Lecture Notes on Sequent Calculus

15-317: Constructive Logic
Frank Pfenning*

Lecture 9
March 11, 2021

1 Introduction

In this lecture we shift to a different presentation style for proof calculi. We develop the sequent calculus as a formal system for proof search in natural deduction. In addition to enabling an understanding of proof search, sequent calculus leads to a more transparent management of the scope of assumptions during a proof than two-dimensional natural deduction, and also allows us more proof theory, so proofs about properties of proofs.

Sequent calculus was originally introduced by Gentzen [Gen35], primarily as a technical device for proving consistency of predicate logic, but turned out to be quite influential for other purposes subsequently. 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 5. Recall that (following the directions of their arrows) 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 a proposition we have deduced from assumptions may be used as a verification. In sequent calculus, both steps work bottom-up, which will ultimately allows us to prove the *global* versions of local soundness and completeness properties that we wished we could show for verifications/uses.

*Edits by André Platzer

2 Sequents

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

$$\begin{array}{ccc} A_1{\downarrow} & \cdots & A_n{\downarrow} \\ & \vdots & \\ & C{\uparrow} & \end{array}$$

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

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

where the judgments $A\ \mathsf{left}$ and $C\ \mathsf{right}$ 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 we are trying to verify called the *succedent*. Sequent calculus is explicit about the assumptions that are available for use (antecedent) and about the proposition to be verified (succedent).

The rules that define the $A\ \mathsf{left}$ and $A\ \mathsf{right}$ 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:

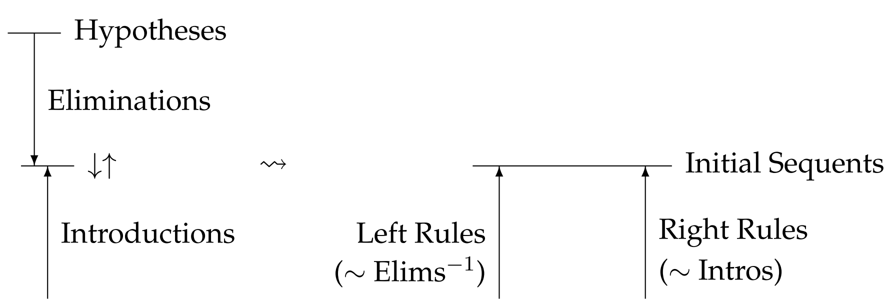

We now proceed connective by connective, constructing the right and left rules from the introduction and elimination rules. When writing a sequent, we can always tell which propositions are on the left of $\Longrightarrow$ and which are on the right of $\Longrightarrow$, so we omit the judgments left and right for brevity. Also, we abbreviate a collection of a natural number of antecedents

A_1 left, ..., A_n left by Γ. The order of the antecedents does not matter, because it does not matter in which order we write down our available assumptions, so we will allow them to be implicitly reordered.

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

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

The only difference is that the antecedents Γ are made explicit. Unlike $\wedge I$, rule $\wedge R$ is explicit about the fact that both premises have the same antecedents, because any assumption can be used in both subdeductions. A proof of $A \wedge B$ from assumptions Γ consists of a proof of A from assumptions Γ together with a proof of B from the same assumptions Γ.

There are two elimination rules, so also two corresponding left rules.

$$\frac{A \wedge B \downarrow}{A \downarrow} \wedge E_1 \qquad \frac{\Gamma, A \wedge B, A \Longrightarrow C}{\Gamma, A \wedge B \Longrightarrow C} \wedge L_1$$

$$\frac{A \wedge B \downarrow}{B \downarrow} \wedge E_2 \qquad \frac{\Gamma, A \wedge B, B \Longrightarrow C}{\Gamma, A \wedge B \Longrightarrow C} \wedge L_2$$

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. But keeping $A \wedge B$ in $\wedge L_i$ is important, because dropping it would make it impossible to read off both A and B from an assumption $A \wedge B$.

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

$$\dfrac{\begin{array}{c}\overline{A{\downarrow}}\ u\\ \vdots\\ B{\uparrow}\end{array}}{A \supset B{\uparrow}}\ {\supset}^u \qquad \dfrac{\Gamma, A \Longrightarrow B}{\Gamma \Longrightarrow A \supset B}\ {\supset}R$$

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 rules so that antecedents are *persistent*: they are always propagated from the conclusion to all premises, because all assumptions are still available. A proof of $A \supset B$ from assumptions Γ consists of a proof of B from Γ, B.

