APPLICATIONS OF METRIC COINDUCTION

DEXTER KOZEN AND NICHOLAS RUOZZI

Computer Science Department, Cornell University, Ithaca, NY 14853-7501, USA
e-mail address: kozen@cs.cornell.edu

Computer Science Department, Yale University, New Haven, CT 06520-8285, USA
e-mail address: Nicholas.Ruozzi@yale.edu

Abstract. Metric coinduction is a form of coinduction that can be used to establish properties of objects constructed as a limit of finite approximations. One can prove a coinduction step showing that some property is preserved by one step of the approximation process, then automatically infer by the coinduction principle that the property holds of the limit object. This can often be used to avoid complicated analytic arguments involving limits and convergence, replacing them with simpler algebraic arguments. This paper examines the application of this principle in a variety of areas, including infinite streams, Markov chains, Markov decision processes, and non-well-founded sets. These results point to the usefulness of coinduction as a general proof technique.

1. Introduction

Mathematical induction is firmly entrenched as a fundamental and ubiquitous proof principle for proving properties of inductively defined objects. Mathematics and computer science abound with such objects, and mathematical induction is certainly one of the most important tools, if not the most important, at our disposal.

Perhaps less well entrenched is the notion of coinduction. Despite recent interest, coinduction is still not fully established in our collective mathematical consciousness. A contributing factor is that coinduction is often presented in a relatively restricted form. Coinduction is often considered synonymous with bisimulation and is used to establish equality or other relations on infinite data objects such as streams [20] or recursive types [11].

In reality, coinduction is far more general. For example, it has been recently been observed [14] that coinductive reasoning can be used to avoid complicated ε-δ arguments

1998 ACM Subject Classification: F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—proof theory; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs—logics of programs; I.1.3 [Symbolic and Algebraic Manipulation]: Languages and Systems—special-purpose algebraic systems; I.2.3 [Artificial Intelligence]: Deduction and Theorem Proving—mathematical induction, probabilistic reasoning.

Key words and phrases: coinduction, coalgebra, logic in computer science, probabilistic logic.

A preliminary version of this paper appeared as [15].

LOGICAL METHODS
IN COMPUTER SCIENCE

DOI:10.2168/LMCS-???

involving the limiting behavior of a stochastic process, replacing them with simpler algebraic arguments that establish a *coinduction hypothesis* as an invariant of the process, then automatically deriving the property in the limit by application of a coinduction principle. The notion of bisimulation is a special case of this: establishing that a certain relation is a bisimulation is tantamount to showing that a certain coinduction hypothesis is an invariant of some process.

Coinduction, as a proof principle, can handle properties other than equality and inequality and extends to other domains. The goal of this paper is to explore some of these applications. We focus on four areas: infinite streams, Markov chains, Markov decision processes, and non-well-founded sets. In Section 2, we present the metric coinduction principle. In Section 3, we illustrate the use of the principle in the context of infinite streams as an alternative to traditional methods involving bisimulation. In Sections 4 and 5, we rederive some basic results of the theories of Markov chains and Markov decision processes, showing how metric coinduction can simplify arguments. Finally, in Section 6, we use metric coinduction to derive a new characterization of the hereditarily finite non-well-founded sets.

2. Coinduction in Complete Metric Spaces

2.1. **Contractive Maps and Fixpoints.**

Let (V, d) be a complete metric space. A function $H : V \to V$ is *contractive* if there exists $0 \leq c < 1$ such that for all $u, v \in V$, $d(H(u), H(v)) \leq c \cdot d(u, v)$. The value c is called the *constant of contraction.* A continuous function H is said to be *eventually contractive* if H^n is contractive for some $n \geq 1$. Contractive maps are uniformly continuous, and by the Banach fixpoint theorem, any such map has a unique fixpoint in V.

The fixpoint of a contractive map H can be constructed explicitly as the limit of a Cauchy sequence $u, H(u), H^2(u), \ldots$ starting at any point $u \in V$. The sequence is Cauchy; one can show by elementary arguments that

$$d(H^{n+m}(u), H^n(u)) \quad \leq \quad c^n(1 - c^m)(1 - c)^{-1} \cdot d(H(u), u).$$

Since V is complete, the sequence has a limit u^*, which by continuity must be a fixpoint of H. Moreover, u^* is unique: if $H(u) = u$ and $H(v) = v$, then

$$d(u, v) = d(H(u), H(v)) \leq c \cdot d(u, v) \quad \Rightarrow \quad d(u, v) = 0,$$

therefore $u = v$.

Eventually contractive maps also have unique fixpoints. If H^n is contractive, let u^* be the unique fixpoint of H^n. Then $H(u^*)$ is also a fixpoint of H^n. But then $d(u^*, H(u^*)) = d(H^n(u^*), H^{n+1}(u^*)) \leq c \cdot d(u^*, H(u^*))$, which implies that u^* is also a fixpoint of H.

2.2. **The Coinduction Rule.**

In the applications we will consider, the coinduction rule takes the following simple form: If φ is a closed nonempty subset of a complete metric space V, and if H is an eventually contractive map on V that preserves φ, then the unique fixpoint u^* of H is in φ. Expressed as a proof rule, this says for φ a closed property,

$$\frac{\exists u \ \varphi(u) \qquad \forall u \ \varphi(u) \Rightarrow \varphi(H(u))}{\varphi(u^*)}. \tag{2.1}$$

In [14], the rule was used in the special form in which V was a Banach space (normed linear space) and H was an eventually contractive linear affine map on V.

2.3. **Why Is This Coinduction?** We have called (2.1) a coinduction rule. To justify this terminology, we must exhibit a category of coalgebras and show that the rule (2.1) is equivalent to the assertion that a certain coalgebra is final in the category. This construction was given in [14], but we repeat it here for completeness.

Say we have a contractive map H on a metric space V and a nonempty closed subset $\varphi \subseteq V$ preserved by H. Define $H(\varphi) = \{H(s) \mid s \in \varphi\}$. Consider the category C whose objects are the nonempty closed subsets of V and whose arrows are the reverse set inclusions; thus there is a unique arrow $\varphi_1 \rightarrow \varphi_2$ iff $\varphi_1 \supseteq \varphi_2$. The map $\bar{H}$ defined by $\bar{H}(\varphi) = \mathrm{cl}(H(\varphi))$, where cl denotes closure in the metric topology, is an endofunctor on C, since $\bar{H}(\varphi)$ is a nonempty closed set, and $\varphi_1 \supseteq \varphi_2$ implies $\bar{H}(\varphi_1) \supseteq \bar{H}(\varphi_2)$. An $\bar{H}$-coalgebra is then a nonempty closed set φ such that $\varphi \supseteq \bar{H}(\varphi)$; equivalently, such that $\varphi \supseteq H(\varphi)$. The final coalgebra is $\{u^*\}$, where u^* is the unique fixpoint of H. The coinduction rule (2.1) says that $\varphi \supseteq H(\varphi) \Rightarrow \varphi \supseteq \{u^*\}$, which is equivalent to the statement that $\{u^*\}$ is final in the category of $\bar{H}$-coalgebras.

3. Streams

Infinite streams have been a very successful source of application of coinductive techniques. The space $\mathcal{S}_\Sigma = (\Sigma^\omega, \mathsf{head}, \mathsf{tail})$ of infinite streams over Σ is the final coalgebra in the category of *simple transition systems* over Σ, whose objects are $(X, \mathsf{obs}, \mathsf{cont})$, where X is a set, $\mathsf{obs} : X \rightarrow \Sigma$ gives an *observation* at each state, and $\mathsf{cont} : X \rightarrow X$ gives a *continuation* (next state) for each state. The unique morphism $(X, \mathsf{obs}, \mathsf{cont}) \rightarrow (\Sigma^\omega, \mathsf{head}, \mathsf{tail})$ maps a state $s \in X$ to the stream $\mathsf{obs}(s), \mathsf{obs}(\mathsf{cont}(s)), \mathsf{obs}(\mathsf{cont}^2(s)), \ldots \in \Sigma^\omega$.

