Coinductive Models of Finite Computing Agents
Peter Wegner and Dina Goldin, Draft, January 25, 1999

Abstract: This paper explores the role of coinductive methods in modeling finite interactive computing agents. The computational extension of computing agents from algorithms to interaction parallels the mathematical extension of set theory and algebra from inductive to coinductive models. Maximal fixed points are shown to play a role in models of observation that parallels minimal fixed points in inductive mathematics. The impact of interactive (coinductive) models on Church's thesis and the connection between incompleteness and greater expressiveness are examined. A final section shows that actual software systems are interactive rather than algorithmic. Coinductive models could become as important as inductive models for software technology as computer applications become increasingly interactive.

Contents

1. Extensions of Expressiveness

The robust equivalence of Turing machines, the lambda calculus, and recursively enumerable sets has led to widespread acceptance of the thesis that Turing machines express the intuitive notion of computation. Because researchers assumed that questions of expressiveness of finite computing agents had been settled once and for all, they did not explore alternative models of computability, focusing instead on questions of complexity, performance, and design of algorithms. The hypothesis that interactive finite computing agents are more expressive than algorithms opens up a research area that had been considered closed, and requires reexamination of fundamental assumptions about models of computation.

We have shown [We1, WG] that interaction machines that provide services over time are more expressive than Turing machines that compute functions. The extension of expressiveness from algorithms to interaction is formalized by recent extensions of set theory and algebra [Ac, BM, JR]. Non-well-founded set theory models sequential interaction by formalizing the semantics of streams [BM]:

  extension of computational expressiveness: algorithms -> interaction
  extension of set theoretic expressiveness: well-founded sets -> non-well-founded sets.
  extension of algebraic expressiveness: algebras -> coalgebras

Models of interaction, just like algorithmic models, can be described in complementary ways by sets and algebras. Non-well-founded sets are the interactive analog of recursively enumerable sets for Turing machines, while coalgebras play the role of the lambda calculus. Coinduction underlies non-well-founded set theory, coalgebra, and interaction as an abstract principle of definition, reasoning, and modeling. Coinduction is related to *abduction*, which infers inner system properties from their observed behavior [We3].

  extension of mathematical expressiveness: inductive models -> coinductive (abductive) models

The equivalent expressiveness of TMs, algorithms, computable functions, and formal systems is due largely to their common use of induction as a basis for system specification.

The Church-Turing thesis has the form "the formally definable notion X corresponds to an intuitive notion Y". It equates the intuitive notion of algorithmic computation with the formal notion of Turing computable functions from integers to integers:

Church-Turing Thesis: The intuitive notion of algorithms is formally expressed by Turing machines.
X = algorithms, Y = Turing machines
In section 5 we extend the Church-Turing thesis from algorithms to sequential and multi-agent interaction by extending both the intuitive notion X of computation and the associated formal notion Y:
E1: The intuitive notion of sequential interaction is formally expressed by non-well-founded sets.
X = sequential (single-stream) interaction, Y = non-well-founded sets
E2: General interactive computing is formally expressed by coinductive models (coalgebras).
X = general (multi-stream) interaction, Y = coinductive models (coalgebras)
The extensions E1, E2 of the Church-Turing thesis provide mathematical legitimacy for models of interaction. Non-well-founded set theory and coalgebras are coinductive extensions of inductive formal models of computation that express extension of the intuitive algorithmic model of computation.

Section 2 introduces interaction machines as stream-processing agents and observation equivalence (bisimulation) as a metric for interactive expressiveness. Section 3 introduces non-well-founded set theory and coalgebras and indicates how they model streams and labeled transition systems. Section 4 shows that eliminating initiality and replaceming minimality by maximality explain the greater expressiveness of coinduction over induction. Section 5 examines metamathematical implications of computing, including the extension of Church's thesis, ontological commitment in mathematical models, and the role of incompleteness theorems. Section 6 examines software engineering, showing that interactive models provide a unifying framework for object-oriented and agent-oriented programming.

2. From Turing to Interaction Machines

Turing machines (TMs) are finite agents that noninteractively transform input into output strings by sequences of actions. TM computations for a given input are history-independent and reproducible, since TMs always start in the same initial state and shut out the world during computation.

