The Knaster–Tarski Theorem

Transfinite Ordinals

Everyone is familiar with the set $\omega = \{0, 1, 2, \ldots\}$ of *finite ordinals*, also known as the *natural numbers*. An essential mathematical tool is the *induction principle* on this set, which states that if a property is true of zero and is preserved by the successor operation, then it is true of all elements of ω.

In theoretical computer science, we often run into inductive definitions that take longer than ω to close, and it is useful to have an induction principle that applies to these objects. Cantor recognized the value of such a principle in his theory of infinite sets. Any modern account of the foundations of mathematics will include a chapter on ordinals and transfinite induction.

Unfortunately, a complete understanding of ordinals and transfinite induction is impossible outside the context of set theory, since many issues impact the very foundations of the subject. Here we will only give a cursory account of the basic facts, tools, and techniques we will need.

Set-Theoretic Definition of Ordinals

Ordinals are defined as certain sets of sets. The key facts we will need about ordinals, succinctly stated, are:

(i) There are two kinds: *successors* and *limits*.

(ii) They are well-ordered.

(iii) There are a lot of them.

(iv) We can do induction on them.

We will explain each of these statements in more detail below.

A set C of sets is said to be *transitive* if $A \in C$ whenever $A \in B$ and $B \in C$. Equivalently, C is transitive if every element of C is a subset of C; that is, $C \subseteq 2^C$. Formally, an *ordinal* is defined to be a set A such that

- A is transitive
- all elements of A are transitive.

It follows that any element of an ordinal is an ordinal. We use $\alpha, \beta, \gamma, \ldots$ to refer to ordinals. The class of all ordinals is denoted **Ord**. It is not a set, but a proper class.

This neat but rather obscure definition of ordinals has some far reaching consequences that are not at all obvious. For ordinals α, β, define $\alpha < \beta$ if $\alpha \in \beta$. The relation $<$ is a strict partial order. As usual, there is an associated nonstrict partial order $\leq$ defined by $\alpha \leq \beta$ if $\alpha \in \beta$ or $\alpha = \beta$.

It follows from the axioms of set theory that the relation $<$ on ordinals is a linear order. That is, if α and β are any two ordinals, then either $\alpha < \beta$, $\alpha = \beta$, or $\alpha > \beta$. This is most easily proved by induction on the well-founded relation

$$(\alpha, \beta) \leq (\alpha', \beta') \stackrel{\text{def}}{\iff} \alpha \leq \alpha' \text{ and } \beta \leq \beta'.$$

Then every ordinal is equal to the set of all smaller ordinals (in the sense of $<$). The class of ordinals is well-founded in the sense that any nonempty set of ordinals has a least element.

If α is an ordinal, then so is $\alpha \cup \{\alpha\}$. The latter is called the *successor* of α and is denoted $\alpha + 1$. Also, if A is any set of ordinals, then $\bigcup A$ is an ordinal, and is the supremum of the ordinals in A under the relation $\leq$.

The smallest few ordinals are

$$\begin{aligned} 0 &\stackrel{\text{def}}{=} \varnothing \\ 1 &\stackrel{\text{def}}{=} \{0\} = \{\varnothing\} \\ 2 &\stackrel{\text{def}}{=} \{0,1\} = \{\varnothing, \{\varnothing\}\} \\ 3 &\stackrel{\text{def}}{=} \{0,1,2\} = \{\varnothing, \{\varnothing\}, \{\varnothing, \{\varnothing\}\}\} \end{aligned}$$

Pictorially,

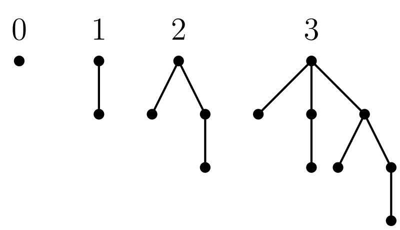

The first infinite ordinal is

$$\omega \stackrel{\text{def}}{=} \{0, 1, 2, 3, \ldots\}.$$

An ordinal is called a *successor ordinal* if it is of the form $\alpha + 1$ for some ordinal α, otherwise it is called a *limit ordinal*. The smallest limit ordinal is 0 and the next smallest is ω. Of course, $\omega + 1 = \omega \cup \{\omega\}$ is an ordinal, so it doesn't stop there.

