

# Lecture Notes on Bounded Model Checking and Symbolic Execution

Matt Fredrikson

Carnegie Mellon University  
Lecture 5

## 1 Introduction & Recap

In the previous lecture we looked into proving that programs satisfy safety properties given as formulas in the first-order dynamic logic. In particular, we can write contract properties with precondition  $P$  and postconditions  $Q$  for a program  $\alpha$  as:

$$P \rightarrow [\alpha]Q \tag{1}$$

If this formula is valid, then it means that in every state  $\omega$ , if  $\omega \models P$  then after all terminating runs of  $\alpha$  starting in  $\omega$  the final state  $\nu \models Q$ .

We then defined the semantics of dynamic logic formulas, so that we can actually prove the validity of such formulas. However, this is not always such an easy thing to do because there are an infinite number of initial states to reason about. To address this we derived a set of axioms that can be used in sequent calculus proofs to reason about the validity of DL formulas.

$$([:=]) \quad [x := e]p(x) \leftrightarrow p(e)$$

$$([\text{assert}]) \quad [\text{assert}(Q)]P \leftrightarrow (Q \wedge P)$$

$$([\text{if}]) \quad [\text{if}(Q) \alpha \text{ else } \beta]P \leftrightarrow (Q \rightarrow [\alpha]P) \wedge (\neg Q \rightarrow [\beta]P)$$

$$([\text{;}]) \quad [\alpha; \beta]P \leftrightarrow [\alpha][\beta]P$$

$$([\text{unwind}]) \quad [\text{while}(Q) \alpha]P \leftrightarrow [\text{if}(Q) \{ \alpha; \text{while}(Q) \alpha \}]P$$

$$([\text{unfold}]) \quad [\text{while}(Q) \alpha]P \leftrightarrow (Q \rightarrow [\alpha][\text{while}(Q) \alpha]P) \wedge (\neg Q \rightarrow P)$$

Each of these axioms reduces reasoning about formulas involving the box modality to reasoning about a series of simpler formulas, sometimes not involving modalities at all. Moreover, the axioms can be “implemented” as a syntactic transformation on formulas, and so automated by tools. When a sequence of syntactic transformations results in formulas that don’t contain any box modalities, then the original DL validity question becomes one of validity of arithmetic formulas. It again becomes possible to automate by invoking a decision procedure for first-order arithmetic, and many good ones exist [?].

This is all very good, but not entirely true in all cases. In particular the axioms [unwind] and [unfold] for dealing with loops do not actually simplify matters at all. Applying either of these axioms results in an equivalent DL formula that reasons about the safety of the first iteration of the loop, and then a DL formula that is a carbon copy of the original for the remaining iterations.

The way to deal with this properly is to reason about loop invariants, introducing an axiom that requires us to prove the correctness of the invariant and that it implies the postcondition. But loop invariants cannot in general be derived automatically, and we are interested in techniques that can be automated and implemented in tools.

Another way to deal with this is to apply [unwind] or [unfold] repeatedly for a while, and simply cut the proof off at a certain point. This is equivalent to unrolling the loop some bounded number of times, and assuming that it will terminate before hitting that bound. If we can prove safety for all iterations up to the bound, then we know that any safety violation must occur on longer executions. This is not an ideal solution because we know nothing about the program’s behavior past the unwinding bound, but it will certainly give us more assurance than we would have gotten by running a few test cases.

## 2 Bounded Model Checking

The principle behind bounded model checking is straightforward. First, pick a bound  $N$  on the execution depth of the program. This bound can refer to the total number of commands that are executed, or to the number of times loops are unrolled. We will adopt the latter convention, as it frees us from the potential awkwardness of splitting compound commands like conditionals and loops partway through their bodies.