Turing machine: TMs are state-transition machines $M = (S, T, s_0, F)$, with finite sets of states S and tape symbols T, a starting state s_0, and a state-transition relation F: SxT -> SxT. TMs transform finite input strings $x \in T^*$ to outputs $y = M(x)$ by a finite sequence of computation steps that read a symbol i from the tape, perform a state transition (s,i) -> (s',o), write a symbol o, and/or move the reading head one position right or left [Tu1].

Interaction machines (IMs) play a role in modeling finite interactive agents comparable to *Turing machines (TMs)* for algorithms. S*equential interaction machines (SIMs)* model sequential interaction, transducing input into output streams.

Sequential interaction machine: SIMs are state-transition machines $M = (S, A, m)$ where S is an enumerable set of states, A is an enumerable set of dynamically bound actions (stream elements), and the transition mapping m: SxA -> S maps state-action pairs into new states. Dynamic binding of inputs is explicitly indicated by rewriting m as a mapping S -> P(AxS) from states to action-state pairs that expresses *input nondeterminism* of actions A.

Expressiveness of finite agents is specified by *observation equivalence* and is measured by the ability of observers to distinguish agent behavior [WG]. TM behaviors consist of sets of input-output pairs, while SIM behaviors consist of sets of sequences of observations. Given two finite agents M1, M2, their behavior $B(M_1)$, $B(M_2)$ is *distinguishable* relative to a class of observers (testers) if the exclusive-or B(M1)⊕B(M2) is nonempty. Members of this set are called *distinguishability certificate*s. Finite agents are *equivalent* if they cannot be distinguished. Equivalence of two SIMs can be falsified by a finite distinguishability certificate (observation sequence) but cannot be finitely verified, illustrating Popper's point that empirical laws of physics can be empirically falsified but not verified [Po]. Interactive equivalence expressed by *bisimulation* [Mi] has greater distinguishing power than off-line equivalence.

Inputs of TMs and computable functions are completely specified at the start of a computation. Com-

putable functions f: X -> Y have a domain X of input strings (integers) and determine a unique $y = f(x)$ for each $x \in X$. Interactive behavior cannot be specified as transformation from a domain X to a range Y because streams are dynamically generated.

Interaction streams have the form $(a_1, o_1), (a_2, o_2), \dots$, where output o_k is computed from action a_k but precedes and can influence a_{k+1}. This *input-output coupling* violates the separation of domains and ranges, causing dynamic dependence of inputs on prior outputs that is characteristic of interactive question-answering, dialog, and control processes. Late (lazy) binding of interactive inputs (input nondeterminism) is more expressive than early (static) binding not because interaction has greater function transformation power but because the richer input domains cannot be inductively specified by strings.

Persistent Turing machines (PTMs) are a canonical model for sequential interaction [GW].

Persistent Turing Machines: A PTM = (S, A, f) is a multitape TM with a persistent work tape whose content is preserved between interactions. S, the set of PTM's *states*, consists of the states of the persistent worktape (i.e., its contents); A is the set of PTM's *actions*, or inputs; f is a *mapping*, f: AxS -> SxO, computed by the underlying TM. Elements of both S and A are finite at any given time but unbounded.

To understand the extra power of SIMs and PTMs over TMs, it is helpful to consider the extra power of TMs over finite lookup tables (FLTs), which return a value when given the corresponding key. They may be considered as a primitive finite computing agent whose expressiveness is that of propositional calculus. A single-step transition for a TM computation consists of performing table look-up of the new TM state and tape action based on the current state and tape character. Besides table look-up, the only additional work performed by a TM during each step consists of 'remembering' the state and performing tape I/O. A TM computation as a whole is more expressive than an FLT computation because its behavior is no longer finite. Similarly, Turing-computable single-step transitions are at the heart of SIM computations. SIM computations are more expressive than those of TMs, since their behavior is no longer enumerable.

The informal argument that SIMs are not expressible by algorithms is surprisingly simple, since it depends only on the notion that stream elements are supplied dynamically.

Proposition: Computations of SIMs are not expressible by or reducible to computations of TMs.

Informal proof: TMs cannot be represented by finite inputs because any finite input can always be interactively extended. Transduction of streams cannot be modeled by transformation of strings. Streams do not have a last element, can be dynamically extended by unpredictable adversaries, and can use previous outputs in determining the next input. TM input strings have a predetermined finite length that cannot be changed once the computation has started.