Since the ordinals form a proper class, there can be no one-to-one function $\mathbf{Ord} \to A$ into a set A. This is what we mean above by, "There are a lot of ordinals." In practice, this will come up when we construct functions $f : \mathbf{Ord} \to A$ from $\mathbf{Ord}$ into a set A by induction. Such an f, regarded as a collection of ordered pairs, is necessarily a class and not a set. We will always be able to conclude that there exist distinct ordinals α, β with $f(\alpha) = f(\beta)$.

Transfinite Induction

The *transfinite induction principle* is a method of establishing that a particular property is true of all ordinals (or of all elements of a class of objects indexed by ordinals). It states that in order to prove that the property is true of all ordinals, it suffices to show that the property is true of an arbitrary ordinal α whenever it is true of all ordinals $\beta < \alpha$. Proofs by transfinite induction typically contain two cases, one for successor ordinals and one for limit ordinals. The basis of the induction is often a special case of the case for limit ordinals, since $0 = \varnothing$ is a limit ordinal; here the premise that the property is true of all ordinals $\beta < \alpha$ is vacuously true.

The validity of this principle ultimately follows from the well-foundedness of the set containment relation $\in$. This is an axiom of set theory.

Zorn's Lemma and the Axiom of Choice

Related to the ordinals and transfinite induction are the *axiom of choice* and *Zorn's lemma*.

The *axiom of choice* states that for every set A of nonempty sets, there exists a function f with domain A that picks an element out of each set in A; *i.e.*, for every $B \in A$, $f(B) \in B$. Equivalently, any Cartesian product of nonempty sets is nonempty.

Zorn's lemma states that every set of sets closed under unions of chains contains a $\subseteq$-maximal element. Here a *chain* is a family of sets linearly ordered by the inclusion relation $\subseteq$, and to say that a set C of sets is closed under unions of chains means that if $B \subseteq C$ and B is a chain, then $\bigcup B \in C$. An element $B \in C$ is $\subseteq$*-maximal* if it is not properly included in any $B' \in C$.

The *well ordering principle* states that every set is in one-to-one correspondence with some ordinal. A set is *countable* if it is either finite or in one-to-one correspondence with ω.

The axiom of choice, Zorn's lemma, and the well ordering principle are equivalent to one another and *independent* of Zermelo–Fraenkel (ZF) set theory in the sense that if ZF is consistent, then neither they nor their negations can be proven from the axioms of ZF.

Complete Partial Orders

A *complete partial order* is a set U with a distinguished partial ordering relation $\leq$ defined on it (reflexive, antisymmetric, transitive) such that every subset of U has a *supremum* or *least upper bound* with respect to $\leq$. That is, for every subset $A \subseteq U$, there is an element $\sup A \in U$ with the properties

(i) for all $x \in A$, $x \leq \sup A$ ($\sup A$ is an upper bound for A),

(ii) if for all $x \in A$, $x \leq y$, then $\sup A \leq y$ ($\sup A$ is the least upper bound).

It follows from (i) and (ii) that $\sup A$ is unique. We abbreviate $\sup\{x, y\}$ by $x \vee y$.

Any complete partial order U has a maximum element $\top \stackrel{\text{def}}{=} \sup U$ and a minimum element $\bot \stackrel{\text{def}}{=} \sup \varnothing$. Also, every subset $A \subseteq U$ has an *infimum* or *greatest lower bound* $\inf A \stackrel{\text{def}}{=} \sup\{y \mid \forall z \in A\ y \leq z\}$. One can show that $\inf A$ is unique and satisfies the properties

(i) for all $y \in A$, $\inf A \leq y$ ($\inf A$ is a lower bound for A),

(ii) if for all $y \in A$, $x \leq y$, then $x \leq \inf A$ ($\inf A$ is the greatest lower bound).

(Miscellaneous Exercise **??**).

