## **OPERATION TRANSFORMATION**

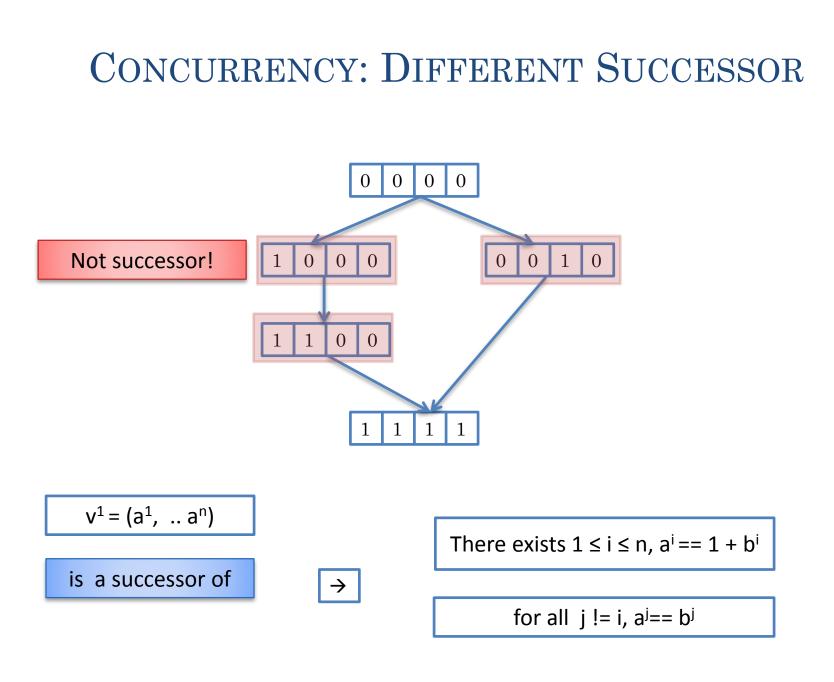
Prasun Dewan Department of Computer Science University of North Carolina at Chapel Hill <u>dewan@cs.unc.edu</u>



#### IMMEDIATELY DELIVERING CONCURRENT MESSAGES

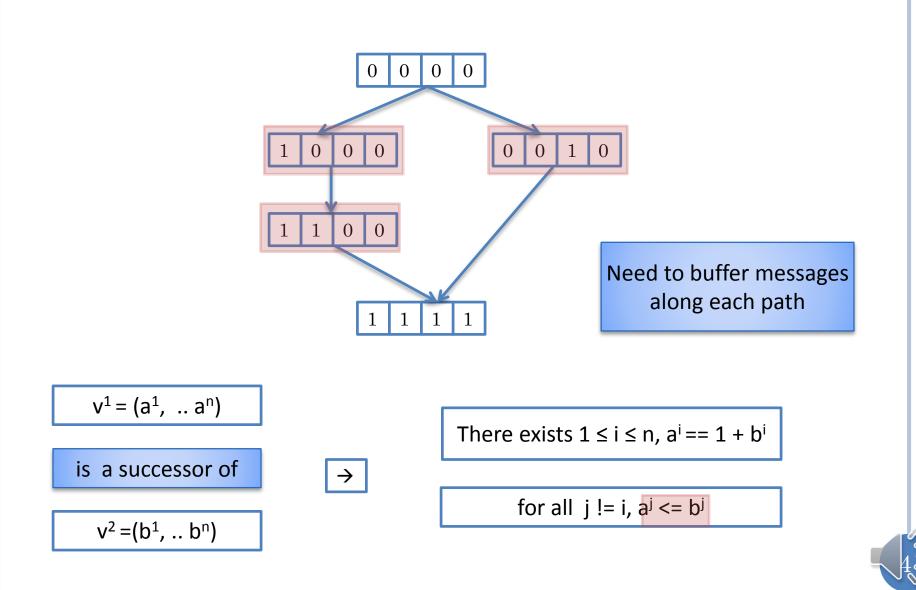
- A tree of message paths exist
- Create vector time stamp and buffer for each leaf in the path
- When a message arrives see if its vector time stamp > one of the vector time stamps, put in the buffer for that vector time stamp
- Otherwise create a new vector time stamp and buffer

(VectorTimeStampCopiedAndNewBufferCreated) deliver the message after flagging concurrency



Bo

# CONCURRENT SUCCESSORS



## Delivering Concurrent Messages

- A DAG of message paths exist
- Create buffer for each leaf in the DAG
- When a message arrives find a buffer
  - See if a buffer exists in which the head message is not concurrent with this message
  - Otherwise create a new buffer
- Insert message at appropriate position in buffer
- Process buffer as before but now use a different successor function
  - $(b^1, b^2, \dots, b^n)$  is successor of  $(a^1, a^2, \dots, a^n)$  if for some j,  $b^j = a^j + 1$  and for all i!= j,  $b^i \le a^i$
  - Takes into account that a message on a different path from an ancestor of the current time stamp node arrived

#### Delivering Concurrent Messages – Full Search

- A DAG of message paths exist
- Create a single buffer for all messages
- When a message arrives put it in the buffer at some position
- Process buffer as before
  - but now use a different successor function
    - (b<sup>1</sup>, b<sup>2</sup>, ..., b<sup>n</sup>) is successor of (a<sup>1</sup>, a<sup>2</sup>, ... a<sup>n</sup>) if for some j, b<sup>j</sup> = a<sup>j</sup> + 1 and for all i != j, b<sup>i</sup> <= a<sup>i</sup>
    - Takes into account that a message on a different path from an ancestor of the current time stamp node arrived
  - Search the entire buffer rather than look at the head of the buffer

#### **OPERATION TRANSFORMATION**

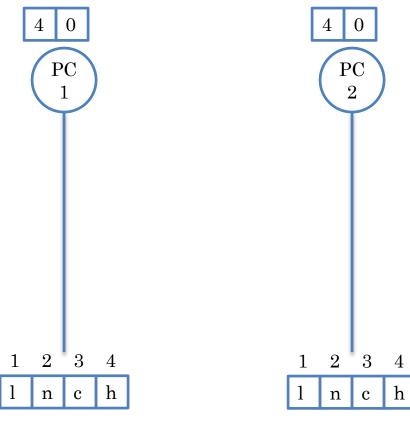


#### Delivering Concurrent Messages – Full Search

- A DAG of message paths exist
- Create a single buffer for all messages
- When a message arrives put it in the buffer at some position
- Process buffer as before
  - but now use a different successor function
    - (b<sup>1</sup>, b<sup>2</sup>, ..., b<sup>n</sup>) is successor of (a<sup>1</sup>, a<sup>2</sup>, ..., a<sup>n</sup>) if for some j, b<sup>j</sup> = a<sup>j</sup> + 1 and for all i != j, b<sup>i</sup> <= a<sup>i</sup>
    - Takes into account that a message on a different path from an ancestor of the current time stamp node arrived
  - Search the entire buffer rather than look at the head of the buffer

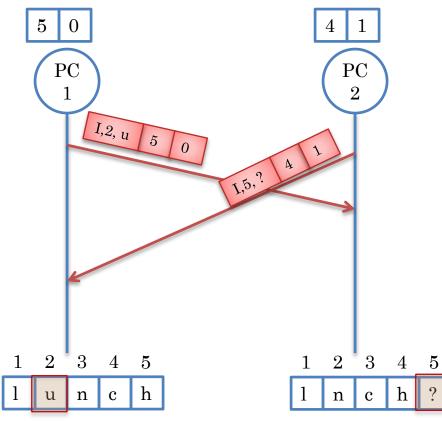
Are causality guarantees among concurrent paths enough?

## CONCURRENT EDITING: INITIAL STATE



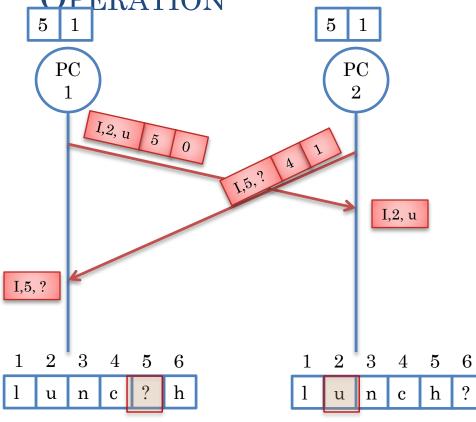


## **CONCURRENT INSERTIONS**





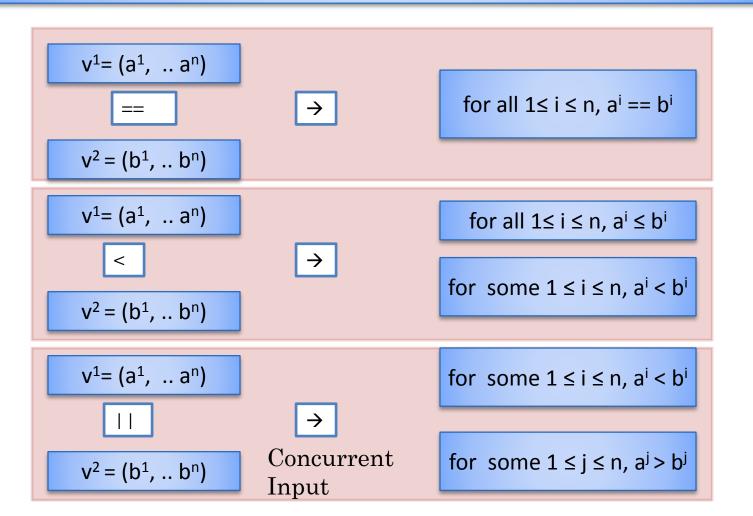
#### IMMEDIATELY DELIVERING CONFLICTING REMOTE OPERATION

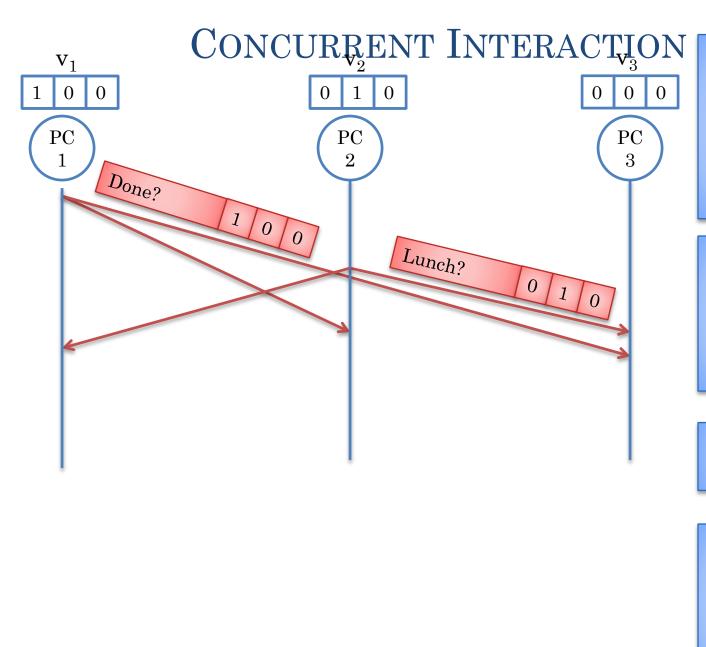




# PARTIALLY ORDERED VECTOR TIME STAMPS

 $v = (x^1, ... x^n)$  at Site S<sup>j</sup>  $\rightarrow$  Site S<sup>j</sup> has received x<sup>i</sup> messages from Site S<sup>i</sup> for all  $1 \le i \le n$ 



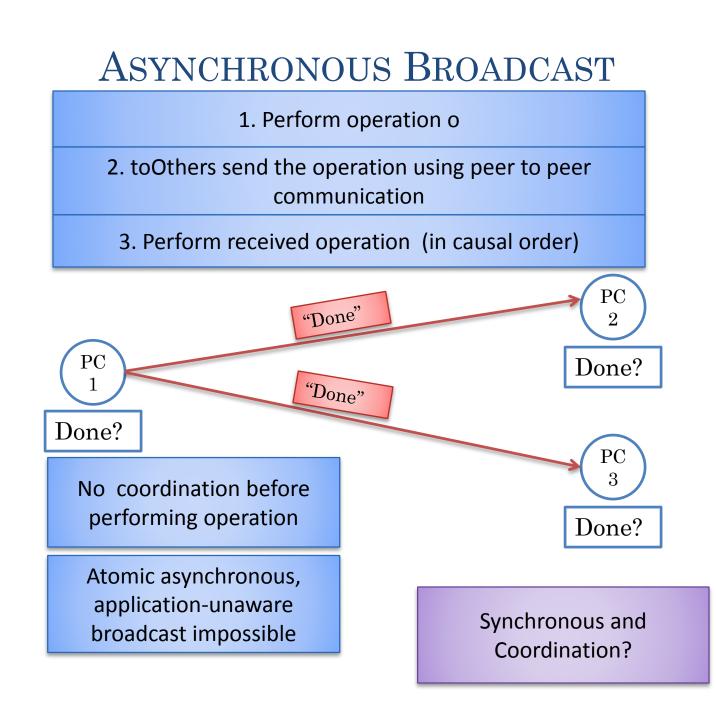


Causal time stamps allow computer to determine concurrent actions

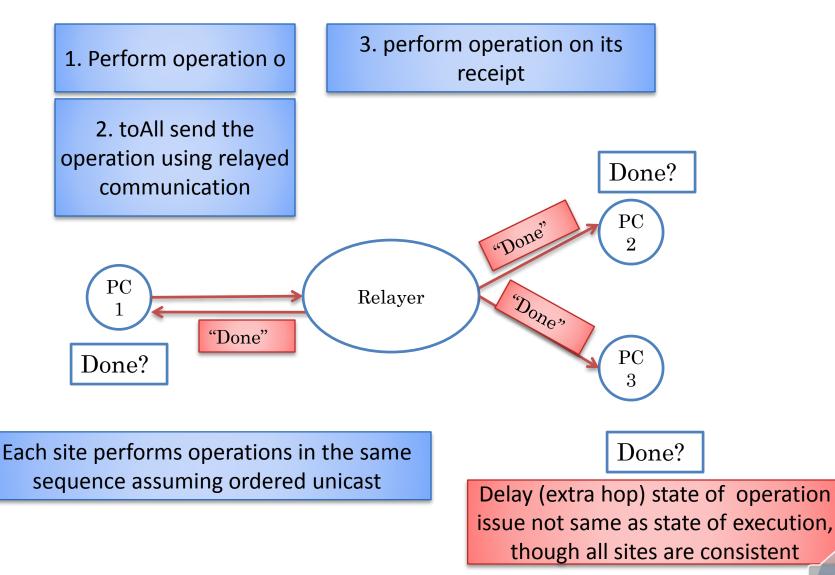
Do not impose a common or total order, which inherently does not exist.

Total order often important

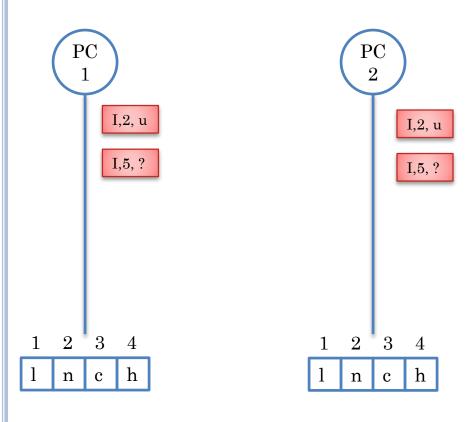
Broadcast supporting total order called atomic broadcast



## Synchronous Relayed Broadcast

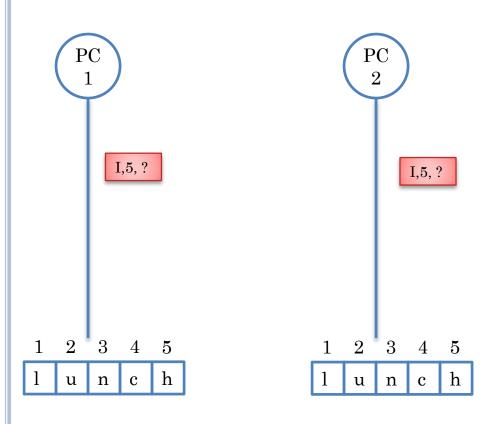


## Ordering with Atomic Broadcast





## FIRST OPERATION EXECUTES

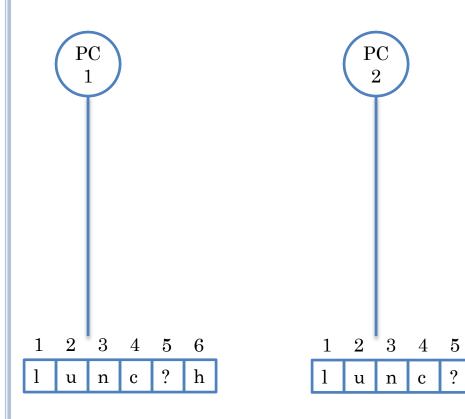




## SECOND OPERATION EXECUTES

6

h

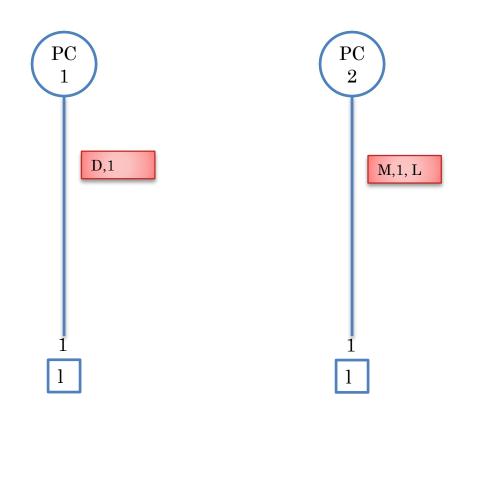


Common state but "intention" violation

Context of operation not the same as when it was issued

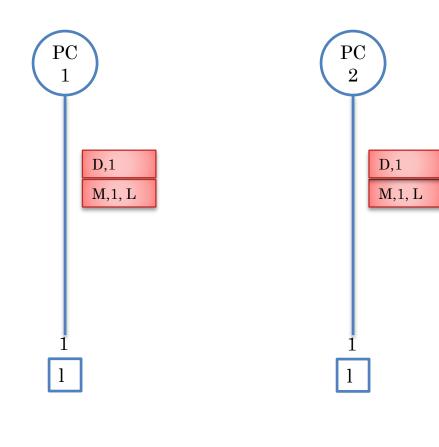
Worse outcome than intention violation?

#### CONCURRENT INTERACTION: DELETE, MODIFY



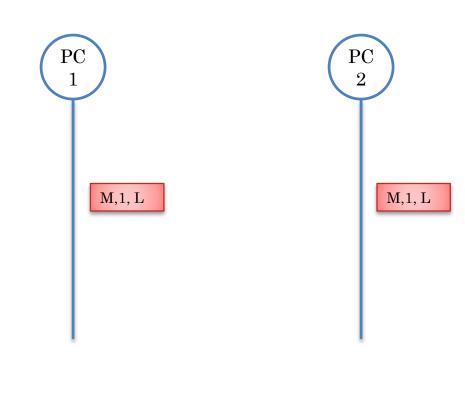


## ATOMIC BROADCAST ORDERING





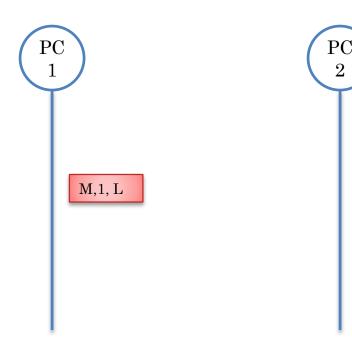
## FIRST OPERATION EXECUTES



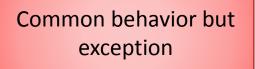


## SECOND OPERATION CAUSES EXCEPTION

M,1, L



#### ArrayIndexOutOfBounds

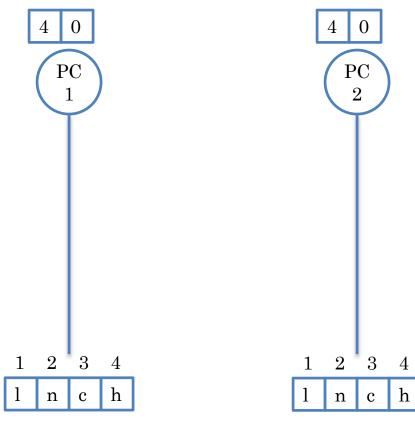


Context of operation not the same as when it was issued

Can concurrency control explain or fix this?

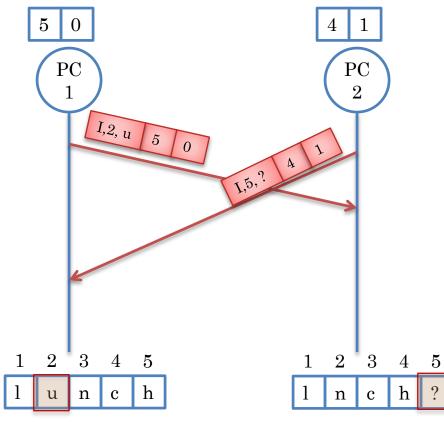


#### CONCURRENT EDITING: INITIAL STATE (REVIEW)



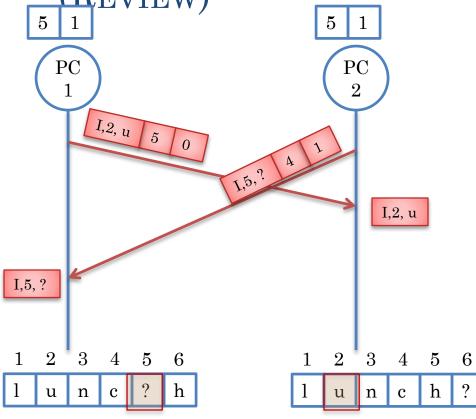


## **CONCURRENT INSERTIONS (REVIEW)**



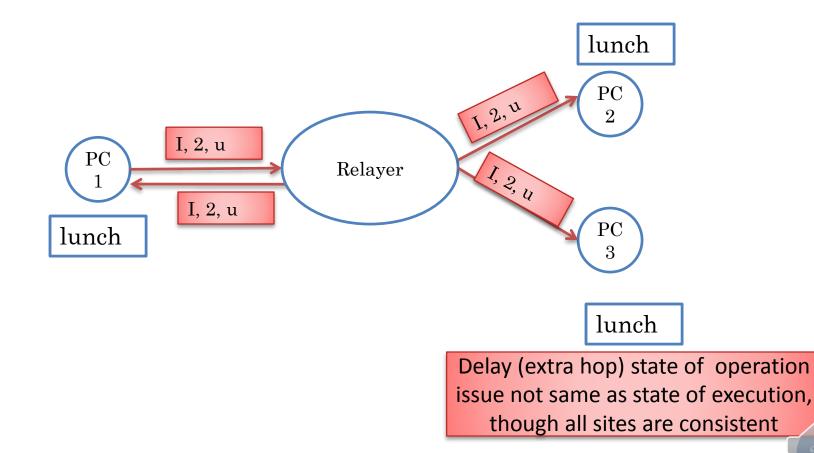


#### DELIVERING CONFLICTING REMOTE OPERATION (REVIEW)

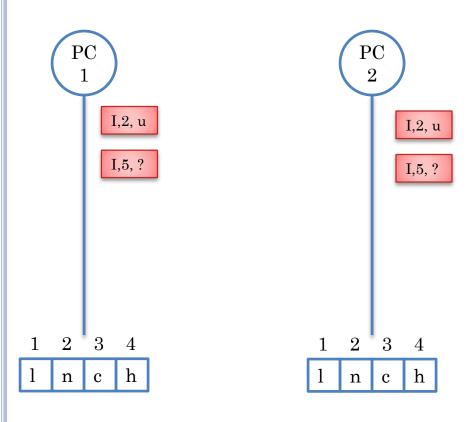




#### SYNCHRONOUS RELAYED BROADCAST (REVIEW)

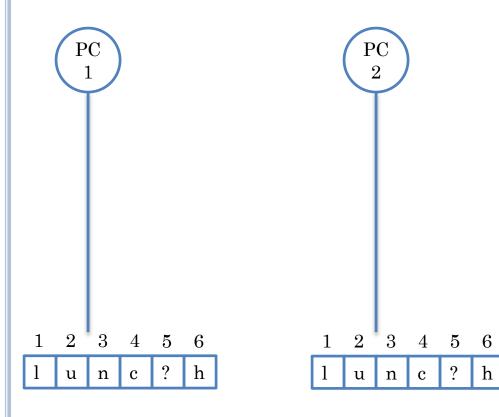


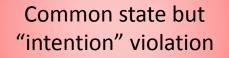
#### ORDERING WITH ATOMIC BROADCAST (REVIEW)





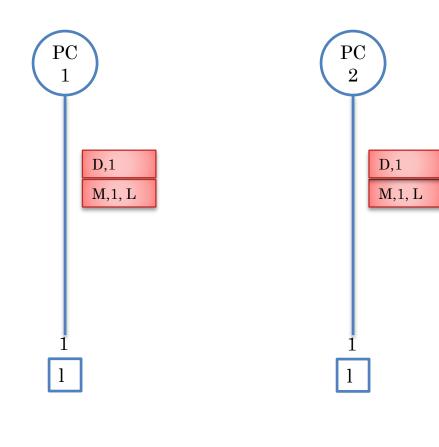
# SECOND OPERATION EXECUTES (REVIEW)





Context of operation not the same as when it was issued

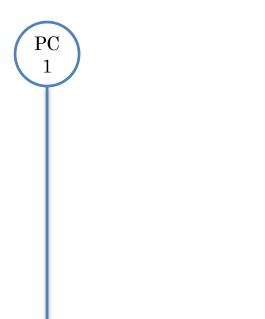
## ATOMIC BROADCAST ORDERING

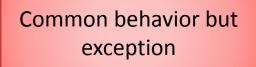




## SECOND OPERATION CAUSES EXCEPTION

 $\frac{PC}{2}$ 





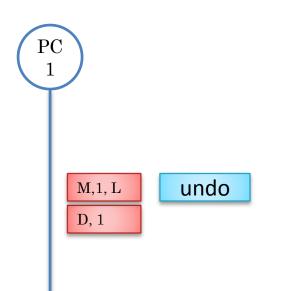
Context of operation not the same as when it was issued

ArrayIndexOutOfBounds

ArrayIndexOutOfBounds



## SINGLE-USER CASE: SELECTIVE UNDO

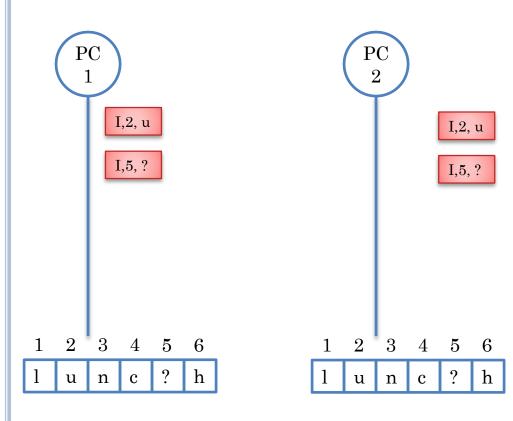


L

Undoing non last operation is done in a context different from the one in which it was executed



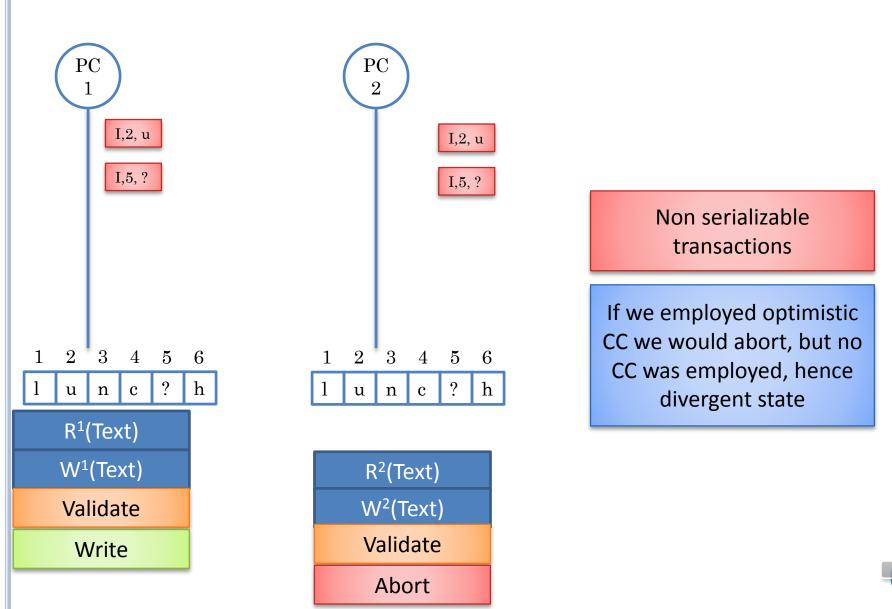
## CONCURRENCY CONTROL?