A more refined argument shows that SIMs are more expressive than TMs for finite computations. Let k be the number of interaction steps (length) of an interactive computation, and let B_k be the class of behaviors that can be observed by observers who can make k or fewer interactions. We can show that, for all $k>0$, $B_{k+1} > B_k$, while the behavior of TMs corresponds to B_1 [WG]. Computations are more expressive for k+1 interactions than for k interactions in the sense that they can make finer observational distinctions. Bisimulation [Mi] captures the greater observational distinguishing power of k+1 over k inputs.

SIMs that buffer the output of their first k interactions can be distinguished by k+1 but not by k interactions. Interactive questioners (Ken Starr) can elicit more information about a questioned subject (Clinton) by k+1 than by k follow-up questions. The game 20 Questions further illustrates the power of follow-up questions in gaining information. TMs have the expressiveness of PTMs with k=1: their behavior is limited to answering a single question (or a predetermined noninteractive sequence of questions). The expressiveness of sequential interaction has been examined in greater detail in [WG], both for question-answering and for the Turing test [Tu2].

Multistream interaction machines (MIMs) are finite agents that interact with multiple autonomous agents, like airline reservation systems or distributed databases that provide services to multiple autonomous clients. MIMs are shown to be more expressive than and not reducible to SIMs [WG], contrasting

with the fact that multitape TMs are no more expressive than single-tape TMs. The proof that MIMs are more expressive than SIMs shows that input-output coupling of streams cannot be preserved under merging. We show that input-output coupling (serializability) cannot be preserved for delegation of subtasks and other forms of nonserializable behavior:

Input-output coupling of autonomous streams cannot be preserved under merging

MIMs allow collaboration and coordination to be precisely modeled. The greater expressiveness of MIMs over SIMs captures the greater problem-solving power of senior managers and CEOs than workers who perform sequences of interactively assigned tasks. Adding autonomous streams (observation channels) to a finite agent increases expressiveness, while adding noninteractive tapes simply increases the structural complexity of predetermined inputs, and does not affect expressive power.

Collaboration and coordination cannot be modeled by SIMs or PTMs

The behavior of MIMs is harder to formalize than that of SIMs, and greater expressiveness of MIMs than SIMs is harder to prove than greater expressiveness of SIMs than TMs. MIMs support the behavior of nonserializable transactions and true concurrency [Pr], while SIMs support only serializable transactions and interleaving concurrency. The metric of distinguishability (observation nonequivalence) determines an expressiveness hierarchy for sequential interaction and greater expressiveness of MIMs than SIMs.

3. Non-Well-Founded Sets and Coalgebras

Though questions of expressiveness can be formulated and proved entirely in terms of machine-based notions, reformulation in terms of non-well-founded set theory and coalgebras lends mathematical legitimacy to interactive models.

A binary relation R on a set S is *non-well-founded* if there is an infinite sequence $b_i \in S$ for $i = 0,1,2,...,$ such that $b_{i+1} R b_i$ for all i, and is *well-founded* otherwise. Non-well-founded sets are defined in terms of non-well-foundedness of the set membership relation $\in$. A set is well-founded if the set membership relation over its structure is well-founded and is non-well-founded otherwise. *Well-founded sets* have only finite set-membership chains, while *non-well-founded sets* may have infinite set-membership chains.

Sets in traditional set theory are inductively constructible from atomic elements by union, intersection, complement, and cross-product constructors that construct sets with atomic elements. The powerset constructor constructs sets that contain sets as elements. Zermelo-Frankel set theory provides an axiomatic specification ZFC- of set theory that can be supplemented by the *foundation axiom* to obtain the inductively specifiable sets ZFC.

foundation axiom (FA): all sets are well-founded (inductively definable).

FA excludes non-well-founded sets, which are defined by systems of flat set equations.

flat set equations: A system of flat set equations has the form (S, A, m), where S is a set of variables (states), A is a set of constants (actions, observations) and m: S -> P(A$\cup$S) is a mapping that specifies for each $s \in S$ a set of constants in A and variables in S. A *non-well-founded set* can be defined as a set that is the solution of a system of flat set equations.

Equations of the form m: S -> P(AxS) whose right hand sides are sets of ordered pairs can be expressed as flat equations by the transformation (a,b) -> {a, {a,b}} [BM].

