Lecture Notes on
Sequent Calculus

15-317: Constructive Logic
Frank Pfenning

Lecture 9
September 24, 2009

1 Introduction

In this lecture we develop the sequent calculus as a formal system for proof search in natural deduction. The sequent calculus was originally introduced by Gentzen [Gen35], primarily as a technical device for proving consistency of predicate logic. Our goal of describing a proof search procedure for natural deduction predisposes us to a formulation due to Kleene [Kle52] called G_3.

Our sequent calculus is designed to *exactly* capture the notion of a *verification*, introduced in Lecture 3. Recall that verifications are constructed bottom-up, from the conclusion to the premises using introduction rules, while uses are constructed top-down, from hypotheses to conclusions using elimination rules. They meet in the middle, where an assumption may be used as a verification for an atomic formula. In the sequent calculus, both steps work bottom-up, which will allows us to prove global versions of the local soundness and completeness properties foreshadowed in Lecture 3.

2 Sequents

When constructing a verification, we are generally in a state of the following form

$$A_1\downarrow \quad \cdots \quad A_n\downarrow$$
$$\vdots$$
$$C\uparrow$$

where $A_1, \ldots, A_n$ are all assumptions we may *use*, while C is the conclusion we are trying to *verify*. A *sequent* is just a local notation for such a partially complete verification. We write

$$A_1 \text{ left}, \ldots, A_n \text{ left} \Longrightarrow C \text{ right}$$

where the judgments A left and C right correspond to $A\downarrow$ and $C\uparrow$, respectively. The judgments on the left are assumptions called *antecedents*, the judgment on the right is the conclusion called the *succedent*.

The rules that define the A left and A right judgment are systematically constructed from the introduction and elimination rules, keeping in mind their directions in terms of verifications and uses. Introduction rules are translated to corresponding *right rules*. Since introduction rules already work from the conclusion to the premises, this mapping is straightforward. Elimination rules work top-down, so they have to be flipped upside-down in order to work as sequent rules, and are turned into *left rules*. Pictorially:
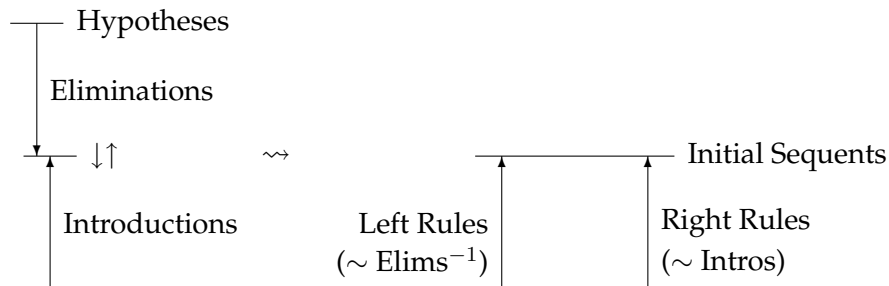
We now proceed connective by connective, constructing the right and left rules from the introduction and elimination rules. When writing a sequent, we can always tell which propositions are on the left and which are on the right, so we omit the judgments left and right for brevity. Also, we abbreviate a collection of antecedents $A_1 \text{ left}, \ldots, A_n \text{ left}$ by Γ. The order of the antecedents does not matter, so we will allow them to be implicitly reordered.

Conjunction. We recall the introduction rule first and show the corresponding right rule.

$$\frac{A\uparrow \quad B\uparrow}{A \wedge B\uparrow} \wedge I \qquad\qquad \frac{\Gamma \Longrightarrow A \quad \Gamma \Longrightarrow B}{\Gamma \Longrightarrow A \wedge B} \wedge R$$

The only difference is that the antecedents Γ are made explicit. Both premises have the same antecedents, because any assumption can be used in both subdeductions.