Having fixed an upper bound on the execution depth, proceed to check the safety property  $P \rightarrow [\alpha]Q$  by repeatedly applying [;], [:=], [assert], and [if], as well as the axioms from the propositional sequent calculus when necessary. Whenever the proof reaches a point at which no further progress can be made because all box modalities contain outermost  $\text{while}(Q)\alpha$  commands, then apply [unfold] and repeat the above process on the loop body  $\alpha$ . When [unfold] has been applied to each loop  $N$  times, then replace each occurrence of  $[\text{while}(Q)\alpha]P$  with  $P$ . Then proceed to close out the proof by reducing any remaining obligations to arithmetic formulas and applying the rule  $\mathbb{Z}$ .

The best way to understand how this works is to see it in action. Let’s start off simple,

supposing that we wish to verify the following up to a bound of  $N = 1$ :

$$x \neq 0 \rightarrow [z := 0; \text{while}(y > 0) \{ \text{if}(y \% 2 = 1) \{ z := z + x \} x := 2 * x; y := y / 2 \}] z \neq 0$$

To keep things easier to read, we will let  $\alpha$  denote the body of the while loop. We proceed as follows.

$$\begin{array}{c} \frac{x \neq 0, z = 0 \vdash [\text{while}(y > 0) \alpha] z \neq 0}{\text{[:=]} = \frac{x \neq 0 \vdash [z := 0][\text{while}(y > 0) \alpha] z \neq 0}{\text{[;]} \frac{x \neq 0 \vdash [z := 0; \text{while}(y > 0) \alpha] z \neq 0}{\rightarrow R \vdash x \neq 0 \rightarrow [z := 0; \text{while}(y > 0) \alpha] z \neq 0}} \end{array}$$

There are two things to point out here.

**New assignment axiom.** First, we have used a new rule that hasn't been mentioned before,  $[\text{:=}] =$ .

$$([\text{:=}] =) \frac{\Gamma, y = e \vdash p(y), \Delta}{\Gamma \vdash [x := e] p(x), \Delta} \text{ (} y \text{ new)}$$

The  $[\text{:=}] =$  rule says that we can move an assignment on the right of a sequent to an equality on the left. However, just as we needed to be careful about variable capture using the basic assignment axiom  $[\text{:=}]$ , we must also be careful here. The way to ensure that problems don't arise because of conflicting assumptions or proof goals is to rename the variable to something completely new, and rewrite all unbound instances of the old variable appearing in the formula to the right of the box modality to the new fresh variable name. In the partial proof above, we optimized slightly because  $z$  does not occur anywhere in the assumptions, and all its instances in the rest of the formula are bound by the assignment itself.

Consider what would have happened if such conditions failed to hold, and we weren't careful about renaming and used the rule incorrectly.

$$[\text{:=}] = \frac{x = 0, x = 1 \vdash x = 0}{x = 0 \vdash [x := 1] x = 0}$$

We end up with a contradiction  $x = 0 \wedge x = 1$  in our assumptions, from which we can prove anything. This is obviously not sound, because we should not be able to prove that  $x = 0$  immediately after assigning  $x := 1$ ! The correct application of  $[\text{:=}] =$  in this case would be as follows.

$$[\text{:=}] = \frac{x = 0, y = 1 \vdash y = 0}{x = 0 \vdash [x := 1] x = 0}$$

Notice that we replaced  $x$  with the fresh variable  $y$  before moving it to the left of the sequent, and replaced the remaining unbound instance of  $x$  after the box with  $y$ .

**Unfolding bugs.** The second thing to notice is that we can proceed no further in the proof without applying [unfold]. So we proceed to unfold the loop, and apply non-loop axioms afterwards until we get stuck again.

$$\begin{array}{c} \frac{x \neq 0, z = 0, y > 0 \vdash [\alpha][\mathbf{while}(y > 0) \alpha]z \neq 0 \quad x \neq 0, z = 0, y \leq 0 \vdash z \neq 0}{\frac{\rightarrow R}{x \neq 0, z = 0 \vdash y > 0 \rightarrow [\alpha][\mathbf{while}(y > 0) \alpha]z \neq 0} \quad \frac{\rightarrow R}{x \neq 0, z = 0 \vdash y \leq 0 \rightarrow z \neq 0}}{\wedge R} \\ \frac{\wedge R}{x \neq 0, z = 0 \vdash (y > 0 \rightarrow [\alpha][\mathbf{while}(y > 0) \alpha]z \neq 0) \wedge (y \leq 0 \rightarrow z \neq 0)} \\ \text{[unfold]} \frac{}{x \neq 0, z = 0 \vdash [\mathbf{while}(y > 0) \alpha]z \neq 0} \end{array}$$