Example: The solution of the set equation s = {(0,s), (1,s)} specifies the nonenumerable infinite binary strings, illustrating that non-well-founded sets may be nonenumerable.

Non-well-founded set theory augments ZFC- with the *anti-foundation axiom (AFA)* to obtain ZFA. [BM] specifies AFA by the strong condition that all flat systems of equations have unique solutions.

AFA: Every flat system of equations has a unique solution (ZFC- + AFA = ZFA).

ZFA is consistent if ZFC- is consistent and admits a larger class of sets (models a larger class of systems) than ZFC. The solution of the flat equation s = {(0,s), (1,s)} does not satisfy FA and cannot be inductively generated. It is not a member of ZFC, illustrating that the class ZFA is strictly larger than ZFC.

Non-well-founded sets can model **labeled transition systems** (LTSs). An LTS with n states can be

specified by a system of n simultaneous set equations (S, A, m), where m:S -> P(AxS) maps states into subsets of action-state pairs. For example, a state with outgoing edges labeled a to s1, b to s2, and c to s3 is represented by the mapping s -> {(a,s1), (b,s2), (c,s3)}.

The mapping function m for systems of flat set equations is coalgebraic. Coalgebras are a model for interactive computing that parallels algebras as a modeling framework for algorithms. Algebraic computing steps specify incremental expression evaluation whose goal is to compute a value, while coalgebraic computing steps specify incremental system observation whose goal is to determine system behavior.

Algebras are structures A = (S, m: F(S) -> S), where S is the carrier set and F is a functor that determines the signature of A. We usually interpret S as a set of values and m: F(S) -> S is a value-preserving homomorphic mapping m from a syntactically specified set of expressions into a value set. Algebras determine reduction processes (mappings) from an inductively defined *initial algebra* of syntactically specified expressions to a quotient algebra that determines the value set.

Coalgebras are structures CA = (S, m: S -> Γ(S)), where S is a carrier set and Γ is a functor that determines the signature of CA. We usually interpret S as a set of states of an observed system and m: S -> Γ(S) as a behavior-preserving homomorphism of observed systems with an unknown state into unfolded set of behaviors (observation sequences). Coalgebras whose mappings are one-to-one are called *final coalgebras*. States of final coalgebras represent unfolded system behaviors and are specified by non-well-founded sets.

Initial algebras are an inductively specified set of expressions, while final coalgebras are coinductively specified non-well-founded sets. The greater richness of non-well-founded over well-founded sets translates into greater computation power of coalgebras over algebras. Coalgebraic homomorphisms that map systems into their behaviors are stronger than algebraic homomorphisms that map expressions into values because equivalence classes of systems with the same behavior (for example, programs that compute the same function) are nonenumerable, while classes of expressions with the same value are enumerable.

Coalgebras that model non-well-founded sets have mappings of the form m: S -> P(A∪S), corresponding to the AFA. It is shown in [BM] that mappings such as m: S -> P(AxS) of LTSs or mappings m: S -> OxP(AxS) of automata are expressible as flat systems of equations. These mappings are "Markovian" in the sense that the new state depends only on the previous state. We have suggested that transaction systems (MIMs) may require more complex non-Markovian mappings of the form Γ(S) -> Γ'(S) such that the next state depends on several previous states. Interaction of overlapping operations with duration or multi-agent interaction may require still more complex mappings. Moreover, we conjecture that consistent extensions of set theory that admit larger sets than the non-well-founded sets are possible, corresponding to nonsequential and multistream finite agents [WG]. In principle, any finite agent can be modeled by a consistent extension of set theory that expresses the class of behaviors computable by the agent and may have a coalgebraic mapping function more complex than that of non-well-founded sets.

4. From Induction to Coinduction

Induction determines a construction paradigm for definition, reasoning, and modeling characterized by an *initiality condition*, an *iteration condition* that allows new elements to be constructed (derived) from initial elements, and a *minimality condition* that only elements so constructed are definable.

inductive definition: initiality (base) condition, constructive iteration condition, minimality condition
construction paradigm: generate structures inductively from base elements

Induction is the basis for defining sets of strings, languages, formal systems, and computable functions. The set A* of all strings over an alphabet A is inductively defined by an iterative concatenation rule:
Strings: (1) the empty string e∈ A*; (2) if x∈ A* then ax∈ A*; (3) A* contains no other elements.
Grammars G define languages L(G) over terminal symbols T with the aid of nonterminal symbols N as sets of strings inductively generated from an initial nonterminal S by generating rules (productions) P.