There are two elimination rules, so we two corresponding left rules. Since the letters L and R are used to denote the type of rule in the sequent calculus, we index the rules as first and second conjunction left rule.

$$\frac{A \wedge B \downarrow}{A \downarrow} \wedge E_L \qquad \frac{\Gamma, A \wedge B, A \Longrightarrow C}{\Gamma, A \wedge B \Longrightarrow C} \wedge L_1$$

$$\frac{A \wedge B \downarrow}{B \downarrow} \wedge E_R \qquad \frac{\Gamma, A \wedge B, B \Longrightarrow C}{\Gamma, A \wedge B \Longrightarrow C} \wedge L_2$$

We preserve the *principal formula* $A \wedge B$ of the left rule in the premise. This is because we are trying to model proof construction in natural deduction where assumptions can be used multiple times. If we temporarily ignore the copy of $A \wedge B$ in the premise, it is easier to see how the rules correspond.

Truth. Truth is defined just by an introduction rule and has no elimination rule. Consequently, there is only a right rule in the sequent calculus and no left rule.

$$\frac{}{\top \uparrow} \top I \qquad \frac{}{\Gamma \Longrightarrow \top} \top R$$

Implication. Again, the right rule for implication is quite straightforward, because it models the introduction rule directly.

$$\frac{\overline{A \downarrow}^{\,u} \atop \vdots \atop B \uparrow}{A \supset B \uparrow} \supset^u \qquad \frac{\Gamma, A \Longrightarrow B}{\Gamma \Longrightarrow A \supset B} \supset R$$

We see here one advantage of the sequent calculus over natural deduction: the scoping for additional assumptions is simple. The new antecedent A left is available anywhere in the deduction of the premise, because in the sequent calculus we only work bottom-up. Moreover, we arrange all the rule so that antecedents are *persistent*: they are always propagated from the conclusion to all premises.

The elimination rule is trickier, because it involves a more complicated combination of verifications and uses.

$$\frac{A \supset B \downarrow \quad A \uparrow}{B \downarrow} \supset E \qquad \frac{\Gamma, A \supset B \Longrightarrow A \quad \Gamma, A \supset B, B \Longrightarrow C}{\Gamma, A \supset B \Longrightarrow C} \supset L$$

In words: in order to use $A \supset B$ we have to produce a verification of A, in which case we can use B. The antecedent $A \supset B$ is carried over to both premises to maintain persistence. Note that the premises of the left rule are reversed, when compared to the elimination rule to indicate that we do not want to make the assumption B unless we have already established A.

In terms of provability, there is some redundancy in the $\supset L$ rule. For example, once we know B, we no longer need $A \supset B$, because B is a stronger assumption. As stressed above, we try to maintain the correspondence to natural deductions and postpone these kinds of optimization until later.

Disjunction.   The right rules correspond directly to the introduction rules, as usual.

$$\frac{A\uparrow}{A \vee B\uparrow} \vee I_L \qquad \frac{\Gamma \Longrightarrow A}{\Gamma \Longrightarrow A \vee B} \vee R_1$$

$$\frac{B\uparrow}{A \vee B\uparrow} \vee I_R \qquad \frac{\Gamma \Longrightarrow B}{\Gamma \Longrightarrow A \vee B} \vee R_2$$

The disjunction elimination rule was somewhat odd, because it introduced two new assumptions, one for each case of the disjunction. The left rule for disjunction actually has a simpler form that is more consistent with all the other rules we have shown so far.

$$\frac{A \vee B\downarrow \quad \begin{array}{c} \overline{A\downarrow}\ u \\ \vdots \\ C\uparrow \end{array} \quad \begin{array}{c} \overline{B\downarrow}\ w \\ \vdots \\ C\uparrow \end{array}}{C\uparrow} \vee E^{u,w} \qquad \frac{\Gamma, A \vee B, A \Longrightarrow C \quad \Gamma, A \vee B, B \Longrightarrow C}{\Gamma, A \vee B \Longrightarrow C} \vee L$$

As for implication, scoping issues are simplified because the new assumptions A and B in the first and second premise, respectively, are available anywhere in the deduction above.

Falsehood.   Falsehood has no introduction rule, and therefore no right rule in the sequent calculus. To arrive at the left rule, we need to pay attention to the distinction between uses and verifications, or we can construct the 0-ary case of disjunction from above.

$$\frac{\bot\downarrow}{C\uparrow} \bot E \qquad \frac{}{\Gamma, \bot \Longrightarrow C} \bot L$$

Atomic propositions. Recall that we cannot use an introduction rule to verify atomic propositions P because they cannot be broken down further. The only possible verification of P is directly via a use of P. This turns into a so-called *initial sequent*.

$$\frac{P\downarrow}{P\uparrow} \;\downarrow\uparrow \qquad \frac{}{\Gamma, P \Longrightarrow P} \;\text{init}$$

This rule has a special status in that it does not break down any proposition, but establishes a connection between two judgments. In natural deduction, it is the connection between uses and verifications; in sequent calculus, it is the connection between the left and right judgments.