Can concurrency control explain or fix this?



## **OPTIMISTIC TRANSACTION**



# VALIDATION/CHECKING TIME (REVIEW)

#### • Early

- Pessimistic
- Late
  - Optimistic
- Merging

## EARLY VS. LATE VALIDATION (REVIEW)

- Per-operation checking and communication overhead
- No compression possible.
- Prevents inconsistency.
- Tight coupling: incremental results shared
- Not functional if disconnected
  - Unless we lock very conservatively, limiting concurrency.

- No per-operation checking, communication overhead
- Compression possible.
- Inconsistency possible resulting in lost work.
- Allows parallel development.
- Functional when disconnected.

# MERGING (REVIEW)

• Like optimistic

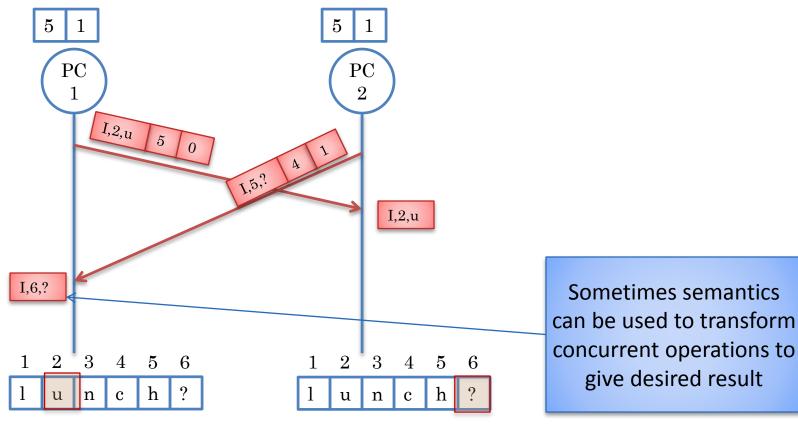
• Allow operation to execute without local checks

#### • But no aborts

- Merge conflicting operations
- E.g. insert 1,a || insert 2, b = insert 1, a; insert 3, b
   || insert 2, b; insert 1, a
- Serializability not guaranteed
  - Ignore reads
  - New transaction to replace conflicting transactions
  - Strange results possible
    - E.g. concurrent dragging of an object in whiteboard

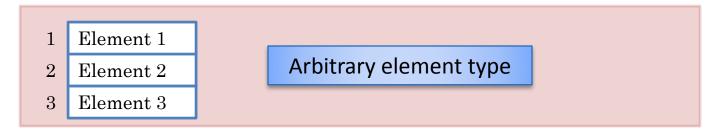
# • App-specific

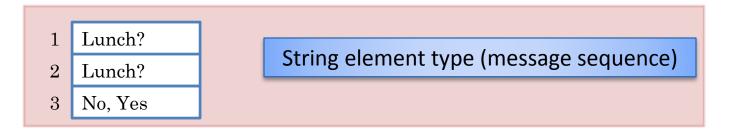
#### TRANSFORMING REMOTE OPERATION

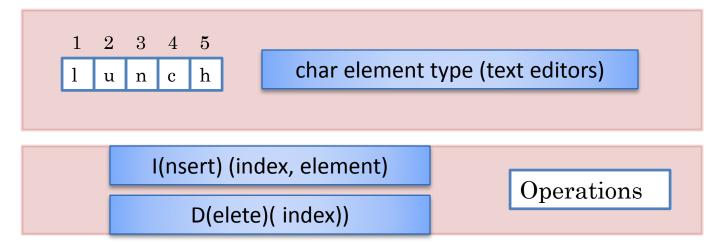


Bound to text buffer?

# ASSUME INDEXED SEQUENCE DATA TYPE

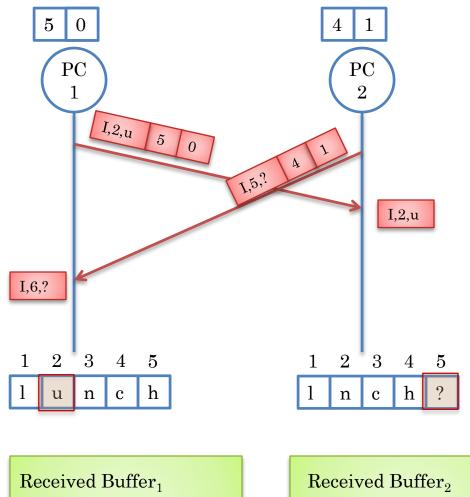








# **CONCURRENT INORDER INTERACTION**



Assume only two sites

Assume messages are received in order from other site

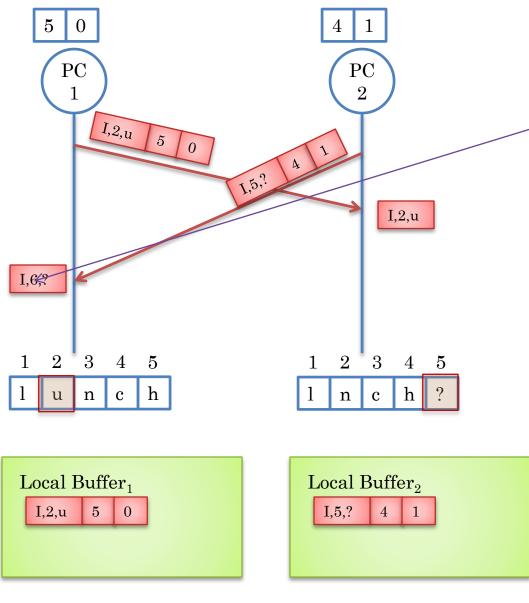
No need for received buffer

Still need time stamp to discover concurrency

General rule for transforming operation



#### TRANSFORMATION FUNCTION



Increment index of remote operation if it has higher index before processing it

Remote operation, R transformed!

Based on index of concurrent local operation, L

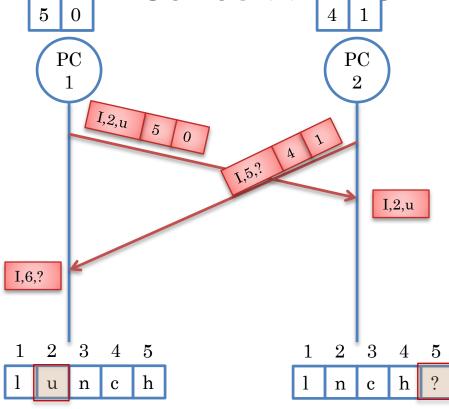
 $R^{T}$  = Transform (R, L)

Apply R<sup>T</sup> instead of R at local site

R Transformed if  $R^T != R$ 

Need local buffer to store L rather than remote buffer

# CONTROL ALGORITHM: SINGLE LOCAL CONCURRENT OPERATION



Given Remote op R, concurrent with exactly one local op L

 $R^{T}$  = Transform (R, L)

Execute R<sup>T</sup>

Site.TimeStamp.increment(R.site)

# OT System Components

Transformation function: Handles single local and concurrent operations

Control algorithm: Calls transformation function, processed buffer and local time stamps

Both must be correct.



# INCLUSION TRANSFORMATION

```
Operation Transform (Operation R, Operation L) {
    if (R.type == Insert && L.type == Insert)
        return TransformInsertInsert (R,L);
    else ....
```

Transform includes effect of second operand on first operand

Other uses in which first and second operands are not remote and local operations

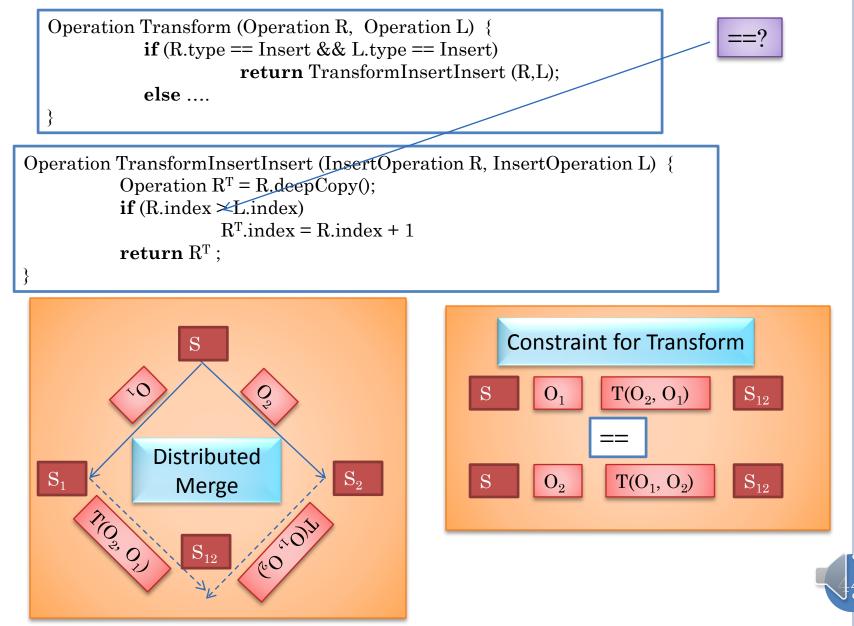
Names indicate we include effect of earlier executed local operation on later received concurrent remote operation Called inclusion transformation

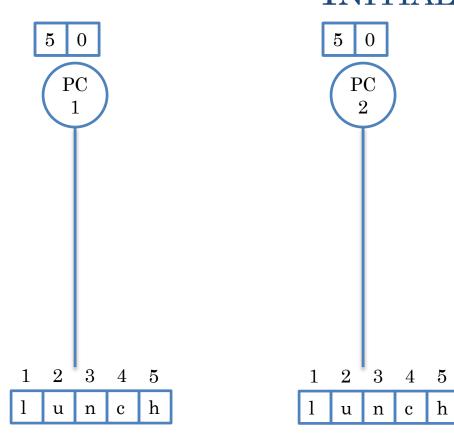
Correct?

Correctness criterion?



#### **CORRECTNESS CRITERION**

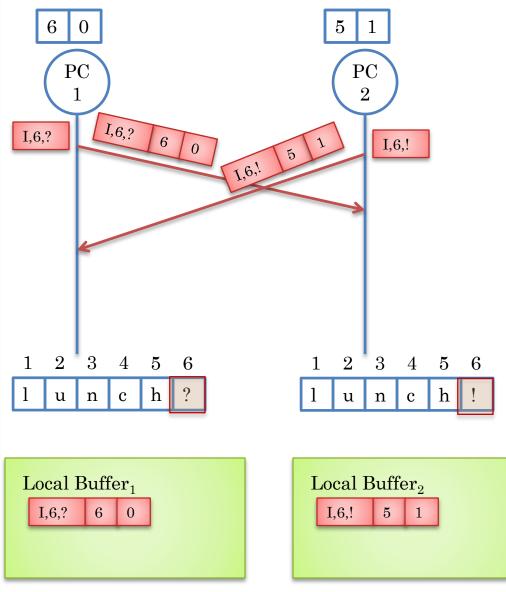




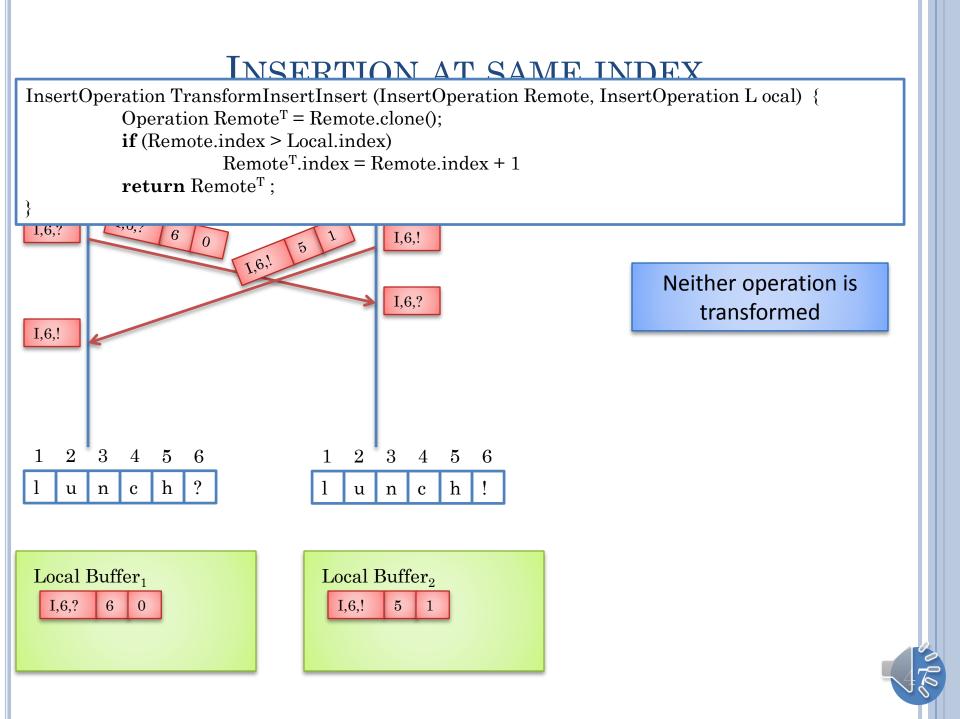


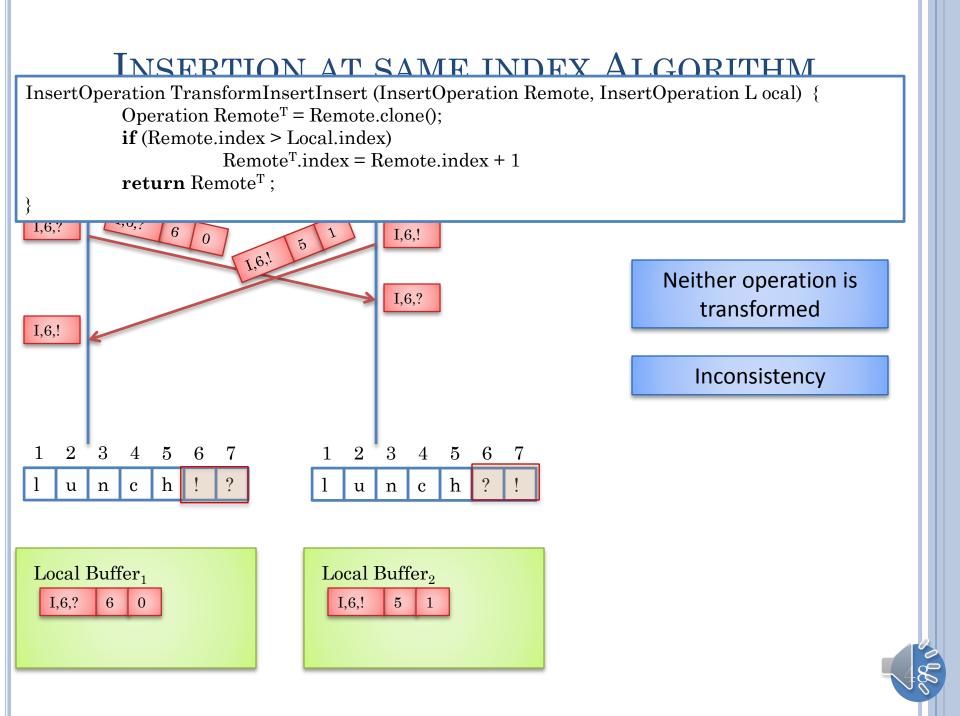
#### INITIAL STATE

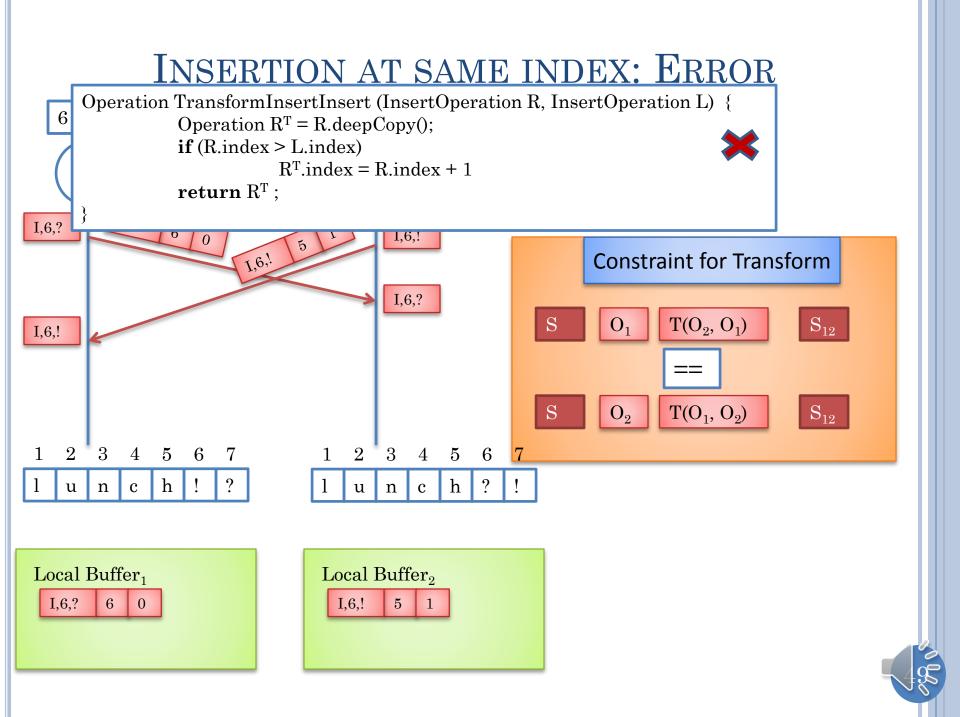
#### **INSERTION AT SAME INDEX**

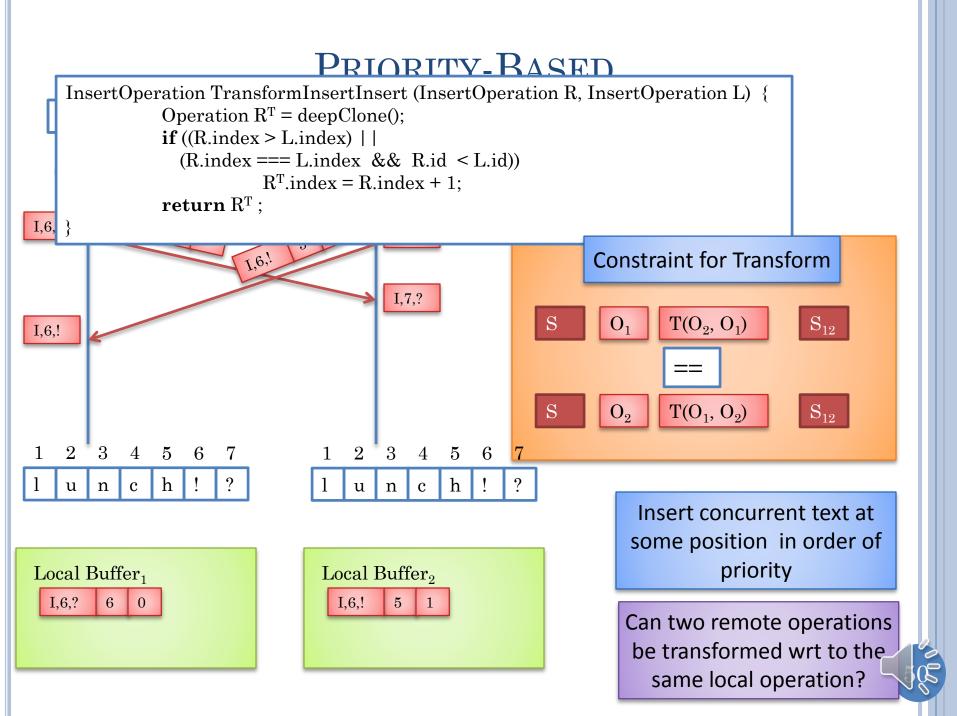




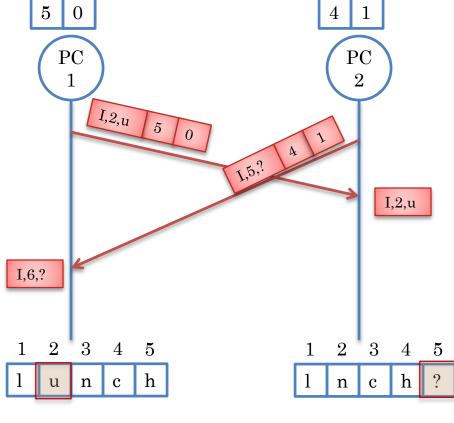




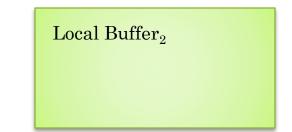




# CONTROL ALGORITHM: SINGLE LOCAL CONCURRENT OPERATION



Local Buffer<sub>1</sub>



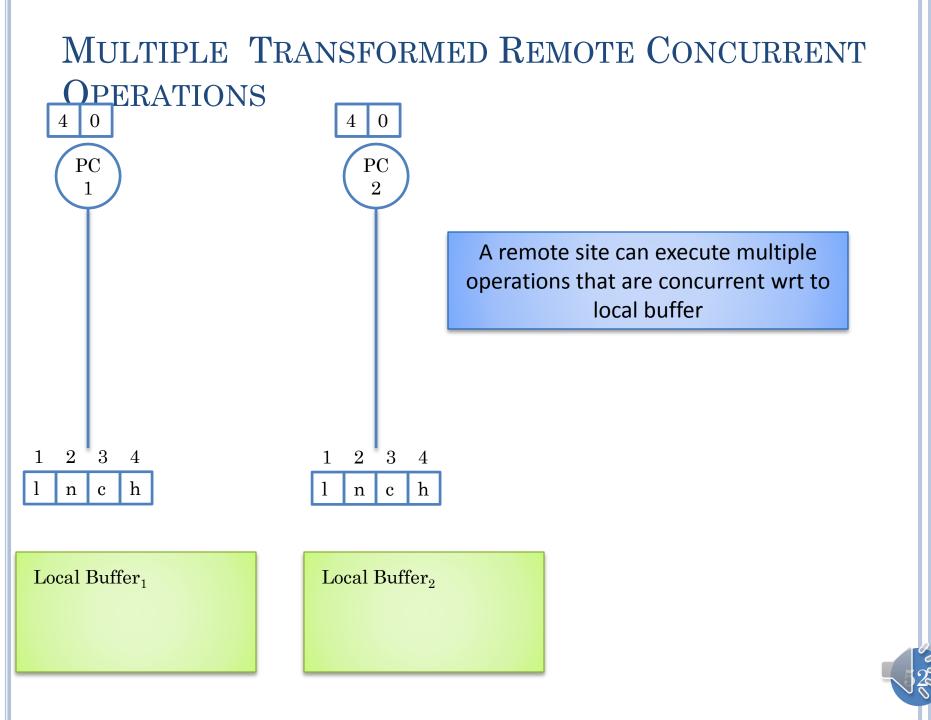
Given Remote op, R, concurrent with exactly one local op L

 $R^{T}$  = Transform (R, L)

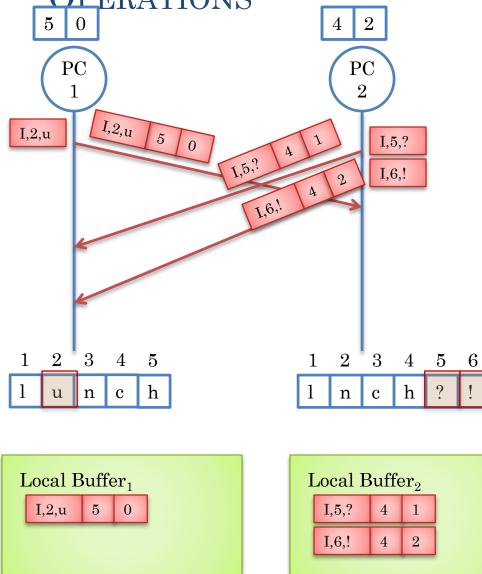
Execute R<sup>T</sup>

Site.TimeStamp.increment(R.site)





#### MULTIPLE TRANSFORMED REMOTE CONCURRENT OPERATIONS





#### SITE 1 OPERATION ARRIVES AND IS NOT TRANSFORMED $\mathbf{5}$ $\mathbf{5}$ $\mathbf{2}$ 0 $\mathbf{PC}$ $\mathbf{PC}$ 21 I,2,u I,2,u 5 I,5,? 1 0 A 1,5,? I,6,! A 1,6,! I,2,u

1 2 3 4 5 1 u n c h 1 u n

Local Buffer<sub>1</sub>

I,2,u 5 0

Ι	Local Buffer <sub>2</sub>						
	I,5,?	4	1				
	I,6,!	4	2				

4

С

 $\mathbf{5}$ 

h

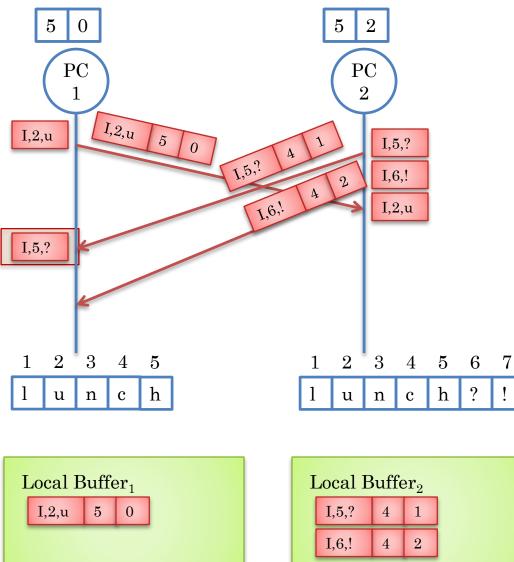
6 7

!