We begin by illustrating the use of the metric coinduction principle in this context as an alternative to traditional methods involving bisimulation. It is well known that $\mathcal{S}_\Sigma$ forms a complete metric space under the distance function $d(\sigma, \tau) \stackrel{\mathrm{def}}{=} 2^{-n}$, where n is the first position at which σ and τ differ. The metric d satisfies the property

$$d(x :: \sigma, y :: \tau) \quad = \quad \begin{cases} \frac{1}{2} d(\sigma, \tau), & \text{if } x = y \\ 1, & \text{if } x \neq y. \end{cases}$$

One can also form the product space $\mathcal{S}_\Sigma^2$ with metric

$$d((\sigma_1, \sigma_2), (\tau_1, \tau_2)) \quad \stackrel{\mathrm{def}}{=} \quad \max d(\sigma_1, \tau_1),\, d(\sigma_2, \tau_2).$$

Since distances are bounded, the spaces of continuous operators $\mathcal{S}_\Sigma^2 \rightarrow \mathcal{S}_\Sigma$ and $\mathcal{S}_\Sigma \rightarrow \mathcal{S}_\Sigma^2$ are also complete metric spaces under the sup metric

$$d(E, F) \quad \stackrel{\mathrm{def}}{=} \quad \sup_x d(E(x), F(x)).$$

Consider the operators $\mathsf{merge} : \mathcal{S}_\Sigma^2 \rightarrow \mathcal{S}_\Sigma$ and $\mathsf{split} : \mathcal{S}_\Sigma \rightarrow \mathcal{S}_\Sigma^2$ defined informally by

$$\begin{aligned} \mathsf{merge}\,(a_0 a_1 a_2 \cdots,\, b_0 b_1 b_2 \cdots) \quad &= \quad a_0 b_0 a_1 b_1 a_2 b_2 \cdots \\ \mathsf{split}\,(a_0 a_1 a_2 \cdots) \quad &= \quad (a_0 a_2 a_4 \cdots,\, a_1 a_3 a_5 \cdots). \end{aligned}$$

Thus merge forms a single stream from two streams by taking elements alternately, and split separates a single stream into two streams consisting of the even and odd elements, respectively.

Formally, one would define merge and split coinductively as follows:

$$\begin{aligned}\mathsf{merge}\,(x :: \sigma,\, \tau) &\stackrel{\text{def}}{=} x :: \mathsf{merge}\,(\tau,\, \sigma)\\ \mathsf{split}\,(x :: y :: \sigma) &\stackrel{\text{def}}{=} (x :: \mathsf{split}\,(\sigma)_1,\, y :: \mathsf{split}\,(\sigma)_2).\end{aligned}$$

These functions exist and are unique, since they are the unique fixpoints of the eventually contractive maps

$$\alpha : (\mathcal{S}_\Sigma^2 \to \mathcal{S}_\Sigma) \;\to\; (\mathcal{S}_\Sigma^2 \to \mathcal{S}_\Sigma) \qquad\qquad \beta : (\mathcal{S}_\Sigma \to \mathcal{S}_\Sigma^2) \;\to\; (\mathcal{S}_\Sigma \to \mathcal{S}_\Sigma^2)$$

defined by

$$\begin{aligned}\alpha(M)(x :: \sigma,\, \tau) &\stackrel{\text{def}}{=} x :: M(\tau,\, \sigma)\\ \beta(S)(x :: y :: \sigma) &\stackrel{\text{def}}{=} (x :: S(\sigma)_1,\, y :: S(\sigma)_2).\end{aligned}$$

We would like to show that merge and split are inverses. Traditionally, one would do this by exhibiting a bisimulation between $\mathsf{merge}\,(\mathsf{split}\,(\sigma))$ and σ, thus concluding that $\mathsf{merge}\,(\mathsf{split}\,(\sigma)) = \sigma$, and another bisimulation between $\mathsf{split}\,(\mathsf{merge}\,(\sigma,\, \tau))$ and $(\sigma,\, \tau)$, thus concluding that $\mathsf{split}\,(\mathsf{merge}\,(\sigma,\, \tau)) = (\sigma,\, \tau)$.

Here is how we would prove this result using the metric coinduction rule (2.1). Let $M : \mathcal{S}_\Sigma^2 \to \mathcal{S}_\Sigma$ and $S : \mathcal{S}_\Sigma \to \mathcal{S}_\Sigma^2$. If M is a left inverse of S, then $\alpha^2(M)$ is a left inverse of $\beta(S)$:

$$\begin{aligned}\alpha^2(M)(\beta(S)(x :: y :: \sigma)) &= \alpha(\alpha(M))(x :: S(\sigma)_1,\, y :: S(\sigma)_2)\\ &= x :: \alpha(M)(y :: S(\sigma)_2,\, S(\sigma)_1)\\ &= x :: y :: M(S(\sigma)_1,\, S(\sigma)_2)\\ &= x :: y :: M(S(\sigma))\\ &= x :: y :: \sigma.\end{aligned}$$

Similarly, if M is a right inverse of S, then $\alpha^2(M)$ is a right inverse of $\beta(S)$:

$$\begin{aligned}\beta(S)(\alpha^2(M)(x :: \sigma,\, y :: \tau)) &= \beta(S)(\alpha(\alpha(M))(x :: \sigma,\, y :: \tau))\\ &= \beta(S)(x :: \alpha(M)(y :: \tau,\, \sigma))\\ &= \beta(S)(x :: y :: M(\sigma,\, \tau))\\ &= (x :: S(M(\sigma,\, \tau))_1,\, y :: S(M(\sigma,\, \tau))_2)\\ &= (x :: (\sigma,\, \tau)_1,\, y :: (\sigma,\, \tau)_2)\\ &= (x :: \sigma,\, y :: \tau).\end{aligned}$$

We conclude that if M and S are inverses, then so are $\alpha^2(M)$ and $\beta(S)$.

The property

$$\varphi(M, S) \;\stackrel{\text{def}}{\Longleftrightarrow}\; M \text{ and } S \text{ are inverses} \tag{3.1}$$

is a nonempty closed property of $(\mathcal{S}_\Sigma^2 \to \mathcal{S}_\Sigma) \times (\mathcal{S}_\Sigma \to \mathcal{S}_\Sigma^2)$ which, as we have just shown, is preserved by the contractive map $(M, S) \mapsto (\alpha^2(M), \beta(S))$. By (2.1), φ holds of the unique fixpoint (merge, split).

That φ is nonempty and closed requires an argument, but these conditions typically follow from general topological considerations. For example, (3.1) is nonempty because the spaces $\mathcal{S}_\Sigma$ and $\mathcal{S}_\Sigma^2$ are both homeomorphic to the topological product of countably many copies of the discrete space Σ.

4. Markov Chains

A finite Markov chain is a finite state space, say $\{1,\ldots,n\}$, together with a stochastic matrix $P \in \mathbb{R}^{n\times n}$ of transition probabilities, with P_{st} representing the probability of a transition from state s to state t in one step. The value P^m_{st} is the probability that the system is in state t after m steps, starting in state s.

A fundamental result in the theory of Markov chains is that if P is irreducible and aperiodic (definitions given below), then P^m_{st} tends to $1/\mu_t$ as $m \to \infty$, where μ_t is the *mean first recurrence time* of state t, the expected time of first reentry into state t after leaving state t. Intuitively, if we expect to be in state t about every μ_t steps, then in the long run we expect to be in state t about $1/\mu_t$ of the time.

The proof of this result as given in Feller [10] is rather lengthy, involving a complicated argument to establish the uniform convergence of a certain countable sequence of countable sequences. The complete proof runs to several pages. Introductory texts devote entire chapters to it (e.g. [12]) or omit the proof entirely (e.g. [17]). In this section we show that, assuming some basic spectral properties of stochastic matrices, the coinduction rule can be used to give a simpler alternative proof.

4.1. Spectral Properties.

Recall that P is *irreducible* if its underlying support graph is strongly connected. The *support graph* has vertices $\{1,\ldots,n\}$ and directed edges $\{(s,t) \mid P_{st} > 0\}$. A directed graph is *strongly connected* if there is a directed path from any vertex to any other vertex. The matrix P is *aperiodic* if in addition, the gcd of the set $\{m \mid P^m_{ss} > 0\}$ is 1 for all states s. By the Perron–Frobenius theorem (see [5, 16]), if P is irreducible and aperiodic, then P has eigenvalue 1 with multiplicity 1 and all other eigenvalues have norm strictly less than 1.

The matrix P is itself not contractive, since 1 is an eigenvalue. However, consider the matrix

$$P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}},$$

where $\mathbf{1}$ is the column vector of all 1's and ${}^{\mathrm{T}}$ denotes matrix transpose. The matrix $\frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}$ is the $n \times n$ matrix all of whose entries are $1/n$.

The spectra of P and $P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}$ are closely related, as shown in the following lemma.