At this point we can't help but notice that the branch of our proof with  $y \leq 0$  in the assumptions has no path forward. The formula:

$$x \neq 0 \wedge z = 0 \wedge y \leq 0 \rightarrow z \neq 0 \quad (2)$$

simply is not a valid formula of arithmetic. So, we've found a bug. What's more, examining the sequent that we are unable to prove:

$$x \neq 0, z = 0, y \leq 0 \vdash z \neq 0 \quad (3)$$

we can extract more useful information about the bug. In particular the context provided by our assumptions tells us exactly what conditions of the initial state need to hold in order for the program to produce a trace that violates the safety property. So if we take any values of  $x, y, z$  that satisfy the assumptions  $x \neq 0 \wedge z = 0 \wedge y \leq 0$  then we are guaranteed to "exercise" the bug. Such a set of inputs and its corresponding trace is called a *counterexample* to the safety property, and is a useful artifact of model checkers when debugging programs in practice.

**Back to basic axioms.** In hindsight, perhaps this is not so impressive because one of our assumptions is  $z = 0$ . We probably should have seen this coming, because the program begins by initializing  $z$  in this way, and will only update it when the loop body is executed. Perhaps there are more bugs to uncover if we continue with the other branch of the loop. In the following, let  $\beta$  denote the program  $x := 2 * x; y := y/2$ .

$$\frac{\textcircled{1} \quad \textcircled{2}}{\frac{\text{[if]} \quad x \neq 0, z = 0, y > 0 \vdash [\mathbf{if}(y \% 2 = 1) \{z := z + x\}][\beta][\mathbf{while}(y > 0) \alpha]z \neq 0}{\text{[;]} \quad x \neq 0, z = 0, y > 0 \vdash [\mathbf{if}(y \% 2 = 1) \{z := z + x\}; \beta][\mathbf{while}(y > 0) \alpha]z \neq 0}}$$

The branch of the proof marked  $\textcircled{1}$  continues below.

$$\begin{array}{c} \frac{x \neq 0, z = 0, y > 0, y \% 2 = 1, z_1 = z + x, x_1 = 2 * x, y_1 = y/2 \vdash [\mathbf{while}(y_1 > 0) \alpha]z_1 \neq 0}{\text{[:=]}=} \\ \frac{\text{[:=]}=} \frac{x \neq 0, z = 0, y > 0, y \% 2 = 1, z_1 = z + x, x_1 = 2 * x \vdash [y := y/2][\mathbf{while}(y > 0) \alpha]z_1 \neq 0}{x \neq 0, z = 0, y > 0, y \% 2 = 1, z_1 = z + x \vdash [x := 2 * x][y := y/2][\mathbf{while}(y > 0) \alpha]z_1 \neq 0} \\ \frac{\text{[:=]}=} \frac{x \neq 0, z = 0, y > 0, y \% 2 = 1 \vdash [z := z + x][x := 2 * x][y := y/2][\mathbf{while}(y > 0) \alpha]z \neq 0}{\text{[;]}=} \\ \frac{\text{[;]}=} \frac{x \neq 0, z = 0, y > 0, y \% 2 = 1 \vdash [z := z + x][\beta][\mathbf{while}(y > 0) \alpha]z \neq 0}{\rightarrow R} \\ \frac{}{x \neq 0, z = 0, y > 0 \vdash y \% 2 = 1 \rightarrow [z := z + x][\beta][\mathbf{while}(y > 0) \alpha]z \neq 0} \end{array}$$

And the branch of the proof marked ② continues here.

