

Guide to Proofs on Discrete Structures

In Problem Set One, you got practice with the art of proofwriting in general (as applied to numbers, sets, puzzles, etc.) Problem Set Two introduced first-order logic and gave you some practice writing more intricate proofs than before. Now that you've hit Problem Set Three, you'll be combining these ideas together. Specifically, you'll be writing a lot of proofs, as before, but those proofs will often be about terms and definitions that are specified rigorously in the language of first-order logic.

The good news is that all the general proofwriting rules and principles you've been honing over the first two problem sets still apply. The new skill you'll need to sharpen for this problem set is determining how to start with a statement like "prove that R is an equivalence relation" and to figure out what exactly it is that you're supposed to be doing. This takes a little practice to get used to, but once you've gotten a handle on things we think you'll find that it's not nearly as tricky as it might seem.

This handout illustrates how to set up proofs of a number of different types of results. While you're welcome to just steal these proof setups for your own use, we recommend that you focus more on the process by which these templates were developed and the context for determining how to proceed. As you'll see when we get to Problem Set Four and beyond, you'll often be given new first-order definitions and asked to prove things about them.

The Big Table

The table given below summarizes each of the first-order logic quantifiers and connectives and how to approach proving each of them. We'll use the concepts here to derive the templates included elsewhere in this handout.

<i>Statement Form</i>	<i>Proof Approach</i>
$\forall x. P$	<p>Direct proof: Pick an arbitrary x, then prove P is true for that choice of x.</p> <p>By contradiction: Suppose for the sake of contradiction that there is some x where P is false. Then derive a contradiction.</p>
$\exists x. P$	<p>Direct proof: Do some exploring and find a choice of x where P is true. Then, write a proof explaining why P is true in that case.</p> <p>By contradiction: Suppose for the sake of contradiction that P is always false and derive a contradiction.</p>
$\neg P$	<p>Direct proof: Simplify your formula by pushing the negation deeper, then apply the appropriate rule.</p> <p>By contradiction: Suppose for the sake of contradiction that P is true, then derive a contradiction.</p>
$P \wedge Q$	<p>Direct proof: Prove each of P and Q independently.</p> <p>By contradiction: Assume $\neg P \vee \neg Q$. Then, try to derive a contradiction.</p>
$P \vee Q$	<p>Direct proof: Prove that $\neg P \rightarrow Q$, or prove that $\neg Q \rightarrow P$. (<i>Great question to ponder: why does this work?</i>)</p> <p>By contradiction: Assume $\neg P \wedge \neg Q$. Then, try to derive a contradiction.</p>
$P \rightarrow Q$	<p>Direct proof: Assume P is true, then prove Q.</p> <p>By contradiction: Assume P is true and Q is false, then derive a contradiction.</p> <p>By contrapositive: Assume $\neg Q$, then prove $\neg P$.</p>
$P \leftrightarrow Q$	Prove both $P \rightarrow Q$ and $Q \rightarrow P$.

Proof Template: Equivalence Relations

Equivalence relations are one of the more common classes of binary relations, and there's a good chance that going forward, you're going to find equivalence relations "in the wild."

Let's imagine that you have a binary relation R over a set A and you want to prove that R is an equivalence relation. How exactly should you go about doing this? As is almost always the case when writing formal proofs, let's return back to the definition to see what we find. Formally speaking, an equivalence relation is a binary relation that is reflexive, symmetric, and transitive. That means that if we want to prove that R is an equivalence relation, we'd likely write something like this:

Proof: We will prove that R is reflexive, symmetric, and transitive.

First, we'll prove that R is reflexive. [...]

Next, we'll prove that R is symmetric. [...]

Finally, we'll prove that R is transitive. [...] ■

Each of those subsequent steps – proving reflexivity, symmetry, and transitivity – is essentially a mini-proof in of itself. In lecture, when we worked with the example of the relation \sim over the set \mathbb{Z} , we split this apart into three separate lemmas. You can do that if you'd like, though it's not strictly necessary, and for some shorter proofs bundling everything together into one larger proof with clearly-demarcated sections is perfectly acceptable.

The above proof template shows how you'd proceed if you found yourself getting to a point in a proof where you needed to show that a particular binary relation is an equivalence relation. There are several different ways that this could happen. First, you might be given a specific relation and then be asked to prove something about it. For example, you might come across a problem like this one:

Let R be a binary relation over \mathbb{Z} where xRy holds if $x^2 = y^2$. Prove that R is an equivalence relation.

In this case, you've been handed a concrete relation R , and the task is to show that it's an equivalence relation. A proof of this result might look like this:

Proof: We will prove that R is reflexive, symmetric, and transitive.

First, we'll prove that R is reflexive. [...]

Next, we'll prove that R is symmetric. [...]

Finally, we'll prove that R is transitive. [...] ■

Something to note about the above proof: notice that we didn't repeat the definition of the relation R at the top of the proof before then proceeding to show that R is reflexive, symmetric, and transitive. This follows a general convention: if you're proving something about an object that's explicitly defined somewhere, you should not repeat the definition of that object in the proof. Remember – don't *restate* definitions; *use them* instead (see the Discrete Structures Proofwriting Checklist for more information).

