

Lecture 9: Critical Sections revisited, and Reasoning about Programs

K. V. S. Prasad
Dept of Computer Science
Chalmers University
Monday 6 and 13 Oct 2014

Plan for today

Chap 2, 3 recap and complete

Chap 4 intro to logic

If time permits, some Linda programming.

REMINDER: I'm away Thu 16 Oct. Maybe guest.

REMINDER: Class rep meeting Thu 9 Oct.

Recap – state diagrams

- (Discrete) computation = states + transitions
 - Both sequential and concurrent
 - Can two frogs move at the same time?
 - We use labelled or unlabelled transitions
 - According to what we are modelling
 - Chess games are recorded by transitions alone (moves)
 - States used occasionally for illustration or as checks
 - In message passing, the (labelled) transitions
 - Are what we see of a (sub)system
 - So they matter more than the states

How to program multiple processes

- Concurrent vs. sequential
 - Concurrent has more states due to interleaving
- But a concurrent sort program should sort
 - No matter which interleaving
 - So cut out unwanted interleavings
 - through synchronisation (waits)

What is interleaved?

Atomic statements

- The thing that happens without interruption
 - Can be implemented as high priority
- We must say what the atomic statements are
 - In the book, assignments and boolean conditions
 - How to implement these as atomic?

Correctness - safety

- A safety property must always hold
 - In every state of every computation
- = “nothing bad ever happens”
 - Typically, partial correctness
 - Program is correct if it terminates
 - E.g., “loop until head, toss”
 - sure to produce a toss if it terminates
 - But not sure it will terminate
 - » Will do so with increasing probability the longer we go on
 - How about “loop until sorted, shuffle deck”?
 - Sure to produce sorted deck if it terminates
 - Needs much longer expected run to terminate

Correctness - Liveness

- A liveness property must eventually hold
 - Every computation has a state where it holds
- = a good thing happens eventually
 - Termination
 - Progress = get from one step to the next
 - Non-starvation of individual process

Safety and Liveness are duals

- Let P be a safety property
 - Then $\text{not } P$ is a liveness property
- Let P be a liveness property
 - Then $\text{not } P$ is a safety property

(Weak) Fairness assumption

- If at any state in the scenario, a statement is continuously enabled, that statement will eventually appear in the scenario.
- So an unfair version of coin tossing cannot guarantee we will eventually see a head.
- We usually assume fairness

What is the critical section problem?

- Specification
 - Both p and q cannot be in their CS at once (mutex)
 - If p and q both wish to enter their CS, one must succeed eventually (no deadlock)
 - If p tries to enter its CS, it will succeed eventually (no starvation)
- GIVEN THAT
 - A process in its CS will leave eventually (progress)
 - Progress in non-CS optional

Different kinds of requirement

- Safety:
 - Nothing bad ever happens on any path
 - Example: mutex
 - In no state are p and q in CS at the same time
 - If state diagram is being generated incrementally, we see more clearly that this says "in every path, mutex"
- Liveness
 - A good thing happens eventually on every path
 - Example: no starvation
 - If p tries to enter its CS, it will succeed eventually
 - Often bound up with fairness
 - We can see a path that starves, but see it is unfair

Deadlock?

- With higher level of process
 - Processes can have a blocked state
 - If all processes are blocked, deadlock
 - So require: no path leads to such a state
- With independent machines (always running)
 - Can have livelock
 - Everyone runs but no one can enter critical section
 - So require: no path leads to such a situation

Language, logic and machines

- Evolution
 - Language fits life – why?
 - What is language?
- What is logic?
 - Special language
- What are machines?
 - Why does logic work with them?
- What kind of logic?