$$\begin{array}{l}
 \frac{x \neq 0, z = 0, y > 0, y \% 2 \neq 1, x_1 = 2 * x, y_1 = y / 2 \vdash [\text{while}(y_1 > 0) \alpha] z \neq 0}{[:=]=} \\
 \frac{x \neq 0, z = 0, y > 0, y \% 2 \neq 1, x_1 = 2 * x \vdash [y := y / 2][\text{while}(y > 0) \alpha] z \neq 0}{[:=]=} \\
 \frac{x \neq 0, z = 0, y > 0, y \% 2 \neq 1 \vdash [x := 2 * x][y := y / 2][\text{while}(y > 0) \alpha] z \neq 0}{[;]} \\
 \frac{[;]}{\rightarrow R} \frac{x \neq 0, z = 0, y > 0, y \% 2 \neq 1 \vdash [x := 2 * x; y := y / 2][\text{while}(y > 0) \alpha] z \neq 0}{x \neq 0, z = 0, y > 0 \vdash y \% 2 \neq 1 \rightarrow [x := 2 * x; y := y / 2][\text{while}(y > 0) \alpha] z \neq 0}
 \end{array}$$

Now in both branches of the proof, we can't go any further without applying [\[unfold\]](#). We initially set our bound to  $N = 1$ , and we've unrolled the loop exactly one time.

**Verification conditions.** So we proceed to replace the formula  $[\text{while}(y > 0) \alpha] z \neq 0$  with  $z \neq 0$  on both branches, yielding the sequents:

$$x \neq 0, z = 0, y > 0, y \% 2 = 1, z_1 = z + x, x_1 = 2 * x, y_1 = y / 2 \vdash z_1 \neq 0 \quad (4)$$

$$x \neq 0, z = 0, y > 0, y \% 2 \neq 1, x_1 = 2 * x, y_1 = y / 2 \vdash z \neq 0 \quad (5)$$

The sequents shown in 4 and 5 are our remaining proof obligations: if they are valid, then we can conclude that the safety property holds on the program traces corresponding to the paths that generated these obligations.

- In the case of 4, the corresponding path enters the while loop (reflected by the assumption  $y > 0$  in the sequent), enters the body of the conditional (reflected by the assumption  $y \% 2 = 1$ ), and executes the remainder of the loop body stopping just before iterating again.
- In the case of 5, the corresponding path enters the while loop ( $y > 0$  is still in the assumptions), skips over the body of the conditional (reflected by  $y \% 2 \neq 1$ ), and executes the rest of the loop body stopping prior to another iteration.

Observe that in both cases the proof obligations involve nothing but arithmetic. Recalling the meaning of sequents, we can derive arithmetic formulas whose validity implies the correctness of these paths.

$$x \neq 0 \wedge z = 0 \wedge y > 0 \wedge y \% 2 = 1 \wedge z_1 = z + x \wedge x_1 = 2 * x \wedge y_1 = y / 2 \rightarrow z_1 \neq 0 \quad (6)$$

$$x \neq 0 \wedge z = 0 \wedge y > 0 \wedge y \% 2 \neq 1 \wedge x_1 = 2 * x \wedge y_1 = y / 2 \rightarrow z \neq 0 \quad (7)$$

Equations 6 and 7 are called *verification conditions*. Recall from earlier Equation 2, the formula whose invalidity told us that the program contains a bug whenever  $x \neq 0 \wedge z = 0 \wedge y \leq 0$ . This was also a verification condition, corresponding to the program path where the body of the while loop is skipped over immediately leading to termination.

The primary job of a bounded model checker is to generate verification conditions for each program path within the execution depth bound. This can be done fully automatically, because there is nothing particularly difficult about applying the axioms [\[;\]](#), [\[:=\]=](#), [\[assert\]](#), and [\[if\]](#). As each verification condition is derived, the bounded model

checker consults an automated decision procedure for arithmetic. This often involves exploiting the duality between satisfiability and validity covered in the second lecture, as most decision procedures are designed to answer satisfiability queries rather than validity. But this is not a practical hurdle, as it merely involves negating the verification condition.