?

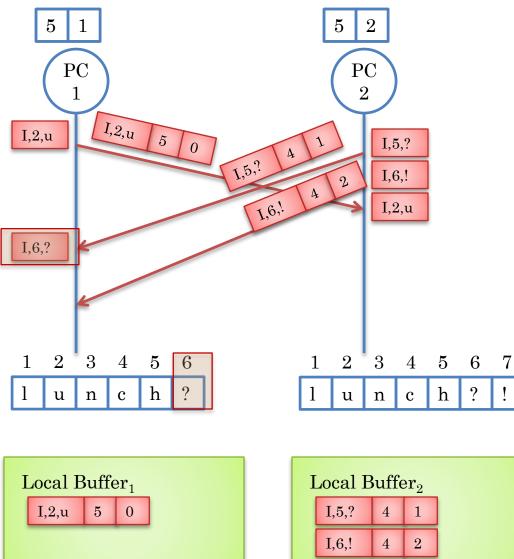


#### FIRST SITE 2 OPERATION ARRIVES



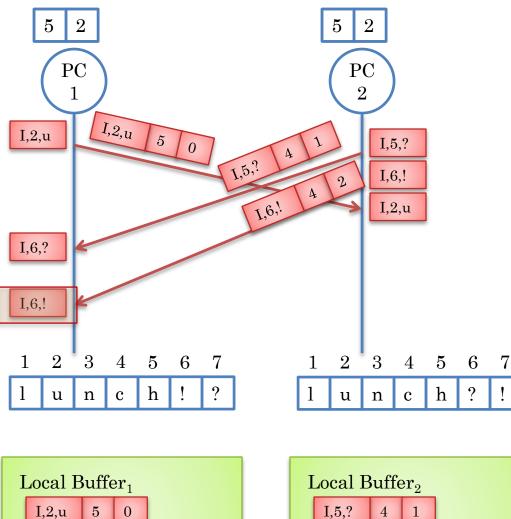


#### FIRST OPERATION TRANSFORMED





#### SECOND OPERATION ARRIVES



I,6,!

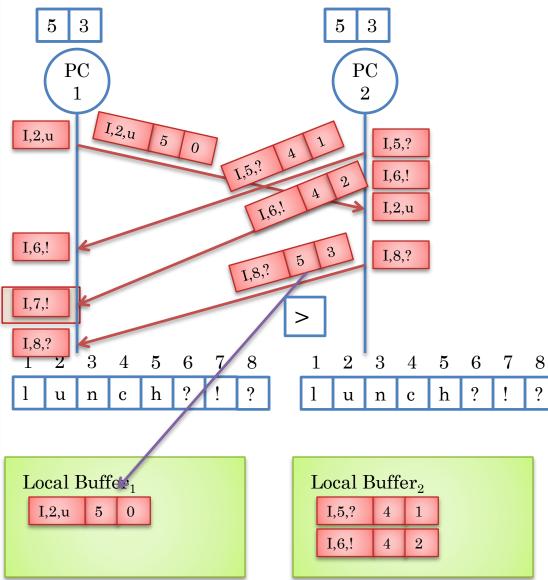
2

4

Multiple remote operations transformed with respect to same local operation



# BUFFER CLEANUP



Multiple remote operations transformed with respect to same local operation

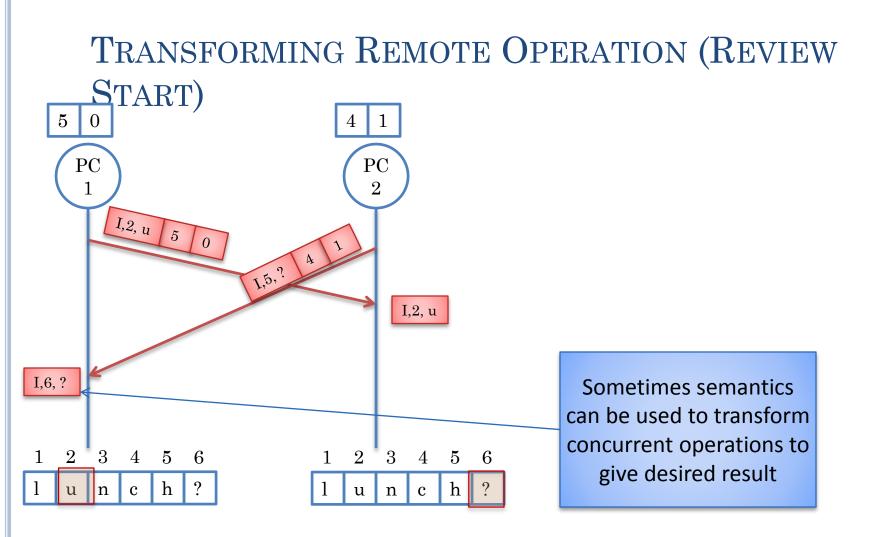
How long should local buffer be kept?

No need to transform

If local.timestamp < msg.timestamp

Each subsequent message has larger time stamp

Remove all locals from buffer with time stamp smaller than time stamp of received message





# OT System Components

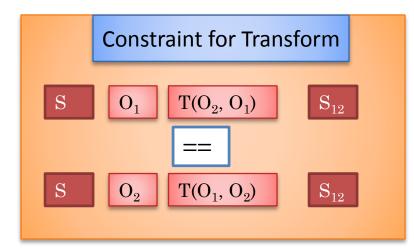
Transformation function: Handles single local and concurrent operations

Control algorithm: Calls transformation function, processed buffer and local time stamps

Both must be correct.



# PRIORITY-BASED TRANSFORMATION



```
 \begin{array}{l} \mbox{InsertOperation TransformInsertInsert (InsertOperation R, InsertOperation L) } \\ \mbox{Operation } R^T = deepClone(); \\ \mbox{if } ((R.index > L.index) \mid | \\ (R.index === L.index &\& R.id < L.id)) \\ R^T.index = R.index + 1; \\ \mbox{return } R^T; \end{array}
```



#### CONTROL ALGORITHM: SINGLE LOCAL CONCURRENT OPERATION

Given Remote op R, concurrent with exactly one local op L

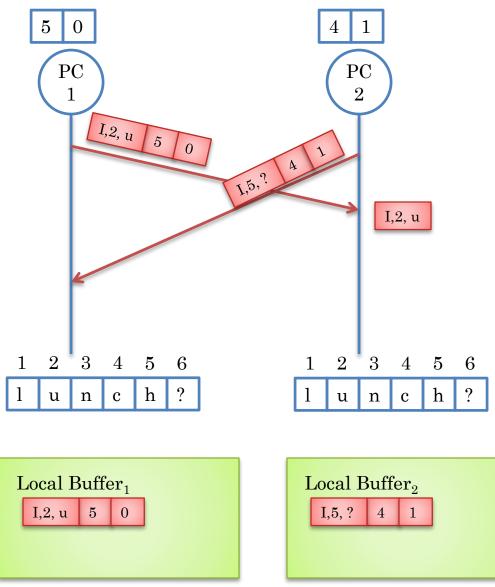
 $R^{T}$  = Transform (R, L)

Execute R<sup>T</sup>

Site.TimeStamp.increment(R.site)



#### NEED FOR LOCAL BUFFER

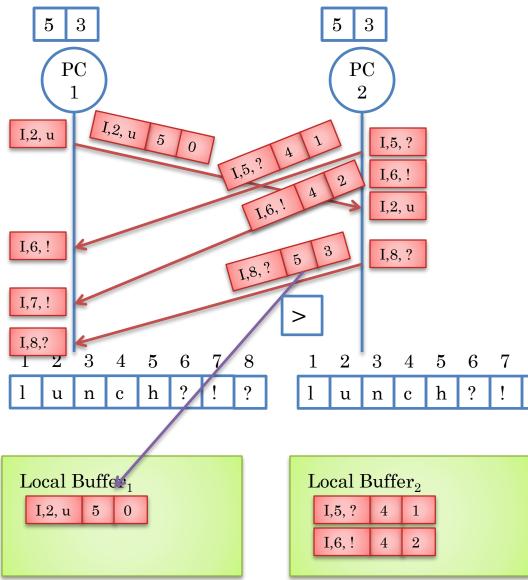




# BUFFER CLEANUP (REVIEW END)

8

?



Multiple remote operations transformed with respect to same local operation

How long should local buffer be kept?

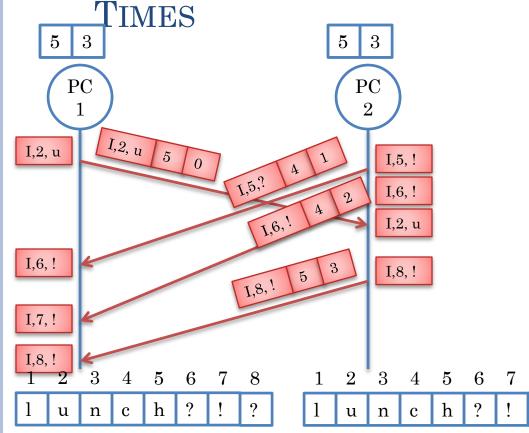
No need to transform

If local.timestamp < msg.timestamp

Each subsequent message has larger time stamp

Remove all locals from buffer with time stamp smaller than time stamp of received message

# DUAL: OPERATION TRANSFORMED MULTIPLE



#### Local Buffer<sub>1</sub>

I,2, u 5 0

Local Buffer <sub>2</sub>						
	I,5, ?	4	1			
	I,6, !	4	2			

8

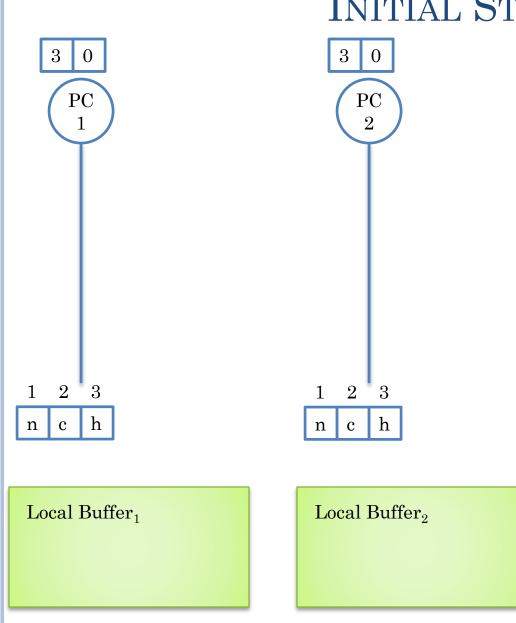
?

Multiple remote operations transformed with respect to same local operation

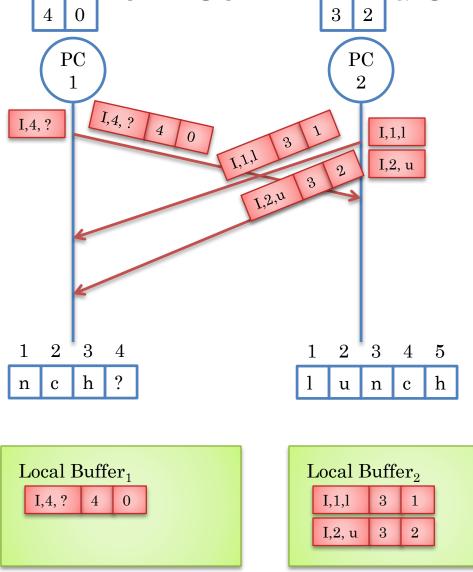
#### Dual?

A remote operation transformed with respect to multiple local operations

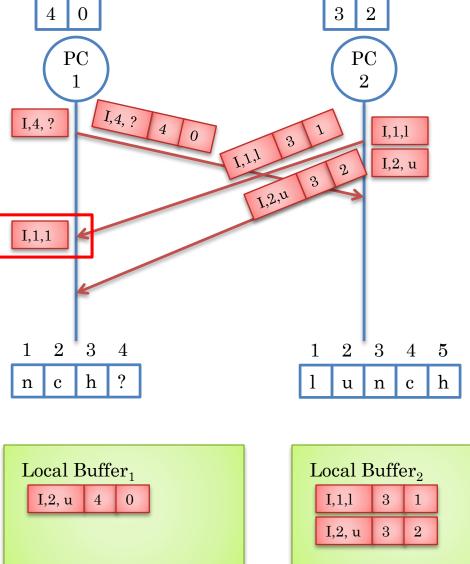




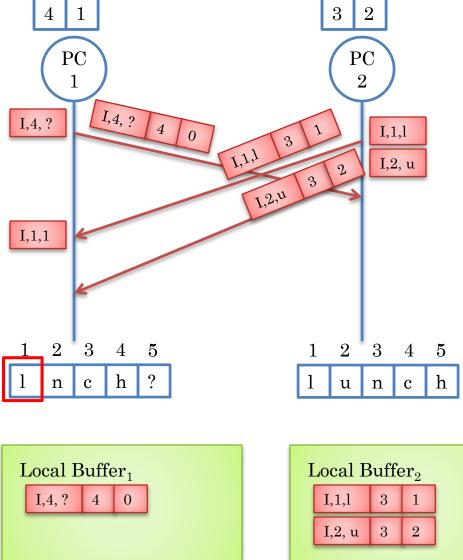
#### INITIAL STATE



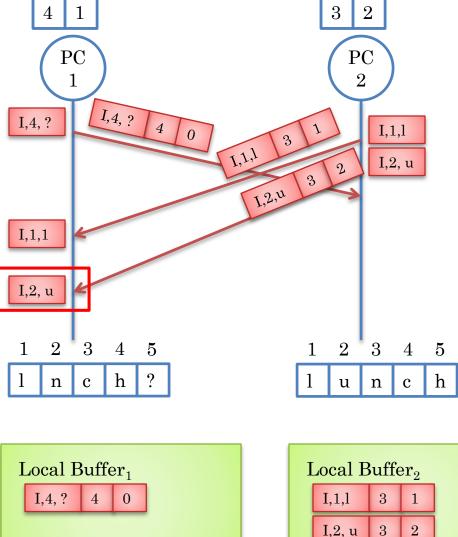




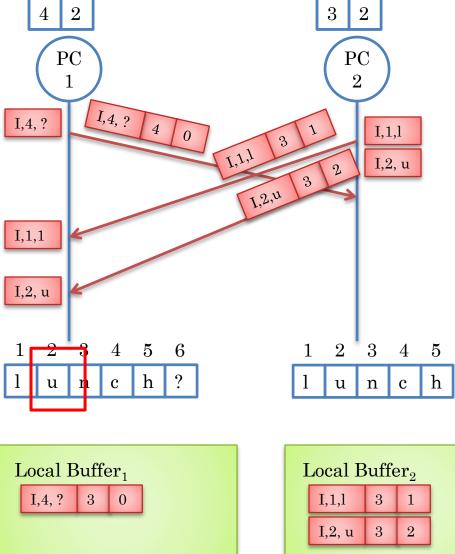




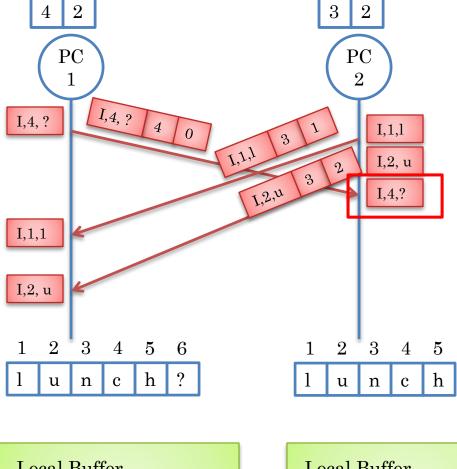














I,4, ? 3 0

Ι	Local Buffer $_2$						
	I,1,1	3	1				
	I,2, u	3	2				



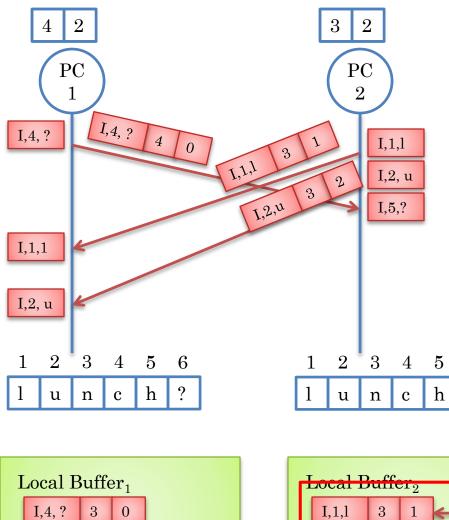
### FIRST TRANSFORMATION

1

 $\mathbf{2}$ 

3

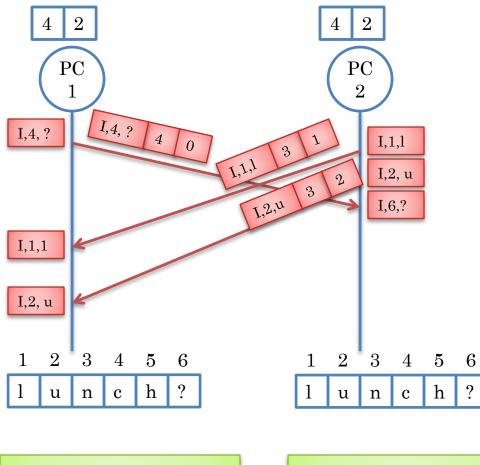
I,2, u



Transform wrt to first concurrent local operation

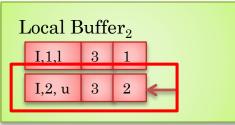


### SECOND TRANSFORMATION/CONTROL ALGORITHM



Local Buffer<sub>1</sub>

I,4,? 3 0



Transform wrt to second concurrent local operation

A remote operation transformed with respect to multiple local operations

Run transform function with respect to all concurrent operations in the local log: Transform (Transform (Transform (R, L1), L2) ...LN)

Control algorithm now handles multiple concurrent/local operations using Transform function addressing single concurrent remote/local operation

### CONTROL ALGORITHM: SINGLE LOCAL CONCURRENT OPERATION

Given Remote op, R, concurrent with exactly one local op L

 $R^{T}$  = Transform (R, L)

Execute R<sup>T</sup>

Site.TimeStamp.increment(R.site)

Algorithm for multiple local concurrent operations?



### CONTROL ALGORITHM: MULTIPLE CONCURRENT LOCAL OPERATIONS

Given Remote op, R, concurrent with local ops  $L^1$ ,  $L^2$ , ...  $L^N$ 

For each L

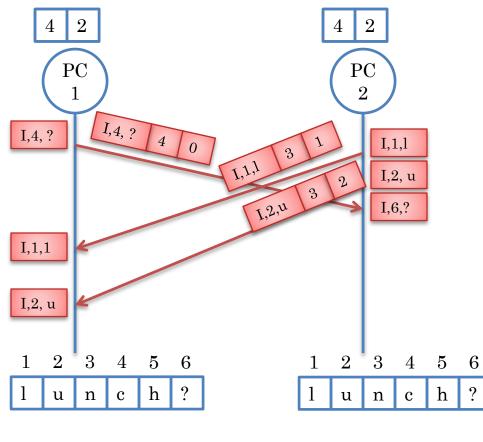
R = Transform (R, L)

Execute R

Site.TimeStamp.increment(R.site)



# SINGLE SITE TRANSFORMATION

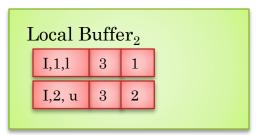


All examples so far involved transformation(s) at one site

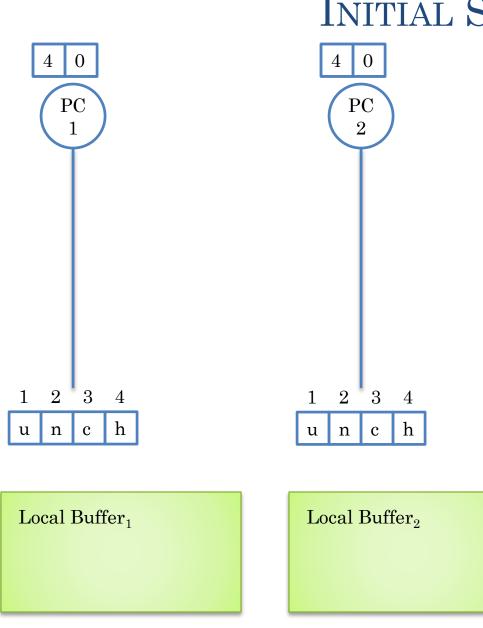
Transformation at both sites?

Local Buffer<sub>1</sub>

I,4,? 3 0

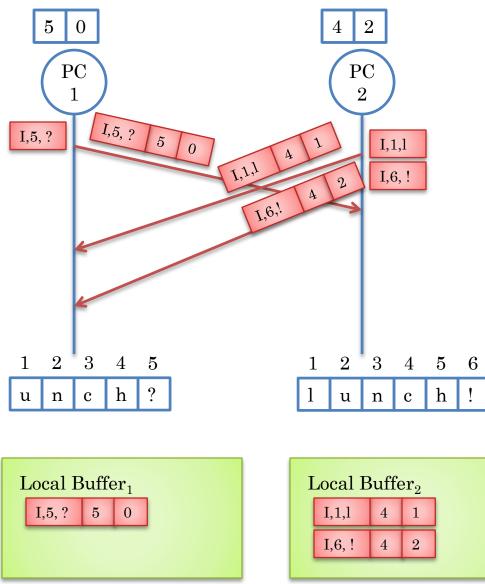






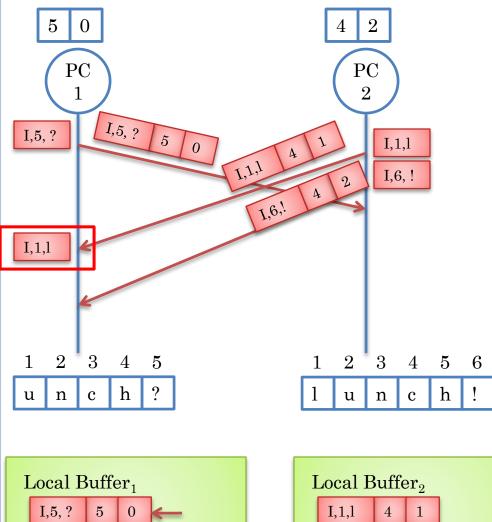
# INITIAL STATE

### MULTIPLE CONCURRENT REMOTE OPERATIONS





# ARRIVAL AT SITE 1



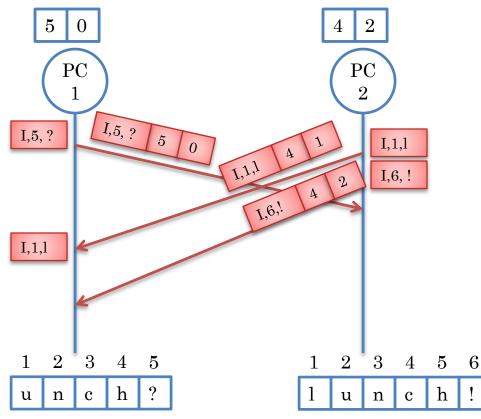
I,6, !