A **language** L(G) over T is the set L(G) = {x | S => x and x∈ T*}, where G = (N, T, S, P) is a grammar with nonterminals N, terminals T, initial nonterminal symbol S and productions P. The derivation relation => is the multistep transitive closure of the single-step derivation relation defined by P.

Theorems of a formal system are inductively derived from axioms by rules of inference.

A **formal system** FS = (AX, TH, RI) is an inductive specification of a set TH of theorems, where (1) axioms AX are theorems; (2) formulae derived from theorems by rules of inference RI are theorems; and (3) no other formulae are theorems.

Computable functions f: X -> Y are inductive at two distinct levels. They map an inductively defined domain X to a range Y by inductively defined computing processes [Tu1]. TMs transform inductively-defined strings by inductively defined state-transition steps. This condition is sufficient as well as necessary: computable functions can be defined as inductive computations over inductively defined domains.

Inductively defined domain (statics): The domain X is inductively defined.

Inductively defined computation (dynamics): The process of computation is inductively defined.

Formal systems likewise make use of two levels of inductive definition. Well-formed formulae and axioms are defined by static induction, while processes of proof are defined by dynamic induction. The set of provable theorems is defined by inductively derivable formulae from inductively defined axioms.

Coinduction determines a *deconstruction paradigm* that deconstructs composite structures into progressively more primitive ones. Non-well-founded sets are coinductively defined by a process of progressive decomposition into subsets that terminates finitely for well-founded (inductive) sets. Coinduction reverses the direction of iteration of an associated inductive process and replaces initiality with finality. Non-well-founded sets introduce a larger class of sets by eliminating finite termination, thereby eliminating initiality of dual inductive processes.

Coinduction models processes of *observation*. Elimination of finite termination corresponds to the fact that processes of observation can continue to reveal new knowledge about the observed objects indefinitely: hidden information is progressively approximated by processes that do not terminate. Coinduction reverses the direction of iteration, eliminates the initiality/finality condition, and replaces the minimality condition by a maximality condition.

coinductive definition: deconstructive iteration condition, maximality condition.

observation paradigm: observe and transduce already existing constructed elements.

A *stream* over A is an ordered pair s = (a, t), where $a \in A$ and t is another stream. The set S of all streams over A is the maximal solution (maximal fixed point) of the equation S = AxS. Sets that are solutions of recursive stream equations do not exist in traditional set theory, where stream equations have the trivial empty set as their solution. The minimal fixed point of the equation S = AxS is the empty set, while the maximal fixed point is the set of all streams over S.

By expressing induction and coinduction in terms of more primitive concepts, we can separately consider the effect of reversing iteration, eliminating initiality, and replacing minimality by maximality in definition and reasoning processes. Finite coinductive termination, like initiality, is a closed-system requirement whose elimination models open systems. Elimination of initiality is related to shedding of the initial-state (and final-state) requirement in extending Turing machines to interaction machines. It removes the requirement of complete environment specification before the computation starts.

Minimality, modeled by least fixed points, is a property of constructive processes of computation and constructive mathematics, while maximality, modeled by greatest fixed points, is a property of empirical observation paradigms for describing observed behavior in an already constructed (already existing) world. Minimality of behavior is associated with maximality of constraints on behavior, while maximality of behavior is associated with minimality of constraints. Maximal fixed points (minimal constraints) provide a mathematical framework for the empirical paradigm that any behavior (possible world) consistent with observation is admissible. Specifications that admit any possible world consistent with a specification are maximal fixed points (minimal constraints on behavior). The distinction between possible-world semantics of Kripke [Kr] and traditional model theory is precisely that of maximality versus minimality. Distinctions between restrictive and permissive social organization, such as those between totalitarian and democratic societies, are modeled by minimality versus maximality. Minimality models centrally controlled structures while maximality admits distributed control.

totalitarianism embodies the minimality principle: **everything is forbidden that is not allowed.**

democracy supports the maximality principle: **everything is allowed that is not forbidden.**

Non-well-founded set theory extends traditional set theory with sets that are solutions to recursive set equations. It provides a framework for formalizing sets with a coinductive (non-well-founded) structure, and correspondingly extends the class of models that sets can formalize to include interactive models of computation. Coinduction extends the definition and reasoning ability of finite computing agents so they can model nonenumerable sets defined by streams.