**Other first-order theories.** So far in this course we have assumed that our programs operate over “real” unbounded integers. This means that the verification conditions that we generate are formulas in the first-order theory of integer arithmetic. This theory is defined by the interpretation given to the constants (0,1,2,..), functions (+, −, ×, ...), and predicates ( $\leq$ , =, ...). We expect that claims of validity for these formulas assume the usual interpretation for such entities that we are familiar with from arithmetic over the integers.

Programs written in languages like C do not operate over such integers. Rather, they operate over machine integers that have bounded with (i.e., 32 or 64 bits), and thus can only take values from a finite set that can be represented as binary numbers of the according width. Likewise, functions like addition and subtraction have a different interpretation over the machine integers, which for example manifest when the result of an operation is too large or too small to be represented by the width of the architecture (i.e., overflow and underflow). Machine integers have additional functions such as bitwise & and |, and shift operators >> and <<.

Just as one can define a first-order theory of “real” integers, it is possible to define a first-order theory of machine integers, or perhaps floating-point decimal numbers, by assigning the appropriate interpretations to the constants, functions, and predicates pertinent to machine integers. Decision procedure developers of course realize this, and have indeed built support for such theories in widely-used tools. Bounded model checkers exploit this to faithfully model the semantics of machine arithmetic for languages like C by generating verification conditions for the first-order theory of machine integers. Importantly, *doing so does not generally entail changes to the verification condition generator itself*, as this is a purely syntactic analysis that does not depend on the interpretation of the underlying term constructors!

Let’s look at a quick example involving machine integer arithmetic to get a sense of the differences that might arise when reasoning about correctness and safety. Consider the following program, which diligently checks that the denominator is non-zero before using it.

```

1 if(a>0 ∨ b>0) {
2   assert(a+b ≠ 0);
3   x := c/(a+b);
4 }
```

We can check that the safety property corresponding to the assertion holds by reasoning about the validity of  $[\text{if}(a > 0 \vee b > 0) \{ \text{assert}(a + b \neq 0); x := c/(a + b) \}] \top$ ,

as follows.

$$\frac{\frac{\frac{\frac{\vdash a > 0 \vee b > 0 \rightarrow a + b \neq 0}{[\text{assert}]} \vdash a > 0 \vee b > 0 \rightarrow [\text{assert}(a + b \neq 0)]\top}{[:=]} \vdash a > 0 \vee b > 0 \rightarrow [\text{assert}(a + b \neq 0)][x := c/(a + b)]\top}{[;]} \vdash a > 0 \vee b > 0 \rightarrow [\text{assert}(a + b \neq 0); x := c/(a + b)]\top}{[\text{if}]} \vdash [\text{if}(a > 0 \vee b > 0)] \{ \text{assert}(a + b \neq 0); x := c/(a + b) \} \top$$

The last step follows because  $a + b \neq 0 \wedge \top$  is equivalent to  $a + b \neq 0$ . Now in the theory of integer arithmetic  $\mathbb{Z}$ , our verification condition

$$a > 0 \vee b > 0 \rightarrow a + b \neq 0 \tag{8}$$

is perfectly valid, and we could close out the proof by simply applying the rule  $\mathbb{Z}$ . But in the theory of 32-bit machine integer arithmetic, if  $a$  and  $b$  are unsigned then Eq. 8 is not valid. Consider the counterexample  $a = 2^{32} - 1 = 0xFFFFFFFF = 4294967295$  and  $b = 1$ . Then  $a + b = 0$  because the result  $4294967296 = 2^{32}$  is too large to fit in a 32-bit unsigned integer representation, so the result of the addition wraps around to zero.

From now on, when we want to use the theory of machine integer arithmetic in our proofs, we will use the rule  $\mathbb{Z}_M$ .

**Unwinding assertions.** From what we've seen so far, bounded model checking gives us a certain limited kind of assurance about the safety of a program. Let's break it down into cases.

1. If the bounded model checker finds a bug, it can report a counterexample. If the verification condition generator was implemented correctly and the correct theory was used by the decision procedure to discharge the proof obligation leading to the bug, then we can be sure that there is actually a bug in our program.