As a simple example, we consider the proof of $(A \vee B) \supset (B \vee A)$.

$$\frac{\dfrac{\overline{A \vee B, A \Longrightarrow A}\;\text{init}}{A \vee B, A \Longrightarrow B \vee A}\;\vee R_2 \quad \dfrac{\overline{A \vee B, B \Longrightarrow B}\;\text{init}}{A \vee B, B \Longrightarrow B \vee A}\;\vee R_1}{\dfrac{\dfrac{A \vee B \Longrightarrow B \vee A}{\Longrightarrow (A \vee B) \supset (B \vee A)}\;\supset R}{}}\;\vee L$$

3 Observations on Sequent Proofs

We have already mentioned that antecedents in sequent proofs are *persistent*: once an assumption is made, it is henceforth usable above the inference that introduces it. Sequent proofs also obey the important *subformula property*: if we examine the complete or partial proof above a sequent, we observe that all sequents are made up of subformulas of the sequent itself. This is consistent with the design criteria for the verifications: the verification of a proposition A may only contain subformulas of A. This is important from multiple perspectives. Foundationally, we think of verifications as defining the meaning of the propositions, so a verification of a proposition should only depend on its constituents. For proof search, it means we do not have to try to resort to some unknown formula, but can concentrate on subformulas of our proof goal.

If we trust for the moment that a proposition A is true if and only if it has a deduction in the sequent calculus (as $\Longrightarrow A$), we can use the sequent calculus to formally prove that some proposition can *not* be true in general. For example, we can prove that intuitionistic logic is *consistent*.

Theorem 1 (Consistency) *It is not the case that $\implies \perp$.*

Proof: No left rule is applicable, since there is no antecedent. No right rule is applicable, because there is no right rule for falsehood. Therefore, there cannot be a proof of $\implies \perp$. $\qquad \square$

Theorem 2 (Disjunction Property) *If $\implies A \lor B$ then either $\implies A$ or $\implies B$.*

Proof: No left rule is applicable, since there is no antecedent. The only right rules that are applicable are $\lor R_1$ and $\lor R_2$. In the first case, we have $\implies A$, in the second $\implies B$. $\qquad \square$

Theorem 3 (Failure of Excluded Middle) *It is not the case that $\implies A \lor \neg A$ for arbitrary A.*

Proof: From the disjunction property, either $\implies A$ or $\implies \neg A$. For the first sequent, no rule applies. For the second sequent, only $\supset R$ applies and we would have to have a deduction of $A \implies \perp$. But for this sequent no rule applies. $\qquad \square$

There are other simple observations that are important for some applications. The first is called *weakening*, which means that we can add an arbitrary proposition to a derivable sequent and get another derivable sequent with a proof that has the same structure.

Theorem 4 (Weakening) *If $\Gamma \implies C$ then $\Gamma, A \implies C$ with a structurally identical deduction.*

Proof: Add A to every sequent in the given deduction of $\Gamma \implies C$, but never use it. The result is a structurally identical deduction of $\Gamma, A \implies C$. $\square$

Theorem 5 (Contraction) *If $\Gamma, A, A \implies C$ then $\Gamma, A \implies C$ with a structurally identical deduction.*

Proof: Pick one copy of A. Wherever the other copy of A is used in the given deduction, use the first copy of A instead. The result is a structurally identical deduction with one fewer copy of A. $\qquad \square$

The proof of contraction actually exposes an imprecision in our presentation of the sequent calculus. When there are two occurrences of a proposition A among the antecedents, we have no way to distinguish which one is being used, either as the principal formula of a left rule or in an initial sequent. It would be more precise to label each antecedent with a unique label and then track labels in the inferences. We may make this precise at a later stage in this course; for now we assume that occurrences of antecedents can be tracked somehow so that the proof above, while not formal, is at least somewhat rigorous.