Logic Review

- How to check that our programs are correct?
 - Testing
 - Can show the presence of errors, but never absence
 - Unless we test every path, usually impractical
 - How do you show math theorems?
 - For **every** triangle, ... (wow!)
 - For **every** run
 - Nothing bad ever happens (safety)
 - Something good eventually happens (liveness)

Propositional logic

- Assignment – atomic props mapped to T or F
 - Extended to interpretation of formulae (B.1)
- Satisfiable – f is true in some interpretation
- Valid - f is true in all interpretations
- Logically equal
 - same value for all interpretations
 - $P \rightarrow q$ is equivalent to $(\text{not } p) \text{ or } q$
- Material implication
 - $p \rightarrow q$ is true if p is false

Proof methods

- State diagram
 - Large scale: "model checking"
 - A logical formula is true of a set of states
- Deductive proofs
 - Including inductive proofs
 - Mixture of English and formulae
 - Like most mathematics
 - But can be formalised
 - Theorem provers
 - Proof checkers

Algorithm 3.1

- We can prove mutex by checking all the states
- How to prove absence of deadlock?
 - Do we have to look at all scenarios?
 - Would be even worse; all paths through the state diagram
 - Fortunately, only paths from states (p_2, q_2, i)
 - Then we can argue from the text of the program
 - Don't need state diagram
- Starvation?
 - Sadly, a loop in the state diagram can starve a process

Algorithms 3.2 – 3.4

- 3.2 : No mutex
 - easy to see scenario that leads to this
 - just do both processes in step
- 3.3 : Deadlocks
 - Again, by doing both processes in step
- 3.4 : Both can starve
 - Again, by doing both processes in step

Dekker's algorithm (3.10)

- Each process in turn has the right to insist on entering its critical section
- Perhaps surprisingly, this does the trick!
- Can prove by state diagram
 - but will do by logic
- Dekker, Peterson, etc.
 - These algorithms are now only nice examples
- Actual mutex etc. achieved by hardware
 - instructions such as test-and-set, compare-and-swap.

Complex atomic hardware instructions

- Show correctness of 3.11 and 3.12
- Test-and-set(common, local) is
 local := common
 common := 1

Exchange(a, b) is
 local temp
 temp := a
 a := b
 b := temp

Atomic Propositions (true in a state)

- *wantp* is true in a state
 - iff (boolean) var *wantp* has value true
- *p4* is true iff the program counter is at *p4*
 - *p4* is the command about to be executed
 - Then *pj* is false for all $j \neq 4$
- *turn=2* is true iff integer var *turn* has value 2
- *not (p4 and q4)* in alg 4.1, slide 4.1
 - Should be true in all states to ensure mutex

Mutex for Alg 4.1

- Invariant Inv1: $(p3 \text{ or } p4 \text{ or } p5) \rightarrow \text{wantp}$
 - Base: $p1$, so antecedent is false, so Inv1 holds.
 - Step: Process q changes neither wantp nor Inv1.
 - Neither $p1$ nor $p3$ nor $p4$ change Inv1.
 - $p2$ makes both $p3$ and wantp true.
 - $p5$ makes antecedent false, so keeps Inv1.

So by induction, Inv1 is always true.

Mutex for Alg 4.1 (contd.)

- Invariant Inv2: $wantp \rightarrow (p3 \text{ or } p4 \text{ or } p5)$
 - Base: $wantp$ is initialised to false, so Inv2 holds.
 - Step: Process q changes neither $wantp$ nor Inv1.
Neither $p1$ nor $p3$ nor $p4$ change Inv1.
 $p2$ makes both $p3$ and $wantp$ true.
 $p5$ makes antecedent false, so keeps Inv1.

So by induction, Inv2 is always true.

Inv2 is the converse of Inv1.

Combining the two, we have

Inv3: $wantp \leftrightarrow (p3 \text{ or } p4 \text{ or } p5)$ and
 $wantq \leftrightarrow (q3 \text{ or } q4 \text{ or } q5)$

Mutex for Alg 4.1 (concluded)

- Invariant Inv4: not (p4 and q4)
 - Base: p4 and q4 is false at the start.
 - Step: Only p3 or q3 can change Inv4.
 - p3 is "await (not wantq)". But at q4, wantq is true by Inv3, so p3 cannot execute at q4.
 - Similarly for q3.