Algorithmic denotational semantics is specified by lattice-structured approximations to least fixed points that determine computable functions [Sc], while traditional operational semantics expresses the state-transition structure of algorithm execution steps. Non-well-founded sets express richer denotations than well-founded sets, specifying maximal fixed points. Non-well-founded set theory provides a denotational semantics for streams, while coalgebras provide a framework for operational semantics of interaction machines that are transducers of streams [WG].

5. Metamathematics of Coinduction

The extension from inductive to coinductive models requires reexamination of the role of formalism in models of both mathematics and computing. The views of Church and Turing in the early 1930s were strongly influenced by the intense debate among logicist, formalist, and empirical schools of mathematical modeling in the early 20th century, which were in turn influenced by Kant's earlier analysis of the distinction between necessary truths of mathematics and contingent truths of physics. We examine the relation between intuitive and formal notions of computing exemplified by the Church-Turing thesis, reexamine notions of constructive, formalist, and realist mathematics, and consider the role of Godel-style incompleteness theorems for coinductive models of computation.

Church's interest in connections between formal models and intuitive concepts of computing was due at least in part to the debate between Brouwer and Hilbert [De] on the role of intuition in logic and mathematical modeling. Brouwer, following Kant, believed that mathematical reasoning was based on inner intuitions, while Hilbert believed that mathematical reasoning was based on formal (objective) rules of logic. However, Brouwer's intuitionism was weaker than Kant's because non-Euclidean geometry had discredited the idea of geometric (spatial) intuition, leaving only a weaker form of arithmetic intuition that excluded nonconstructive reasoning or proof and became increasingly identified with constructive mathematics [De]. Hilbert's formalism was likewise weak in the sense that it excluded consistent circular (coinductive) reasoning [BM] as an overreaction to the paradoxes of set theory [WG].

Coinductive models can be viewed either as a new paradigm that abandons intuitionism and formalism or as a strengthening of intuitionism and formalism that reconciles the role of intuition and formalism in mathematics by allowing these notions to coincide. Coinductive models support a larger class of mathematical models than traditional intuitionist or formal models because of a stronger ontological commitment to the existence of mathematical objects.

5.1 Coinduction and Church's Thesis

Theses that relate intuitive to formal models of computing can be motivated either by the desire to formalize a given intuitive notion or by the goal of providing intuition for a given formal concept. The Church-Turing thesis aimed to provide intuition for a robust formal concept expressible by Turing machines, the lambda calculus, or partial recursive functions. Church recognized that the answer to such questions could not be definitive, but the equivalent expressiveness of alternative formalisms for computability appeared to provide strong evidence for the Church-Turing thesis.

Church-Turing thesis: Formal effective computability by the lambda calculus (Church) or TMs (Turing) expresses the intuitive notion of effective computability of functions (over positive integers).

However, the robustness of Church-Turing models is due to their common inductive foundation rather than their ability to completely express all forms of computation. Coinductive models uniformly extend the expressiveness of machines, algebras, and set theory. Coinduction provides a more expressive mental tool

for definition, reasoning, and modeling that changes the status of Turing machines from maximally to weakly expressive models limited by induction. When the intuitive notion of effective computation is broadened to include interaction, the formal notion of computability must be correspondingly broadened.

The Church-Turing thesis answers the question "What is the intuitive notion of computing that is expressed by TMs?", but not the question "What is the formal model that expresses the intuitive notion of computing". In the early years of computing the intuitive notion of computing was identified with algorithms, and the two above questions were considered to have the same answer. As technology became increasingly interactive, and it was realized that algorithms could not express interaction, the intuitive notion of computing continued to be formalized by TMs because no formal model beyond that of TMs or well-founded set theory was available. Non-well-founded set theory and SIMs provide well-defined mathematical and machine models that go beyond algorithms.

Thesis for sequential interaction: The intuitive notion of effective computability for sequential interaction corresponds to specifiability by non-well-founded sets, SIMs, or PTMs.