Now we can show that double negation elimination does not hold in general

Theorem 6 (Failure of Double Negation Elimination) *It is not the case that* $\implies \neg\neg A \supset A$ *for arbitrary* A.

Proof: Assume we have the shortest proof of $\implies \neg\neg A \supset A$. There is only one rule that could have been applied ($\supset R$), so we must also have a proof of $\neg\neg A \implies A$. Again, only one rule could have been applied,

$$\frac{\neg\neg A \implies \neg A \quad \neg\neg A, \bot \implies A}{\neg\neg A \implies A} \supset L$$

We can prove the second premise, but not the first. If such a proof existed, it must end either with the $\supset R$ or $\supset L$ rules.

Case: The proof proceeds with $\supset R$.

$$\frac{\neg\neg A, A \implies \bot}{\neg\neg A \implies \neg A} \supset R$$

Now only $\supset L$ could have been applied, and it premises must have been

$$\frac{\neg\neg A, A \implies \neg A \quad \neg\neg A, A, \bot \implies \bot}{\neg\neg A, A \implies \bot} \supset L$$

Again, the second premise could have been deduced, but not the first. If it had been inferred with $\supset R$ and, due to contraction, we would end up with another proof of a sequent we have already seen, and similarly if $\supset L$ had been used. In either case, it would contradict the assumption of starting with a shortest proof.

Case: The proof proceeded with $\supset L$.

$$\frac{\neg\neg A \Longrightarrow \neg A \quad \neg\neg A, \bot \Longrightarrow \neg A}{\neg\neg A \Longrightarrow \neg A} \supset L$$

The first premise is identical to the conclusion, so if there were a deduction of that, there would be one without this rule, which is a contradiction to the assumption that we started with the shortest deduction.

$\square$

4 Identity

We permit the init rule only for atomic propositions. However, the version of this rule with arbitrary propositions A is *admissible*, that is, each instance of the rule can be deduced. We call this the *identity theorem* because it shows that from an assumption A we can prove the identical conclusion A.

Theorem 7 (Identity) *For any proposition A, we have $A \Longrightarrow A$.*

Proof: By induction on the structure of A. We show several representative cases and leave the remaining ones to the reader.

Case: $A = P$. Then

$$\frac{}{P \Longrightarrow P} \text{ init}$$

Case: $A = A_1 \wedge A_2$. Then

By i.h. on A_1 and weakening $\quad$ By i.h. on A_2 and weakening

$$\frac{\dfrac{A_1 \wedge A_2, A_1 \Longrightarrow A_1}{A_1 \wedge A_2 \Longrightarrow A_1} \wedge L_1 \qquad \dfrac{A_1 \wedge A_2, A_2 \Longrightarrow A_2}{A_1 \wedge A_2 \Longrightarrow A_2} \wedge L_2}{A_1 \wedge A_2 \Longrightarrow A_1 \wedge A_2} \wedge R$$

Case: $A = A_1 \supset A_2$. Then

By i.h. on A_1 and weakening $\quad$ By i.h. on A_2 and weakening

$$\frac{\dfrac{\dfrac{A_1 \supset A_2, A_1 \Longrightarrow A_1 \qquad A_1 \supset A_2, A_1, A_2 \Longrightarrow A_2}{A_1 \supset A_2, A_1 \Longrightarrow A_2} \supset L}{A_1 \supset A_2 \Longrightarrow A_1 \supset A_2} \supset R}{}$$

Case: $A = \perp$. Then

$$\overline{\perp \Longrightarrow \perp} \; {\perp}L$$

$\square$

The identity theorem is the global version of the local completeness property for each individual connective. One can recognize the local expansion as embodied in each case of the inductive proof of identity.

In the next lecture we will see a new theorem, called *cut*, which is the global analogue of local soundness.

References

[Gen35] Gerhard Gentzen. Untersuchungen über das logische Schließen. *Mathematische Zeitschrift*, 39:176–210, 405–431, 1935. English translation in M. E. Szabo, editor, *The Collected Papers of Gerhard Gentzen*, pages 68–131, North-Holland, 1969.

[Kle52] Stephen Cole Kleene. *Introduction to Metamathematics*. North-Holland, 1952.