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

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

In words: in order to use $A \supset B$ to verify C we have to produce a verification of A, in which case we can then use B in the verification of C. The antecedent $A \supset B$ is carried over to both premises to maintain persistence. While irrelevant (since all premises need to be proved for a proof to succeed), the order of premises emphasizes 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.

$$\dfrac{A{\uparrow}}{A \vee B{\uparrow}}\ {\vee}I_1 \qquad \dfrac{\Gamma \Longrightarrow A}{\Gamma \Longrightarrow A \vee B}\ {\vee}R_1$$

$$\dfrac{B{\uparrow}}{A \vee B{\uparrow}}\ {\vee}I_2 \qquad \dfrac{\Gamma \Longrightarrow B}{\Gamma \Longrightarrow A \vee B}\ {\vee}R_2$$

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 in sequent calculus actually has a simpler form that is more consistent with all the other rules we have shown so far.

$$\dfrac{A \lor B{\downarrow} \quad \begin{matrix} \overline{A{\downarrow}}^{\,u} \\ \vdots \\ C{\uparrow} \end{matrix} \quad \begin{matrix} \overline{B{\downarrow}}^{\,w} \\ \vdots \\ C{\uparrow} \end{matrix}}{C{\uparrow}} \lor E^{u,w} \qquad \dfrac{\Gamma, A \lor B, A \Longrightarrow C \quad \Gamma, A \lor B, B \Longrightarrow C}{\Gamma, A \lor B \Longrightarrow C} \lor L$$

As for implication, scoping issues are more explicit and simplified because the new assumptions A and B in the first and second premise, respectively, are available anywhere in their respective deductions above. But the assumption A is only available in the deduction for the left premise, while assumption B is only available in the right premise. Sequent calculus is explicit about that. The sequent calculus formulation also makes it more transparent what the appropriate verification/uses assignment is.

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 take inspiration from the 0-ary case of disjunction above.

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

Completing verifications. 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 atomic P is directly via a use of P. In the version of verifications we have presented, we can complete the construction of a verification whenever $A\downarrow$ is available to conclude $A\uparrow$.[1] This turns into a so-called *initial sequent* or application of the *identity rule*.

$$\dfrac{A \downarrow}{A \uparrow} {\downarrow\uparrow} \qquad \dfrac{}{\Gamma, A \Longrightarrow A} \,\mathsf{id}$$

This rule has a special status in that it does not decompose any proposition, but establishes a connection between two judgments. In natural deduction,

[1] A stricter version of this rule only allows the use of $\downarrow\uparrow$ rule for atomic propositions P. We use the name init for the sequent rule id restricted to atomic A. See next lectures.

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 $(A \vee B) \supset (B \vee A)$.

$$\dfrac{\dfrac{\dfrac{}{A \vee B, A \Longrightarrow A}\ \mathsf{id}}{A \vee B, A \Longrightarrow B \vee A}\ \vee R_2 \qquad \dfrac{\dfrac{}{A \vee B, B \Longrightarrow B}\ \mathsf{id}}{A \vee B, B \Longrightarrow B \vee A}\ \vee R_1}{\dfrac{A \vee B \Longrightarrow B \vee A}{\Longrightarrow (A \vee B) \supset (B \vee A)}\ \supset R}\ \vee L$$

Unlike in natural deduction, sequent calculus proofs are always *constructed bottom-up*, with the desired conclusion at the bottom, working upwards using the respective left or right proof rules in the antecedent or succedent.

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* $\Longrightarrow \bot$.

Proof: No left rule is applicable, since there is no antecedent. No right rule is applicable, because there is no right rule for falsehood. The identity rule is not applicable either. Therefore, there cannot be a proof of $\Longrightarrow \bot$, because this exhausted all proof rules matching on the shape $\Longrightarrow \bot$. □

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

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

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

Proof: From the disjunction property, either $\Longrightarrow A$ or $\Longrightarrow \neg A$. For the first sequent, no rule applies. For the second sequent, only $\supset R$ applies and we would have to have a deduction of $A \Longrightarrow \bot$. But for this sequent no rule applies, because it does not have the form of id and $\bot$ has no right rule. □

Of course, there are still specific formulas A for which $\Longrightarrow A \vee \neg A$ will be provable, such as $\Longrightarrow \top \vee \neg\top$ or $\Longrightarrow \bot \vee \neg\bot$, but not generally for any A.

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* $\Gamma \Longrightarrow C$ *then* $\Gamma, A \Longrightarrow C$ *with a structurally identical deduction.*

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

Theorem 5 (Contraction) *If* $\Gamma, A, A \Longrightarrow C$ *then* $\Gamma, A \Longrightarrow C$ *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. □

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. Theorem 5 shows that there is no need to keep multiple copies of the same assumption, so labels are not crucial.

