

CS3110 Spring 2016 Lecture 25: Event Logic

Robert Constable

Topics

1. Reasoning about 2/3 consensus.

We will prove three critical properties of consensus protocols. We discussed these properties in Lecture 22. They are:

P1: Unanimous inputs of value v cause the protocol to decide on v .
We call this *unanimity*.

P2: All decided values are inputs. We call this *validity*.

P3: Any two decisions have the same value. We call this *agreement*.

2. The Logic of Events.

We can axiomatize and formally prove the above properties of protocols as well as others. The reasoning is based on events in the (general process) computing model. We can axiomatize this logic with a few intuitively clear axioms that capture features of the *message sequence diagrams*.

3. We will review the *enduring concepts* from the course next lecture. The great programming languages leave their mark on computing theory as well as on decades of important applications. The earliest major languages were Fortran, Algol, and Lisp – from the 1960s approximately. Fortran and Lisp are still well used, and Algol left its mark on the theory and methodology of creating and implementing programming languages.

1 Reasoning about 2/3 (simple) consensus

The Synthesized 2/3 Protocol

Begin

$r : \text{Nat}$
 $\text{decided}_i, \text{vote}_i : \text{Bool}$

$r = 0, L = \text{nil}$
 $\text{decided}_i = \text{false}$
 $v_i = \text{input to } P_i$
 $\text{vote}_i = v_i$

Until decided_i **do:**

1. $r := r + 1$
2. **Broadcast** $\text{vote} \langle r, \text{vote}_i \rangle$ to group G
3. **Collect** $2f + 1$ round r votes in list L
4. $\text{vote}_i := \text{majority}(L)$
5. **If** $\text{unanimous}(L)$ **then** $\text{decided}_i := \text{true}$

End**Broadcast decision**

Note: Collect means to add any received votes to L , and check whether there are $2f + 1$ votes with the same round number.

We will prove the three properties of the protocol that we discussed previously: *unanimity*, *validity*, and *agreement*.

P1: If all inputs propose value v , then any decision must have value v .

We wrote this symbolically as:

$$\forall v : T. (\forall e : \text{Event}. \text{input}(e) = v) \Rightarrow (\forall e : \text{Event}. \text{decide}(e) = v)$$

We prove this by induction on rounds. In round 1 only inputs are broadcast to the replicas. Hence only inputs are collected at step 3. Thus only inputs are counted in the majority. So only inputs are broadcast in round 2. This argument also shows that if only inputs are broadcast in

round r , then only inputs can be decided. If the only input is v , then this is the only possible decision. QED

P2: All decided values are input values.

$$\begin{aligned} \forall e : \text{Event. } E(\text{decide}(e)) \Rightarrow \\ \exists e' : \text{Event. } (e' < e \ \& \ \text{decide}(e) = \text{input}(e')) \end{aligned}$$

We can also prove this by induction on rounds. If the protocol is in round 1, then all collected values are inputs. For any round r , if all collected values are inputs, then only inputs will be broadcast. Hence any unanimous value will be an input. QED

P3. Agreement, any two decisions have the same value.

$$\begin{aligned} \forall e_1, e_2 : \text{Event. } E(\text{decide}(e_1)) \ \& \ E(\text{decide}(e_2)) \Rightarrow \\ \text{decide}(e_1) = \text{decide}(e_2) \end{aligned}$$


Suppose that two processes execute *decide* in the same round r , e.g. set *decided_i* := true and *decided_j* := true. This means that they have collected $2f + 1$ votes and they are unanimous at process i say for a value v_i . (For example, for $f = 1$, 3 votes are collected, and all three are for v_i .) In process j , 3 values are collected, suppose they are all for $v_j \neq v_i$. This is impossible since there are only 4 votes to collect and three are for v_i .

In general, we see that the requirement that a decision is made only when $2f + 1$ votes are the same, disallows two different collections of $2f + 1$ votes since there are only $3f + 1$ possibilities, not $4f + 2$. QED

If they decide in different rounds, say i in r and j in $r' > r$, then we need a more subtle argument. **Optional exercise:** Finish the proof for this case.