Lemma 4.1. *Let $P \in \mathbb{R}^{n\times n}$ be a stochastic matrix. Any (left) eigenvector x^{T} of $P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}$ that lies in the hyperplane $x^{\mathrm{T}}\mathbf{1} = 0$ is also an eigenvector of P with the same eigenvalue, and vice-versa. The only other eigenvalue of P is 1 and the only other eigenvalue of $P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}$ is 0.*

Proof. For any eigenvalue λ of P and corresponding eigenvector x^{T},

$$\lambda x^{\mathrm{T}}\mathbf{1} \;=\; x^{\mathrm{T}}P\mathbf{1} = x^{\mathrm{T}}\mathbf{1}$$

since $P\mathbf{1} = \mathbf{1}$, so either $\lambda = 1$ or $x^{\mathrm{T}}\mathbf{1} = 0$. Similarly, for any eigenvalue λ of $P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}$ and corresponding eigenvector x^{T},

$$\lambda x^{\mathrm{T}}\mathbf{1} \;=\; x^{\mathrm{T}}(P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}})\mathbf{1} = x^{\mathrm{T}}\mathbf{1} - x^{\mathrm{T}}\mathbf{1} = 0,$$

so either $\lambda = 0$ or $x^{\mathrm{T}}\mathbf{1} = 0$. But if $x^{\mathrm{T}}\mathbf{1} = 0$, then

$$x^{\mathrm{T}}(P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}) = x^{\mathrm{T}}P - \frac{1}{n}x^{\mathrm{T}}\mathbf{1}\mathbf{1}^{\mathrm{T}} = x^{\mathrm{T}}P,$$

so in this case x^{T} is an eigenvector of P iff it is an eigenvector of $P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}$ with the same eigenvalue. □

4.2. **Coinduction and the Convergence of P^m.** If P is irreducible and aperiodic, then $P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}$ is eventually contractive, since $\inf_n \sqrt[n]{\|(P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}})^n\|}$ is equal to the *spectral radius* or norm of the largest eigenvalue of $P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}$ (see [9]), which by Lemma 4.1 is less than 1. Thus the map

$$x^{\mathrm{T}} \quad \mapsto \quad x^{\mathrm{T}}(P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}) + \frac{1}{n}\mathbf{1}^{\mathrm{T}} \tag{4.1}$$

is of the proper form to be used with the metric coinduction rule (2.1) to establish the convergence of P^m.

Since $P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}$ is eventually contractive, the map (4.1) has a unique fixpoint u^{T}. The set of stochastic vectors

$$S \quad = \quad \{x^{\mathrm{T}} \mid x^{\mathrm{T}} \geq 0,\ x^{\mathrm{T}}\mathbf{1} = 1\}$$

is closed and preserved by the map (4.1), since

$$x^{\mathrm{T}}\mathbf{1} = 1 \quad \Rightarrow \quad x^{\mathrm{T}}(P - \frac{1}{n}\mathbf{1}\mathbf{1}^{\mathrm{T}}) + \frac{1}{n}\mathbf{1}^{\mathrm{T}} = x^{\mathrm{T}}P,$$

and S is preserved by P. By the metric coinduction rule (2.1), the unique fixpoint u^{T} is contained in S. By Lemma 4.1, it is also an eigenvector of 1, and $y^{\mathrm{T}}P^m$ tends to u^{T} for any $y^{\mathrm{T}} \in S$. Applying this to the rows of any stochastic matrix E, we have that EP^m converges to the matrix $\mathbf{1}u^{\mathrm{T}}$.

4.3. **Recurrence Statistics.** Once we have established the convergence of P^m, we can give a much shorter argument than those of [10, 12] that the actual limit of P^m_{st} is $1/\mu_t$. We follow the notation of [10].

Fix a state t, and let $\mu = \mu_t$. Let f_m be the probability that after leaving state t, the system first returns to state t at time m. Let $u_m = P^m_{tt}$ be the probability that the system is in state t at time m after starting in state t. By irreducibility, $\sum_{m=1}^{\infty} f_m = 1$ and $\mu = \sum_{m=1}^{\infty} mf_m < \infty$. Let $\rho_m \stackrel{\mathrm{def}}{=} \sum_{k=m+1}^{\infty} f_k$, and consider the generating functions

$$f(x) \stackrel{\mathrm{def}}{=} \sum_{m=1}^{\infty} f_m x^m \qquad u(x) \stackrel{\mathrm{def}}{=} \sum_{m=0}^{\infty} u_m x^m$$

$$\rho(x) \stackrel{\mathrm{def}}{=} \sum_{m=0}^{\infty} \rho_m x^m \qquad \sigma(x) \stackrel{\mathrm{def}}{=} u_0 + \sum_{m=0}^{\infty} (u_{m+1} - u_m)x^{m+1}.$$

The probabilities u_n obey the recurrence

$$u_0 = 1 \qquad\qquad u_n = \sum_{m=0}^{n-1} u_m f_{n-m},$$

which implies that $f(x)u(x) = u(x) - 1$. Elementary algebraic reasoning gives

$$\sigma(x)\rho(x) \;=\; 1. \tag{4.2}$$

Now we claim that both $\sigma(1)$ and $\rho(1)$ converge. The sequence $\rho(1)$ converges to $\mu > 0$, since

$$\rho(1) \;=\; \sum_{m=1}^{\infty} \rho_m = \sum_{m=1}^{\infty} m f_m = \mu, \tag{4.3}$$

and the latter sequence in (4.3) converges absolutely. For $\sigma(1)$, we have

$$\sigma(1) \;=\; u_0 + \sum_{m=0}^{\infty} (u_{m+1} - u_m),$$

which converges by the results of Section 4.2. By (4.2), $\sigma(1)\rho(1) = 1$, therefore $\sigma(1) = 1/\mu$. But the mth partial sum of $\sigma(1)$ is just $u_0 + \sum_{k=0}^{m-1}(u_{k+1} - u_k) = u_m$, so the sequence u_m converges to $1/\mu$.

5. Markov Decision Processes

In this section, we rederive some fundamental results on Markov decision processes using the metric coinduction principle. A fairly general treatment of this theory is given in [8], and we follow the notation of that paper. However, the strategic use of metric coinduction allows a more streamlined presentation.

5.1. **Existence of Optimal Strategies.**

Let V be the space of bounded real-valued functions on a set of states Ω with the sup norm $\|v\| \overset{\text{def}}{=} \sup_{x\in\Omega} |v(x)|$. The space V is a complete metric space with metric $\|v - u\|$.

For each state $x \in \Omega$, say we have a set Δ_x of *actions*. A *deterministic strategy* is an element of $\Delta \overset{\text{def}}{=} \prod_{x\in\Omega} \Delta_x$, thus a selection of actions, one for each state $x \in \Omega$. More generally, if Δ_x is a measurable space, let $\mathbb{M}(\Delta_x)$ denote the space of probability measures on Δ_x. A *probabilistic strategy* is an element of $\prod_{x\in\Omega} \mathbb{M}(\Delta_x)$, thus a selection of probability measures, one for each $x \in \Omega$. A deterministic strategy can be viewed as a probabilistic strategy in which all the measures are point masses.

Now suppose we have a *utility function* $h : \prod_{x\in\Omega}(\Delta_x \to V \to \mathbb{R})$ with the three properties listed below.[1] The function h induces a function H such that $H_\delta(u)(x) = h(x, \delta_x, u) \in \mathbb{R}$, where $x \in \Omega$, $\delta \in \Delta$, and $u \in V$.

(i) The function H is uniformly bounded as a function of δ and x. That is, $H_\delta : V \to V$, and for any fixed $u \in V$, $\sup_{\delta\in\Delta} \|H_\delta(u)\|$ is finite.

(ii) The functions H_δ are uniformly contractive with constant of contraction $c < 1$. That is, for all $\delta \in \Delta$ and $u, v \in V$, $\|H_\delta(v) - H_\delta(u)\| \le c \cdot \|v - u\|$. Thus H_δ has a unique fixpoint, which we denote by v_δ.

(iii) Every H_δ is *monotone*: if $u \le v$, then $H_\delta(u) \le H_\delta(v)$. The order $\le$ on V is the pointwise order.

Lemma 5.1. *Define $A : V \to V$ by $A(u)(x) \overset{\text{def}}{=} \sup_{d\in\Delta_x} h(x, d, u)$. The supremum exists since the H_δ are uniformly bounded. Then A is contractive with constant of contraction c.*

[1] We write $h(x, \delta_x, u)$ instead of $h(x)(\delta_x)(u)$ for readability.