 $\mathbf{2}$ 

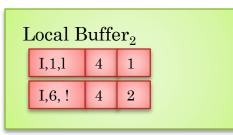
4



### **Result of Transform**



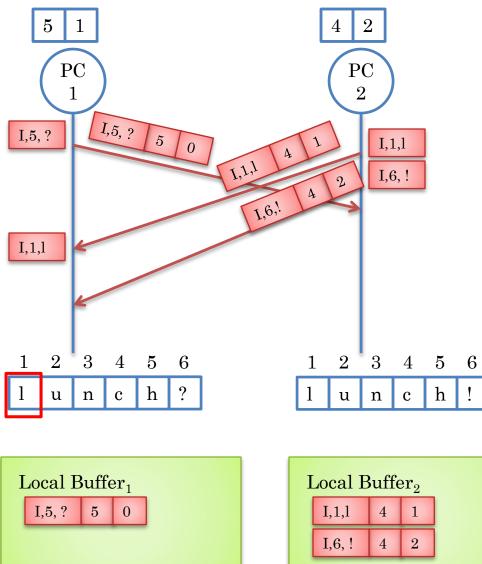
Local Buffer<sub>1</sub>



Remote operation not transformed

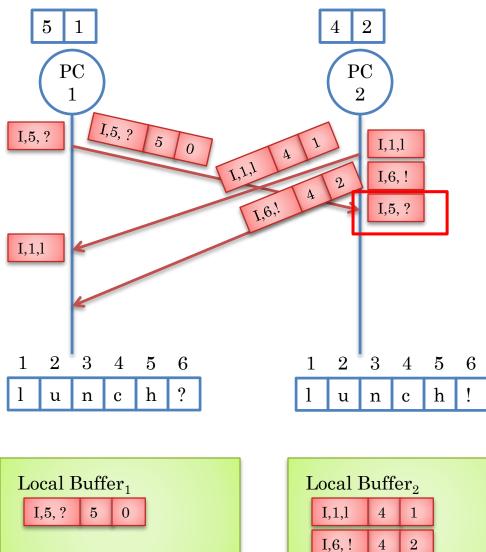


### **UNTRANSFORMED APPLICATION**



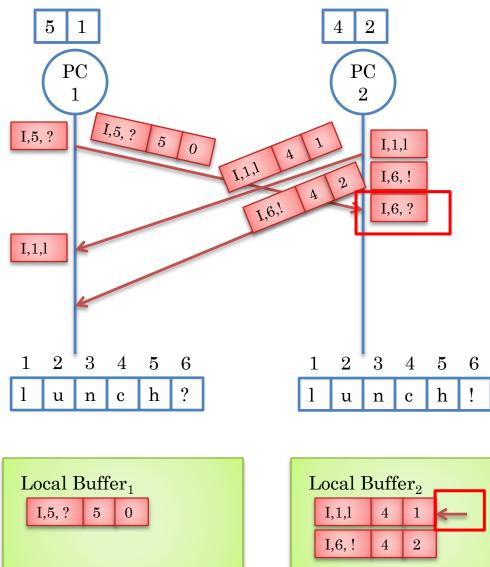


### ARRIVAL AT SECOND SITE



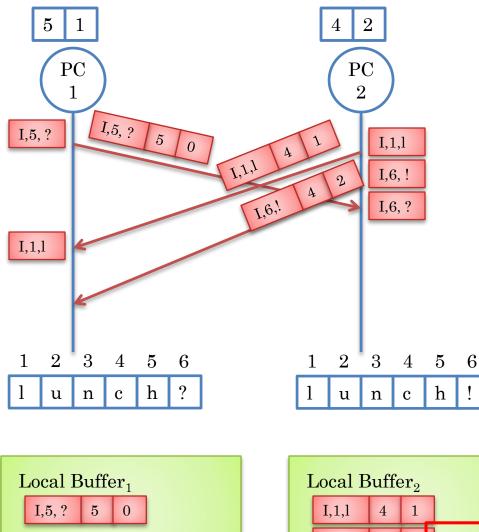


### TRANSFORMED WRT TO FIRST LOCAL OPERATION





# EXAMINING SECOND LOCAL OPERATION



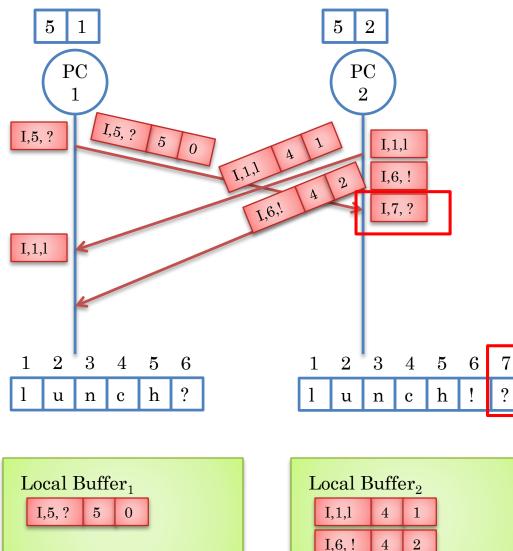
I,6, !

2

4



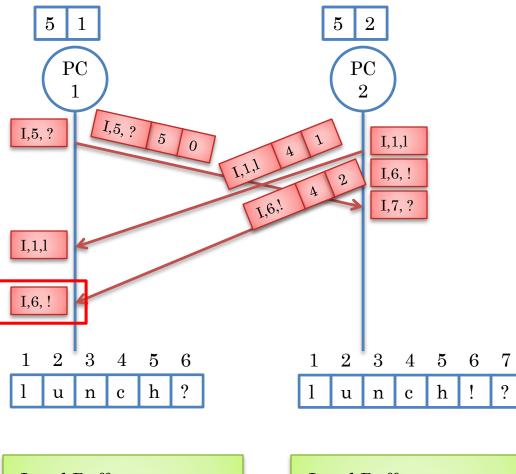
### SECOND TRANSFORMATION AND APPLICATION



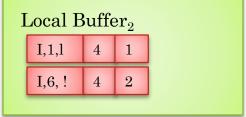
User 1 < User 2



### SECOND ARRIVAL AT SITE 1

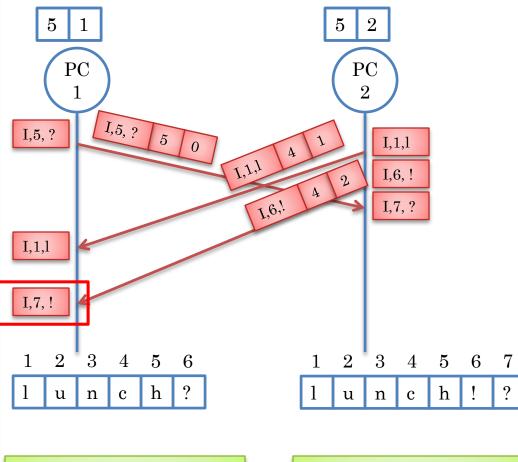


Local Buffer<sub>1</sub>





### **OPERATION TRANSFORMED**





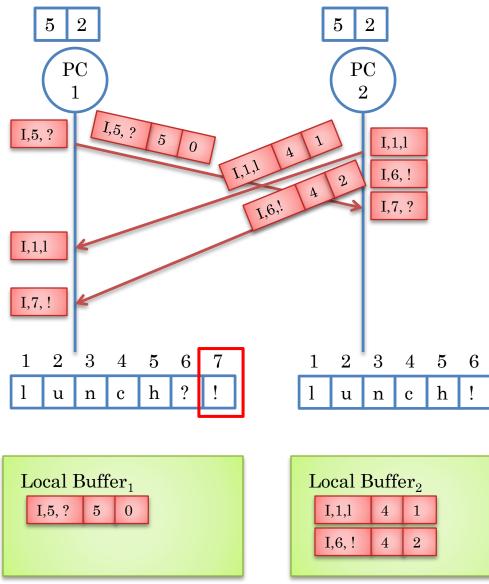




# APPLICATION OF TRANSFORMED OPERATION

7

?



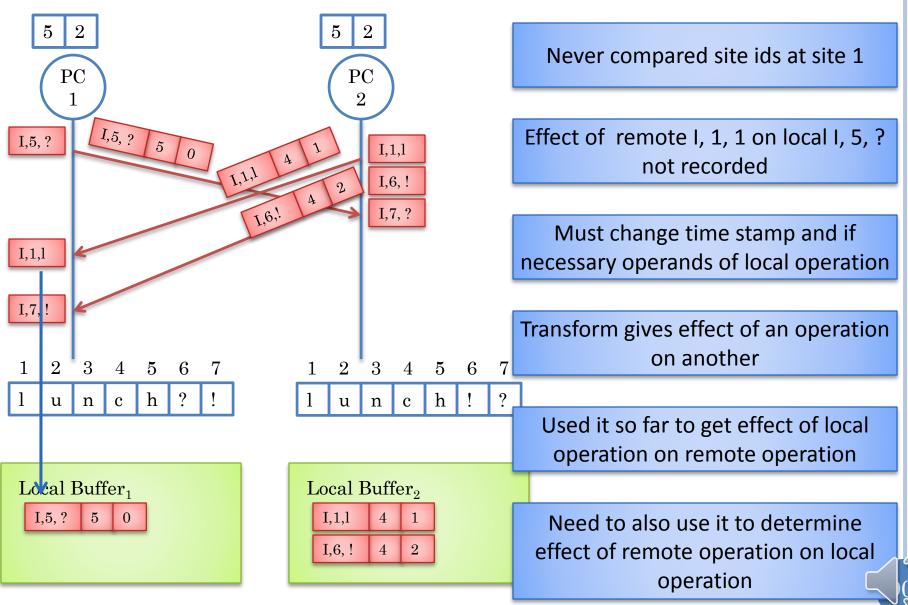
### Inconsistency!

Never compared user ids at site 1

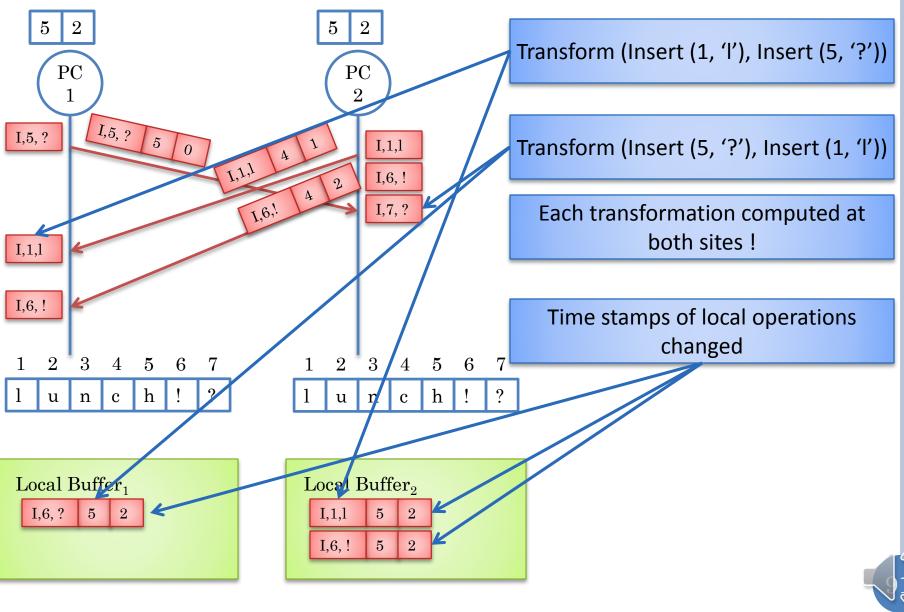
#### What went wrong?



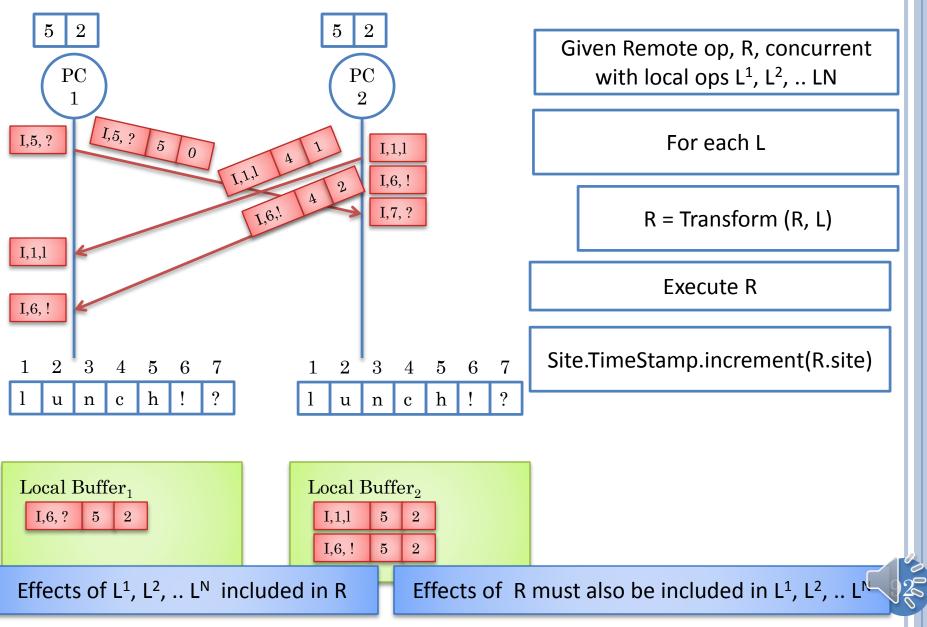
# WHAT WENT WRONG?



