

15-414: [Bug Catching: Automated Program Verification](#)

Lecture Notes on Propositional Logic and Proofs

Matt Fredrikson

Carnegie Mellon University
Lecture 2

1 Introduction

The purpose of this lecture is to investigate the most basic of all logics: propositional logic, which is the logic of elementary logical connectives such as and/or etc. The primary motivation in this course for the study of propositional logic is in order to get a more precise understanding of the logical conditions typically used in program contracts. Today's lecture will not be enough to fully understand contracts, nor will propositional logic be sufficient for that purpose. But it is the most elementary preparation regardless.

We will get into the habit of thoroughly understanding all sides of the objects we deal with. Since we will be dealing with logical combinations such as ands and ors in contracts, today's lecture will right away explore the syntax of the language of propositional logic as well as its semantics and proof principles that it provides.

2 A Stroll Down Memory Lane: Recalling Contracts

Thinking back of the [15-122 Principles of Imperative Computation](#) course, we recall that contracts have served a valuable role in understanding programs. The experience in that particular course emphasized their role in imperative C0 programs and focused on informal proofs and dynamic checking of contracts. For example, Dijkstra's algorithm for computing the greatest common divisor of x and y needs a loop invariant and a precondition, because `Dijkstra(5, 0)` would not work in this C0 program:

```

int Dijkstra(int x, int y)
//@requires x>0 && y>0;
//@ensures  \result>0 && x % \result == 0 && y % \result == 0;
{
  int a=x;
  int b=y;
  int u=b;
  int v=a;
  while (x!=y)
  //@loop_invariant 2*a*b == u*x + v*y;
  {
    if (x>y) {
      x=x-y; v=v+u;
    } else {
      y=y-x; u=u+v;
    }
  }
  return x;
}

```

This algorithm uses contracts, which is a good thing. Are they all correct? Are they easy to follow? Is it enough to show $x \% \backslash\text{result} == 0 \ \&\& \ y \% \backslash\text{result} == 0$ holds at the return statement to show the postcondition? Are x and y the right variables to use in the `@ensures` clause or should we have used a and b instead? Does the postcondition follow easily from the loop invariant?

This is all quite exciting. But the purpose of today's lecture is not actually to get us back into specifying or checking contracts of programs, because that is what the entire next lecture is good for.

Instead of understanding any particular program or the meaning or effect that a contract has in a particular program, we, instead, zoom in on the formulation of the conditions in the contract themselves and try to understand what exactly they are.

What kind of expression is $x>0 \ \&\& \ y>0$ in the `@requires` precondition and what does it mean? Our layman's reading in the 15-122 course was that the C0 contracts `@requires`, `@ensures`, `@loop_invariant` and `@assert` just expect ordinary C0 expressions of type `bool` that are being evaluated and need to come back with value `true` to successfully pass.

Well, what exactly does the expression `\result` mean in the `@ensures` postcondition? What if the C0 expression in a contract calls a function that has the side effect of changing a data structure? Are side effects even allowed during contract checking? What does a recursive function call mean during a contract? What exactly is the meaning of the `&&` operator itself? What should its meaning be? Some form of logical and. Does it perform short-circuit evaluation? When exactly and how are the contracts evaluated? What if an expression crashes during contract evaluation? How do we know that the contracts are correct for a C0 program?

These are quite a number of subtle questions for something that we thought we had already mastered as well as the contracts from Principles of Imperative Computation. Maybe we should first take a step back and give the expressions within a contract a more careful look to see how they can best be understood.

3 Propositional Logic

Definition 1 (Syntax of propositional logic). The formulas F, G of propositional logic are defined by the following grammar (where p is an atomic proposition):

$$F ::= p \mid \neg F \mid F \wedge G \mid F \vee G \mid F \rightarrow G \mid F \leftrightarrow G$$

The way to read such a grammar is that whenever F and G are formulas then the conjunction $F \wedge G$ also is a formula and so is the disjunction $F \vee G$ as well as implication $F \rightarrow G$ and bisubjunction $F \leftrightarrow G$. And whenever F is a formula then the negation $\neg F$ is a formula, too. Finally, any atomic proposition, usually written p, q, r , is a formula. For example, this is a propositional formula:

$$(p \wedge q \rightarrow r) \wedge (p \rightarrow q) \rightarrow (p \rightarrow r) \quad (1)$$