How do we know this? Recall the axioms used to generate the verification condition:  $[;]$ ,  $[:=]$ ,  $[\text{assert}]$ ,  $[\text{if}]$ , and  $[\text{unfold}]$ . Each of these is an equivalence reducing one formula to another, so that if the resulting verification condition is valid, then the original formula is as well. Likewise, if the VC is *not* valid, then the original formula was not either. If this isn't convincing enough, then most decision procedures will produce a counterexample to the VC, that we can construct an input to the program with and actually run to observe the bug.

2. If the bounded model checker does *not* find a bug, then all that we can say with confidence is that there are no bugs on paths up to the execution depth bound. At least we can say this, again due to the fact that we have proved the axioms of dynamic logic to be valid equivalences. But importantly, we can't say that there are *no* safety bugs in the program when the model checker fails to find one, as there could be a bug on some path past the depth bound. Likewise, it could just as well be the case that there are indeed no such bugs on longer paths, but we shouldn't take this view without solid evidence to back it up.

The second case is unfortunate. If the model checker finds no bugs then we will probably want to follow up to see if we can convince ourselves that there are no safety violations, but how?

One simple approach that sometimes works in practice is to use an *unwinding assertion*. As the name suggests, an unwinding assertion is an assertion command that is added to the program as the bounded model checker applies the `[unfold]` axiom. Whenever the depth bound is reached, rather than replacing `[while(Q) α]P` with `P`, `[while(Q) α]P` is replaced with `[assert(¬Q)]P`. This way, verification will only succeed if the program would have terminated anyway after the bound was reached because `¬Q` is true at that point. Let's look at a short example to illustrate the idea. We will set a depth bound of  $N = 2$  in the following proof.

$$\begin{array}{c}
 \vdots \\
 \hline
 \text{[:=]} \frac{x = 2, 0 < x, x_1 = x - 1 \vdash [\text{while}(0 < x_1) x_1 := x_1 - 1]x_1 = 0}{x = 2, 0 < x \vdash [x := x - 1][\text{while}(0 < x) x := x - 1]x = 0} \quad * \\
 \text{→R} \frac{x = 2 \vdash 0 \leq x \rightarrow [x := x - 1][\text{while}(0 < x) x := x - 1]x = 0}{x = 2 \vdash (0 \leq x \rightarrow x := x - 1; [\text{while}(0 < x) x := x - 1]x = 0) \wedge (0 \geq x \rightarrow x = 0)} \quad \mathbb{Z}_M \\
 \text{∧R} \frac{x = 2 \vdash (0 \leq x \rightarrow x := x - 1; [\text{while}(0 < x) x := x - 1]x = 0) \wedge (0 \geq x \rightarrow x = 0)}{x = 2 \vdash [\text{while}(0 < x) x := x - 1]x = 0} \\
 \text{[unfold]}
 \end{array}$$

We continue with the proof below, as we have run out of space. Let  $P$  denote our assumptions so far  $x = 2, 0 < x, x_1 = x - 1$ . This time when we apply `[unfold]`, we will hit the execution bound immediately, and at that point insert the unwinding assertion.

$$\begin{array}{c}
 \text{[assert]} \frac{P, 0 < x_1, x_2 = x_1 - 1 \vdash 0 \geq x_2 \wedge x_2 = 0}{P, 0 < x_1, x_2 = x_1 - 1 \vdash [\text{assert}(0 \geq x_2)]x_2 = 0} \\
 \text{[:=]} \frac{P, 0 < x_1 \vdash [x_1 := x_1 - 1][\text{assert}(0 \geq x_1)]x_1 = 0}{P \vdash 0 < x_1 \rightarrow [x_1 := x_1 - 1][\text{assert}(0 \geq x_1)]x_1 = 0} \quad * \\
 \text{→R} \frac{P \vdash 0 < x_1 \rightarrow [x_1 := x_1 - 1][\text{assert}(0 \geq x_1)]x_1 = 0}{P \vdash (0 < x \rightarrow [x_1 := x_1 - 1][\text{assert}(0 \geq x_1)]x_1 = 0) \wedge (0 \geq x_1 \rightarrow x_1 = 0)} \quad \mathbb{Z}_M \\
 \text{∧R} \frac{P \vdash (0 < x \rightarrow [x_1 := x_1 - 1][\text{assert}(0 \geq x_1)]x_1 = 0) \wedge (0 \geq x_1 \rightarrow x_1 = 0)}{P \vdash [\text{while}(0 < x_1) x_1 := x_1 - 1]x_1 = 0} \\
 \text{[unfold]}
 \end{array}$$