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* $\Longrightarrow \neg\neg A \supset A$ *for arbitrary* A.

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

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

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, because these are the only applicable rules.

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

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

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

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

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.

□

4 Optimizations

We will devote significant attention to "optimizations" of the sequent calculus where we eliminate redundancies while preserving the same set of theorems. One form of redundancy arises if one antecedent of the premise of the rules is never needed to prove the succedent. For example, the rule

$$\frac{\Gamma, A \vee B, A \Longrightarrow C \quad \Gamma, A \vee B, B \Longrightarrow C}{\Gamma, A \vee B \Longrightarrow C} \vee L$$

has a redundant antecedent $A \vee B$ in both premises. In the first premise, for example, we also have A and A is stronger than $A \vee B$ (in the sense that $A \supset (A \vee B)$). In the second premise, it is B which is stronger than $A \vee B$.

In the following summary of sequent calculus we put [brackets] around the antecedents that could be considered redundant *after optimization*.

$$\frac{}{\Gamma, A \Longrightarrow A} \mathsf{id}$$

$$\frac{\Gamma \Longrightarrow A \quad \Gamma \Longrightarrow B}{\Gamma \Longrightarrow A \wedge B} \wedge R \qquad \frac{\Gamma, A \wedge B, A \Longrightarrow C}{\Gamma, A \wedge B \Longrightarrow C} \wedge L_1 \quad \frac{\Gamma, A \wedge B, B \Longrightarrow C}{\Gamma, A \wedge B \Longrightarrow C} \wedge L_2$$

$$\frac{}{\Gamma \Longrightarrow \top} \top R \qquad \text{no rule } \top L$$

$$\frac{\Gamma, A \Longrightarrow B}{\Gamma \Longrightarrow A \supset B} \supset R \qquad \frac{\Gamma, A \supset B \Longrightarrow A \quad \Gamma, [A \supset B], B \Longrightarrow C}{\Gamma, A \supset B \Longrightarrow C} \supset L$$

$$\frac{\Gamma \Longrightarrow A}{\Gamma \Longrightarrow A \vee B} \vee R_1 \quad \frac{\Gamma \Longrightarrow B}{\Gamma \Longrightarrow A \vee B} \vee R_2 \qquad \frac{\Gamma, [A \vee B], A \Longrightarrow C \quad \Gamma, [A \vee B], B \Longrightarrow C}{\Gamma, A \vee B \Longrightarrow C} \vee L$$

$$\text{no rule } \bot R \qquad \frac{}{\Gamma, \bot \Longrightarrow C} \bot L$$

We could also replace the two left rules $\wedge L_1$ and $\wedge L_2$ for conjunction with

$$\frac{\Gamma, [A \wedge B], A, B \Longrightarrow C}{\Gamma, A \wedge B \Longrightarrow C} \wedge L$$

5 Classical Sequent Calculus

One of Gentzen's remarkable discoveries was the encoding of *classical logic* in sequent calculus. In natural deduction, classical logic can be incorporated by the law of excluded middle, by double negation elimination, or by the rule of indirect proof. All of these are clearly outside the simple beauty of the natural deduction rules as defined by introductions and eliminations.

How do we obtain sequent calculus for classical logic? Simply by allowing a sequent to have multiple conclusions! A sequent then has the form $\Gamma \stackrel{\mathsf{CL}}{\Longrightarrow} \Delta$, where Δ is also a collection of propositions. Now succedents as well as antecedents in the rules are *persistent* in all the rules. Remarkably, this is all we need to do!

We can then prove the law of excluded middle as follows, remembering that $\neg A \triangleq A \supset \bot$:

$$\dfrac{\dfrac{\dfrac{\dfrac{}{A \stackrel{\mathsf{CL}}{\Longrightarrow} A \vee \neg A, A, \neg A}\,\mathsf{id}}{\stackrel{\mathsf{CL}}{\Longrightarrow} A \vee \neg A, A, \neg A}\,{\supset}R}{\stackrel{\mathsf{CL}}{\Longrightarrow} A \vee \neg A, A}\,{\vee}R_2}{\stackrel{\mathsf{CL}}{\Longrightarrow} A \vee \neg A}\,{\vee}R_1$$

Somehow, by allowing us to "hedge our bets" about which disjunct is true (first we say "A" using $\vee R_1$, then we later say "$\neg A$" using $\vee R_2$) and then using the second possibility to establish the first we have circumvented the usual constructive nature of the disjunction.

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.