4 Semantics of Propositional Logic

Writing down logical formulas that fit to the syntax of propositional logic is one thing, but not particularly useful unless we also know whether the formulas are actually true or not. Or, in fact, under which circumstances they are true or false. We cannot generally know whether the atomic propositions in a propositional logical formula are true or false, because they are just called p, q, r , which does not tell us much about their intention. But we can ask somewhere. Let's fix a function ω , called *interpretation*, that tells us the truth-value for each atomic proposition. So $\omega(p) = \text{true}$ iff atomic proposition p is interpreted as true in interpretation ω . For example, we could fix the following interpretation when interpreting formula (1):

$$\omega = \{q, r\} \quad (2)$$

By this common notation, we mean the interpretation that satisfies $\omega(q) = \text{true}$ and $\omega(r) = \text{true}$ and interprets all other atomic propositions such as p as *false*.

Having fixed an interpretation ω for the atomic proposition, we can now easily evaluate all propositional formulas to see whether they are true or false in that interpretation ω of atomic propositions, because the logical operators $\wedge, \vee, \neg, \rightarrow, \leftrightarrow$ always have exactly the same meaning.

Definition 2 (Semantics of propositional logic). The propositional formula F is true in interpretation ω , written $\omega \models F$, as inductively defined by distinguishing the shape of formula F :

1. $\omega \models p$ iff $\omega(p) = \text{true}$ for atomic propositions p
2. $\omega \models F \wedge G$ iff $\omega \models F$ and $\omega \models G$.
3. $\omega \models F \vee G$ iff $\omega \models F$ or $\omega \models G$.
4. $\omega \models \neg F$ iff $\omega \not\models F$, i.e. it is not the case that $\omega \models F$.
5. $\omega \models F \rightarrow G$ iff $\omega \not\models F$ or $\omega \models G$.
6. $\omega \models F \leftrightarrow G$ iff both are true or both false, i.e., it is either the case that both $\omega \models F$ and $\omega \models G$ or it is the case that $\omega \not\models F$ and $\omega \not\models G$.

With this definition, it is easy to establish that formula (1) is true in interpretation (2):

$$\omega \models (p \wedge q \rightarrow r) \wedge (p \rightarrow q) \rightarrow (p \rightarrow r)$$

For example, the evaluation of the right-hand side formula after the implication \rightarrow proceeds as follows:

$$\omega \models p \rightarrow r \text{ because } \omega \models r \text{ because } \omega(r) = \text{true}$$

Was this a coincidence? Is formula (1) only true in this particular interpretation (2) or what happens with other interpretations of the atomic propositions?

The most exciting formulas are those that are true no matter what the interpretation of the atomic propositions is. Such a formula is called *valid* and very helpful, because it expresses a true property no matter what specific interpretation of the atomic propositions we had in mind.

Definition 3 (Validity). A formula F is called *valid* iff it is true in all interpretations, i.e. $\omega \models F$ for all interpretations ω . Because any interpretation makes valid formulas true, we also write $\models F$ iff formula F is valid. A formula F is called *satisfiable* iff there is an interpretation ω in which it is true, i.e. $\omega \models F$. Otherwise it is called *unsatisfiable*.

Indeed, if we try *all* other interpretations to evaluate formula (1) we will find that it is always true. Let's tabulate our results by writing down each combination of truth-values for all atomic propositions and evaluating all subformulas of (1) according to their semantics.

p	q	r	$p \wedge q$	$p \wedge q \rightarrow r$	$p \rightarrow q$	$p \rightarrow r$	$(p \wedge q \rightarrow r) \wedge (p \rightarrow q)$	(1)
true	true	true	true	true	true	true	true	true
false	true	true	false	true	true	true	true	true
true	false	true	false	true	false	true	false	true
false	false	true	false	true	true	true	true	true
true	true	false	true	false	true	false	false	true
false	true	false	false	true	true	true	true	true
true	false	false	false	true	false	false	false	true
false	false	false	false	true	true	true	true	true

Indeed, the truth-value of the formula (1) is *true* in all interpretations, thus, (1) is valid:

$$\models (p \wedge q \rightarrow r) \wedge (p \rightarrow q) \rightarrow (p \rightarrow r)$$

The only downside is all this busywork to evaluate all interpretations, which is exponential in the number of variables and incredibly boring on top of that.