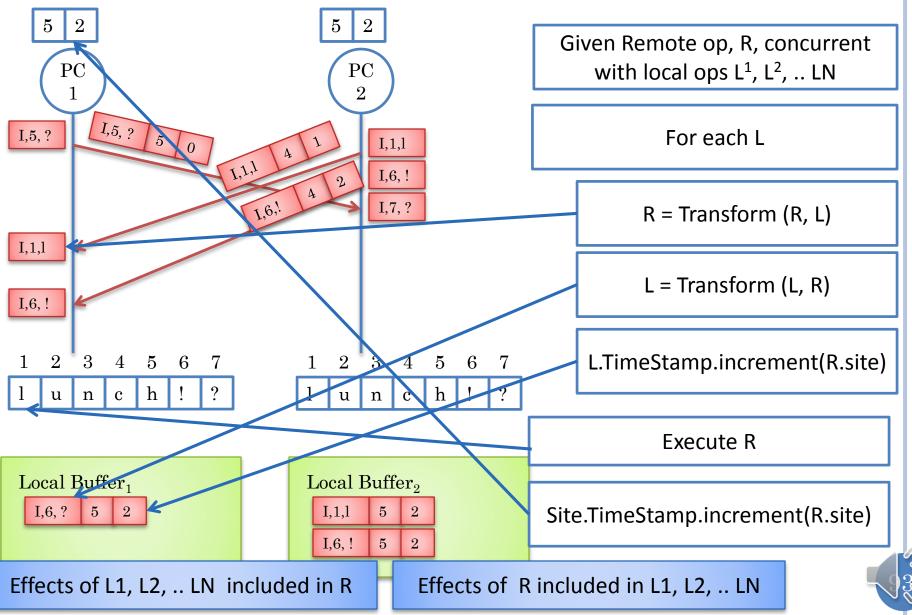
# **RUNNING TRANSFORM IN PAIRS**



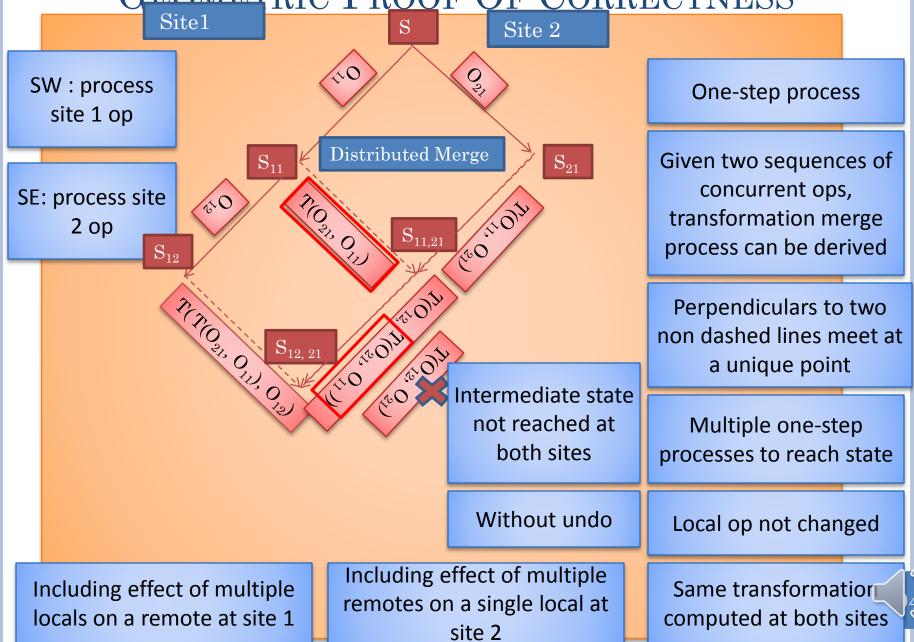
### PREVIOUS CONTROL ALGORITHM

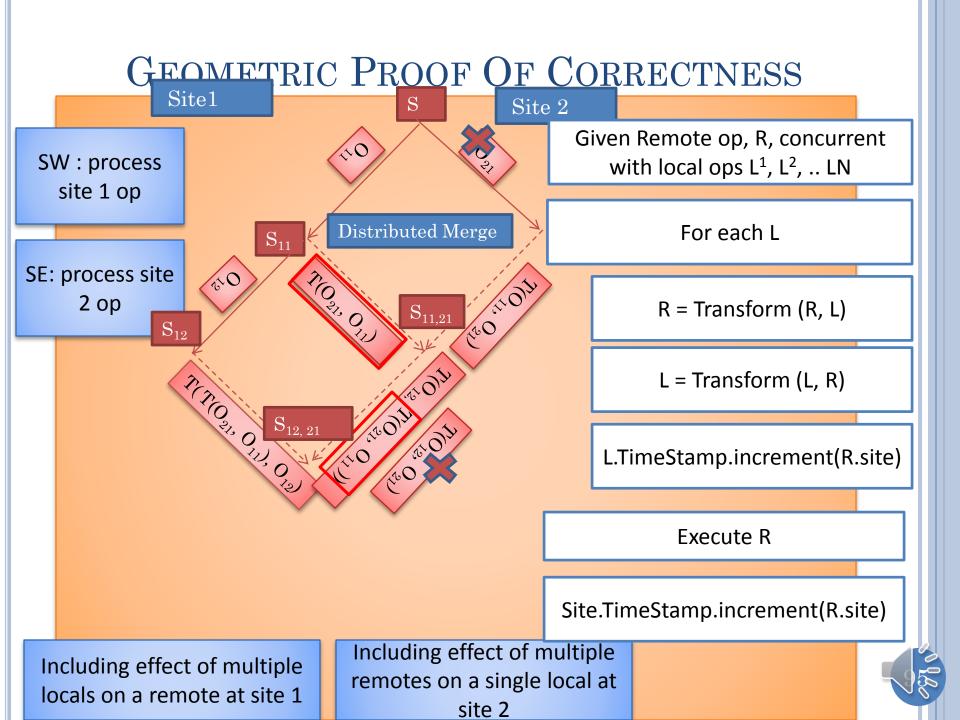


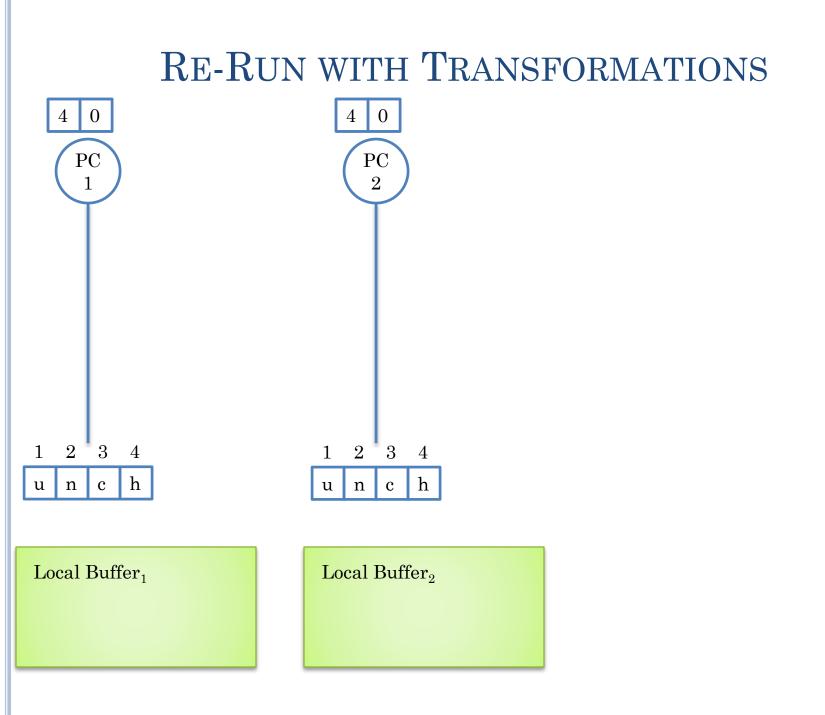
# NEW CONTROL ALGORITHM



### GEOMETRIC PROOF OF CORRECTNESS

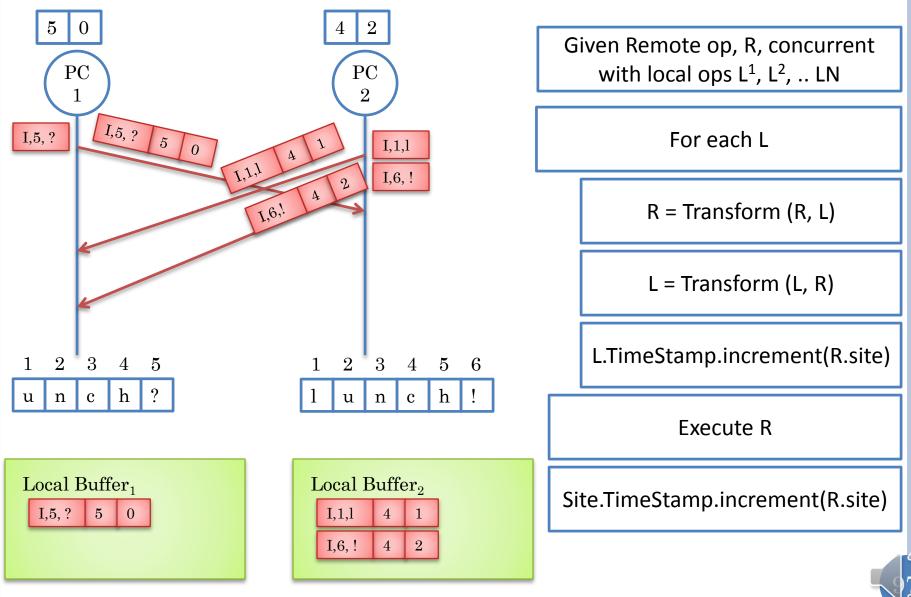




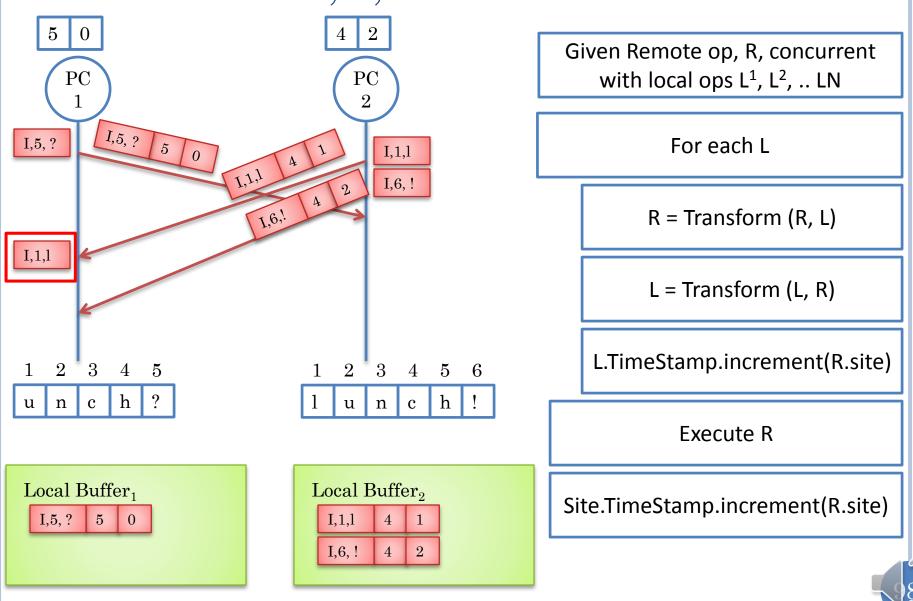




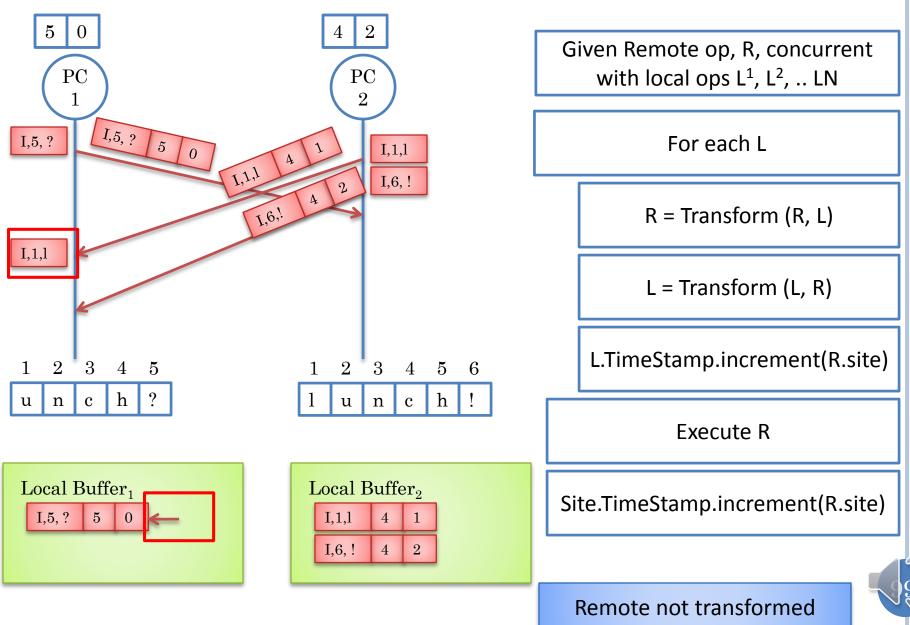
# **CONCURRENT INTERACTION**



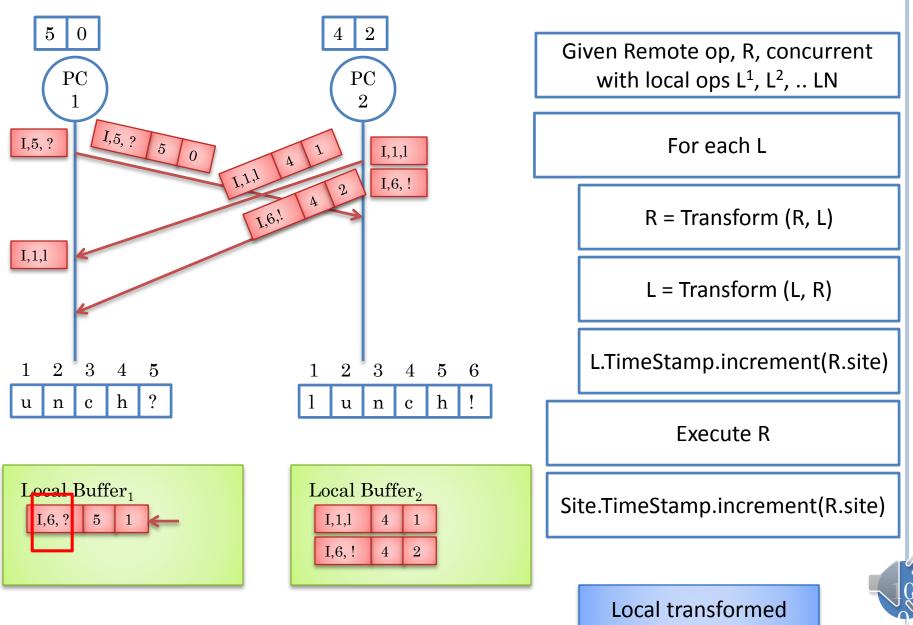
# REMOTE I, 1, L ARRIVES AT SITE 1



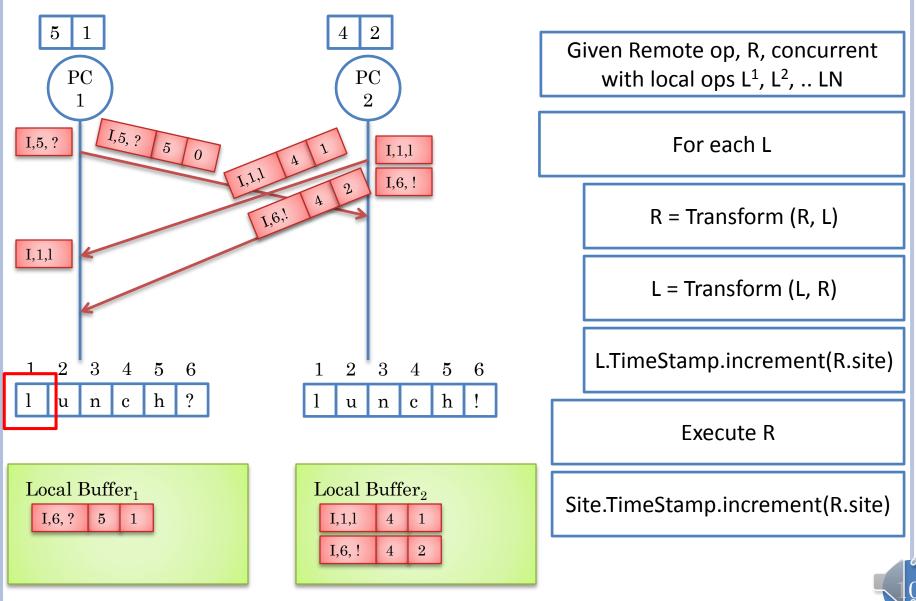
### **REMOTE NOT TRANSFORMED**



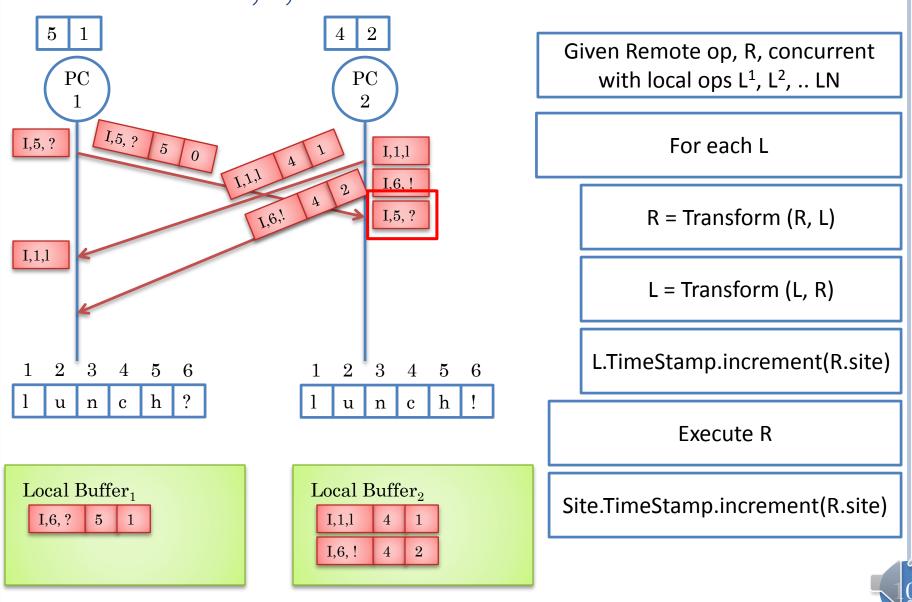
# LOCAL TRANSFORMED



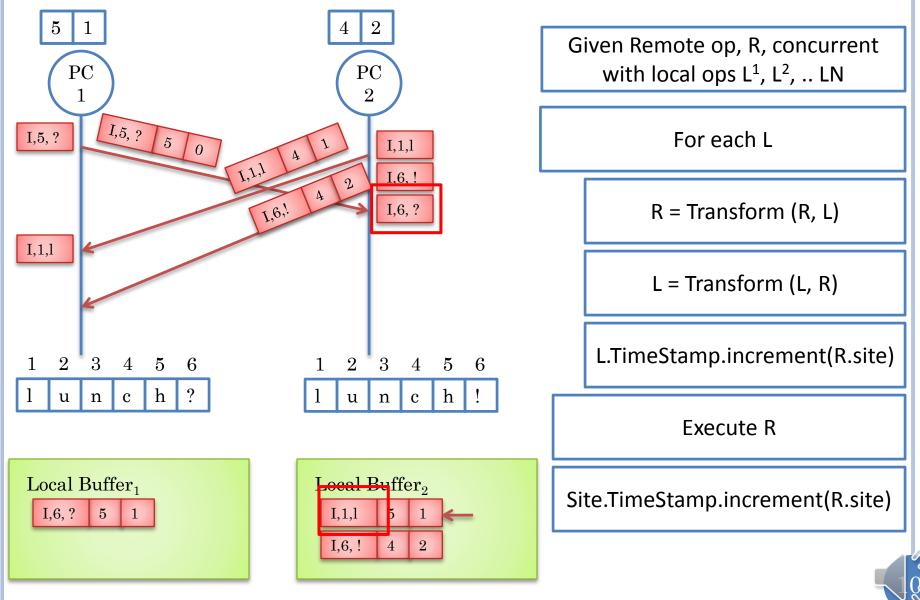
# REMOTE APPLIED



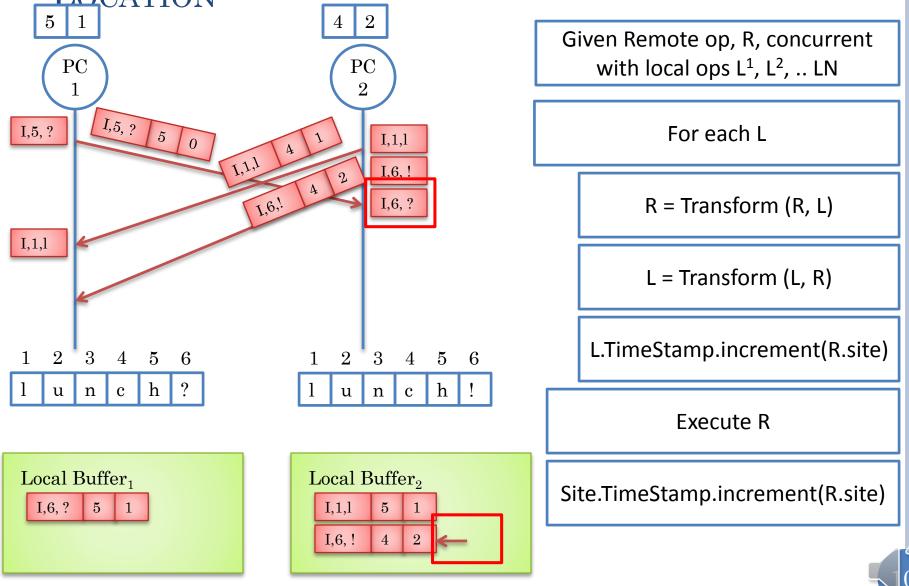
# I,5, ? Arrives at Site 2



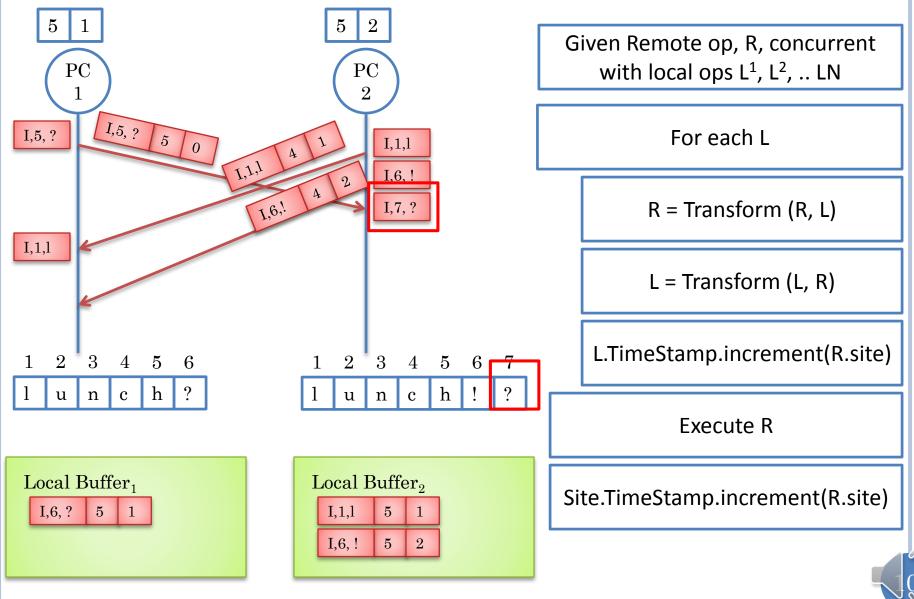
### REMOTE TRANSFORMED BUT NOT FIRST LOCAL



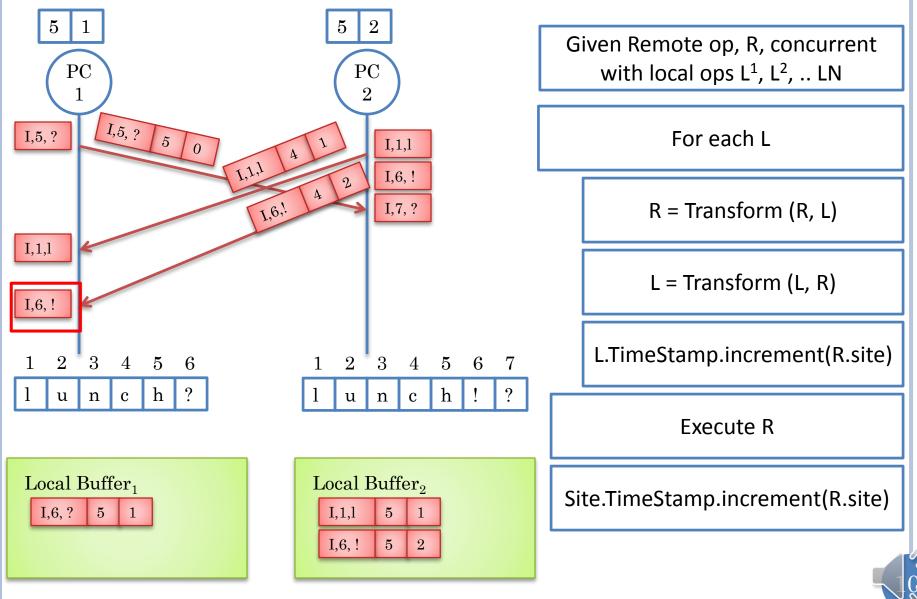
# COMPARED WITH SECOND LOCAL (I. 6, !) AT SAME LOCATION

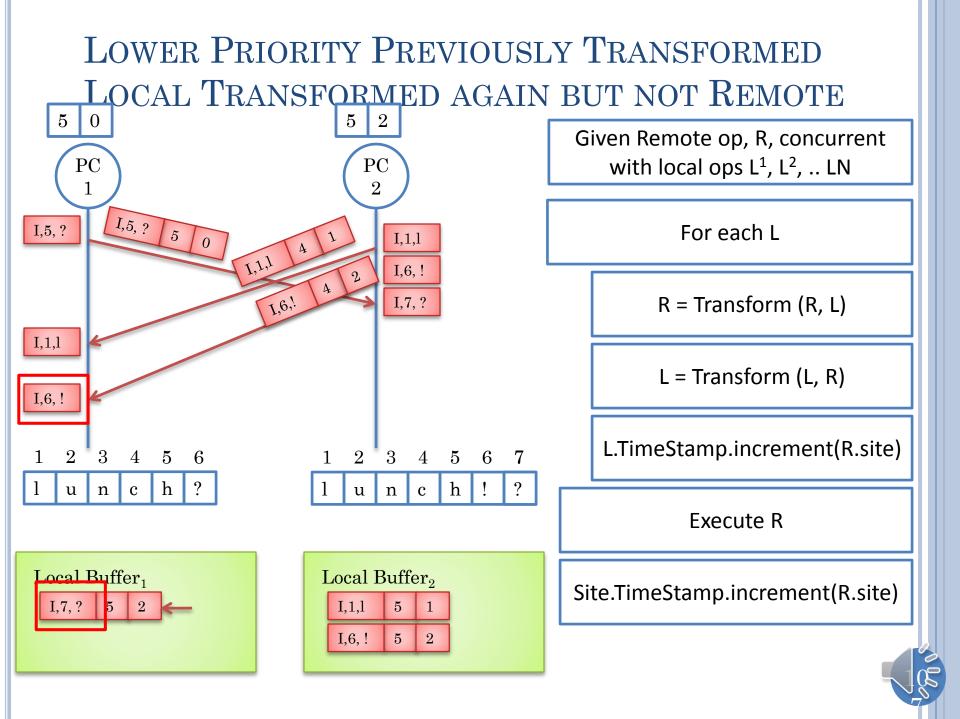


### TRANSFORMATION AND APPLICATION

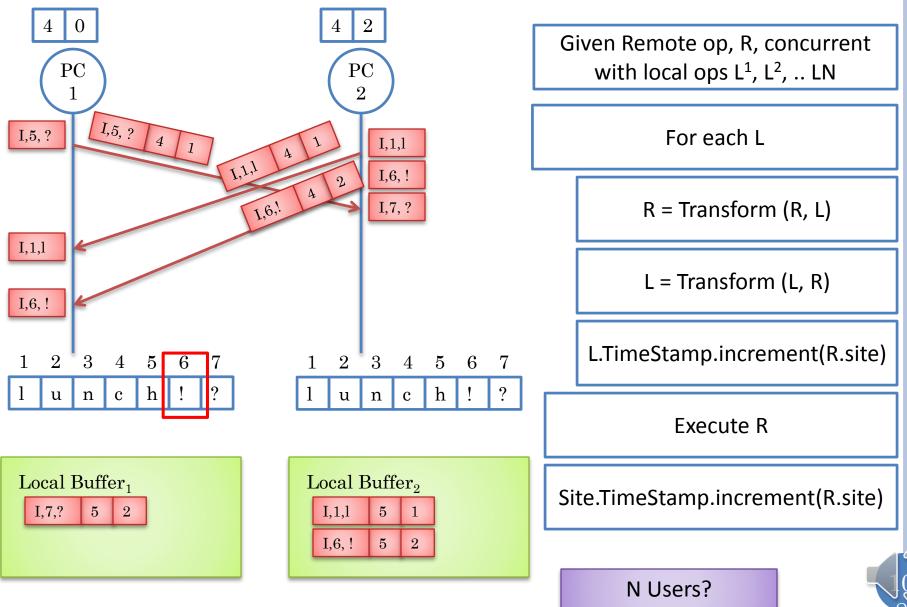


### SECOND REMOTE OPERATION ARRIVES AT SITE 1

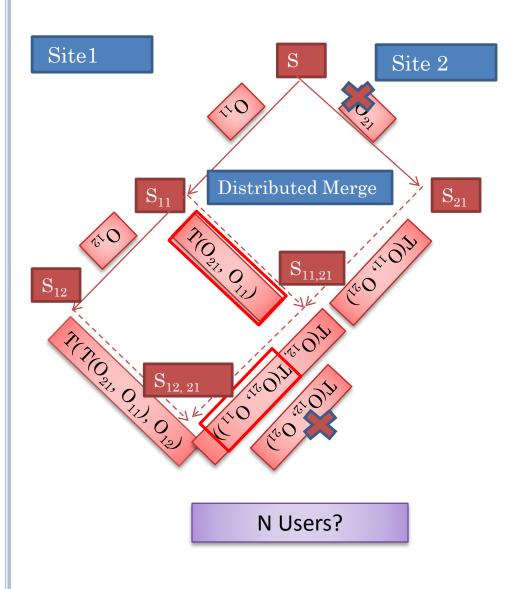