So we have mutex for Alg 4.1

Proof of Dekker's Algorithm (outline)

- Invariant Inv2: $(\text{turn} = 1) \text{ or } (\text{turn} = 2)$
- Invariant Inv3: $\text{wantp} \leftrightarrow p3..5 \text{ or } p8..10$
- Invariant Inv4: $\text{wantq} \leftrightarrow q3..5 \text{ or } q8..10$
- Mutex follows as for Algorithm 4.1
 - NB: "turn" alone won't prove it.
- Will show neither p nor q starves
 - Effectively shows absence of livelock

Liveness via Progress

- Invariants can prove safety properties
 - Something good is always true
 - Something bad is always false
- But invariants cannot state liveness
 - Something good happens eventually
- Progress A to B
 - if we are in state A, we will progress to state B.
- Weak fairness assumed
 - to rule out trivial starvation because process never scheduled.
 - A scenario is weakly fair if
 - B is continually enabled at state A in scenario ->
B will eventually appear in the scenario

Box and Diamond

- A request is eventually granted
 - For all t . $\text{req}(t) \rightarrow \text{exists } t'. (t' \geq t) \text{ and } \text{grant}(t')$
 - New operators indicate time relationship implicitly
 - $\text{box} (\text{req} \rightarrow \text{diam grant})$
- If "successor state" is reflexive,
 - $\text{box } A \rightarrow A$ (if it holds indefinitely, it holds now)
 - $A \rightarrow \text{diam } A$ (if it holds now, it holds eventually)
- If "successor state" is transitive,
 - $\text{box } A \rightarrow \text{box box } A$
 - if not transitive, A might hold in the next state, but not beyond
 - $\text{diam diam } A \rightarrow \text{diam } A$
- See Wikipedia page on LTL

Formalising Fairness

- Absolute Fairness: every process should be executed infinitely often:
 - **for all** i : $GF\ ex_i$
 - But a process might not be enabled.
- Strong Fairness: a process that is infinitely often enabled executes infinitely often when enabled:
 - **for all** i : $(GF\ en_i) \Rightarrow (GF(en_i \wedge ex_i))$:
- Weak Fairness: a process that is ultimately always enabled should execute infinitely often:
 - **for all** i : $(FG\ en_i) \Rightarrow (GF\ ex_i)$

Progress in (non-)critical section

- The following are notes re Ben-Ari's proof, but I prefer the liveness proof in the Utwente notes.
- Progress in critical section
 - box (p8 \rightarrow diam p9)
 - It is always true that if we are at p8, we will eventually progress to p9
- Non-progress in non-critical section
 - diam (box p1)
 - It is possible that we will stay at p1 indefinitely

Progress through control statements

- For "p1: if A then s" to progress to s, need
 - p1 and box A
 - p1 and A is not enough
 - does not guarantee A holds by the time p1 is scheduled
- So in Dekker's algorithm
 - p4 and box (turn = 2) -> diam p5
 - But turn = 2 is not true forever!
 - It doesn't have to be. Only as long as p4.

Lemma 4.11

- box want_p and $\text{box}(\text{turn} = 1) \rightarrow$
 $\text{diam box}(\text{not want}_q)$
 - If it is p 's turn, and it wants to enter its CS, q
will eventually defer
- Note that at q_1 , want_q is always false
 - Both at init and on looping
- q will progress through $q_2..q_5$ and wait at q_6
 - Inv4: $\text{want}_q \leftrightarrow q_3..5$ or $q_8..10$
 - Implies $\text{box}(\text{not want}_q)$ at q
- Lemma follows

Progress to CS in Dekker's algorithm

- Suppose p_2 and box ($\text{turn}=2$)
 - If p_3 and not want_q then $\text{diam } p_8$
 - p_2 and box ($\text{turn}=2$ and want_q) \rightarrow
 $\text{diam } \text{box } p_6 \leftrightarrow \text{diam } \text{box}$ (not want_p)
 - p_6 and box ($\text{turn}=2$ and not want_p) \rightarrow $\text{diam } q_9$
 - p_2 and box ($\text{turn}=2$) \rightarrow $\text{diam } \text{box}$ (p_6 and $\text{turn}=1$)
 - Lemma 4.11 now yields $\text{diam } p_8$