5 Proofs for Propositional Logic

Literally evaluating a formula in all possible interpretations is certainly one way of establishing that a propositional logical formula is valid, but it always requires exponential effort and is quite uninteresting, because it does not even provide a comprehensible reason for the validity of the formula. The only way to check that a truth-table is constructed correctly for a formula is to check that it enumerates all cases of interpretations and all its computations of truth-values are according to the semantics and that, indeed, *true* is the outcome in all cases. Possible but incredibly dull. Besides, this finite enumeration principle cannot work for the significantly more interesting and expressive logics that we will be pursuing to understand programs in subsequent lectures.

The semantics considered one operator at a time. Let's try to make the same thing happen for proofs as well. What about a proof of a conjunction $F \wedge G$? How could that work?

A proof of a conjunction $F \wedge G$ should consist of a proof of the left conjunct F together with a proof of the right conjunct G , because both proofs together prove the conjunction $F \wedge G$. So stapling a proof of F together with a proof of G will give us a proof of $F \wedge G$. That was easy enough.

But what does a proof of an implication $F \rightarrow G$ consist of? It certainly isn't a proof of F together with a proof of G anymore. A proof of G would constitute a proof of $F \rightarrow G$, but such a proof is missing out on an important power. It would have been allowed to assume F , because the formula $F \rightarrow G$ only says that F implies G , so that G is true in case F is. If F isn't true, then the implication $F \rightarrow G$ doesn't say anything about whether G is true or not. (Check back with Def. 2 if you don't believe this). Consequently, an unconditional proof of G certainly does establish $F \rightarrow G$, but is a bit much to ask for. The proof of $F \rightarrow G$ should, instead, consist of a proof of G that is allowed to assume F . This requires the capability to manage assumptions in a proof, which, retrospectively, should not actually come as a surprise.

For managing assumptions in a structured way, we will follow in the footsteps of Gerhard Gentzen [Gen35], who introduced sequent calculus for the study of logic. But it turns out that sequent calculi are also immensely useful not just for understanding logical reasoning, but also for organizing and conducting proofs without risking to lose track of assumptions.

5.1 Simple Sequents

The first kind of *sequent* that we will consider (and subsequently generalize) is of the form

$$\Gamma \vdash F$$

with the available assumptions as a list of formulas Γ as *antecedent* and with the formula we want to prove from it as F . The symbol \vdash is called *sequent turnstile* and separates the available assumptions from what we try to prove from them.

There are some sequents where we are obviously done with a proof. For example when literally the same formula F is in the antecedent and the succedent, because F easily follows when assuming F . So the sequent $\Gamma, F \vdash F$ has a trivial proof. We will later capture this thought with a proof rule [id](#), but first consider proofs for the operators we already started considering.

Coming back to conjunctions, proving a conjunction $P \wedge Q$ requires proving P and proving Q . This fact does not change when working from a list of assumptions Γ .

$$(\wedge R) \frac{\Gamma \vdash P \quad \Gamma \vdash Q}{\Gamma \vdash P \wedge Q}$$

This proof rule $\wedge R$ expresses that all it takes to prove the *conclusion* $\Gamma \vdash P \wedge Q$ below the rule bar is to prove all the *premises* $\Gamma \vdash P$ and $\Gamma \vdash Q$ above the rule bar. In the proof of the left premise $\Gamma \vdash P$, the same assumptions Γ will still be available that were available in the conclusion $\Gamma \vdash P \wedge Q$. And likewise for the right premise.

Proving an implication $P \rightarrow Q$, with which we had difficulties before, now simply allows us to add the assumption P to the antecedent with the list of all available assumptions and continue a proof of Q from this augmented list of assumptions:

$$(\rightarrow R) \frac{\Gamma, P \vdash Q}{\Gamma \vdash P \rightarrow Q}$$

Reading the rule $\rightarrow R$ from bottom to top means that a proof of an implication $P \rightarrow Q$ from a list of assumptions Γ requires us to prove Q from the assumptions Γ together with P . If we keep on applying rule $\rightarrow R$ (and the other rules) then all our available assumptions will ultimately land in the antecedent.