The assumption that computable functions f: X -> Y from integers to integers, with inductively definable domains X of integers, model effective computability, excludes interaction streams (section 2) whose domains are not expressible by integers and whose domains and ranges are not separable. Expressing domains as integers requires complete specification of arguments before the computation starts and separability of domains and ranges: conditions that are assumed by Church and Turing and violated by streams. Stream computation cannot be represented in terms of functions from integers to integers. Whether f: stream -> stream is a noncomputable function or a computable nonfunction [WG] is a matter of definition.

The Church-Turing thesis should be viewed as a conjecture about the restricted notion of algorithmic computation rather than about the broader intuitive notion of computing. The thesis for sequential interaction is likewise a conjecture about a restricted notion of computing. Non-well-founded sets are not the most expressive class of consistent set theoretic models. We conjecture that more powerful models, possibly defined by replacing the AFA by a consistent axiom that permits a larger class of sets, will determine a more expressive class of computing mechanisms. We further conjecture that there is no well-defined maximal class of consistent sets, and that there is consequently no well-defined formal notion of maximal effective computability. This implies that effective computability is a relative rather than an absolute notion that can be defined only relative to assumptions about the form of interaction.

5.2 Ontological Commitment and Expressiveness

Questions of the existential (ontological) status of mathematical objects were central to both Kant's attempt to reconcile mathematical and empirical models and to the metamathematical models of Russell, Brouwer, Hilbert, and Godel in the early twentieth century [De]. Coinductive models can be viewed as a strong form of intuitionism based on intuitions about processes of observation and a strong kind of formalism based on formally consistent reasoning about incompletely observable systems [WG]. Because non-well-founded set theory accords existence to a larger class of sets than well-founded set theory, it can model larger classes of objects, situations, and computational problems.

The ontological strength of mathematical paradigms determines the expressiveness of their models

Logicians were reluctant to accept coinductive models based on circular reasoning because they failed to distinguish between inconsistent and consistent forms of circular reasoning, following Russell in overreacting to the paradoxes of set theory [BM]. An additional reason for excluding coinductive thinking from mainstream mathematics was the influence on logic of Brouwer's intuitionism (constructive mathematics), which accords existence to mathematical entities only if they can be effectively constructed.

Hilbert in principle accepted the stronger Platonic position that consistency implies existence. But in practice Hilbert restricted the consistency principle to inductive definition and reasoning, excluding consistent as well as inconsistent forms of circular reasoning. When Godel proved incompleteness of arithmetic by showing, through diagonalization, that true assertions of arithmetic cannot be inductively defined, the mainstream mathematical community accepted incompleteness as an absolute result rather than a relative result about a limited (inductive) form of logic.

Finsler's prescient work on set theory in the 1920s [Fi] showed the consistency of circular reasoning and anticipated Godel's incompleteness result, but was largely ignored because it did not conform to the mainstream paradigm of formalist mathematics. Finsler, influenced by Cantor's model of the real numbers, took to its logical conclusion the viewpoint that concepts exist independently of formalisms in which they are expressed. He went beyond Hilbert's formalism in applying the principle "consistency implies existence" to accept the existence of consistent sets of concepts independently of whether they are formalized.

We call this viewpoint a *realist ontology* of mathematics, since it accords existence to all consistent possible worlds, including the real world. Realism traditionally refers to models that accord objects an existence independently of whether they are perceived [EB]. By analogy, mathematical realism, as we define it, accords objects an existence independently of whether they are formalized. There is a connection between realism in mathematics and empiricism in physics and computing [We3], since both are concerned with modeling independently existing sets of objects.

Mathematical realism allows finite agents to model both nonenumerable real-number domains and physical domains like the real world. It relates mathematical and physical meanings of the term "real" because coinduction models both mathematical nonenumerability and physical interaction. Realist models provide a mathematical foundation for empirical computer science and a connection between empirical models of computing and physics [We3].

Intuitionism, inductive formalism, and realism can be classified by their degree of commitment to the existence of mathematical objects, embodying progressively stronger forms of ontological commitment that can model progressively larger classes of applications.

intuitionism (Brouwer): existence requires (inductive) constructibility (minimal ontology)
inductive formalism (Hilbert): consistency of (inductive) formal systems implies existence of models
realism (Cantor, Finsler): consistency (of a specification) implies existence (maximal ontology)

Hilbert accepted the principle "consistency implies existence" (C->E) [Fe], but in practice limited the application of this principle to formal systems with inductively specifiable behavior. We show in section 5.3 that Hilbert's limited application of the C->E principle is a primary cause of Godel incompleteness and conflicts with realism. But an extension of formalism to include coinduction is compatible with realism.