Now we have reduced the problem to machine arithmetic, leading to the verification condition

$$x = 2 \wedge 0 < x \wedge x_1 = x - 1 \wedge 0 < x_1 \wedge x_2 = x_1 - 1 \rightarrow 0 \geq x_2 \wedge x_2 = 0 \quad (9)$$

This is of course valid, and because of the unwinding condition we know that there are no paths in the program that exceed the execution depth  $N = 2$ . From this we conclude that the program satisfies the safety property on *all* traces.

### 3 Symbolic Execution

Recall that when we discussed verification conditions, we saw that the assumptions in the context of each proof obligation reflect the path covered by that proof obligation. Also worth noting is that because we used `[:=]=` rather than `[:=]`, moving assignments



to the assumptions and renaming variables each time, the context also tracks the intermediate state rather explicitly. Consider for example the following derivations.

$$\frac{x = a, y = b \vdash y = b \wedge x = a}{[:=] x = a, y = b \vdash [z := x]y = b \wedge z = a}$$

$$\frac{[:=] x = a, y = b \vdash [z := x][x := y]x = b \wedge z = a}{[:=] x = a, y = b \vdash [z := x][x := y][y := z]x = b \wedge y = a}$$

$$\frac{[s];[s] x = a, y = b \vdash [z := x; x := y; y := z]x = b \wedge y = a}{[s];[s] x = a, y = b \vdash [z := x; x := y; y := z]x = b \wedge y = a}$$

In this case, the verification condition is simply

$$x = a \wedge y = b \rightarrow y = b \wedge x = a \quad (10)$$

All of the information about the intermediate states that the program entered to achieve its final result is gone from the condition. On the other hand, using  $[:=]_=$  to do a similar derivation:

$$\frac{x = a, y = b, z_1 = x, x_1 = y, y_1 = z_1 \vdash x_1 = b \wedge y_1 = a}{[:=]_= x = a, y = b, z_1 = x, x_1 = y \vdash [y := z_1]x_1 = b \wedge y = a}$$

$$\frac{[:=]_= x = a, y = b, z_1 = x \vdash [x := y][y := z_1]x = b \wedge y = a}{[:=]_= x = a, y = b \vdash [z := x][x := y][y := z]x = b \wedge y = a}$$

$$\frac{[s];[s] x = a, y = b \vdash [z := x; x := y; y := z]x = b \wedge y = a}{[s];[s] x = a, y = b \vdash [z := x; x := y; y := z]x = b \wedge y = a}$$

Now the verification condition is

$$x = a \wedge y = b \wedge z_1 = x \wedge x_1 = y \wedge y_1 = z_1 \rightarrow x_1 = b \wedge y_1 = a \quad (11)$$

In terms of reasoning about correctness, there is no difference here. Equations 10 and 11 are equivalent. But the context tells us that the most recent “version” of  $y$  (i.e.  $y_1$ ) was updated to take the most recent version of  $z$  (i.e.  $z_1$ ), which was in turn updated to take the initial version of  $x$ .

**Path formulas.** The conjunction of the assumptions calculated in this way is called the *path formula* for the corresponding program path behind this derivation. Any feasible path through a program has a corresponding path formula that is satisfiable. *Symbolic execution* is a technique that enumerates program paths that may contain safety violations, generates their path formulas, and checks each formula for satisfiability. Whenever a path formula is satisfiable, it means that there is at least one trace that follows that path.