Proof. Let $\varepsilon > 0$. For $x \in \Omega$, assuming without loss of generality that $A(v)(x) \geq A(u)(x)$,

$$\begin{aligned}
&|A(v)(x) - A(u)(x)| \\
&= \sup_{d\in\Delta_x} h(x,d,v) - \sup_{e\in\Delta_x} h(x,e,u) \\
&\leq \varepsilon + h(x,d,v) - \sup_{e\in\Delta_x} h(x,e,u) \quad \text{for suitably chosen } d \in \Delta_x \\
&\leq \varepsilon + h(x,d,v) - h(x,d,u) \\
&\leq \varepsilon + c \cdot \|v-u\|.
\end{aligned}$$

Since ε was arbitrary, $|A(v)(x) - A(u)(x)| \leq c \cdot \|v-u\|$, thus

$$\|A(v) - A(u)\| \leq \sup_x |A(v)(x) - A(u)(x)| \leq c \cdot \|v-u\|.$$

□

Since A is contractive, it has a unique fixpoint v^*.

Lemma 5.2. *For any δ, $v_\delta \leq v^*$.*

Proof. By the coinduction principle, it suffices to show that $u \leq v$ implies $H_\delta(u) \leq A(v)$. Here the metric space is V^2, the closed property φ is $u \leq v$, and the contractive map is (H_δ, A). But if $u \leq v$, then by monotonicity,

$$H_\delta(u)(x) \leq H_\delta(v)(x) = h(x,\delta_x,v) \leq \sup_{d\in\Delta_x} h(x,d,v) = A(v).$$

□

Lemma 5.3. *The fixpoint v^* can be approximated arbitrarily closely by v_δ for deterministic strategies δ.*

Proof. Let $\varepsilon > 0$. Let δ be such that for all x,

$$\sup_{d\in\Delta_x} h(x,d,v^*) - h(x,\delta_x,v^*) < (1-c)\varepsilon.$$

We will show that $\|v^* - v_\delta\| \leq \varepsilon$. By the coinduction rule (2.1), it suffices to show that $\|v^* - u\| \leq \varepsilon$ implies $\|v^* - H_\delta(u)\| \leq \varepsilon$. Here the metric space is V, the closed property $\varphi(u)$ is $\|v^* - u\| \leq \varepsilon$, and the contractive map is H_δ. But if $\|v^* - u\| \leq \varepsilon$,

$$\begin{aligned}
\|v^* - H_\delta(u)\| &= \sup_x |v^*(x) - H_\delta(u)(x)| = \sup_x |A(v^*)(x) - H_\delta(u)(x)| \\
&= \sup_x |\sup_{d\in\Delta_x} h(x,d,v^*) - h(x,\delta_x,u)| \\
&\leq \sup_x (|\sup_{d\in\Delta_x} h(x,d,v^*) - h(x,\delta_x,v^*)| + |h(x,\delta_x,v^*) - h(x,\delta_x,u)|) \\
&\leq (1-c)\varepsilon + c \cdot \|v^* - u\| \leq (1-c)\varepsilon + c\varepsilon = \varepsilon.
\end{aligned}$$

□

5.2. **Probabilistic Strategies.** We use the metric coinduction rule (2.1) to prove the well-known result that for Markov decision processes, probabilistic strategies are no better than deterministic strategies. If $\sup_{d\in\Delta_x} h(x,d,v^*)$ is attainable for all x, then the deterministic strategy $\delta_x \stackrel{\mathrm{def}}{=} \mathrm{argmax}_{d\in\Delta_x} h(x,d,v^*)$ is optimal, even allowing probabilistic strategies. However, if $\sup_{d\in\Delta_x} h(x,d,v^*)$ is not attainable, then it is not so obvious what to do.

For this argument, we assume that Δ_x is a measurable space and that for all fixed x and u, $h(x,d,u)$ is an integrable function of $d \in \Delta_x$. Given a probabilistic strategy $\mu : \prod_{x\in\Omega} \mathbb{M}(\Delta_x)$, the one-step utility function is $H_\mu : V \to V$ defined by the Lebesgue integral

$$H_\mu(u)(x) \quad \stackrel{\mathrm{def}}{=} \quad \int_{d\in\Delta_x} h(x,d,u)\cdot\mu_x(\triangle d).$$

This integral accumulates the various individual payoffs over all choices of d weighted by the measure μ_x.

The map $H_\mu(u)$ is uniformly bounded in μ, since

$$\begin{aligned}
\|H_\mu(u)\| &= \sup_x \left| \int_{d\in\Delta_x} h(x,d,u)\cdot\mu_x(\triangle d)\right| \le \sup_x \int_{d\in\Delta_x} |h(x,d,u)|\cdot\mu_x(\triangle d) \\
&\le \sup_x \sup_d |h(x,d,u)| \cdot \int_{d\in\Delta_x} \mu_x(\triangle d) = \sup_{x,d} |h(x,d,u)|.
\end{aligned}$$

It is also a contractive map with constant of contraction c, since

$$\begin{aligned}
\|H_\mu(v) - H_\mu(u)\| &= \sup_x |H_\mu(v)(x) - H_\mu(u)(x)| \\
&= \sup_x \left| \int_{d\in\Delta_x} h(x,d,v)\cdot\mu_x(\triangle d) - \int_{d\in\Delta_x} h(x,d,u)\cdot\mu_x(\triangle d)\right| \\
&= \sup_x \left| \int_{d\in\Delta_x} (h(x,d,v) - h(x,d,u))\cdot\mu_x(\triangle d)\right| \\
&\le \sup_x \int_{d\in\Delta_x} |h(x,d,v) - h(x,d,u)|\cdot\mu_x(\triangle d) \\
&\le \sup_x \int_{d\in\Delta_x} c\cdot\|v-u\|\cdot\mu_x(\triangle d) \\
&= c\cdot\|v-u\|\cdot\sup_x \int_{d\in\Delta_x} \mu_x(\triangle d) \\
&= c\cdot\|v-u\|.
\end{aligned}$$

Since it is a contractive map, it has a unique fixpoint v_μ.

Now take any deterministic strategy δ such that $h(x,\delta_x,v_\mu) \ge v_\mu(x)$ for all x. This is always possible, since if $h(x,d,v_\mu) < v_\mu(x)$ for all $d \in \Delta_x$, then

$$v_\mu(x) \quad = \quad H_\mu(v_\mu)(x) = \int_{d\in\Delta_x} h(x,d,v_\mu)\cdot\mu_x(\triangle d) < v_\mu(x),$$

a contradiction. The following lemma says that the deterministic strategy δ is no worse than the probabilistic strategy μ.

Lemma 5.4. *$v_\delta \ge v_\mu$.*

Proof. Assuming $v_\mu \le v$, we have

$$v_\mu(x) \quad \le \quad h(x, \delta_x, v_\mu) \le h(x, \delta_x, v) = H_\delta(v)(x),$$

the second inequality by monotonicity. As x was arbitrary, $v_\mu \le H_\delta(v)$. The result follows from the coinduction principle on the metric space V with $\varphi(v)$ the closed property $v_\mu \le v$ and contractive map H_δ. □

6. Non-Well-Founded Sets

In classical Zermelo–Fraenkel set theory with choice (ZFC), the "element of" relation $\in$ is well-founded, as guaranteed by the axiom of foundation. Aczel [2] developed the theory of *non-well-founded sets*, in which sets with infinitely descending $\in$-chains are permitted in addition to the well-founded sets. These are precisely the sets that are explicitly ruled out of existence by the axiom of foundation.

In the theory of non-well-founded sets, the sets are represented by *accessible pointed graphs* (APGs). An APG is a directed graph with a distinguished node such that every node is reachable by a directed path from the distinguished node. Two APGs represent the same set iff they are bisimilar. The APGs of well-founded sets may be infinite, but may contain no infinite paths or cycles, whereas the APGs of non-well-founded sets may contain infinite paths and cycles. Equality as bisimulation is the natural analog of extensionality in ZFC; essentially, two APGs are declared equal as sets if there is no witness among their descendants that forces them not to be. The class V is the class of sets defined in this way.

Aczel [2] (see also [4, 21]) notes the strong role that coinduction plays in this theory. Since equality between APGs is defined in terms of bisimulation, coinduction becomes a primary proof technique for establishing the equivalence of different APGs representing the same set.