Proving a disjunction $P \vee Q$ is more subtle. How do we prove a disjunction? We could prove a disjunction $P \vee Q$ by proving the left disjunct P :

$$(\vee R_1) \frac{\Gamma \vdash P}{\Gamma \vdash P \vee Q}$$

That works. But then what if the disjunction $P \vee Q$ is true because the right disjunct Q is true? Well, we could adopt yet another proof rule for disjunction that shows the right disjunct instead:

$$(\vee R_2) \frac{\Gamma \vdash Q}{\Gamma \vdash P \vee Q}$$

This would give us a pair of proof rules $\forall R_1$ and $\forall R_2$ to prove disjunctions. But we will have to choose at the time of proving the disjunction $P \vee Q$ whether we prove it by proving its left disjunct P with rule $\forall R_1$ or whether we prove it by proving its right disjunct Q with rule $\forall R_2$. That requires a lot of attention when proving disjunctions. Worse yet: will we always be able to tell which disjunct we will be able to prove?

In many cases, we will be able to predict which disjunct of a disjunction we will be able to prove if we think ahead very carefully. But that is already not particularly helpful and convenient. Worse yet, there are cases where, for principle reasons, we will be unable to predict which disjunct of a disjunction we will prove! Suppose we are trying to prove the formula $p \vee \neg p$, which is certainly valid, because it will evaluate to *true* whether or not the atomic proposition p is interpreted to be *true*. But when trying to prove the law of excluded middle $p \vee \neg p$, neither rule $\forall R_1$ nor rule $\forall R_2$ will succeed because the whole point of the law of excluded middle is that it will evaluate to *true* whether p is *true* or *false* (so $\neg p$ is *true*), but we cannot generally say ahead of time which side will be *true*.

Instead, what we are going to do is to keep our options open. We will record in the sequent the fact that formulas P as well as Q were both available as formulas for us to prove when proving the disjunction $P \vee Q$ by keeping both as a list on the right-hand side of the sequent turnstile \vdash . Of course, we might have already gather other options that we could prove, so the disjunction proof rule is:

$$(\forall R) \frac{\Gamma \vdash P, Q, \Delta}{\Gamma \vdash P \vee Q, \Delta}$$

Proving a disjunction $P \vee Q$ from a list of assumptions Γ with a list of alternatives Δ works by splitting the disjunction into its two options P and Q and continuing with a proof of the alternatives P, Q, Δ from the assumptions Γ .

5.2 Sequent Calculus

To manifest this, let's properly define what a sequent $\Gamma \vdash \Delta$ is and what it means.

Definition 4 (Sequent). A *sequent* $\Gamma \vdash \Delta$ organizes the reasoning into a list Γ of formulas available as assumptions, called *antecedent*, and a list Δ called *succedent*. The semantics of sequent $\Gamma \vdash \Delta$ is the same as that of the formula

$$\left(\bigwedge_{F \in \Gamma} F \right) \rightarrow \left(\bigvee_{G \in \Delta} G \right)$$

In particular, proving a sequent $\Gamma \vdash \Delta$ requires proving that the disjunction of all succedent formulas Δ is implied by the conjunction of all antecedent formulas Γ . For proving a sequent $\Gamma \vdash \Delta$, we can, thus, assume all formulas in Γ and need to show one of the formulas in Δ , or at least show their disjunction.

This list Δ of alternatives to prove is simply preserved in the proof rules we saw so far:

$$(\wedge R) \frac{\Gamma \vdash P, \Delta \quad \Gamma \vdash Q, \Delta}{\Gamma \vdash P \wedge Q, \Delta}$$

$$(\rightarrow R) \frac{\Gamma, P \vdash Q, \Delta}{\Gamma \vdash P \rightarrow Q, \Delta}$$

$$(\vee R) \frac{\Gamma \vdash P, Q, \Delta}{\Gamma \vdash P \vee Q, \Delta}$$

For example in rule $\wedge R$, the same succedent Δ is still available in both premises, because a proof of Δ from the assumptions Γ in either premise would also prove Δ from the assumptions Γ in the conclusion.

When we leave the development of proof rules for the bisubjunction operator \leftrightarrow as an exercise, the only remaining operator to worry about is negation \neg . How do we prove a negation $\neg P$?