A common example of a complete partial order is the powerset 2^X of a set X, or set of all subsets of X, ordered by the subset relation $\subseteq$. The supremum of a set $\mathcal{C}$ of subsets of X is their union $\bigcup \mathcal{C}$ and the infimum of $\mathcal{C}$ is their intersection $\bigcap \mathcal{C}$.

Monotone, Continuous, and Finitary Operators

An *operator* on a complete partial order U is a function $\tau : U \to U$. Here we introduce some special properties of such operators such as monotonicity and closure and discuss some of their consequences. We culminate with a general theorem due to Knaster and Tarski concerning inductive definitions.

In the special case of set-theoretic complete partial orders 2^X ordered by set inclusion $\subseteq$, we call such an operator a *set operator*.

An operator τ is said to be *monotone* if it preserves $\leq$:

$$x \leq y \quad \Rightarrow \quad \tau(x) \leq \tau(y).$$

A *chain* in U is a subset of U totally ordered by $\leq$; that is, for every x and y in the chain, either $x \leq y$ or $y \leq x$. An operator τ is said to be *chain-continuous* if for every chain

A,

$$\tau(\sup A) \;=\; \sup_{x\in A} \tau(x).$$

For set operators $\tau : 2^X \to 2^X$, τ is said to be *finitary* if its action on $A \subseteq X$ depends only on finite subsets of A in the following sense:

$$\tau(A) \;=\; \bigcup_{\substack{B\subseteq A \\ B \text{ finite}}} \tau(B).$$

A set operator is finitary iff it is chain-continuous (Miscellaneous Exercise **??**), and every chain-continuous operator on any complete partial order is monotone (Miscellaneous Exercise **??**). In many applications involving set operators, the operators are finitary.

Example 1.1 For a binary relation R on a set V, define

$$\tau(R) \;=\; \{(a,c) \mid (a,b),\ (b,c) \in R\}.$$

The function τ is a set operator on V^2; that is,

$$\tau : 2^{V^2} \;\to\; 2^{V^2}.$$

The operator τ is finitary, because $\tau(R)$ is determined by the action of τ on two-element subsets of R. □

Prefixpoints and Fixpoints

A *prefixpoint* of an operator τ on U is an element $x \in U$ such that $\tau(x) \leq x$. A *fixpoint* of τ is an element $x \in U$ such that $\tau(x) = x$. Every operator on U has at least one prefixpoint, namely $\sup U$. Monotone operators also have fixpoints, as we shall see.

For set operators $\tau : 2^X \to 2^X$, we often say that a subset $A \subseteq X$ is *closed* under τ if A is a prefixpoint of τ; that is, if $\tau(A) \subseteq A$.

Example 1.2 By definition, a binary relation R on a set V is *transitive* if $(a,c) \in R$ whenever $(a,b) \in R$ and $(b,c) \in R$. Equivalently, R is transitive iff it is closed under the finitary set operator τ defined in Example 1.1. □

Lemma 1.3 *The infimum of any set of prefixpoints of a monotone operator τ is a prefixpoint of τ.*

Proof. Let A be any set of prefixpoints of τ. We wish to show that $\inf A$ is a prefixpoint of τ. For any $x \in A$, we have $\inf A \leq x$, therefore

$$\tau(\inf A) \ \leq \ \tau(x) \ \leq \ x,$$

since τ is monotone and since x is a prefixpoint. Since $x \in A$ was arbitrary, $\tau(\inf A) \leq \inf A$. □

For $x \in U$, define

$$PF_\tau(x) \ \stackrel{\text{def}}{=} \ \{y \in U \mid \tau(y) \leq y,\ x \leq y\}, \tag{A.1}$$

the set of all prefixpoints of τ above x. Note that $PF_\tau(\bot)$ is the set of all prefixpoints of U, and all $PF_\tau(x)$ are nonempty, since $\top$ is in there at least.