In attempting to define a metric on non-well-founded sets, the classical Hausdorff distance suggests itself as a promising candidate. This metric has been previously defined for the hereditarily finite well-founded sets and their completion, the *finitary* non-well-founded sets, by Abramsky [1]. For the more general case of arbitrary non-well-founded sets, there are two complications. One is that we must apply the definition coinductively. Another is that ordinarily, the Hausdorff metric is only defined on compact sets, since otherwise a Hausdorff distance of zero may not imply equality, and that is the case here. However, the definition still makes sense even for non-compact sets and leads to further insights into the structure of non-well-founded sets.

In this section, we define a distance function $d : V^2 \to \mathbb{R}$ based on a coinductive application of the Hausdorff distance function and derive some properties of d. We show that (V, d) forms a compact pseudometric space. Being a pseudometric instead of a metric means that there are sets $s \ne t$ with $d(s, t) = 0$. Nevertheless, we identify a maximal family of sets that includes all the hereditarily finite sets on which d acts as a metric.

We will prove the following results. Define $s \approx t$ if $d(s, t) = 0$. Call a set s *singular* if the only t such that $s \approx t$ is s itself.

- A set is singular if and only if it is hereditarily finite.
- All singular sets are closed in the pseudometric topology. In particular, all hereditarily finite sets are hereditarily closed (but not vice-versa).
- A set is hereditarily closed if and only if it is closed and all elements are singular.

- All hereditarily closed sets are canonical (but not vice-versa), where a set is *canonical* if it is a member of a certain coinductively-defined class of canonical representatives of the $\approx$-classes.
- The map d is a metric on the canonical sets; moreover, the canonical sets are a maximal class for which this is true.

6.1. **Coinductive Definition of Functions.** Just as classical ZFC allows the definition of functions by induction over ordinary well-founded sets, there is a corresponding principle for non-well-founded sets known as the *Solution Lemma* [2, 21]. In particular, the Solution Lemma implies that for any function $H : V \to V$, the equation

$$G(s) \stackrel{\text{def}}{=} \{G(u) \mid u \in H(s)\} \tag{6.1}$$

determines $G : V \to V$ uniquely. This is because if G and G' both satisfy (6.1), then the relation

$$u\ R\ v \quad \stackrel{\text{def}}{\iff} \quad \exists s\ u = G(s) \wedge v = G'(s)$$

is a bisimulation, therefore $G(s) = G'(s)$ for all s. In coalgebraic terms[2], the map G is the unique morphism from the coalgebra $(V, \{(s,t) \mid s \in H(t)\})$ to the final coalgebra $(V, \in)$; see [2, Chp. 7] or [21, Part V].

6.2. **Definition of d.** Let B be the Banach space of bounded real-valued functions $g : \mathrm{APG}^2 \to \mathbb{R}$ with norm

$$\|g\| \stackrel{\text{def}}{=} \sup_{s,t} |g(s,t)|.$$

Define the map $\tau : B \to B$ by

$$\tau(g)(s,t) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } s,t = \varnothing \\ 1 & \text{if } s = \varnothing \Leftrightarrow t \neq \varnothing \\ \frac{1}{2}\max\begin{cases} \sup_{u\in s}\inf_{v \in t} g(u,v) \\ \sup_{v \in t}\inf_{u\in s} g(u,v)\end{cases} & \text{if } s,t \neq \varnothing. \end{cases}$$

It can be shown that $\|\tau(g) - \tau(g')\| \le \frac{1}{2}\|g - g'\|$, thus τ is contractive on B with constant of contraction $1/2$ and has a unique fixpoint $d \in B$. One can therefore use the metric coinduction rule (2.1) to prove properties of d.

To illustrate, let us show that the non-well-founded sets V form a compact (thus complete) pseudometric space with respect to the distance function d. At the outset, it is not immediately clear that d is well-defined on V. We must argue that d is invariant on bisimulation classes; that is, for any bisimulation R, if $s\ R\ s'$ and $t\ R\ t'$, then $d(s,t) = d(s',t')$. We will use the metric coinduction rule (2.1) to prove this.

Consider the following closed property on B, defined with respect to an arbitrary but fixed bisimulation R on the class of APGs:

$$\varphi(g) \quad \stackrel{\text{def}}{\iff} \quad \forall s\ \forall s'\ \forall t\ \forall t'\ \ s\ R\ s' \wedge t\ R\ t' \ \Rightarrow\ g(s,t) = g(s',t').$$

[2] When regarding V as a coalgebra, the notation $(V, \in)$ is a slight but convenient abuse. Formally, these structures are coalgebras with respect to the powerset functor $\mathcal{P}$. To be precise, we should write (V, β), where $\beta : V \to \mathcal{P}V$ and write $s \in \beta(t)$ instead of $s \in t$.

This property is closed in the metric topology on B, since it is an infinite conjunction of closed properties $g(s,t) = g(s',t')$, one for each selection of s, s', t, t' such that $s \mathrel{R} s'$ and $t \mathrel{R} t'$. It is clearly nonempty. We wish to prove that $\varphi(d)$. By the metric coinduction rule (2.1), it suffices to show that φ is closed under τ.

Suppose $\varphi(g)$. We want to show that $\varphi(\tau(g))$, or in other words,

$$\forall s\ \forall s'\ \forall t\ \forall t'\ s \mathrel{R} s' \wedge t \mathrel{R} t' \quad \Rightarrow \quad \tau(g)(s,t) = \tau(g)(s',t').$$

Let s, s', t, t' be such that $s \mathrel{R} s'$ and $t \mathrel{R} t'$. Since R is a bisimulation, we have

$$\begin{array}{ll} \forall u \in s\ \exists u' \in s'\ u \mathrel{R} u' & \qquad \forall u' \in s'\ \exists u \in s\ u \mathrel{R} u' \\ \forall v \in t\ \exists v' \in t'\ v \mathrel{R} v' & \qquad \forall v' \in t'\ \exists v \in t\ v \mathrel{R} v'. \end{array}$$

It follows that $s = \varnothing$ iff $s' = \varnothing$ and $t = \varnothing$ iff $t' = \varnothing$. If $s = s' = \varnothing$, then

$$\tau(g)(s,t) \quad = \quad \left\{ \begin{array}{ll} 0 & \text{if } t, t' = \varnothing \\ 1 & \text{if } t, t' \neq \varnothing \end{array} \right\} \quad = \quad \tau(g)(s',t').$$

A symmetric argument holds if $t = t' = \varnothing$.

Otherwise, all four sets s, s', t, t' are nonempty. In this case,

$$\begin{aligned} \tau(g)(s,t) \quad &= \quad \tfrac{1}{2} \max \left\{ \begin{array}{l} \sup_{u \in s} \inf_{v \in t} g(u,v) \\ \sup_{v \in t} \inf_{u \in s} g(u,v) \end{array} \right. \\ \tau(g)(s',t') \quad &= \quad \tfrac{1}{2} \max \left\{ \begin{array}{l} \sup_{u' \in s'} \inf_{v' \in t'} g(u',v') \\ \sup_{v' \in t'} \inf_{u' \in s'} g(u',v'), \end{array} \right. \end{aligned}$$

so it suffices to show that

$$\sup_{u \in s} \inf_{v \in t} g(u,v) \quad = \quad \sup_{u' \in s'} \inf_{v' \in t'} g(u',v') \tag{6.2}$$

$$\sup_{v \in t} \inf_{u \in s} g(u,v) \quad = \quad \sup_{v' \in t'} \inf_{u' \in s'} g(u',v'). \tag{6.3}$$

We show only (6.2); the argument for (6.3) is symmetric. Also by symmetry, we need only show the inequality in one direction:

$$\sup_{u \in s} \inf_{v \in t} g(u,v) \quad \leq \quad \sup_{u' \in s'} \inf_{v' \in t'} g(u',v').$$

This inequality follows from the property

$$\forall u \in s\ \exists u' \in s'\ \inf_{v \in t} g(u,v) \leq \inf_{v' \in t'} g(u',v'),$$

which in turn follows from

$$\forall u \in s\ \exists u' \in s'\ \forall v' \in t'\ \exists v \in t\ g(u,v) \leq g(u',v').$$

In fact, we have

$$\forall u \in s\ \exists u' \in s'\ \forall v' \in t'\ \exists v \in t\ g(u,v) = g(u',v')$$

by choosing $u' \in s'$ such that $u \mathrel{R} u'$ and $v \in t$ such that $v \mathrel{R} v'$, as guaranteed by the coinduction hypothesis and the fact that R is a bisimulation.