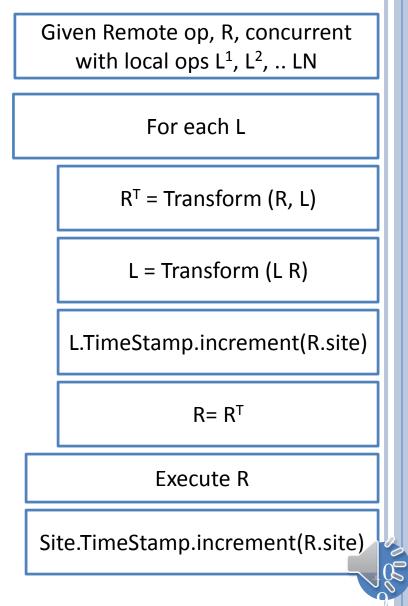


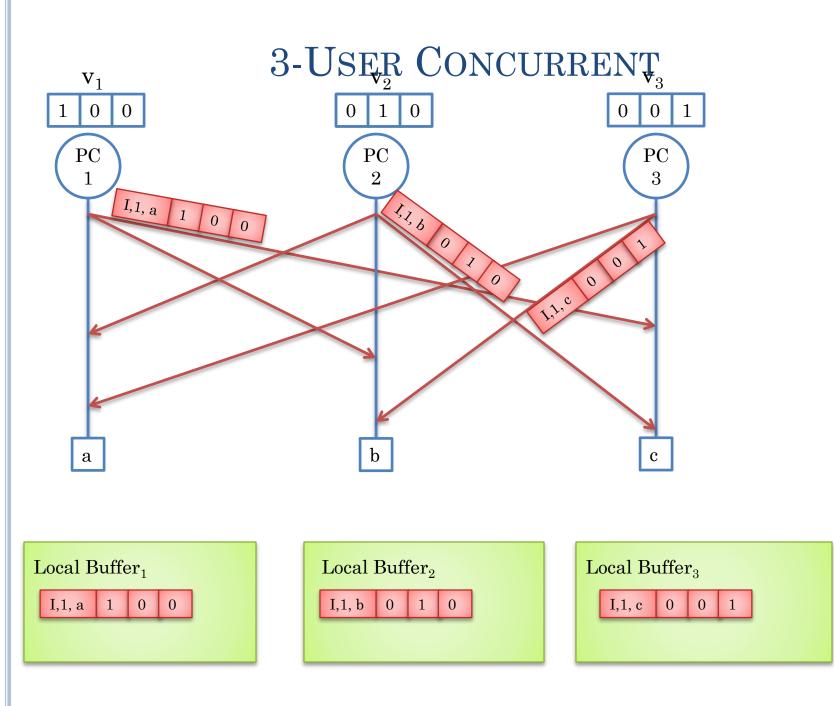
# UNTRANSFORMED REMOTE APPLIED

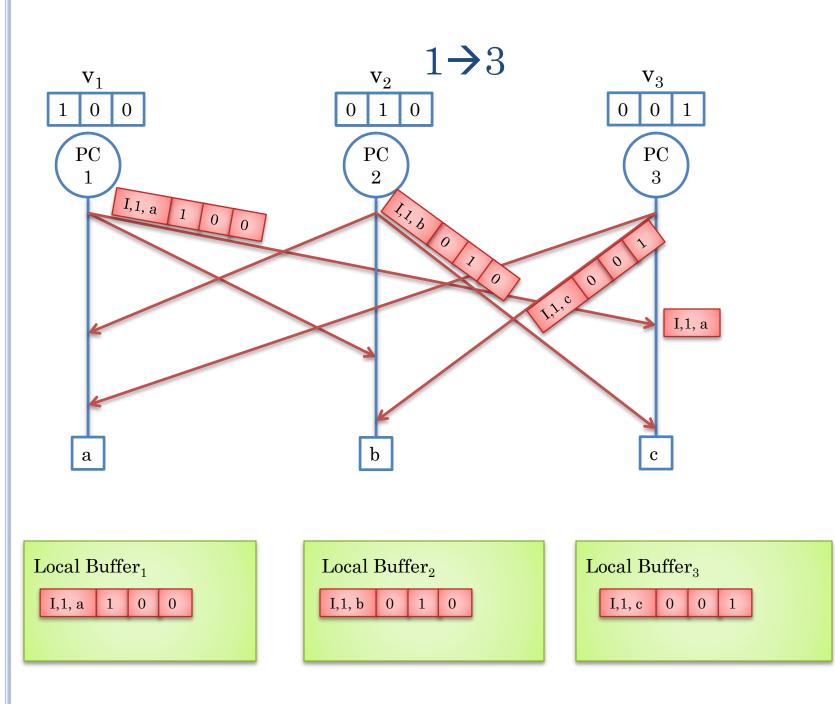


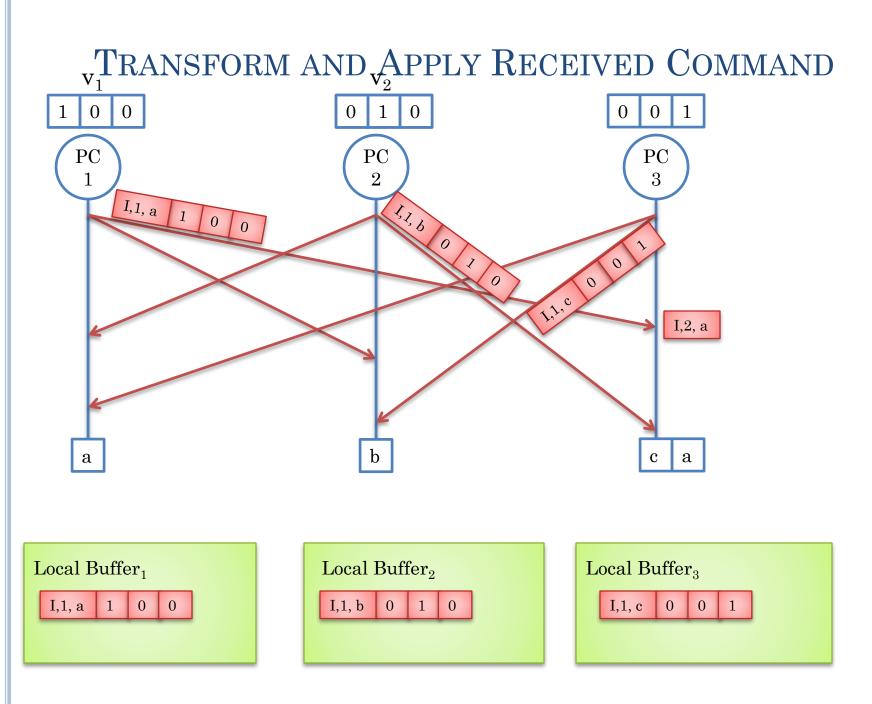
#### Algorithm and Proof Of Correctness

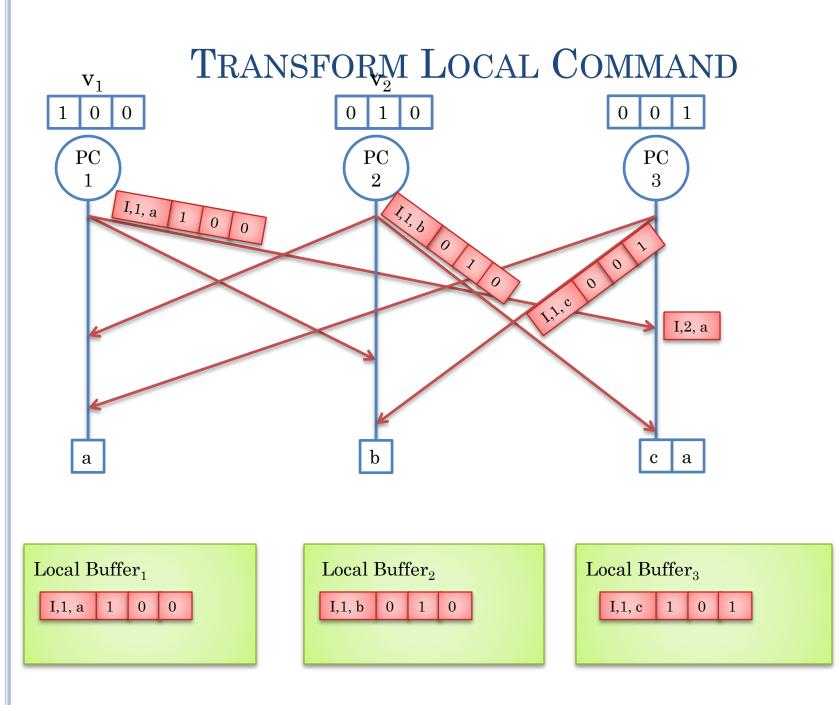


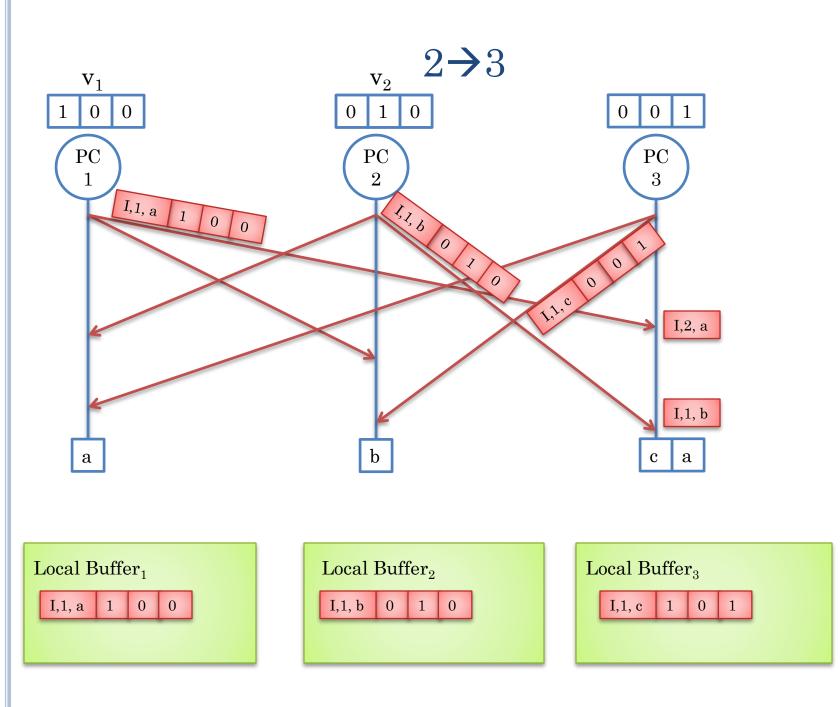




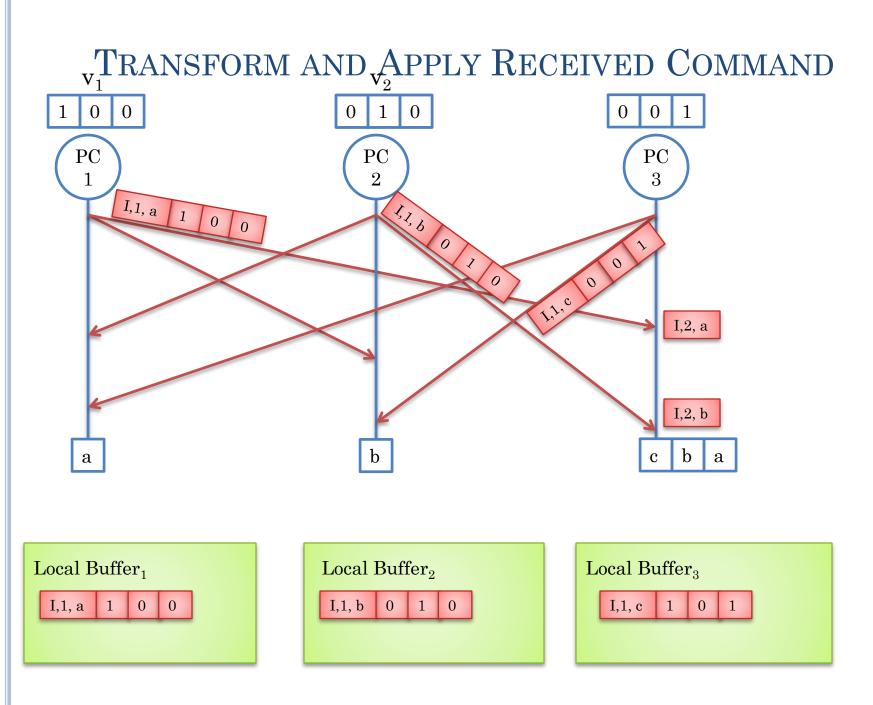


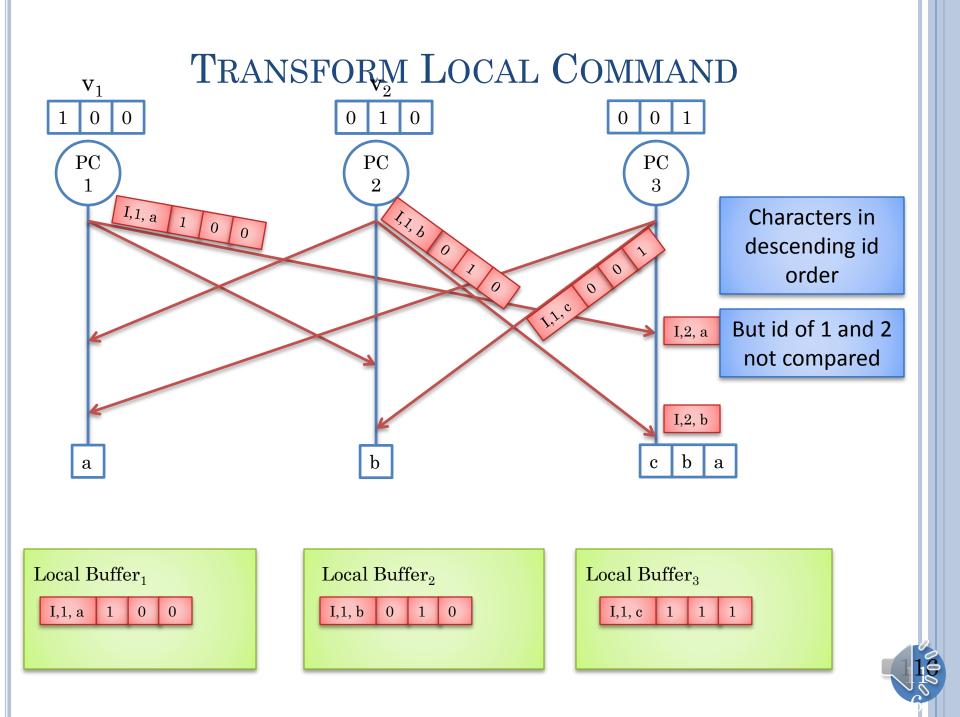


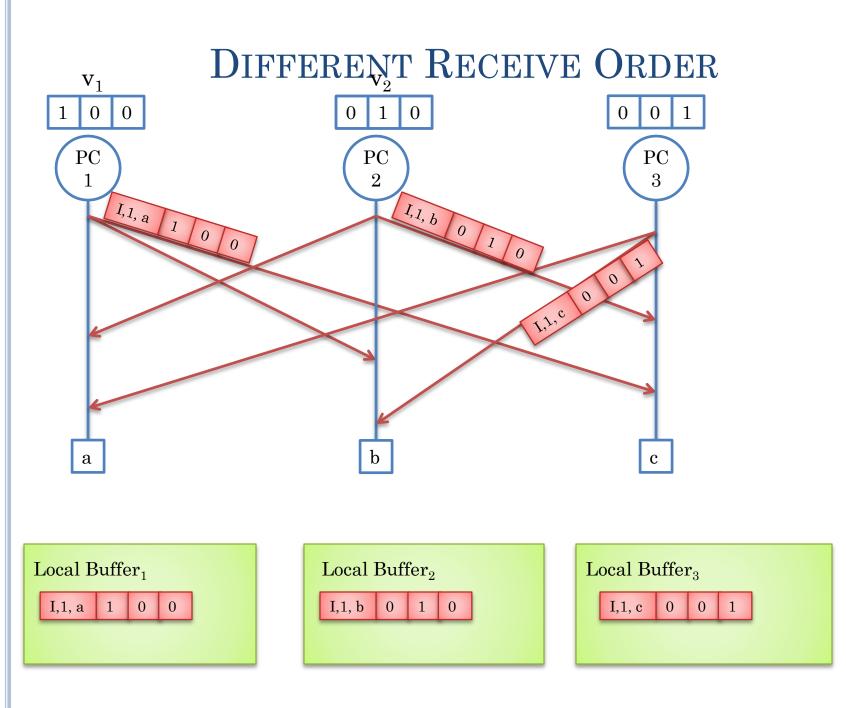




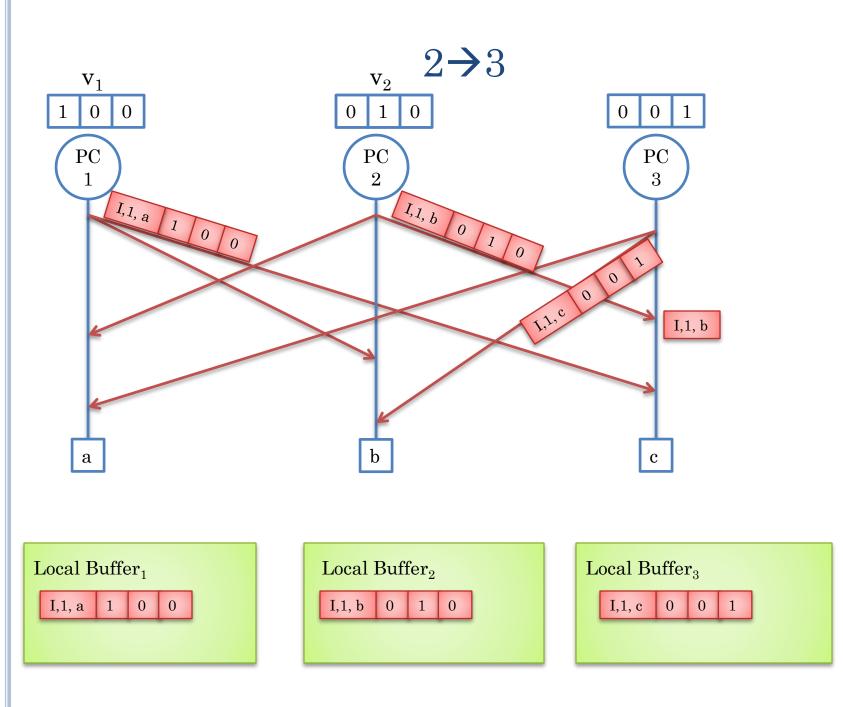


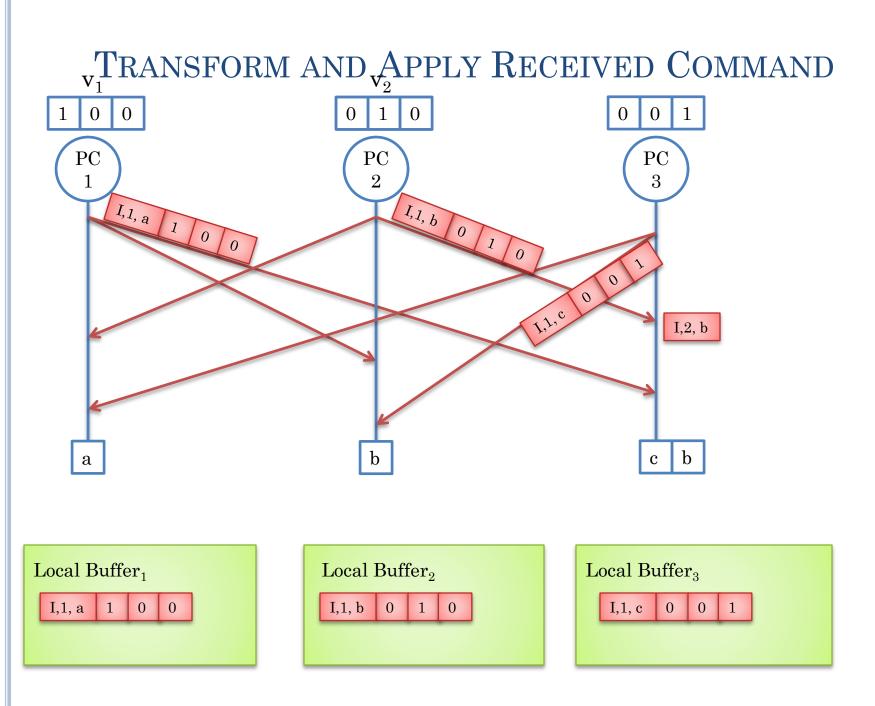


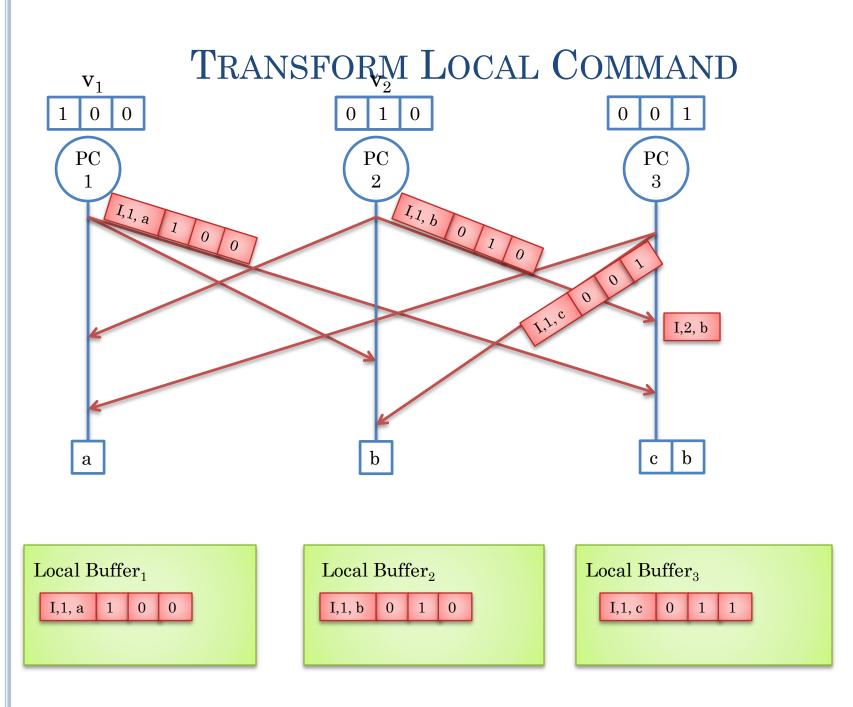




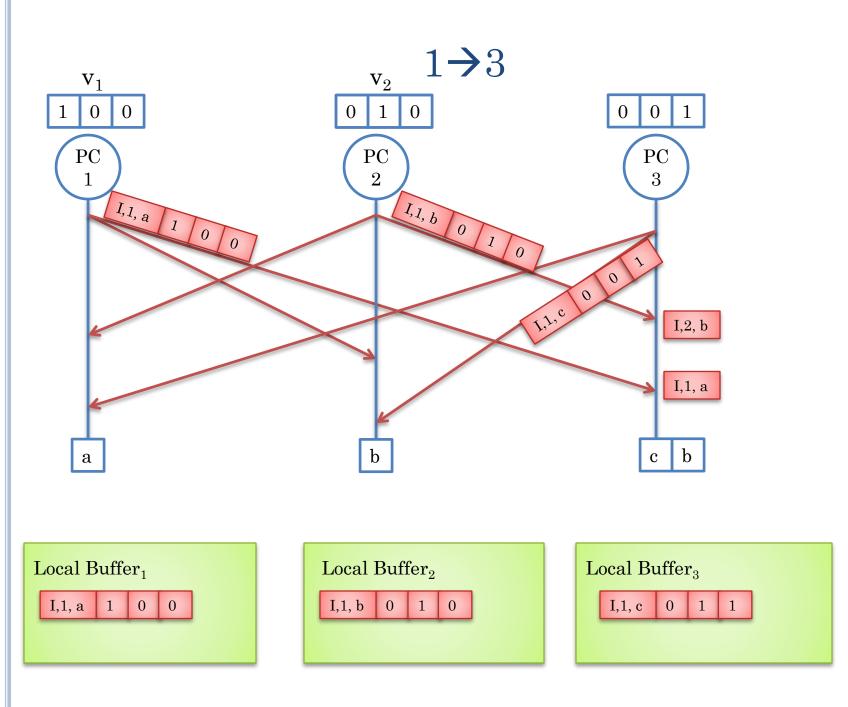


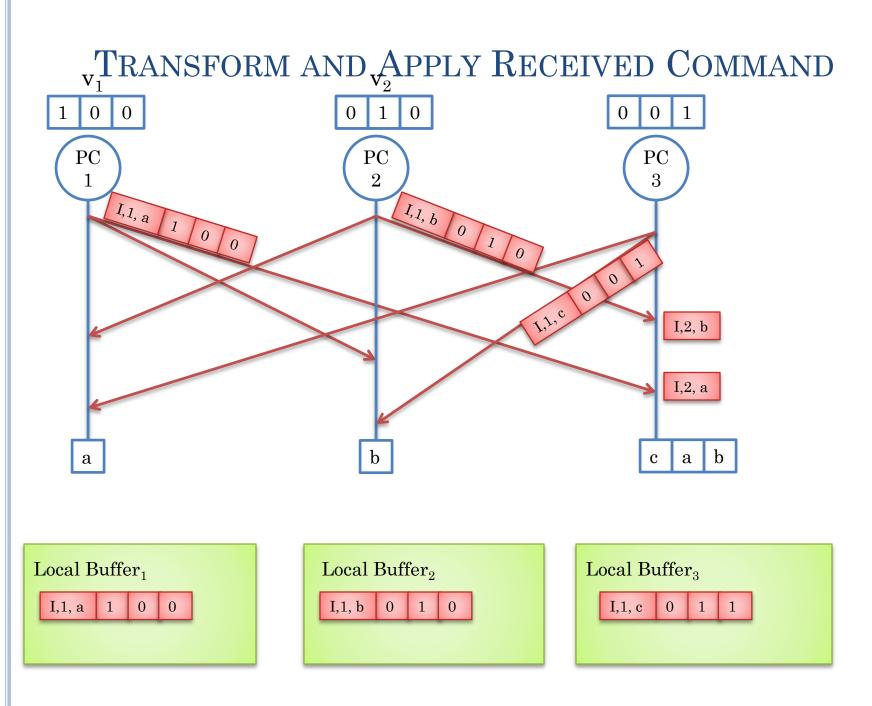


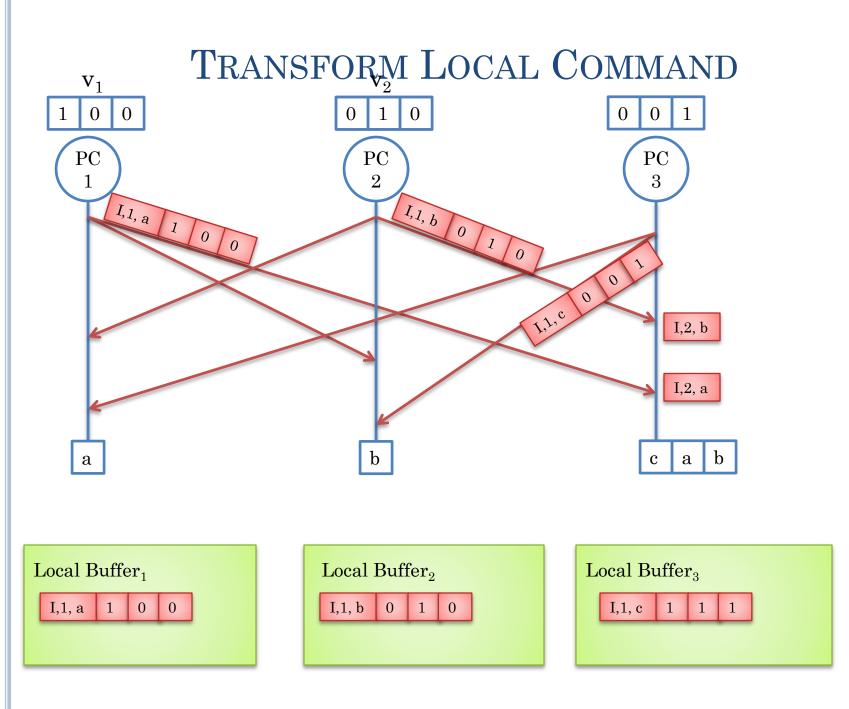








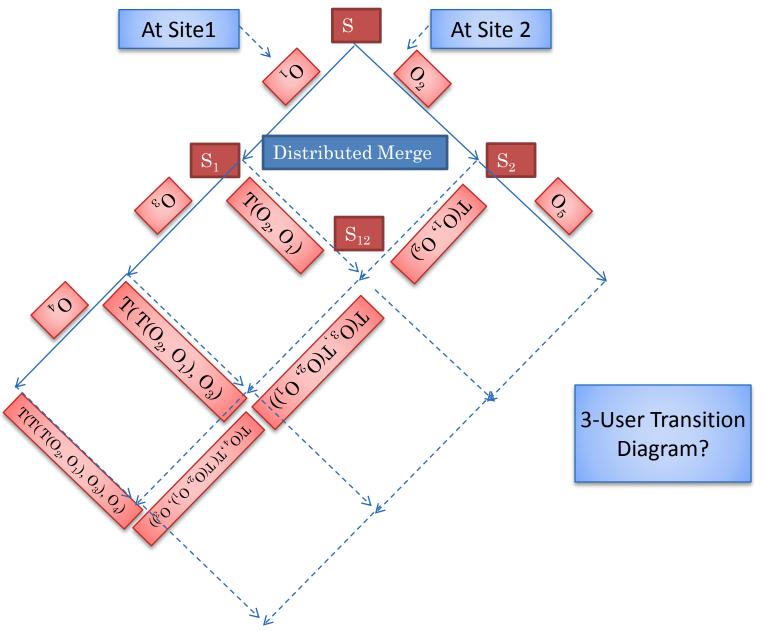


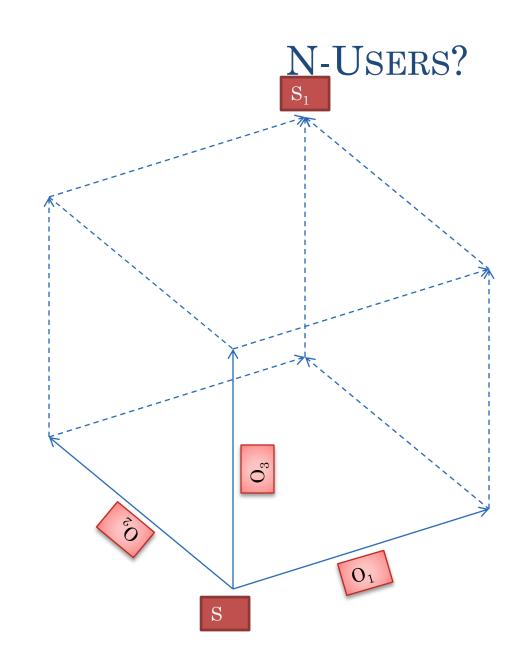


#### PROBLEM WITH 2 USERS

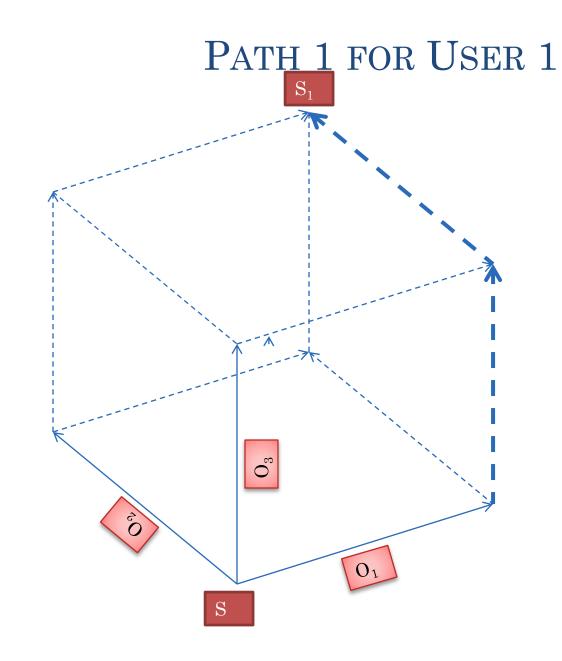
- Order of concurrent messages influences output.
- Same output not guaranteed at a single site.
- Same output not guaranteed at all sites.
- Problem independent of whether local operation is transformed.
- To understand better, need state transition diagram