As a great exercise, go and fill in the blank spots in the above proof!

On the other hand, you might be asked to show that an arbitrary relation R satisfying some criteria is an equivalence relation. For example, consider this problem, which we solved in lecture:

Prove that if R is reflexive and cyclic, then R is an equivalence relation.

In this case, we should structure our proof like this:

Proof: Consider an arbitrary binary relation R over a set A that is reflexive and cyclic. We will prove that R is an equivalence relation. To do so, we will show that R is reflexive, symmetric, and transitive.

First, we'll prove that R is reflexive. [...]

Next, we'll prove that R is symmetric. [...]

Finally, we'll prove that R is transitive. [...] ■

Notice that in this case, we had to explicitly introduce what the relation R is, since we're trying to prove a universally-quantified statement ("if R is reflexive and cyclic, then ...") and there isn't a concrete choice of R lying around for us to use.

Proof Template: Transitivity

Suppose you have a binary relation R over a set A . To prove that R is transitive, you need to show that

$$\forall x \in A. \forall y \in A. \forall z \in A. (xRy \wedge yRz \rightarrow xRz).$$

Remember our first guiding principle: if you want to prove that a statement is true and that statement is specified in first-order logic, look at the structure of that statement to see how to structure the proof. So let's look at the above formula. From the structure of this formula, we can see that

- we *pick* arbitrary elements x, y , and z from A (since these variables are *universally-quantified*);
- we *assume* that xRy and yRz are true (since they're the *antecedent of an implication*); and then
- we *prove* that xRz is true (since it's the *consequent of an implication*).

This means that a proof of transitivity will likely start off like this:

Consider some arbitrary $x, y, z \in A$ where xRy and yRz . We need to prove that xRz . [...]

At this point, the structure of the proof will depend on how R is defined. If you have a concrete definition of R lying around, you would likely go about expanding out that definition. For example, consider the relation S defined over \mathbb{Z} as follows:

$$xSy \quad \text{if} \quad \exists k \in \mathbb{N}. (x + k = y)$$

Suppose we wanted to prove that S is transitive. To do so, we might write something like this:

Proof: Consider some arbitrary integers x, y, z where xSy and ySz . We need to prove that xSz . Since we know that xSy and ySz , we know that there exist natural numbers m and n where $x + m = y$ and $y + n = z$. [...]

Notice that we did *not* repeat the definition of the relation S in this proof. The convention is that if you have some object that's explicitly defined outside of a proof (as S is defined here), then you don't need to – and in fact, shouldn't – redefine it inside the proof. Also, although the relation S is defined in first-order logic, note that there is no first-order logic anywhere in this proof, not even in the part where xSy and ySz are expanded into their longer equivalents.

Alternatively, it might be the case that you don't actually have a concrete definition of R lying around and only know that R happens to satisfy some other properties. For example, let's return to the definition of a cyclic relation that we introduced in lecture. As a reminder, a relation R is called *cyclic* if the following first-order formula is true:

$$\forall x \in A. \forall y \in A. \forall z \in A. (xRy \wedge yRz \rightarrow zRx).$$

So let's imagine that we know that R is cyclic (and possibly satisfies some other properties) and we want to show that R is transitive. In that case, rather than expanding the definition of R (since we don't have one), we might go about applying the existing rules we have. That might look like this, for example:

Consider some arbitrary $x, y, z \in A$ where xRy and yRz . We need to prove that xRz . Since R is cyclic, from xRy and yRz we learn that zRx . [...]

Notice how we invoked the cyclic property: we didn't write out the general definition of a cyclic relation and then argue that it applies here and instead just said how we could use the statements we began with, plus the cyclic property, to learn something new.

Proof Template: Reflexivity

Let's suppose you have a binary relation R over a set A . If you want to prove that R is reflexive, you need to prove that the following statement is true:

$$\forall x \in A. xRx.$$

Following the general rule of "match the proof to the first-order definition," this means that in our proof,

- we *pick* an arbitrary element x from A (since this variable is *universally-quantified*), then
- we *prove* that xRx is true (since it's *the statement that's being quantified*).

A subtle but important point here: notice that in our proof of reflexivity, we don't actually make any assumptions about x other than the fact that it's an element of A . This contrasts with transitivity, where we would make some extra assumptions about the variables we'd chosen because the statement we were interested in proving was an implication. (*Make sure you see why this is!*)

As a result, a proof of reflexivity will often start off with this:

Pick any $x \in A$. We need to prove that xRx . [...]

As before, the next steps in this proof are going to depend entirely about what we know about the relation R . If we have a concrete definition of R in hand (say, if the proof is taking an existing relation that we've defined and then showing that it's an equivalence relation), then we'd need to look at the definition of R to see what to do next. Let's go back to our example of the relation S defined over the set \mathbb{Z} that we defined earlier:

$$xSy \quad \text{if} \quad \exists k \in \mathbb{N}. (x + k = y)$$

If we wanted to prove that S is reflexive, we might do something like this:

Consider some integer x . We need to prove that xSx , meaning that we need to show that there's a natural number k such that $x + k = x$. [...]