We conclude by the metric coinduction principle (2.1) that $\varphi(d)$ holds, thus d is invariant on the equivalence classes of any bisimulation R on APGs, therefore well-defined on V.

To show that d is a pseudometric, we must also show

$$\begin{array}{ll} d(s,t) \geq 0 \text{ (in fact, } d(s,t) \in [0,1]) & \qquad d(s,t) = d(t,s) \\ d(s,u) \leq d(s,t) + d(t,u) & \qquad d(s,s) = 0. \end{array}$$

All these properties can be shown in the same way, by metric coinduction. One need only argue that they are all nonempty closed properties closed under τ.

We will establish compactness (hence completeness) later in section 6.4, but first we introduce the canonical sets.

6.3. **Canonical Sets.** The map d is only a pseudometric and not a metric, since it is possible that $d(s,t) = 0$ even though $s \neq t$. For example, define $\overline{0} = \varnothing$, $\overline{n+1} = \{\overline{n}\}$. Let Ω be the unique non-well-founded set such that $\Omega = \{\Omega\}$. The sets $\{\overline{n} \mid n \geq 0\}$ and $\{\overline{n} \mid n \geq 0\} \cup \Omega$ are distinct, but distance 0 apart (Fig. 1). This follows from the observation that $d(\overline{n}, \Omega) = 2^{-n}$, so

$$\sup_{v \in \{\overline{n} \mid n \geq 0\} \cup \Omega} \; \inf_{u \in \{\overline{n} \mid n \geq 0\}} d(u,v) \quad = \quad \inf_{n \geq 0} d(\overline{n}, \Omega) \quad = \quad 0.$$

Nevertheless, it is possible to relate this map to the coalgebraic structure of V.

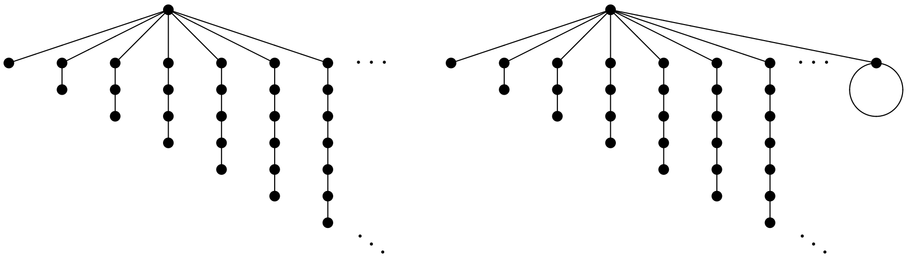
$\{\overline{n} \mid n \geq 0\}$
$\{\overline{n} \mid n \geq 0\} \cup \Omega$

Figure 1: Distinct sets of distance 0

The map d defines a pseudometric topology with basic open neighborhoods $\{t \mid d(s,t) < \varepsilon\}$ for each set s and $\varepsilon > 0$, but because d is only a pseudometric, the topology does not have nice separation properties. However, if we define $s \approx t \overset{\text{def}}{\Longleftrightarrow} d(s,t) = 0$, then d is well-defined on $\approx$-equivalence classes and is a metric on the quotient space.

More interestingly, we can identify a natural class of canonical elements, one in each $\approx$-class, such that d, restricted to canonical elements, is a metric; moreover, the canonical elements are a maximal class for which this is true. Thus the quotient space is isometric to the subspace of canonical elements. The canonical elements include all the hereditarily finite sets.

The canonical elements are defined as the images of the function $F : V \to V$, defined coinductively as follows:

$$F(s) \quad \overset{\text{def}}{=} \quad \{F(u) \mid u \in \text{cl}(s)\}, \tag{6.4}$$

where cl denotes closure in the pseudometric topology. The equation (6.4) determines F uniquely, as with (6.1). A set s is called *canonical* if $s = F(t)$ for some t; equivalently, by Corollary 6.3(ii) below, if s is a fixpoint of F. For example, the right-hand side of Fig. 1 is F applied to the left-hand side, and the set on the right-hand side is canonical.

Lemma 6.1. *$d(s,t) = 0$ iff $\text{cl}(s) = \text{cl}(t)$.*

Proof. If $s = t = \varnothing$, then both sides are true. If exactly one of s, t is $\varnothing$, then both sides are false. Finally, if both $s, t \neq \varnothing$, then

$$\begin{aligned} d(s,t) = 0 \quad &\Leftrightarrow \quad \sup_{u \in s} \inf_{v \in t} d(u,v) = 0 \wedge \sup_{v \in t} \inf_{u \in s} d(u,v) = 0 \\ &\Leftrightarrow \quad \forall u \in s\ \forall \varepsilon > 0\ \exists v \in t\ d(u,v) < \varepsilon \wedge \forall v \in t\ \forall \varepsilon > 0\ \exists u \in s\ d(u,v) < \varepsilon \\ &\Leftrightarrow \quad s \subseteq \mathrm{cl}(t) \wedge t \subseteq \mathrm{cl}(s) \\ &\Leftrightarrow \quad \mathrm{cl}(s) = \mathrm{cl}(t). \end{aligned}$$

□

Theorem 6.2.

(i) *If $d(s,t) = 0$, then $F(s) = F(t)$.*
(ii) *For all s, $d(s, F(s)) = 0$; that is, $s \approx F(s)$.*

Proof. (i) By Lemma 6.1, if $d(s,t) = 0$, then $\mathrm{cl}(s) = \mathrm{cl}(t)$, and the conclusion $F(s) = F(t)$ is immediate from (6.4).

(ii) We proceed by coinduction on the definition of d. We strengthen the coinduction hypothesis $g(s, F(s)) = 0$ with the two extra assertions that $0 \leq g(s,t) \leq d(s,t)$ and that g satisfies the triangle inequality. We wish to show that this combined property holds of $\tau(g)$ under the assumption that it holds of g.

That $0 \leq \tau(g)(s,t) \leq \tau(d)(s,t) = d(s,t)$ is clear from the coinduction hypothesis and the monotonicity of the operators in the definition of τ. The argument that $\tau(g)$ satisfies the triangle inequality is equally straightforward. Thus it remains to show that $\tau(g)(s, F(s)) = 0$.

By definition of F, $s = \varnothing$ iff $F(s) = \varnothing$, and in this case $\tau(g)(s, F(s)) = 0$ by definition of τ. Otherwise $s \neq \varnothing$ and $F(s) \neq \varnothing$. To show $\tau(g)(s, F(s)) = 0$ in this case, we need to show that

$$\begin{aligned} \sup_{u \in s} \inf_{v \in F(s)} g(u,v) = \sup_{u \in s} \inf_{w \in \mathrm{cl}(s)} g(u, F(w)) \quad &= \quad 0, \\ \sup_{v \in F(s)} \inf_{u \in s} g(u,v) = \sup_{w \in \mathrm{cl}(s)} \inf_{u \in s} g(u, F(w)) \quad &= \quad 0. \end{aligned}$$

It suffices to show

$$\forall u \in s \inf_{w \in \mathrm{cl}(s)} g(u, F(w)) = 0, \qquad \forall w \in \mathrm{cl}(s) \inf_{u \in s} g(u, F(w)) = 0.$$

For the former, we can take $w = u$; then the result follows from the coinduction hypothesis $g(u, F(u)) = 0$. For the latter, let $w \in \mathrm{cl}(s)$. Here we use all three clauses of the coinduction hypothesis:

$$\inf_{u \in s} g(u, F(w)) \quad \leq \quad \inf_{u \in s} g(u,w) + g(w, F(w)) \leq \inf_{u \in s} d(u,w) + 0 = 0,$$

the last equation from the fact that $w \in \mathrm{cl}(s)$. □

Corollary 6.3.

(i) *$d(s,t) = 0$ iff $F(s) = F(t)$.*
(ii) *For all s, $F(F(s)) = F(s)$.*
(iii) *Every $\approx$-equivalence class contains exactly one canonical set, and d restricted to canonical sets is a metric. Moreover, the canonical sets are a maximal class for which this is true.*

6.4. **Compactness.** For the results of section 6.5, we need to show that the space of non-well-founded sets is compact under d, thus complete. We will show that every infinite set has a limit point. Define the equivalence relations $\approx_n$ inductively by:

$$s \approx_0 t \text{ for all } s, t \qquad\qquad s \approx_{n+1} t \overset{\text{def}}{\iff} \forall u \in s\ \exists v \in t\ u \approx_n v \\ \wedge\ \forall v \in t\ \exists u \in s\ u \approx_n v.$$

Also define inductively

$$S_0 \overset{\text{def}}{=} \varnothing \qquad\qquad S_{n+1} \overset{\text{def}}{=} 2^{S_n},$$

where 2^A denotes the powerset of A. Each S_n is a well-founded hereditarily finite set. For $n \geq 0$, define the map $f_n : V \to S_{n+1}$ inductively by

$$f_0(s) \overset{\text{def}}{=} \varnothing \qquad\qquad f_{n+1}(s) \overset{\text{def}}{=} \{f_n(u) \mid u \in s\}.$$

The following properties of S_n, $\approx_n$, and f_n are easily established by induction on n.