We can prove a negation $\neg P$ by assuming the converse P and going for a contradiction. In fact, since we may have already gathered a number of other alternatives Δ to prove, all we need to do to prove $\neg P$ from a list of assumptions Γ with a list of alternatives Δ is to prove the remaining alternatives Δ from assuming Γ as well as the opposite P :

$$(\neg R) \frac{\Gamma, P \vdash \Delta}{\Gamma \vdash \neg P, \Delta}$$

Does this list of rules handle all operators? There's one rule per operator, which is a good thing. The catch is that there's really only one rule per operator so far. If the operators occur on the right, so in the succedent, then the respective proof rules tell us what to do. But the implication proof rule $\rightarrow R$ is good about pushing assumptions into the antecedent. What if it pushes a conjunction $P \wedge Q$ into the antecedent? Is there a proof rule to handle what happens then?

Not yet. But there should be a rule for handling the case where there's a conjunction $P \wedge Q$ among the list of assumptions in the antecedent. In fact, for every logical operator, there should be a right proof rule handling the case where it is the top-level operator on the right in the succedent as well as a left proof rule handling when it appears on the left in the antecedent.

5.3 Left Rules

When we find a conjunction $P \wedge Q$ among the list of assumptions in the antecedent, then we can safely split it into two separate assumptions P as well as Q :

$$(\wedge L) \frac{\Gamma, P, Q \vdash \Delta}{\Gamma, P \wedge Q \vdash \Delta}$$

Proving a sequent that has a conjunction $P \wedge Q$ among its assumptions in the antecedent is the same as proving it with two separate assumptions P as well as Q instead.

What happens when we have a disjunction $P \vee Q$ among our assumptions in the antecedent? In that case we have no way of knowing whether P or whether Q is true.

All we know is that either of them is. But we still succeed with a proof if we manage to show the sequent both when assuming P as well as when, instead, assuming Q , because while either are possible, the assumption $P \vee Q$ implies that one of those cases has to apply.

$$(\vee\text{L}) \frac{\Gamma, P \vdash \Delta \quad \Gamma, Q \vdash \Delta}{\Gamma, P \vee Q \vdash \Delta}$$

When an implication $P \rightarrow Q$ is among the assumptions in the antecedent, then we can make use of that assumption by showing its respective assumption P and can then assume Q instead. If we can assume $P \rightarrow Q$ and show P then we can assume Q :

$$(\rightarrow\text{L}) \frac{\Gamma \vdash P, \Delta \quad \Gamma, Q \vdash \Delta}{\Gamma, P \rightarrow Q \vdash \Delta}$$

Wait a moment. The left premise does not actually show P from the assumptions Γ , because it only shows the succedent P, Δ which is interpreted disjunctively. So it is possible that the left premise does not show P but merely Δ . But in that case, the conclusion is justified as well, because it also has the antecedent Δ as the list of alternatives to show.

Since the operator \leftrightarrow is left as an exercise, the only remaining case is to handle a negation $\neg P$ among the assumptions in the antecedent. If we assume $\neg P$ then it is also sufficient if we can show the opposite P (recall the semantics of sequents):

$$(\neg\text{L}) \frac{\Gamma \vdash P, \Delta}{\Gamma, \neg P \vdash \Delta}$$

To understand, we can first pretend there would be no succedent Δ . What happens if there is no succedent? Then the empty disjunction that it means is equivalent to the formula *false* that is never true in any interpretation. In that special case, rule $\neg\text{L}$ says that for proving a contradiction *false* from assumptions Δ and $\neg P$, it is sufficient to prove the opposite P from the remaining assumptions Γ .

5.4 Closing and Forking

The above proof rules excel at splitting operators off of propositional logical formulas. But they never actually prove anything on their own except simplifying all formulas until only atomic propositions are left. What is missing is the observation that a sequent can be proved easily when the same formula P is in the antecedent and succedent with the identity proof rule called [id](#):

$$(\text{id}) \frac{}{\Gamma, P \vdash P, \Delta}$$

Whenever we find the same formula P in the antecedent and succedent, we can use rule [id](#) to prove that sequent without any further questions (no premise, i.e. no more remaining subgoals).