2 The Logic of Events

One way to understand asynchronous distributed protocols is to think in terms of *message sequence diagrams* as we did in explaining the 2/3 consensus protocol. Here is a diagram to remind us for the case of one possible failure, $3f + 1$ will have the value 4. So we show the four participants and the inputs they receive.



The *events* are marked as dots on the “time line” of each process. The logic of events postulates properties of the events. At each location, we see that the events are linearly ordered.

There is another important ordering relation less clearly discernible. Lamport called it out and named it as *causal order*. The red event at P_1 , the 5th event on the P_1 time line, was *caused by* the four previous events at P_1 . Using the message sequence diagram, we can trace out the causal order relation.

The logic of events provides axioms to characterize the local order of events at each process as well as the causal order relation. The logic also depends on assumptions about message delivery and process failures.

The logic of events has been fully implemented by the Nuprl proof assistant and used to prove that protocols have certain behaviors under various assumptions about process failure and message delivery. The current list of axioms and theorems can be examined in the Nuprl mathematics library at <http://www.nuprl.org/wip/EventOrdering/new!event-ordering/index.html> . We discuss two key properties of this logic.

Property 1: The events at any location are totally ordered.

Property 2: If only the events at finitely many locations are considered, then causal order is well founded. *This means that we can prove properties of protocols by induction on causal order.*

Here are some of the other “laws” of the logic of events.

Table 1. Signatures in the Logic of Events

Events with Order

E: \mathbb{D}
pred?: $E \rightarrow (E + Loc)$
sender?: $E \rightarrow (E + Unit)$

and with Values

Kind = $Loc \times Act + Lnk \times Tag$
vtyp: $Kind \rightarrow Type$
kind: $e : E \rightarrow Kind$
val: $e : E \rightarrow vtyp(kind(e))$

and with State

typ: $Id \rightarrow Loc \rightarrow Type$
initially: $x : Id \rightarrow i : Loc \rightarrow typ(x, i)$
when: $x : Id \rightarrow e : E \rightarrow typ(x, loc(e))$
after: $x : Id \rightarrow e : E \rightarrow typ(x, loc(e))$

Definitional extensions

loc: $E \rightarrow Loc$
first: $E \rightarrow Bool$
isrcv: $E \rightarrow Bool$
 $x < y, \quad x <_{loc} y$
sender: $\{e : E | isrcv(e)\} \rightarrow E$
link: $\{e : E | isrcv(e)\} \rightarrow Link$
tag: $\{e : E | isrcv(e)\} \rightarrow Tag$
state(i) = x : $Id \rightarrow typ(x, i)$
state-when: $e : E \rightarrow state(loc(e))$
state-after: $e : E \rightarrow state(loc(e))$

Table 2. Axioms of Basic Event Structures

1. On any link, an event sends boundedly many messages; formally: $\forall l : Link. \forall e : E. \exists e' : E. \forall e'' : E. R(e'', e) \Rightarrow e'' < e' \wedge loc(e') = dst(l)$ where $R(e'', e) \equiv isrcv(e'') \wedge sender(e'') = e \wedge link(e'') = l$
2. The predecessor function is one-to-one; formally: $\forall e_1, e_2 : E. pred?(e_1) = pred?(e_2) \Rightarrow e_1 = e_2$
3. Causal order is (strongly) well-founded; formally: $\exists f : E \rightarrow \mathbb{N}. \forall e_1, e_2 : E. e_1 < e_2 \Rightarrow f(e_1) < f(e_2)$
4. The location of the sender of an event is the source of the link on which the message was received; formally: $\forall e : E. isrcv(e) \Rightarrow loc(sender(e)) = src(link(e))$
5. Links deliver messages in FIFO (first in first out) order; formally: $\forall e_1, e_2 : E. link(e_1) = link(e_2) \Rightarrow sender(e_1) < sender(e_2) \Rightarrow e_1 < e_2$
6. State variables change only at events, so that: $\forall e : E. \neg first(e) \Rightarrow (x \text{ when } e) = (x \text{ after } pred(e))$

Message sequence diagram for 2/3 consensus