#### MULTIPLE REMOTE CONCURRENT OPERATIONS



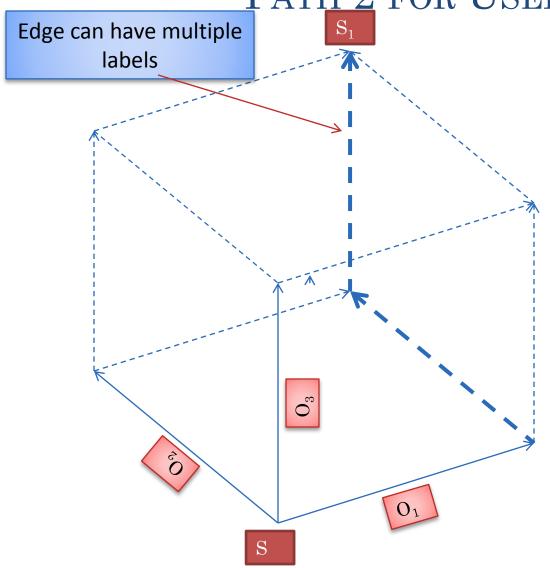








#### PATH 2 FOR USER 1

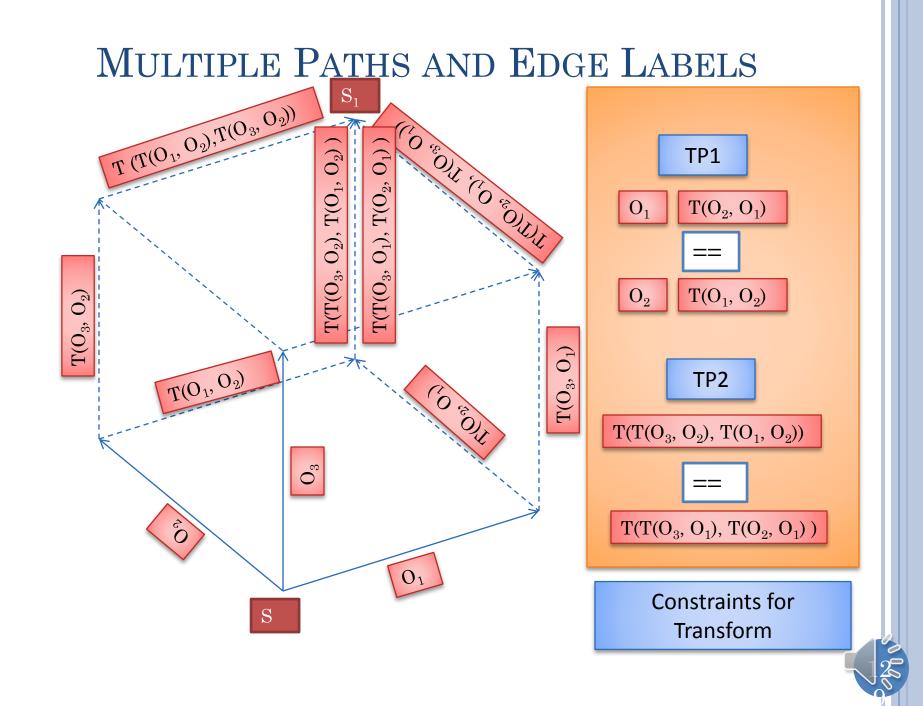


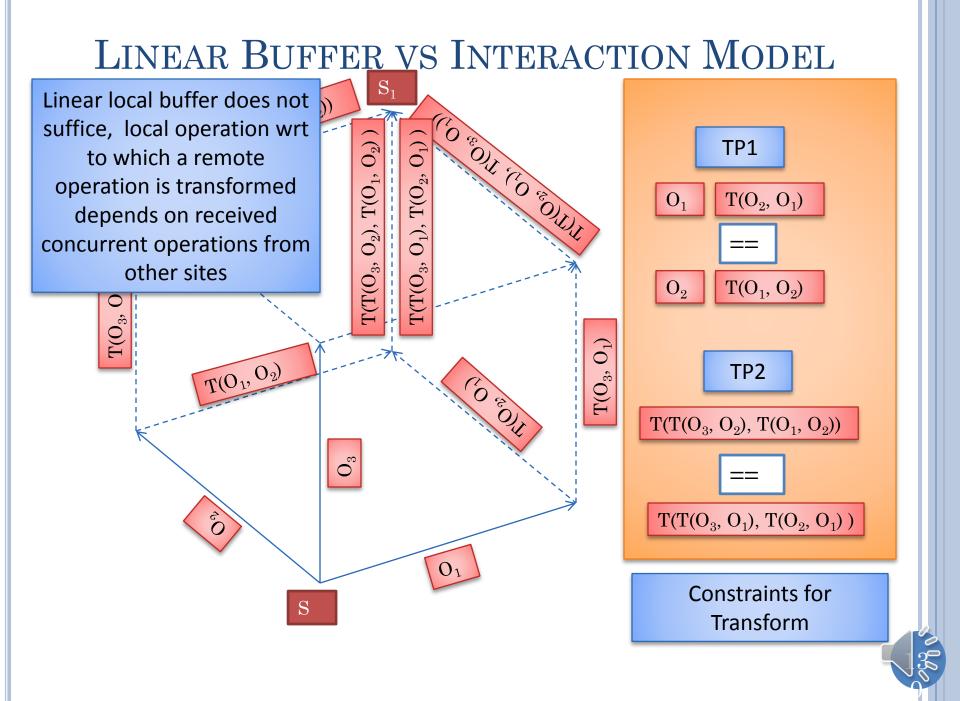
The two paths must give equivalent results.

In our example, our transformation functions did not!

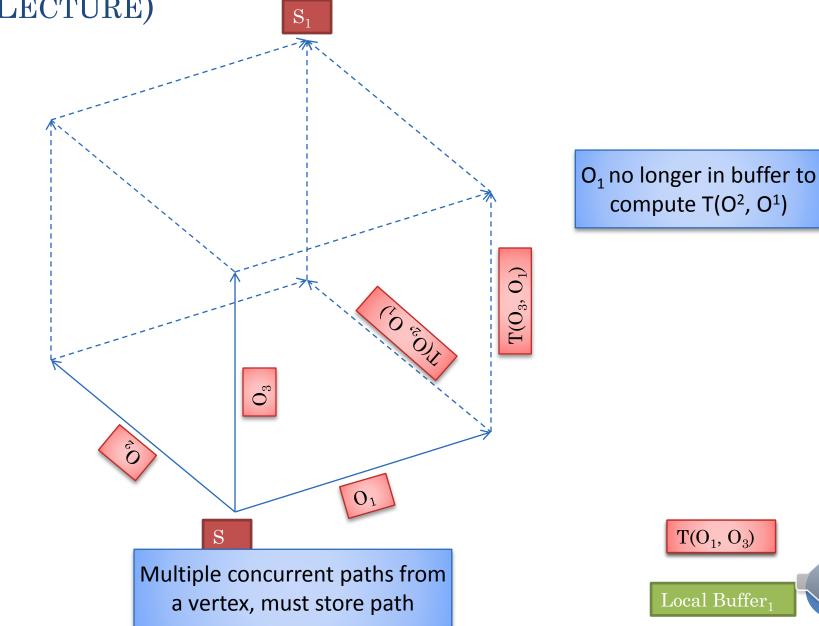
Necessary condition for new transformation functions?

Each edge should have a unique label

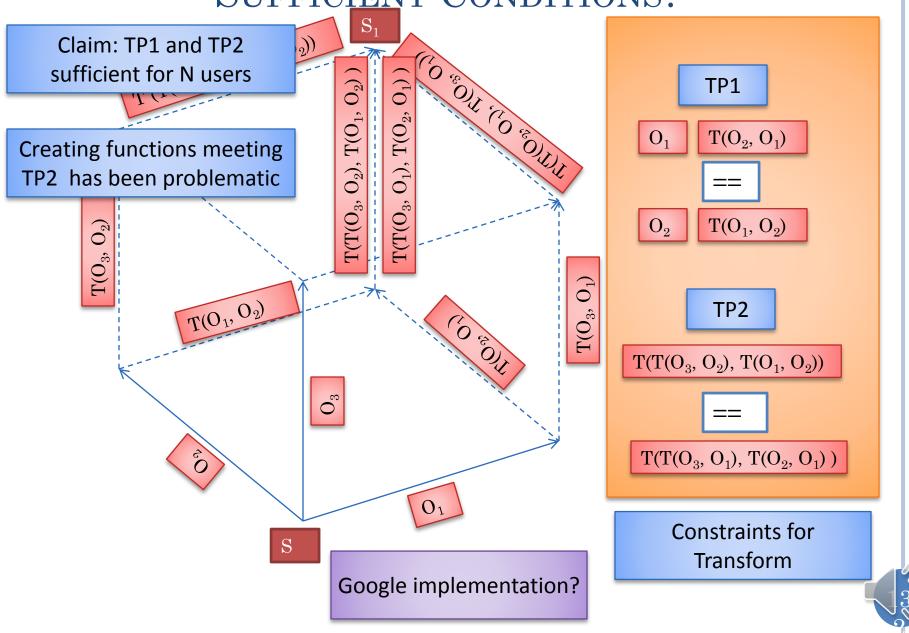




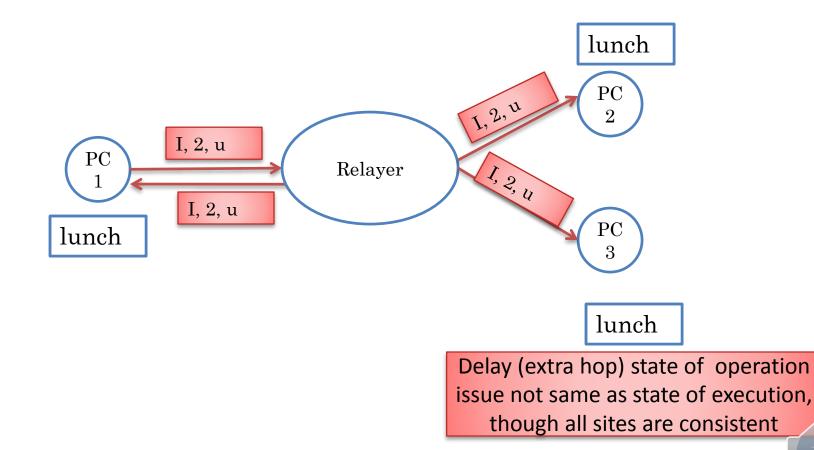
#### LINEAR BUFFER VS INTERACTION MODEL (POST LECTURE) S<sub>1</sub>



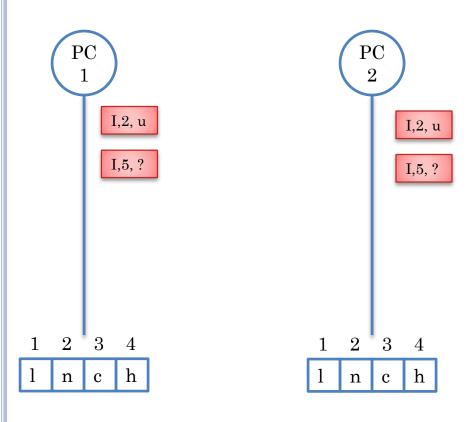
#### SUFFICIENT CONDITIONS?



#### SYNCHRONOUS RELAYED BROADCAST (REVIEW)

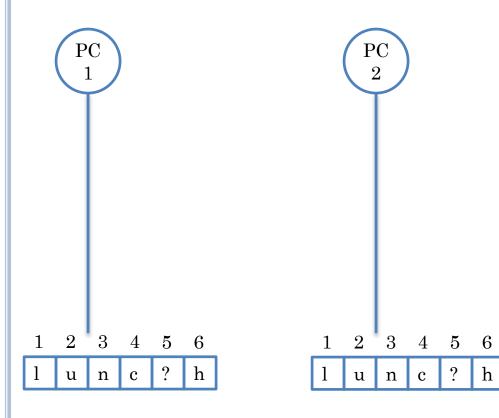


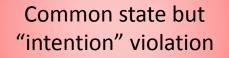
#### ORDERING WITH ATOMIC BROADCAST (REVIEW)



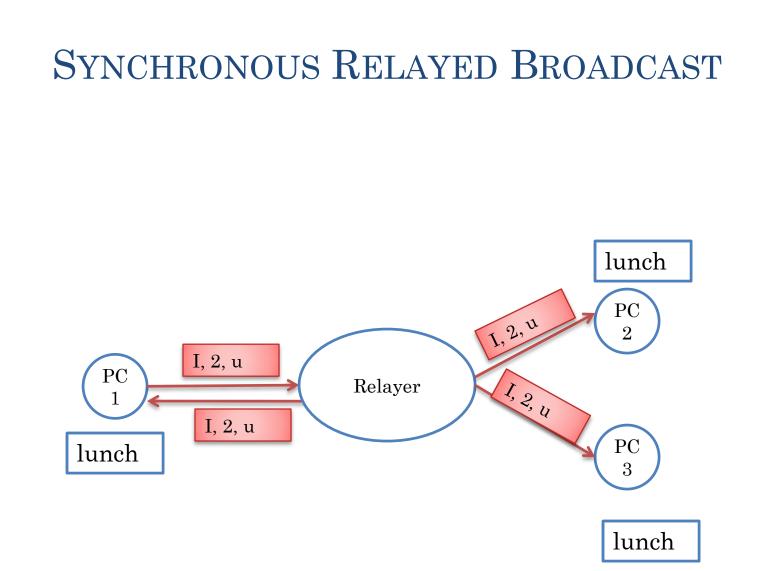


#### SECOND OPERATION EXECUTES (REVIEW)



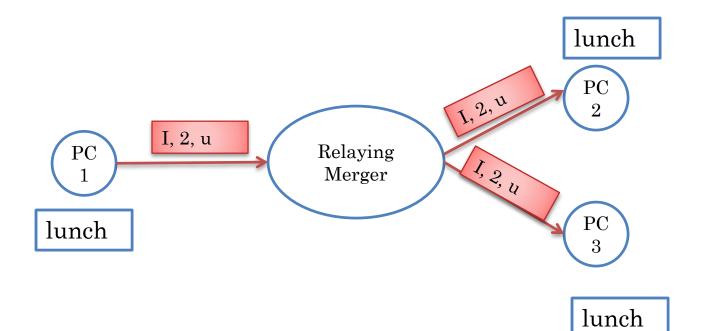


Context of operation not the same as when it was issued





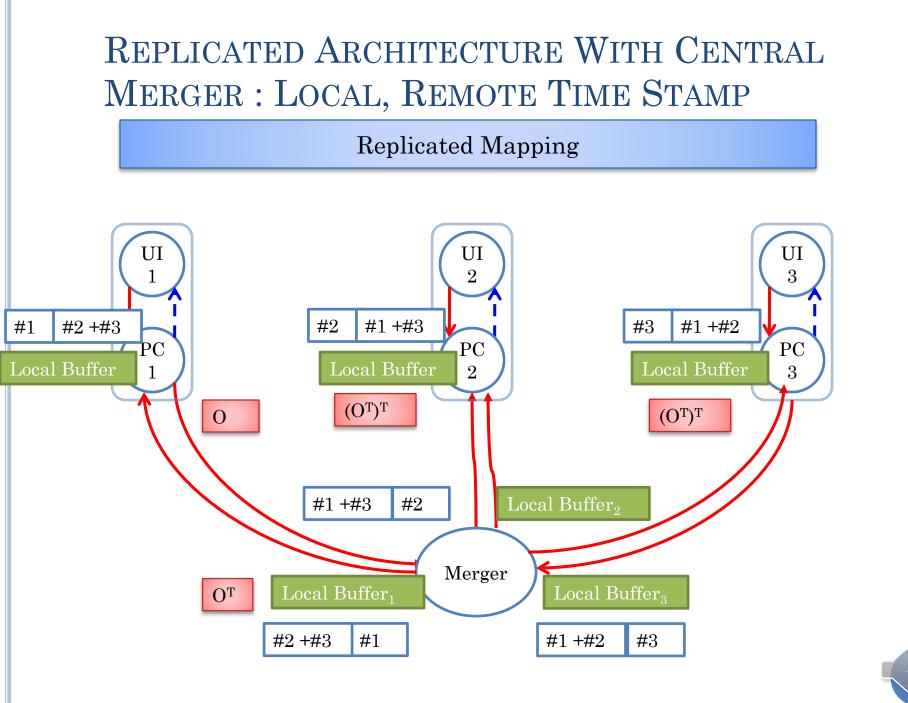
## ASYNCHRONOUS MERGED RELAYED BROADCAST





#### 2 to N Users

- Can do one N-user merge
- Can do N 2-User messages
  - Through a server
  - Each client is consistent with the server
  - Implies each client is consistent with the server
- But server does not issue any operations
  - For each client, server operations are those issued by other clients



#### CLIENT AND SERVER RECEIVE ALGORITHM

Given Remote op, R, concurrent with local ops L<sup>1</sup>, L<sup>2</sup>, .. LN

For each L

 $R^{T}$  = Transform (R, L)

L = Transform (L R)

L.TimeStamp.increment(R.site)

R= R<sup>⊤</sup>

Execute R

Site.TimeStamp.increment(R.site)

#### Client

Given Remote op, R, concurrent with local ops L<sup>1</sup>, L<sup>2</sup>, .. LN

For each L

 $R^{T}$  = Transform (R, L)

L = Transform (L R)

L.TimeStamp.increment(R.site)

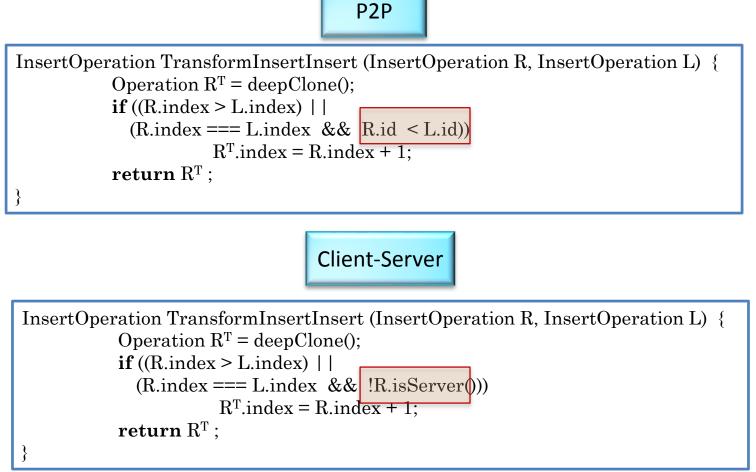
 $R = R^T$ 

For all other sites assume server executed R

Site.TimeStamp.increment(R.site)