Another insightful proof rule is the cut proof rule, which enables us to first prove an arbitrary formula C on the left premise and then assume C on the right premise.

$$\text{(cut)} \frac{\Gamma \vdash C, \Delta \quad \Gamma, C \vdash \Delta}{\Gamma \vdash \Delta}$$

Think of C as a lemma that is proved in the left premise and then assumed to hold in the right premise. The twist is again that the left premise does not necessarily prove C but might also settle for proving another alternative in the remaining succedent Δ , but that also establishes the succedent Δ of the conclusion. The primary purpose of the **cut** rule is for ingenious theoretical studies of reasoning [Gen35] as well as to find clever shortcuts in practical proofs by first proving a lemma C that subsequently helps multiple times in the remaining proof. It plays a crucial role in constructive logics, too.

All these sequent calculus proof rules are *sound*, that is, if all their premises are valid, then their conclusion is valid. Especially if there are no premises any more because we were able to use the identity proof rule **id** on all premises, then the conclusion is valid, which is what we were hoping to achieve with a proof.

5.5 Conducting Sequent Calculus Proofs

As an example, let's prove formula (1). Sequent calculus proofs are conducted in a bit of a funny way by starting with the conjecture at the bottom

$$\vdash (p \wedge q \rightarrow r) \wedge (p \rightarrow q) \rightarrow (p \rightarrow r)$$

and then working our way upwards by applying proof rules to the remaining sequents. The reason why we work like that is that in (sound!) sequent calculus proof rules validity of all premises implies validity of the conclusion. So if we start with our conjecture at the bottom and work our way upwards, then if we are able to prove all premises then the conclusion at the bottom will be valid, too. We apply sequent calculus rules from the bottom to the top but, when a proof is done, their soundness makes validity inherit from the top to the bottom.

Enough said. Let's prove formula (1) in sequent calculus:

$$\begin{array}{c} \text{id}_{q,p \vdash p,r}^* \quad \text{id}_{q,p \vdash q,r}^* \\ \wedge R \frac{}{q,p \vdash p \wedge q,r} \\ \text{id}_{r,q,p \vdash r}^* \\ \rightarrow L \frac{}{p \wedge q \rightarrow r, p \vdash p,r} \\ \rightarrow L \frac{}{p \wedge q \rightarrow r, p \rightarrow q, p \vdash r} \\ \rightarrow R \frac{}{p \wedge q \rightarrow r, p \rightarrow q \vdash p \rightarrow r} \\ \wedge L \frac{}{(p \wedge q \rightarrow r) \wedge (p \rightarrow q) \vdash p \rightarrow r} \\ \rightarrow R \frac{}{\vdash (p \wedge q \rightarrow r) \wedge (p \rightarrow q) \rightarrow (p \rightarrow r)} \end{array}$$

6 Soundness

Having conducted a sequent calculus proof, the most pressing question is what a proof proves. Of course, as we already alluded to before, a proof in a sound proof calculus implies the validity of the conclusion.

Definition 5 (Soundness of a proof rule). A sequent calculus proof rule

$$\frac{\Gamma_1 \vdash \Delta_1 \quad \dots \quad \Gamma_n \vdash \Delta_n}{\Gamma \vdash \Delta}$$

is *sound* iff the validity of all premises implies the validity of the conclusion:

$$\text{if } \models (\Gamma_1 \vdash \Delta_1) \text{ and } \dots \text{ and } \models (\Gamma_n \vdash \Delta_n) \text{ then } \models (\Gamma \vdash \Delta)$$

Recall from Def. 4 that the meaning of the sequent $\Gamma \vdash \Delta$ is the same as that of the formula $(\bigwedge_{F \in \Gamma} F) \rightarrow (\bigvee_{G \in \Delta} G)$.

$$\begin{array}{ll} (\wedge R) \frac{\Gamma \vdash P, \Delta \quad \Gamma \vdash Q, \Delta}{\Gamma \vdash P \wedge Q, \Delta} & (\wedge L) \frac{\Gamma, P, Q \vdash \Delta}{\Gamma, P \wedge Q \vdash \Delta} \\ (\vee R) \frac{\Gamma \vdash P, Q, \Delta}{\Gamma \vdash P \vee Q, \Delta} & (\vee L) \frac{\Gamma, P \vdash \Delta \quad \Gamma, Q \vdash \Delta}{\Gamma, P \vee Q \vdash \Delta} \\ (\rightarrow R) \frac{\Gamma, P \vdash Q, \Delta}{\Gamma \vdash P \rightarrow Q, \Delta} & (\rightarrow L) \frac{\Gamma \vdash P, \Delta \quad \Gamma, Q \vdash \Delta}{\Gamma, P \rightarrow Q \vdash \Delta} \\ (\neg R) \frac{\Gamma, P \vdash \Delta}{\Gamma \vdash \neg P, \Delta} & (\neg L) \frac{\Gamma \vdash P, \Delta}{\Gamma, \neg P \vdash \Delta} \\ (\text{id}) \frac{}{\Gamma, P \vdash P, \Delta} & (\text{cut}) \frac{\Gamma \vdash C, \Delta \quad \Gamma, C \vdash \Delta}{\Gamma \vdash \Delta} \end{array}$$