Lemma 6.4. *For all $s, t \in V$ and $m, n \geq 0$,*

(i) *$f_n(s) \in S_{n+1}$;*
(ii) *if $s \in S_{n+1}$ then $f_n(s) = s$;*
(iii) *$s \approx_n f_n(s)$;*
(iv) *$f_n(f_m(s)) = f_{\min m,n}(s)$;*
(v) *if $s, t \in S_{n+1}$ and $s \approx_n t$, then $s = t$.*

Lemma 6.5. *For all $s, t \in V$ and $n \geq 0$, the following are equivalent:*

(i) *$s \approx_n t$;*
(ii) *$f_n(s) = f_n(t)$;*
(iii) *$d(s, t) \leq 2^{-n}$.*

For each $s \in V$, let $f(s)$ denote the sequence $f_0(s), f_1(s), f_2(s), \ldots$. It follows from Lemma 6.4(iv) that $f_n(f_{n+1}(s)) = f_n(s)$. Moreover, we have the following representation theorem as converse:

Lemma 6.6. *Any sequence $s_0, s_1, s_2, \ldots$ such that $f_n(s_{n+1}) = s_n$ for all $n \geq 0$ is $f(s)$ for some s.*

Proof. Let W be the set of all sequences $s = s_0, s_1, s_2, \ldots$ such that $s_n = f_n(s_{n+1})$, $n \geq 0$. This is a set, since the defining condition implies $s_n \in S_{n+1}$. Consider the system with nodes W and edges N defined by

$$u\ N\ s \quad \overset{\text{def}}{\iff} \quad \forall n \geq 0\ u_n \in s_{n+1}.$$

We claim that $f_n(s) = s_n$. The proof is by induction on n. Certainly $f_0(s) = \varnothing = s_0$, since $s_0 = f_0(s_1) \in S_1$ and $\varnothing$ is the only element of S_1. Now suppose the claim is true for n. Then

$$\begin{aligned} f_{n+1}(s) &= \{f_n(u) \mid u \in W,\ u\ N\ s\} = \{f_n(u) \mid u \in W,\ \forall k\ u_k \in s_{k+1}\} \\ &= \{u_n \mid u \in W,\ \forall k\ u_k \in s_{k+1}\} = s_{n+1}. \end{aligned}$$

The last equation requires that for all $a \in S_{n+1}$, there exists $u \in W$ such that $u_n = a$. The sequence $u = f_0(a), f_1(a), f_2(a), \ldots$ does it. □

Lemma 6.7. *The space V is compact under the pseudometric d, therefore complete.*

Proof. We wish to show that every infinite set s has a limit point p (not necessarily contained in s). Let W be the tree of all sequences $u_0, u_1, u_2, \ldots$ such that $f_n(u_{n+1}) = u_n$ for all $n \geq 0$ as defined in the proof of Lemma 6.6. This is a finitely branching, infinite tree with root $\varnothing$. By König's lemma, there is an infinite path p in W such that for every node p_n on the path, there are infinitely many $u \in s$ such that $f_n(u) = p_n$. The set represented by the path p as given by Lemma 6.6 is the desired limit point, since for all k, there exist infinitely many $u \in s$ such that $f_k(u) = p_k = f_k(p)$, therefore $d(u, p) \leq 2^{-k}$. □

6.5. **Hereditarily Finite Sets Are Canonical.** Let φ be a property of sets. We define a set to be *hereditarily* φ (Hφ) if it has an APG representation in which every node represents a set satisfying φ. Equivalently, Hφ is the largest solution of

$$\mathrm{H}\varphi(s) \quad \overset{\text{def}}{\iff} \quad \varphi(s) \wedge \forall u \in s\ \mathrm{H}\varphi(u).$$

The *hereditarily finite* (HF) sets are those possessing an APG representation in which every node has finite out-degree (not necessarily bounded). Note that this differs from Aczel's definition [2, p. 7]. Aczel defines a set to be hereditarily finite if it has a finite APG, which is a much stronger condition. Aczel's definition and ours coincide for well-founded sets by König's lemma, but not for non-well-founded sets in general. For example, the set $f(0)$, where f is defined coinductively by $f(n) = \{\overline{n}, f(n+1)\}$ (Fig. 2), is hereditarily finite in our sense but not Aczel's. We would prefer the term *regular* or *rational* for sets that are hereditarily finite in Aczel's sense, since they are exactly the sets that have a regular or rational tree representation [6].

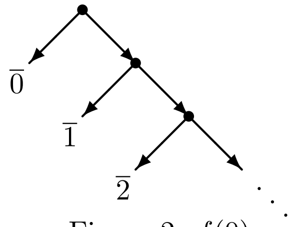

Figure 2: $f(0)$

A set is *hereditarily closed* (HC) if it has an APG representation in which every node represents a closed set in the pseudometric topology. Recall that a set is *singular* if it forms a singleton $\approx$-class.

Lemma 6.8. *If s is singular, then all elements of s are singular. Thus all singular sets are hereditarily singular.*

Proof. Suppose $u \in s$, $v \neq u$, and $d(u, v) = 0$. We claim that (i) if $v \notin s$, then $d(s, s \cup \{v\}) = 0$, and (ii) if $v \in s$, then $d(s, s - \{v\}) = 0$, thus in either case, s is not singular.

In case (i), we must show

$$\sup_{x \in s} \inf_{y \in s \cup \{v\}} d(x, y) = 0, \qquad \sup_{y \in s \cup \{v\}} \inf_{x \in s} d(x, y) = 0.$$

It suffices to show

$$\forall x \in s\ \exists y \in s \cup \{v\}\ d(x, y) = 0, \qquad \forall y \in s \cup \{v\}\ \exists x \in s\ d(x, y) = 0.$$

The former is immediate by picking $y = x$. For the latter, pick $x = y$ if $y \neq v$, otherwise pick $x = u$.

Case (ii) is really the same case as (i), with $s - \{v\}$ in (ii) playing the role of s in (i). □

Lemma 6.9.

(i) *If s is closed and all elements of s are closed, then all elements of s are singular.*
(ii) *Every singular set is closed.*

Proof. (i) Suppose $u \in s$ and $d(u,v) = 0$. Then $v \in s$, since s is closed. By Lemma 6.1, $\mathrm{cl}(u) = \mathrm{cl}(v)$. But u and v are both closed, so $u = v$.

(ii) By Lemma 6.1, $d(\mathrm{cl}(u), u) = 0$, so if u singular then $u = \mathrm{cl}(u)$. □

Theorem 6.10. *A set is hereditarily closed if and only if it is closed and all its elements are singular.*

Proof. This follows directly from Lemmas 6.8 and 6.9. □

Theorem 6.11. *A set is singular if and only if it is hereditarily finite.*

Proof. Suppose first that s is hereditarily finite (HF). Consider the binary relation on sets s, t defined by

$$\mathrm{HF}(s) \wedge d(s,t) = 0. \tag{6.5}$$

We have

$$\begin{aligned} \mathrm{HF}(s) \wedge d(s,t) = 0 \quad &\Rightarrow \quad \forall v \in t\ \forall \varepsilon > 0\ \exists u \in s\ \mathrm{HF}(u) \wedge d(u,v) < \varepsilon \\ &\Rightarrow \quad \forall v \in t\ \exists u \in s\ \mathrm{HF}(u) \wedge d(u,v) = 0, \end{aligned} \tag{6.6}$$

since u is finite. It follows that

$$\begin{aligned} \mathrm{HF}(s) \wedge \mathrm{HF}(t) \wedge d(s,t) = 0 \quad \Rightarrow \quad &\forall u \in s\ \exists v \in t\ \mathrm{HF}(u) \wedge \mathrm{HF}(v) \wedge d(u,v) = 0 \\ &\wedge\ \forall v \in t\ \exists u \in s\ \mathrm{HF}(u) \wedge \mathrm{HF}(v) \wedge d(u,v) = 0, \end{aligned}$$

so the binary relation $\mathrm{HF}(s) \wedge \mathrm{HF}(t) \wedge d(s,t) = 0$ is a bisimulation; thus

$$\mathrm{HF}(s) \wedge \mathrm{HF}(t) \wedge d(s,t) = 0 \quad \Rightarrow \quad s = t.$$