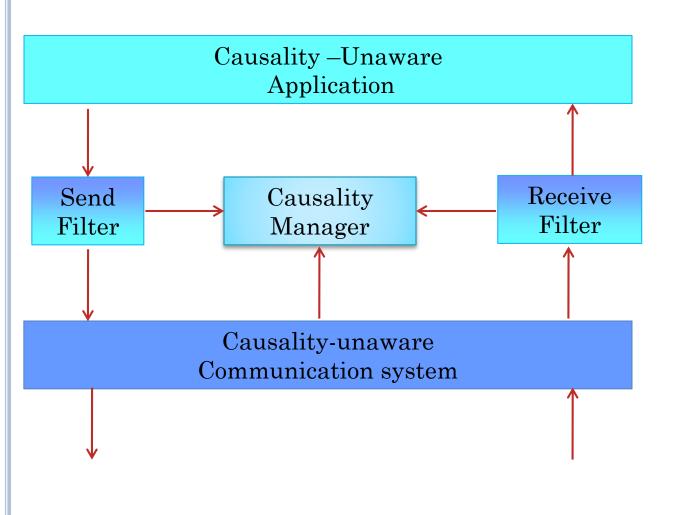
Server

#### TRANSFORM OPERATION FOR CLIENT-SERVER CASE



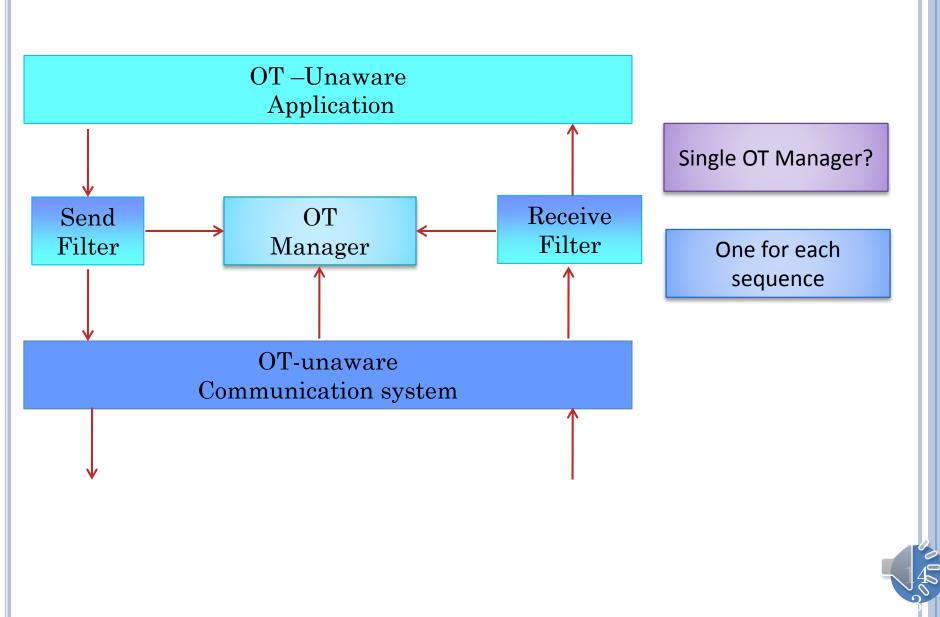


#### CAUSALITY MANAGER

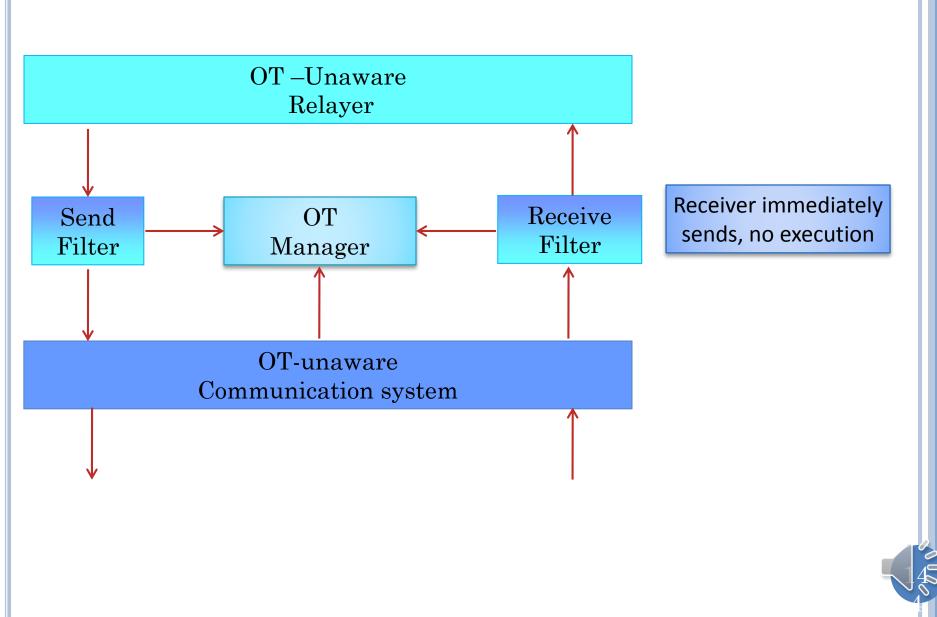




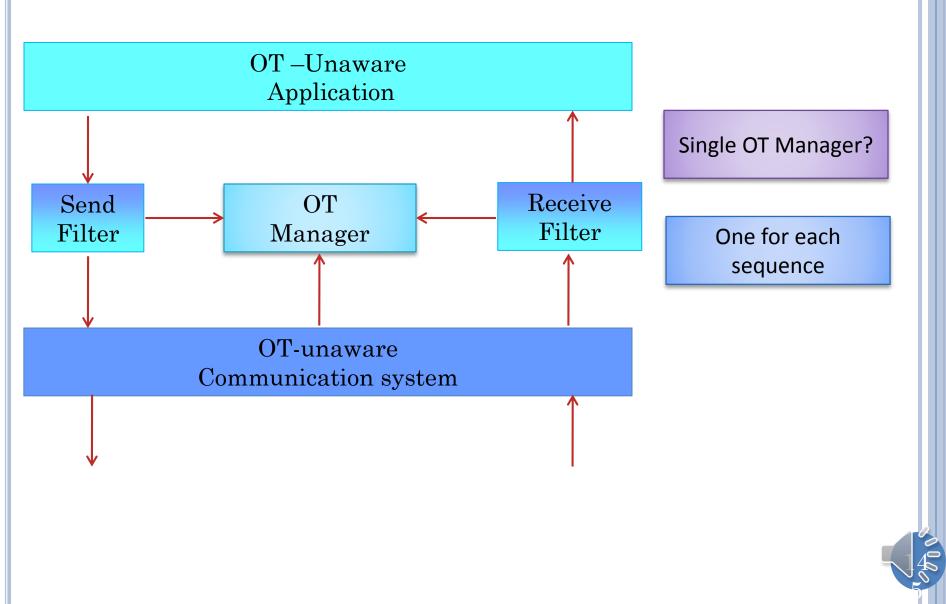
### CLIENT OT MANAGER



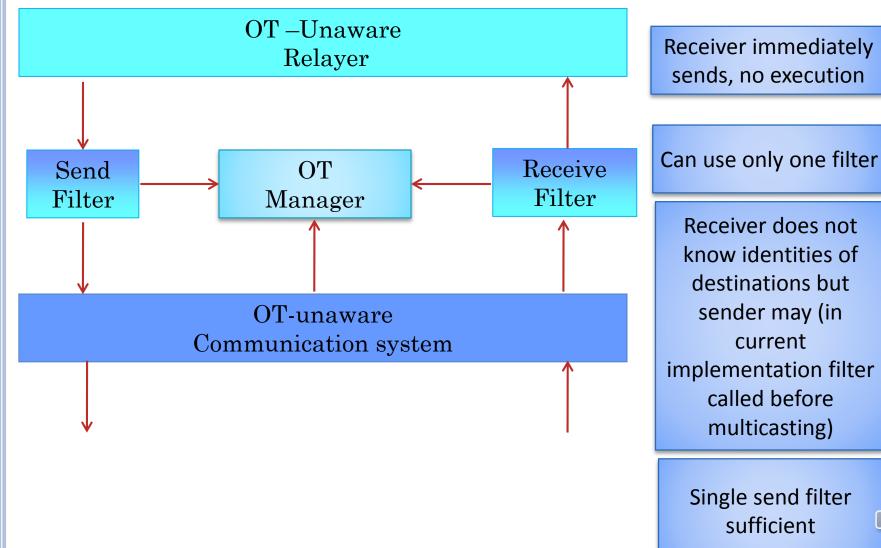
### SERVER OT MANAGER



# CLIENT OT MANAGER (REVIEW)



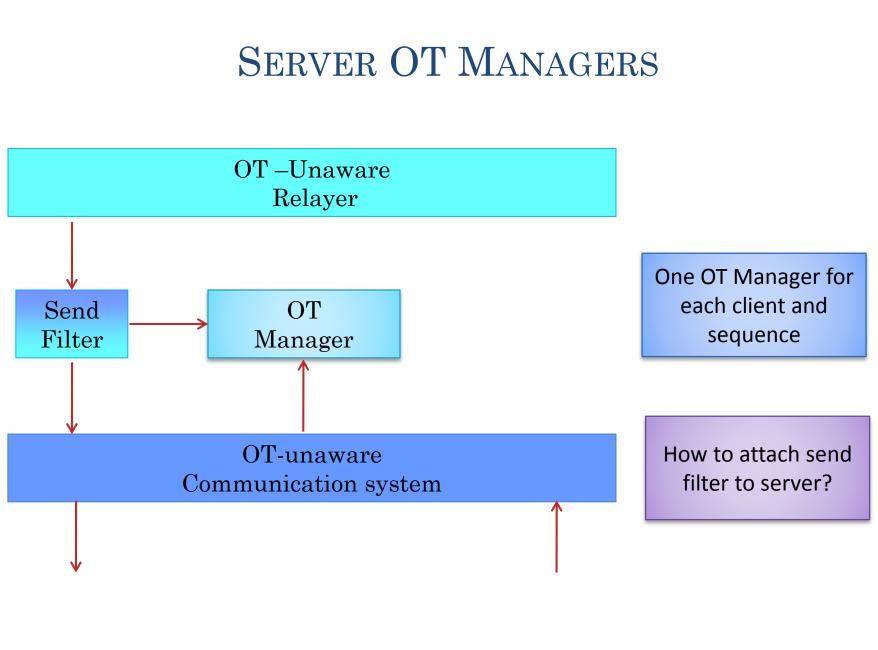
# SINGLE SERVER FILTER



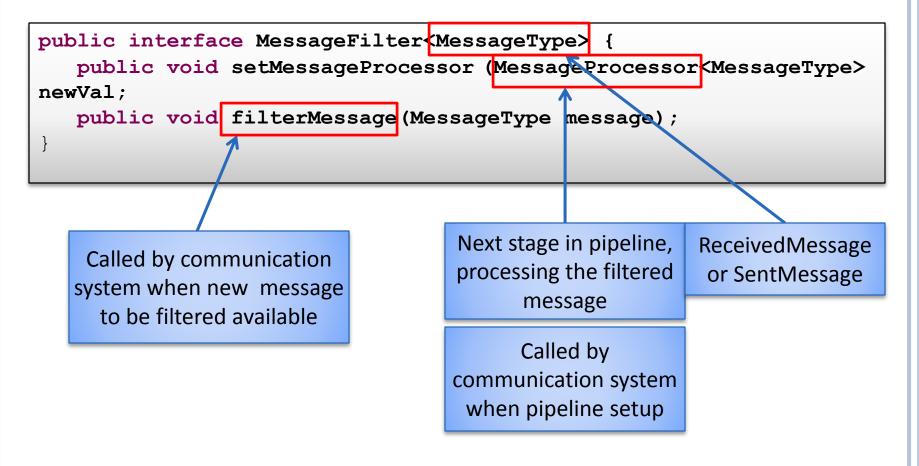
Receiver does not

know identities of destinations but sender may (in current implementation filter called before multicasting)

> Single send filter sufficient



#### MESSAGE FILTER INTERFACE





# SERVER MESSAGE FILTER INTERFACE

public interface ServerMessageFilter extends MessageFilter<SentMessage> {
 public void userJoined(String aSessionName, String anApplicationName,
 String userName);

```
public void userLeft(String aSessionName, String anApplicationName,
String userName);
```

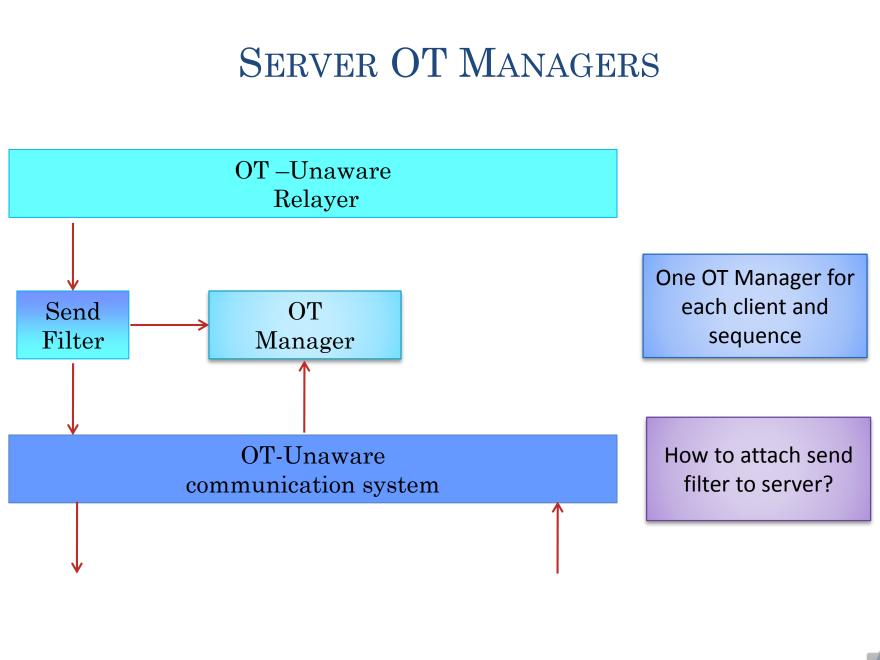
#### SERVER FACTORY INTERFACE

public interface ServerMessageFilterCreator {
 ServerMessageFilter getServerMessageFilter();

#### SEND FILTER (FACTORY) SELECTOR OR ABSTRACT FACTORY

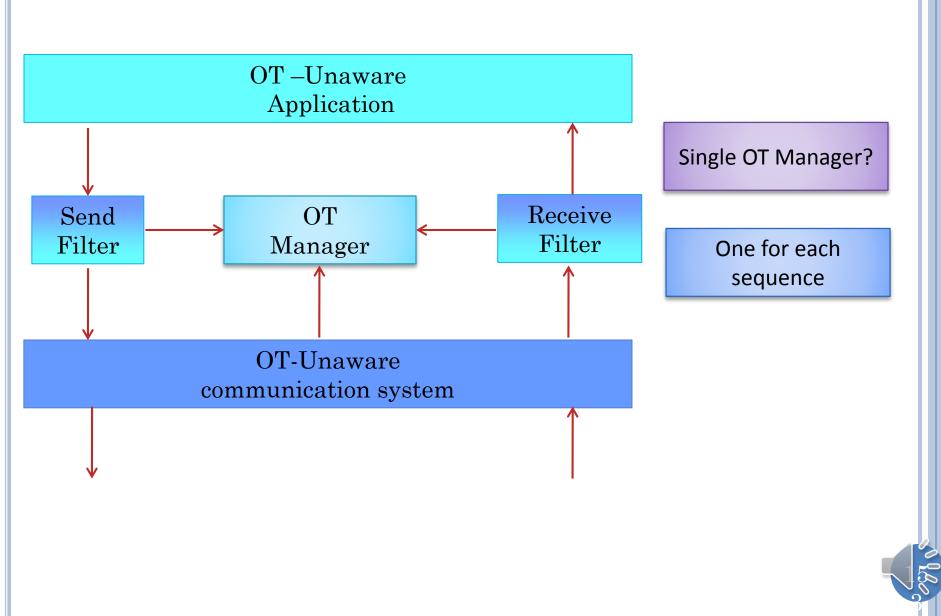
```
public class SentMessageFilterSelector {
   static MessageFilterCreator<SentMessage> filterFactory =
        new AMessageForwarderCreator<SentMessage>();
   public static MessageFilterCreator<SentMessage> getMessageFilterCreator() {
        return filterFactory;
    }
   public static void setMessageFilterCreator(
        MessageFilterCreator<SentMessage> theFactory) {
        filterFactory = theFactory;
    }
}
```





1.A

# CLIENT OT MANAGER



#### CLIENT INITIALIZATION

Init

Create <List, OT Manager> Mapping ListOTManager

For each List, L

ListOTManager(L) ← new OT Manager (ClientName, Not Server)

Create Client Send and Receive Filter Factories, passing them OTManager so they can pass them to the two filters

# SEND FILTER TRACEABLE STEPS

Send Filter

On each user edit about OT List L

Ask ListOTManager(L) to time stamp edit

OTListEditSend the timestamped edit through message processor

Ask ListOTManager(L) to store copy of sent message

As in causality must ensure changing site time stamp does not change message time stamp

### RECEIVE FILTER TRACEABLE STEPS

**Receive Filter** 

On each OTListEditReceived for list L received (through server)

OTListEditFlipped time stamp

Ask OTManager(L) to transform received edit

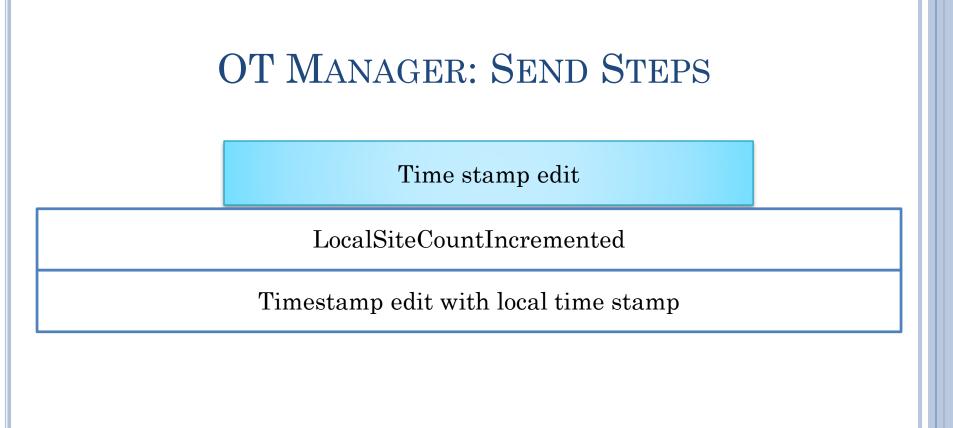
Pass transformed edit to message processor



## OT MANAGER: INIT

Init (User Name, IsServer)

Initial OTT ime Stamp Created



Store Sent Message

MessageBuffered

#### **OT MANAGER: RECEIVE FILTER COMMUNICATION**

Process received timestamped edit

For each local message not concurrent with received edit

Local MessageUnBuffered

For each local buffered concurrent edit L

L= TransformationResult from Transform(L, R)

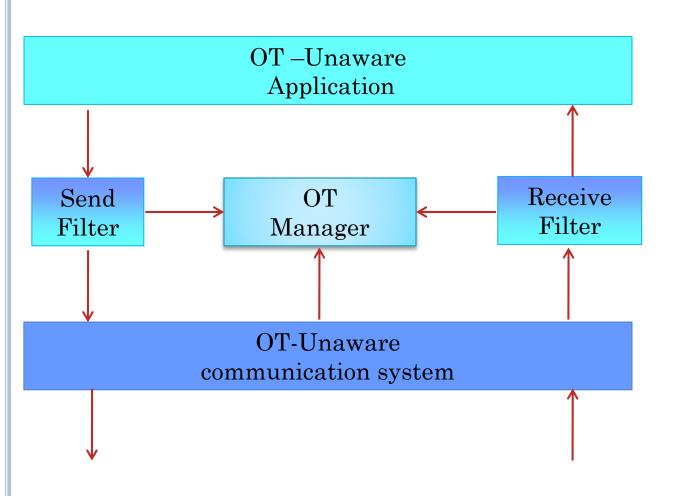
R= TransformationResult from Transform(R, L)

OTListEditRemoteCountIncremented in L

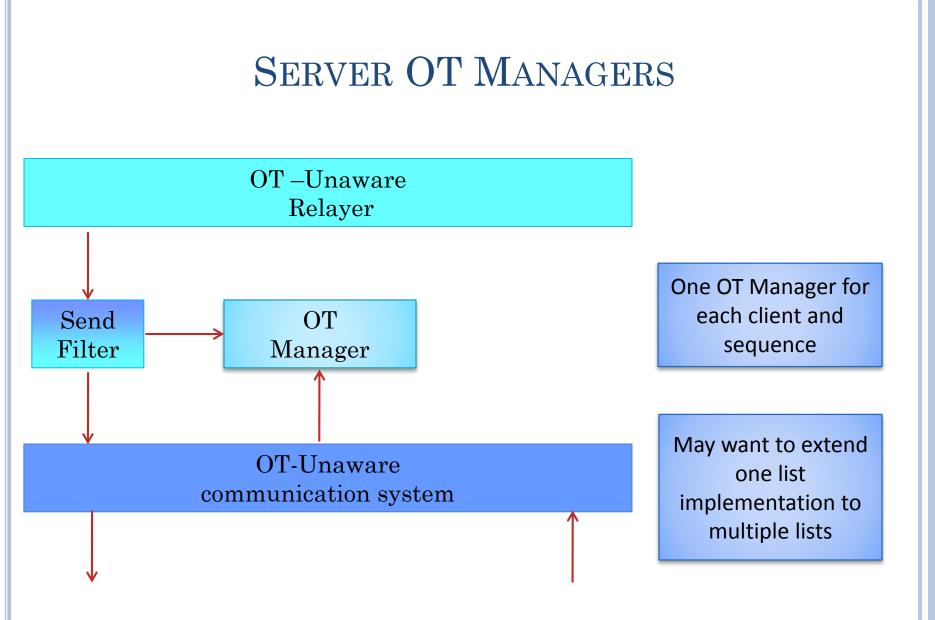
RemoteSiteCountIncremented

User name in trace step is name of user who executed the operation, for local edit, the local user, for remote edit, the remote user (exact name, not server)

# CLIENT OT MANAGER

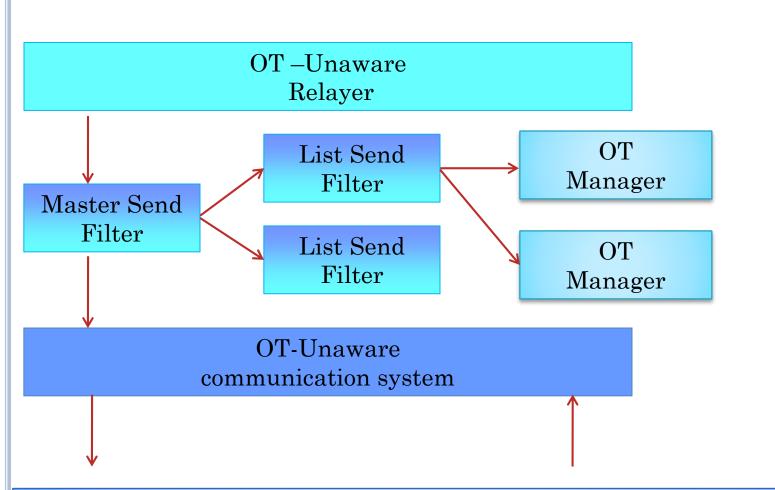








# FILTER COMPOSITION



Could have used this architecture (master/delegate filters) also for clients, but client filters were simple and there were two of them, so it is not clear creating two additional master filters for clients is worth it

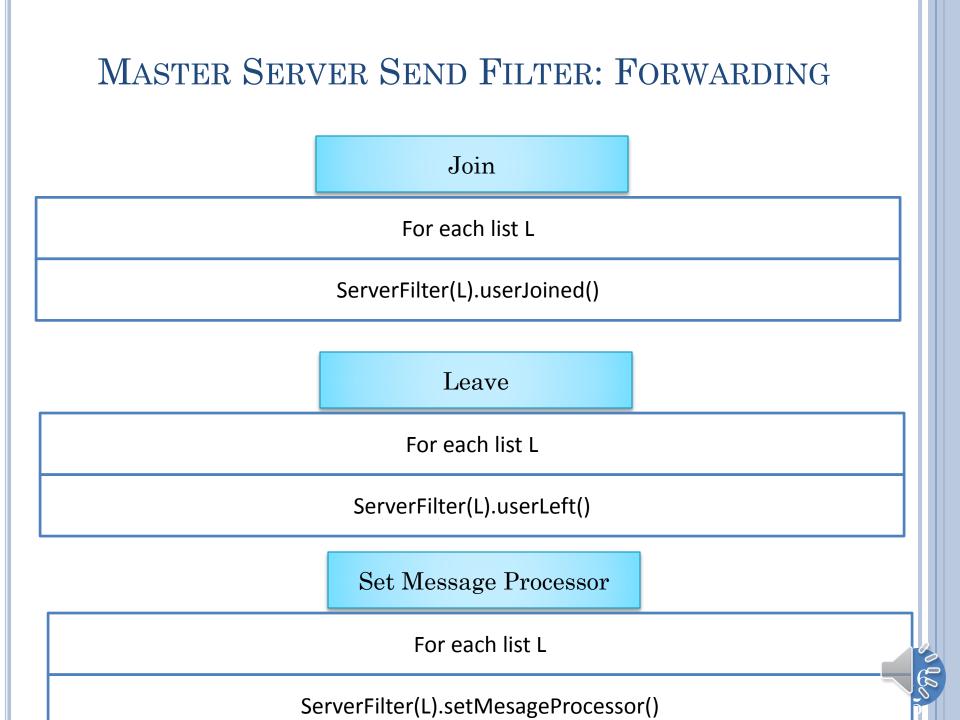
SERVER INITIALIZATION
Init
Create <list, serverfilter=""> Mapping ServerFilter</list,>
For each List, L
ServerFilter(L) ← new ServerFilter()
Create Send Filter Factory, passing it ServerFilterMapping

# SERVER MESSAGE FILTER INTERFACE

public interface ServerMessageFilter extends MessageFilter<SentMessage> {
 public void userJoined(String aSessionName, String anApplicationName,
 String userName);

public void userLeft(String aSessionName, String anApplicationName, String userName);

}



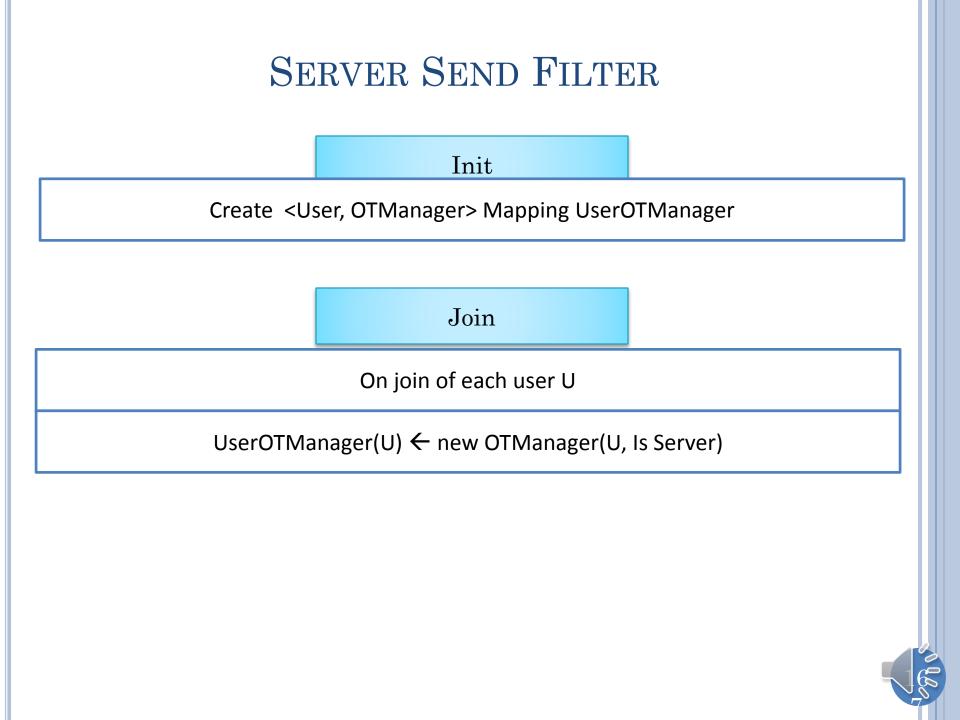
### MASTER SERVER SEND FILTER

New Message

On each client edit of list L

ServerFilter(L).filterMessage()





### SERVER SEND FILTER

New Message

On each OTListEditReceived from Sending User S

OTListEditFlipped time stamp

Ask UserOTManager(S) to create transform received edit, O<sup>T</sup>

For each user R in UserOTManager other than S

Create unicast copy of message containing, O<sup>TR</sup> By calling ASentMessage.toSpecificUser(message, R)

Ask UserOTManager(R) to time stamp O<sup>TR</sup>

Send timestamped edit through message processor

Ask ListOTManager(R) to store copy of sent message

# CENTRALIZED ALGORITHM

• Assume all users merge through central server

- Output is produced locally immediately
- Server keeps local buffer and timestamp for each client
- Each client treats server as second user and sends it each command
- Instead of applying (possibly transformed) command to its local state server sends time-stamped command to each remote client
- Client transforms it further if it has executed concurrently
- Client assumes command executed directly by server
- Each client consistent with server, and thus with each other client
- Unique ordering of all commands from remote machines

# HISTORY

- dOPT (Distributed Operation Transformation) Ellis and Gibbs '89
  - Did not transform local operation
  - Had known problem with multiple users
- adOPTed (Ressel et al '96)
  - Transformed local operation
  - Give conditions for N-user replicated merging
- Jupiter (Nichols, Curtis, et al '95)
  - Centralized merging
  - Inventors of LiveMeeting
  - Implemented in GoogleWave