Figure 1: Sequent calculus proof rules for propositional logic

Lemma 6 (Soundness of propositional logic proof rules). *All propositional logic proof rules (summarized again in Fig. 1), are sound.*

Proof. It is crucial to prove soundness for all proof rules. We will, nevertheless, only prove it for one rule and leave the others as exercises. But we will prove that rule with exceeding care.

$\wedge R$ That proof rule $\wedge R$ is sound can be shown as follows. Assume that both of its premises $\Gamma \vdash P, \Delta$ and $\Gamma \vdash Q, \Delta$ are valid, i.e. both $(\bigwedge_{F \in \Gamma} F) \rightarrow P \vee (\bigvee_{G \in \Delta} G)$ and $(\bigwedge_{F \in \Gamma} F) \rightarrow Q \vee (\bigvee_{G \in \Delta} G)$ are true in all interpretations. We need to show that the conclusion $\Gamma \vdash P \wedge Q, \Delta$ is then also valid, i.e. $\models (\Gamma \vdash P \wedge Q, \Delta)$, which means that $(\bigwedge_{F \in \Gamma} F) \rightarrow (P \wedge Q) \vee (\bigvee_{G \in \Delta} G)$ is true in all interpretations. Consider any interpretation ω and show that $\omega \models (\bigwedge_{F \in \Gamma} F) \rightarrow (P \wedge Q) \vee (\bigvee_{G \in \Delta} G)$. If any of the

antecedent formulas $F \in \Gamma$ is false in ω ($\omega \not\models F$) or any of the remaining succedent formulas $G \in \Delta$ is true ($\omega \models G$), then $\omega \models (\bigwedge_{F \in \Gamma} F) \rightarrow (P \wedge Q) \vee (\bigvee_{G \in \Delta} G)$. Otherwise, all antecedent formulas in Γ are true $\omega \models \bigwedge_{F \in \Gamma} F$ and all Δ formulas are false $\omega \not\models \bigvee_{G \in \Delta} G$.

By premise, $\omega \models (\bigwedge_{F \in \Gamma} F) \rightarrow P \vee (\bigvee_{G \in \Delta} G)$ and $\omega \models (\bigwedge_{F \in \Gamma} F) \rightarrow Q \vee (\bigvee_{G \in \Delta} G)$. Since antecedents in Γ are true and succedents in Δ false in ω , this implies $\omega \models P$ and $\omega \models Q$. By Def. 2, these imply $\omega \models P \wedge Q$, which implies that the conclusion is true in ω , i.e. $\omega \models (\bigwedge_{F \in \Gamma} F) \rightarrow (P \wedge Q) \vee (\bigvee_{G \in \Delta} G)$. \square

In fact, the prelude of the soundness argument is common to all proof rules so that one usually just assumes right away without loss of generality that the common antecedent Γ is true while the common succedent Δ false in the current interpretation ω .

Now that all proof rules of propositional logic are sound it is easy to see that the whole proof calculus is sound, because a proof is entirely built by applying sound proof rules so validity of all premises (of which there are none in a completed proof) implies validity of the conclusion. Because this is so important and we want to practice the important proof principle of induction, we will show this explicitly.

Theorem 7 (Soundness of propositional logic). *The sequent calculus of propositional logic is sound, i.e. it only proves valid formulas. That is, if $\vdash P$ has a proof in the propositional sequent calculus, then P is valid, i.e. $\models P$.*

Proof. What we need to show is that if $\vdash P$ is the conclusion of a completed sequent calculus proof, then P is valid, i.e. $\models P$. A proof of the sequent $\vdash P$ will consist of proofs of sequents of the more general shape $\Gamma \vdash \Delta$. So we instead prove the more general statement that a proof of $\Gamma \vdash \Delta$ implies $\models (\Gamma \vdash \Delta)$. We will prove this by induction on the structure of the proof. That is, we will prove it for the smallest possible proofs. And then, assuming that the proofs of the smaller pieces of a proof have valid conclusions, we will show that one more proof step preserves validity.