To facilitate enumerating path formulas, symbolic execution first constructs a *control flow graph* of the program that reflects all of the paths in the program. Consider the following program as an example.

```

1 if(a > 0) { x := -2; }
2 if(b < 5) {
3   if(a = 0 ∧ c ≠ 0) { y := 1 }
4   z := 2;
5 }
6 assert(x+y+z ≠ 3);

```

The corresponding control flow graph is as follows.



The edges of the control flow graph are labeled to reflect the corresponding program command on that portion of the path. As shorthand, we write  $?P$  to represent conditions that must hold, i.e. in place of  $\text{assert}(P)$ . The node marked in red denotes a path that violates the assertion. Edges labeled with  $\epsilon$  are noops, and don't correspond to any program command. Then the task of symbolic execution is to determine if this node is reachable from an initial state from a feasible path.

Each path through the control flow graph corresponds to a verification condition, which we obtain by listing out a corresponding program for that path and then applying the relevant axioms of dynamic logic. For example, the path that first takes the right edge from the initial state, and then left edges until reaching the red node would be:

$$\vdash [?a > 0; x := -1; ?b < 5; ?a = 0 \wedge c \neq 0; y := 1; z := 2; ?x + y + z = 3] \top \quad (12)$$

Notice that we use the trivial postcondition  $\top$  because we are merely interested in whether there are *any* traces that could follow this path. So we don't care what properties the final state may have, which is described by  $\top$ . Applying axioms  $[\text{:=}]$  and  $[\text{assert}]$ , we derive the following verification condition.

$$a > 0 \wedge x_1 = -1 \wedge b < 5 \wedge a = 0 \wedge c \neq 0 \wedge y_1 = 1 \wedge z_1 = 2 \wedge x_1 + y_1 + z_1 = 3 \quad (13)$$

This formula is not satisfiable, because  $a > 0 \wedge a = 0$  is a contradiction, so there are no feasible traces that follow this path. On the other hand, if we had taken the left branch off of the initial state and then followed the same commands afterwards, it is not hard to check that we would obtain the following verification condition.

$$a \leq 0 \wedge b < 5 \wedge a = 0 \wedge c \neq 0 \wedge y_1 = 1 \wedge z_1 = 2 \wedge x + y_1 + z_1 = 3 \quad (14)$$

Equation 14 is indeed satisfiable, as evidenced by the witness  $x = 0, y_1 = 1, z_1 = 2$ . This means that the path is feasible, which we could show by running the program on any input with  $x = 0$ .

**Checking invariants.** Path formulas contain all of the information that we need to reason about the satisfaction of formulas over state at all points in the execution. For example, suppose we wish to check the invariant  $P$ . Part way through the derivation of

$$x = a, y = b \vdash [z := x; x := y; y := z]x = b \wedge y = a$$

we derived the sequent  $x = a, y = b, z_1 = x \vdash [x := y][y := z]x = b \wedge y = a$ . The context  $x = a, y = b, z_1 = x$  reflects the state after the first assignment  $z := x$ . We need to show that at this point in the execution,  $P$  holds, i.e., we can prove validity of:

$$x = a, y = b, z_1 = x \vdash P$$

Executing the next assignment  $x := y$  lead to the context  $x = a, y = b, z_1 = x, x_1 = y$  and the obligation to prove:

$$x = a, y = b, z_1 = x, x_1 = y \vdash P$$

And finally, executing the last assignment  $y := z$  gives the obligation:

$$x = a, y = b, z_1 = x, x_1 = y, y_1 = z_1 \vdash P$$

This is a useful tactic, and the ability to select paths from the program affords more flexibility than bounded model checking the program:

$$z := x; \text{assert}(P); x := y; \text{assert}(P); y := z; \text{assert}(P)$$

This flexibility is important when the program is too large to perform bounded model checking on to a sufficiently large execution depth. If we can reason efficiently that there is a subset of paths on which the safety property will definitely hold, then the remaining (hopefully much smaller or at least finite) set of paths can be enumerated with symbolic execution and discharged individually. In subsequent lectures, we will see how runtime safety checks can be used to ensure that certain paths will never violate safety, so that targeted techniques like symbolic execution can be used to reason about the rest.