It follows from Lemma 1.3 that $PF_\tau(\bot)$ forms a complete partial order under the induced ordering $\leq$; however, whereas the infimum in $PF_\tau(\bot)$ of any set of prefixpoints A is $\inf A$, the supremum is $\inf PF_\tau(\sup A)$, which is *not* the same as $\sup A$ in general, since $\sup A$ is not necessarily a prefixpoint (Miscellaneous Exercise **??**). Thus we must be careful to say whether we are taking suprema in U or in $PF_\tau(\bot)$.

For $x \in U$, define

$$\tau^\dagger(x) \ \stackrel{\text{def}}{=} \ \inf PF_\tau(x). \tag{A.2}$$

By Lemma 1.3, $\tau^\dagger(x)$ is the least prefixpoint of τ such that $x \leq \tau^\dagger(x)$.

Lemma 1.4 *Any monotone operator τ has a $\leq$-least fixpoint.*

Proof. We show that $\tau^\dagger(\bot)$ is the least fixpoint of τ in U. By Lemma 1.3, it is the least prefixpoint of τ. If it is a fixpoint, then it is the least one, since every fixpoint is a prefixpoint. But if it were not a fixpoint, then by monotonicity, $\tau(\tau^\dagger(\bot))$ would be a smaller prefixpoint, contradicting the fact that $\tau^\dagger(\bot)$ is the smallest. □

Closure Operators

An operator σ on a complete partial order U is called a *closure operator* if it satisfies the following three properties:

(i) σ is monotone

(ii) for all x, $x \leq \sigma(x)$

(iii) for all x, $\sigma(\sigma(x)) = \sigma(x)$.

Because of clause (ii), fixpoints and prefixpoints coincide for closure operators. Thus an element is closed with respect to a closure operator σ iff it is a fixpoint of σ. As shown in Lemma 1.3, the set of closed elements of a closure operator forms complete partial order.

Lemma 1.5 *For any monotone operator τ, the operator $\tau^\dagger$ defined in* (A.2) *is a closure operator.*

Proof. The operator $\tau^\dagger$ is monotone, since

$$x \le y \;\;\Rightarrow\;\; PF_\tau(y) \subseteq PF_\tau(x) \;\Rightarrow\; \inf PF_\tau(x) \le \inf PF_\tau(y),$$

where $PF_\tau(x)$ is the set defined in (A.1).

Property (ii) of closure operators follows directly from the definition of $\tau^\dagger$. Finally, to show property (iii), since $\tau^\dagger(x)$ is a prefixpoint of τ, it suffices to show that any prefixpoint of τ is a fixpoint of $\tau^\dagger$. But

$$\tau(y) \le y \;\;\Leftrightarrow\;\; y \in PF_\tau(y) \;\Leftrightarrow\; y = \inf PF_\tau(y) = \tau^\dagger(y).$$

□

Example 1.6 The *transitive closure* of a binary relation R on a set V is the least transitive relation containing R; that is, it is the least relation containing R and closed under the finitary transitivity operator τ of Example 1.1. The *transitive closure* of R is the relation $\tau^\dagger(R)$. Thus the closure operator $\tau^\dagger$ maps an arbitrary relation R to its transitive closure. □

Example 1.7 The *reflexive transitive closure* of a binary relation R on a set V is the least reflexive and transitive relation containing R; that is, it is the least relation that contains R, is closed under transitivity, and contains the identity relation $\iota = \{(a,a) \mid a \in V\}$. Note that "contains the identity relation" just means closed under the (constant valued) monotone operation $R \mapsto \iota$. Thus the reflexive transitive closure of R is $\sigma^\dagger(R)$, where σ denotes the finitary set operator $R \mapsto \tau(R) \cup \iota$. □

The Knaster–Tarski Theorem

The *Knaster–Tarski theorem* is a useful theorem describing how least fixpoints of monotone operators can be obtained either "from above," as in the proof of Lemma 1.4, or "from below," as a limit of a chain of elements defined by transfinite induction.

Let U be a partially ordered set and let τ be a monotone operator on U. Let $\tau^\dagger$ be the associated closure operator defined in (A.2). We show how to attain $\tau^\dagger(x)$ starting from x and working up. The idea is to start with x and then apply τ repeatedly until achieving closure. In most applications, the operator τ is continuous, in which case this takes only ω iterations; but for monotone operators in general, it can take more.

