

Global State Recording

Global State Collection

- Applications:
 - Checking “stable” properties, checkpoint & recovery
- Issues:
 - Need to capture both node and channel states
 - system cannot be stopped
 - no global clock

Notations

Some notations:

- LS_i : Local state of process i
- $\text{send}(m_{ij})$: Send event of message m_{ij} from process i to process j
- $\text{rec}(m_{ij})$: Similar, receive instead of send
- $\text{time}(x)$: Time at which state x was recorded
- $\text{time}(\text{send}(m))$: Time at which $\text{send}(m)$ occurred

Definitions

- $\text{send}(m_{ij}) \in \text{LS}_i$ iff $\text{time}(\text{send}(m_{ij})) < \text{time}(\text{LS}_i)$
- $\text{rec}(m_{ij}) \in \text{LS}_j$ iff $\text{time}(\text{rec}(m_{ij})) < \text{time}(\text{LS}_j)$
- $\text{transit}(\text{LS}_i, \text{LS}_j)$
 $= \{ m_{ij} \mid \text{send}(m_{ij}) \in \text{LS}_i \text{ and } \text{rec}(m_{ij}) \notin \text{LS}_j \}$
- $\text{inconsistent}(\text{LS}_i, \text{LS}_j)$
 $= \{ m_{ij} \mid \text{send}(m_{ij}) \notin \text{LS}_i \text{ and } \text{rec}(m_{ij}) \in \text{LS}_j \}$

Definitions

- Global state: collection of local states

$$GS = \{LS1, LS2, \dots, LS_n\}$$

- GS is consistent iff

for all $i, j, 1 \leq i, j \leq n$,

$$\text{inconsistent}(LS_i, LS_j) = \emptyset$$

- GS is transitless iff

for all $i, j, 1 \leq i, j \leq n$,

$$\text{transit}(LS_i, LS_j) = \emptyset$$

- GS is strongly consistent if it is consistent and transitless.

Chandy-Lamport's Algorithm

- Uses special marker messages.
- One process acts as initiator, starts the state collection by following the marker sending rule below.
- Marker sending rule for process P:
 - P records its state and
 - For each outgoing channel C from P on which a marker has not been sent already, P sends a marker along C before any further message is sent on C

Chandy Lamport's Algorithm contd..

- When Q receives a marker along a channel C:
 - If Q has not recorded its state then Q records the state of C as empty; Q then follows the marker sending rule
 - If Q has already recorded its state, it records the state of C as the sequence of messages received along C after Q's state was recorded and before Q received the marker along C

Notable Points

- Markers sent on a channel distinguish messages sent on the channel before the sender recorded its states and the messages sent after the sender recorded its state
- The state collected may not be any state that actually happened in reality, rather a state that “could have” happened
- Requires FIFO channels
- Message complexity $O(|E|)$, where E = no. of links