1. The only proofs with just 1 proof step are of the form

$$\text{id} \frac{*}{\Gamma, P \vdash P, \Delta}$$

Its conclusion is valid, because assumption P in the antecedent trivially implies P in the succedent.

2. Consider any proof ending with a proof step of this form:

$$\frac{\Gamma_1 \vdash \Delta_1 \quad \dots \quad \Gamma_n \vdash \Delta_n}{\Gamma \vdash \Delta} \quad (3)$$

By induction hypothesis, we can assume that the (smaller!) proofs of the premises $\Gamma_1 \vdash \Delta_1$ and $\dots \Gamma_n \vdash \Delta_n$ already imply the validity of their respective conclusions so $\models (\Gamma_1 \vdash \Delta_1)$ and $\dots \models (\Gamma_n \vdash \Delta_n)$.

The proof rule used in the proof step (3) must have been one of the proof rules of the sequent calculus of propositional logic. All these sequent calculus proof rules of propositional logic are sound by Lemma 6. Consequently, $\models (\Gamma \vdash \Delta)$, so the conclusion of the proof (3) is valid. \square

Soundness is one thing, and most crucial for any correct reasoning. But since propositional logic is so simple, it enjoys other pleasant properties. It is also the case that every valid propositional logic formula will be provable from the sequent calculus proof rules in Fig. 1, which is called *completeness*.

Theorem 8 (Completeness of propositional logic). *The sequent calculus of propositional logic is complete, i.e. it proves all valid formulas. That is, if P is valid, so $\models P$ then $\vdash P$ has a proof in the propositional sequent calculus.*

In fact, because propositional logic is so simple, it is perfectly decidable whether a propositional logical formula is valid.

Theorem 9 (Decidability of propositional logic). *Propositional logic is decidable, i.e. there is an algorithm that accepts any propositional logical formula as input and correctly outputs "valid" or "not valid" in finite time.*

How could such an algorithm possibly work? Well how to do that as efficiently as possible is the purpose of a SAT solver, which we will learn more about in a later lecture. That it is possible at all, however, is absolutely trivial. All that the algorithm needs to do is write down every interpretation with any true/false assignment for all the (finitely many!) atomic propositions in the logical formula and check whether it evaluates to true according to Def. 2. Easy, but boring. And of inherently exponential effort, because there are exponentially many interpretations to consider (in the number of the variables). This is why SAT solvers try to be a lot more clever about it. Whether SAT solvers have a chance to be inherently faster than exponential in the worst-case is, of course, the exciting open P-vs-NP problem.

Why do SAT solvers have such a funny name? Well, because they solve the question whether a propositional logical formula is satisfiable. What does that have to do with validity? If a formula is satisfiable, what does that tell us about validity? If a formula is valid, what does that tell us about satisfiability?

Of course, if a formula is valid, so true in all interpretations, it is clearly satisfiable so true in at least one interpretation. But the converse is totally wrong. Yet if the negation $\neg P$ of the formula P is satisfiable, then P itself cannot possibly be valid, because there apparently is an interpretation ω in which its negation $\neg P$ is already true. And it is quite impossible for $\omega \models \neg P$ and $\omega \models P$ to hold at the same time. Indeed, the formula P is valid if and only if its negation $\neg P$ is unsatisfiable.

Lemma 10. *A formula P is valid if and only if its negation $\neg P$ is unsatisfiable.*

This lemma would be an incredibly boring observation if it wasn't for the fact that it explains why SAT solvers are useful for checking the validity of propositional logical formulas.

7 Summary

The proof rules for propositional logic that this lecture discussed are summarized in Fig. 1 on p. 11. Other important concepts from this lecture that will be with us in the future are soundness and the principles of structural induction employed in proving it.

References

- [Gen35] Gerhard Gentzen. Untersuchungen über das logische Schließen. I. *Math. Zeit.*, 39(2):176–210, 1935.