Formally, we construct by transfinite induction a chain of elements $\tau^\alpha(x)$ indexed by ordinals α:

$$\begin{aligned}\tau^{\alpha+1}(x) &\stackrel{\text{def}}{=} x \vee \tau(\tau^\alpha(x))\\ \tau^\lambda(x) &\stackrel{\text{def}}{=} \sup_{\alpha<\lambda} \tau^\alpha(x), \quad \lambda \text{ a limit ordinal}\\ \tau^*(x) &\stackrel{\text{def}}{=} \sup_{\alpha\in\mathbf{Ord}} \tau^\alpha(x).\end{aligned}$$

The base case is included in the case for limit ordinals:

$$\tau^0(x) = \bot.$$

Intuitively, $\tau^\alpha(x)$ is the set obtained by applying τ to x α times, reincluding x at successor stages.

Lemma 1.8 *If $\alpha \le \beta$, then $\tau^\alpha(x) \le \tau^\beta(x)$.*

Proof. We proceed by transfinite induction on α. For two successor ordinals $\alpha+1$ and $\beta+1$,

$$\tau^{\alpha+1}(x) = x \vee \tau(\tau^\alpha(x)) \le x \vee \tau(\tau^\beta(x)) = \tau^{\beta+1}(x),$$

where the inequality follows from the induction hypothesis and the monotonicity of τ. For a limit ordinal λ on the left-hand side and any ordinal β on the right-hand side,

$$\tau^\lambda(x) = \sup_{\alpha<\lambda} \tau^\alpha(x) \le \tau^\beta(x),$$

where the inequality follows from the induction hypothesis. Finally, for a limit ordinal λ on the right-hand side, the result is immediate from the definition of $\tau^\lambda(x)$. □

Lemma 1.8 says that the $\tau^\alpha(x)$ form a chain in U. The element $\tau^*(x)$ is the supremum of this chain over all ordinals α.

Now there must exist an ordinal κ such that $\tau^{\kappa+1}(x) = \tau^\kappa(x)$, because there is no one-to-one function from the class of ordinals to the set U. The least such κ is called the *closure ordinal* of τ. If κ is the closure ordinal of τ, then $\tau^\beta(x) = \tau^\kappa(x)$ for all $\beta > \kappa$, therefore $\tau^*(x) = \tau^\kappa(x)$.

If τ is chain-continuous, then its closure ordinal is at most ω, but not for monotone operators in general (Exercise **??**).

Theorem 1.9 (Knaster–Tarski) $\tau^\dagger(x) = \tau^*(x)$.

Proof. First we show the forward inclusion. Let κ be the closure ordinal of τ. Since $\tau^\dagger(x)$ is the least prefixpoint of τ above x, it suffices to show that $\tau^*(x) = \tau^\kappa(x)$ is a prefixpoint of τ. But

$$\tau(\tau^\kappa(x)) \;\leq\; x \vee \tau(\tau^\kappa(x)) \;=\; \tau^{\kappa+1}(x) \;=\; \tau^\kappa(x).$$

Conversely, we show by transfinite induction that for all ordinals α, $\tau^\alpha(x) \leq \tau^\dagger(x)$, therefore $\tau^*(x) \leq \tau^\dagger(x)$. For successor ordinals $\alpha + 1$,

$$\begin{aligned} \tau^{\alpha+1}(x) &= x \vee \tau(\tau^\alpha(x)) \\ &\leq x \vee \tau(\tau^\dagger(x)) && \text{induction hypothesis and monotonicity} \\ &\leq \tau^\dagger(x) && \text{definition of } \tau^\dagger. \end{aligned}$$

For limit ordinals λ, $\tau^\alpha(x) \leq \tau^\dagger(x)$ for all $\alpha < \lambda$ by the induction hypothesis; therefore

$$\tau^\lambda(x) \;=\; \sup_{\alpha<\lambda} \tau^\alpha(x) \;\leq\; \tau^\dagger(x).$$

$\square$