Take a minute to see why we've structured the proof this way.

Alternatively, it might be the case that you don't actually have a concrete definition of R lying around and only know that R happens to satisfy some other properties. In that case, you'd have to use what you knew about the relation R in order to establish that xRx has to hold.

Proof Template: Asymmetry

Suppose you have a binary relation R over a set A and you want to prove that R is asymmetric. This means you need to prove that

$$\forall x \in A. \forall y \in A. (xRy \rightarrow \neg yRx).$$

As usual, we're going to prove this one by looking at the first-order logic statement above to determine what we need to show. The pattern will look like this:

- we **pick** an arbitrary elements x and y from A (since these variables are **universally-quantified**);
- we **assume** that xRy holds (since it's **the antecedent of an implication**), then
- we **prove** that $\neg yRx$ is true (since it's **the consequent of our implication**).

In many ways, this proof structure is similar to that for transitivity – we've picked some arbitrarily-chosen values, made some assumptions about how they relate in some way, and need to show that they relate in some other way as well (here, that yRx holds, or, equivalently, that $\neg yRx$ isn't true.) That would give us a proof template like this one:

Let x and y be arbitrary elements of A where xRy holds. We need to prove that $\neg yRx$ does not hold. [...]

This, of course, isn't the only way that you could prove that R is asymmetric. Another option would be to write this as proof by contradiction. If you wanted to do that, you'd **assume** the negation of the above first-order logic statement, which looks like this:

$$\exists x \in A. \exists y \in A. (xRy \wedge yRx).$$

If we're **assuming** that this statement is true, it means that we're assuming

- there are elements x and $y \in A$ (since we're **assuming an existentially-quantified statement**) where
- both xRy and yRx are true (since we're **assuming a conjunction (and) of two statements**).

That would mean that our proof would start off like this:

Assume for the sake of contradiction that R is not asymmetric. This means that there must be some elements $x, y \in A$ where both xRy and yRx are true. [...]

From here, we'd be off to the races trying to reach a contradiction.

Proof Template: Not Symmetric

In the previous examples we've seen, we've tried to prove that a relation *does* have some particular property. What happens if you want to show that a relation *does not* have some particular property?

For example, let's suppose R is a binary relation over A and we want to prove that R is not symmetric. For starters, remember that “not symmetric” and “asymmetric” do *not* mean the same thing! Proving that a relation is not symmetric is quite different from proving that it's asymmetric. To see why this is, remember that if we want to prove that a relation is not symmetric, we'd need to prove that the following first-order logic formula is not true:

$$\forall x \in A. \forall y \in A. (xRy \rightarrow yRx).$$

Equivalently, we'd want to prove that the negation of the above formula is true. The negation of that formula is

$$\exists x \in A. \exists y \in A. (xRy \wedge y\not Rx).$$

If we want to prove that this formula is true, we would show that

- **there are choices of x and y** (since we're proving an existentially-quantified formula) where
- **xRy is true and yRx is false** (since we're proving the conjunction (and) of these statements).

There are many ways you could write up a proof of this form. For example, you could write something like this:

Let $x = [\dots]$ and $y = [\dots]$. Then xRy is true because $[\dots]$ and yRx is false because $[\dots]$.

This one is pretty simple and to the point. We need to show that a certain x and y exist satisfying some criteria, so we just say what choices we're going to make and justify why those choices are good ones. We could also be even more direct than this and omit the names x and y . For example:

Notice that $[\dots]R[\dots]$ is true because $[\dots]$, but that $[\dots]R[\dots]$ because $[\dots]$.

The hard part of writing a proof like this is probably going to be coming up with what choices of objects you're going to pick in the first place rather than writing down your reasoning.

Some Exercises: Other Types of Proofs

In the preceding sections, we went over some examples of how you'd structure proofs of different properties of binary relations. However, we didn't cover all the different properties you might want to consider. That's intentional – ultimately, it's more important that you understand how to look at a first-order statement and determine what it is that you need to prove than it is to approach these problems as an exercise in pattern-matching.

Here are a few exercises that you may want to work through as a way of seeing whether you've internalized the concepts from this handout.

- What would a good proof template be for proving symmetry? How about reflexivity?
- Suppose $f: A \rightarrow B$ is a function. Give **two** proof templates for showing that f is injective.
- Suppose $f: A \rightarrow B$ is a function. Give a proof template for showing that f is not surjective.
- A binary relation R over a set A is called **functional** if the following is true:

$$\forall x \in A. \forall y \in A. \forall z \in A. (xRy \wedge xRz \rightarrow y = z).$$

Write a proof template for showing that a relation R is functional.

- A function $f: A \rightarrow A$ is called an **involution** if the following is true:

$$\forall a \in A. f(f(a)) = a.$$

Write a proof template for showing that a function f is an involution.