Thus if $\mathrm{HF}(s)$, then there is a positive lower bound $\delta > 0$ on $d(u,v)$ for $u, v \in s$, $u \neq v$. But then

$$\begin{aligned} \mathrm{HF}(s) \wedge d(s,t) = 0 \quad &\Rightarrow \quad \forall u \in s\ \forall \varepsilon > 0\ \exists v \in t\ \mathrm{HF}(u) \wedge d(u,v) < \varepsilon \\ &\Rightarrow \quad \forall u \in s\ \exists v \in t\ \mathrm{HF}(u) \wedge d(u,v) < \delta, \end{aligned}$$

and using (6.6), this gives

$$\begin{aligned} &\mathrm{HF}(s) \wedge d(s,t) = 0 \\ &\quad \Rightarrow \quad \forall u \in s\ \exists v \in t\ \mathrm{HF}(u) \wedge d(u,v) < \delta \wedge \exists w \in s\ d(w,v) = 0 \\ &\quad \Rightarrow \quad \forall u \in s\ \exists v \in t\ \exists w \in s\ \mathrm{HF}(u) \wedge d(w,v) = 0 \wedge d(u,w) < \delta \\ &\quad \Rightarrow \quad \forall u \in s\ \exists v \in t\ \exists w \in s\ \mathrm{HF}(u) \wedge d(w,v) = 0 \wedge u = w \\ &\quad \Rightarrow \quad \forall u \in s\ \exists v \in t\ \mathrm{HF}(u) \wedge d(u,v) = 0. \end{aligned}$$

This combined with (6.6) says that the relation (6.5) itself is a bisimulation. Thus $\mathrm{HF}(s) \wedge d(s,t) = 0$ implies $s = t$; in other words, $\mathrm{HF}(s)$ implies that s is singular.

Now suppose that s is singular. By Lemma 6.8, s is hereditarily singular. We argue that s must be finite. If s is infinite, then by Lemma 6.7, s has a limit point p (not necessarily contained in s). We claim that (i) if $p \notin s$, then $d(s, s \cup \{p\}) = 0$, and (ii) if $p \in s$, then $d(s, s - \{p\}) = 0$, thus in either case s is not singular. For (i),

$$\begin{aligned} d(s, s \cup \{p\}) = 0 \quad \Leftrightarrow \quad & \forall u \in s \; \forall \varepsilon > 0 \; \exists v \in s \cup \{p\} \; d(u,v) < \varepsilon \\ & \wedge \; \forall v \in s \cup \{p\} \; \forall \varepsilon > 0 \; \exists u \in s \; d(u,v) < \varepsilon. \end{aligned}$$

The first clause is true by taking $v = u$. For the second clause, we can take $u = v$ unless $v = p$. But if $v = p$, the condition reduces to

$$\forall \varepsilon > 0 \; \exists u \in s \; d(u,p) < \varepsilon,$$

which is true by Lemma 6.5.

Case (ii) is really the same as case (i), with $s - \{p\}$ in (ii) playing the role of s in (i). □

Theorem 6.12. *Every hereditarily finite set is heretarily closed, and every hereditarily closed set is canonical. Both implications are strict.*

Proof. The first implication $\mathrm{HF}(s) \Rightarrow \mathrm{HC}(s)$ follows directly from Lemma 6.9(ii) and Theorem 6.11.

For the implication $\mathrm{HC}(s) \Rightarrow s = F(s)$, one approach would be to show that the binary relation on sets s, t defined by $\mathrm{HC}(s) \wedge t = F(s)$ is a bisimulation. Alternatively, we can observe that on hereditarily closed sets s, the coinductive definition

$$F(s) \;\; \stackrel{\mathrm{def}}{=} \;\; \{F(u) \mid u \in \mathrm{cl}(s)\}$$

is equivalent to the coinductive definition

$$F(s) \;\; \stackrel{\mathrm{def}}{=} \;\; \{F(u) \mid u \in s\},$$

which uniquely defines the identity function, thus $s = F(s)$ on all such sets.

Both implications are strict. An hereditarily closed set that is not finite is $\{\overline{n} \mid n \geq 0\} \cup \Omega$, and a canonical set that is not closed is $\{\{\overline{n} \mid n \geq 0\} \cup \Omega\}$. □

7. Conclusions and Future Work

We have illustrated the use of the metric coinduction principle in four areas: infinite streams, Markov chains, Markov decision processes, and non-well-founded sets. In all these areas, metric coinduction can be used to simplify proofs or derive new insights.

Other areas are likely to be amenable to such techniques. In particular, iterated function systems seem to be a promising candidate.

Acknowledgements

Thanks to Lars Backstrom, Radha Jagadeesan, Prakash Panangaden, and Jan Rutten for valuable comments. This work was supported in part by ONR Grant N00014-01-1-0968 and by NSF grant CCF-0635028.

References

[1] Samson Abramsky. A cook's tour of the finitary non-well-founded sets. In Sergei Artemov, Howard Barringer, Artur d'Avila Garcez, Luis C. Lamb, and John Woods, editors, *We Will Show Them: Essays in Honour of Dov Gabbay*, volume 1, pages 1–18. College Publications, 2005.

[2] P. Aczel. *Non-Well-Founded Sets*, volume 14 of *CSLI Lecture Notes*. Stanford University, 1988.

[3] M. Barnsley. *Fractals Everywhere*. Academic Press, 1993.

[4] Jon Barwise and Lawrence Moss. *Vicious Circles: On the Mathematics of Non-Wellfounded Phenomena*, volume 60 of *CSLI Lecture Notes*. Center for the Study of Language and Information (CSLI), Stanford University, 1996.

[5] P. Brémaud. *Markov Chains, Gibbs Fields, Monte Carlo Simulation and Queues*. Texts in Applied Mathematics. Springer-Verlag, 1999.

[6] Bruno Courcelle. Fundamental properties of infinite trees. *Theor. Comput. Sci.*, 25:95–169, 1983.

[7] Jaco de Bakker and Erik de Vink. *Control Flow Semantics*. MIT Press, 1996.

[8] Eric V. Denardo. Contraction mappings in the theory underlying dynamic programming. *SIAM Review*, 9(2):165–177, April 1967.

[9] Nelson Dunford and Jacob T. Schwartz. *Linear Operators: Part I: General Theory*. John Wiley, 1957.

[10] W. Feller. *An Introduction to Probability Theory and its Applications*, volume 1. Wiley, 1950.

[11] Marcelo P. Fiore. A coinduction principle for recursive data types based on bisimulation. *Information and Computation*, 127(2):186–198, 1996.

[12] Olle Häggström. *Finite Markov Chains and Algorithmic Applications*. Cambridge University Press, 2002.

[13] D. Isaacson and R. Madsen. *Markov Chains: Theory and Applications*. John Wiley and Sons, 1976.

[14] Dexter Kozen. Coinductive proof principles for stochastic processes. *Logical Methods in Computer Science*, 3(4:8), 2007. DOI: 10.2168/LMCS-3 (4:8) 2007.

[15] Dexter Kozen and Nicholas Ruozzi. Applications of metric coinduction. In T. Mossakowski et al., editor, *Proc. 2nd Conf. Algebra and Coalgebra in Computer Science (CALCO 2007)*, volume 4624 of *Lecture Notes in Computer Science*, pages 327–341. Springer, August 2007.

[16] Henryk Minc. *Nonnegative Matrices*. John Wiley, 1988.

[17] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[18] H. Peitgen and P. Richter. *The Beauty of Fractals: Images of Complex Dynamical Systems*. Springer-Verlag, 1986.

[19] J. J. M. M. Rutten. Behavioural differential equations: a coinductive calculus of streams, automata, and power series. *Theoretical Computer Science*, 308:1–53, 2003.

[20] J.J.M.M. Rutten. Universal coalgebra: A theory of systems. *Theor. Comput. Sci.*, 249:3–80, 2000.

[21] Daniele Turi. *Functorial Operational Semantics and its Denotational Dual*. PhD thesis, Free University, Amsterdam, June 1996.