Coinduction determines an ontological paradigm shift from constructive and formalist to realist modeling principles that express observation of system behavior. Constructive inductively defined ontologies specified by least fixed points are weaker than empiricist coinductively defined ontologies specified by greatest fixed points. Maximal ontology admits interactive and time-dependent behavior for finite agents, while minimal ontology admits only noninteractive behavior. The strong connection between circular reasoning and interactive (realist) models of actual and possible worlds is explored in [BM].

The discovery sixty years after the foundational discussions of the 1920s that non-well-founded set theory models interactive computing provides an incentive for reevaluating Finsler's arguments, and validates the view that conceptually consistent possible worlds exist independently of their formalizability or constructibility. Tensor analysis is another example of an abstract mathematical technique that was later validated by applications in relativity theory, linear logic, and a variety of other areas many years after it was initially developed.

5.3 Incompleteness Theorems

According to Feferman [Fe], Godel credits his epistemological position of philosophical realism, arrived at early in his university years, with his mathematical successes. His Platonic belief in the independent reality of mathematical objects motivated his proof that arithmetic for independently existing integers could not be formalized. Godel's incompleteness result showed the impossibility of Hilbert's program of reducing mathematics to logic [Fe] and by implication the impossibility of reducing the behavior of finite computing agents to first-order logic [WG]. However, Turing's model of computation is "pre-Godelian" in its inductive (constructive) model of the notion of computation. The position that the intuitive notion of computing is completely expressed by TMs can be equated to Hilbert's view that mathematics can be com-

pletely expressed by inductive formalism.

Though it could be argued that the case for constructive theories is stronger in computing than in mathematics, we take the position that computer science is the science of modeling and that models of computation should include models of interactive as well as algorithmic behavior. We conjecture that a completeness theorem for interactive computing can be formulated in terms of coinductive computation, suggesting the following extended Church-Turing thesis:

Extended (coinductive) Church-Turing thesis: A behavior can be expressed by an interactive finite computing agent iff it can be expressed by a coinductive mathematical model.

Godel took advantage of the fact that computable functions are inductively specifiable in his use of Godel numbering to enumerate the TMs. Godel's incompleteness result is a simple corollary of the "folk theorem" that inductively specified enumerable classes of theorems cannot completely model nonenumerable classes of objects or situations [We3]. First-order logic can model only enumerable semantic domains because its number of theorems is enumerable, which is easy to prove. Proving that a particular semantic domain is nonenumerable is generally harder, but diagonalization is a useful tool in proving nonenumerability and was used by Cantor to prove nonenumerability of the reals and by Godel to show nonenumerability of the true assertions about integer arithmetic. Interaction machines can be shown to have a nonenumerable number of computations and their behavior is therefore Godel incomplete.

First-order logic can model only enumerable semantic domains
To prove Godel incompleteness of any domain, show its nonenumerability
Diagonalization is a tool for showing that specific domains are nonenumerable

Godel's incompleteness result implies not only that inductive mathematics is too weak to express formalist models of arithmetic over the integers, but also that interactive models are incomplete. The incompleteness of empiricist (intuitive) models is easier to prove than incompleteness of arithmetic because interactive models are more strongly nonenumerable than recursively enumerable sets.

Feferman indicates [Fe] that Godel concealed his Platonic (realist) beliefs to avoid controversy, using Platonic principles as a basis for research rather than as a topic for analysis. Feferman is surprised at Godel's opposition to C->E and does not adequately explain Godel's reasons for this. Our analysis suggests that Godel's antipathy to C->E may stem from the fact that C->E is incompatible with Godel's belief in realist mathematics and his acceptance of Hilbert's inductive formalism.

Claim: C->E, inductive formalism, and realist mathematics cannot be simultaneously true.

C->E, inductive models, and realist mathematics are three alternative starting points for a mathematical *Weltanschauung* that are incompatible. Hilbert, Godel, and Cantor (Finsler) epistemologies of mathematical modeling each accept two and reject one of these principles. Hilbert's belief in C->E and inductive formalism was shown by Godel to be incompatible with realist mathematics, Godel's belief in realist mathematics and inductive formalism caused him to reject C->E, while Cantor and Finsler's belief in C->E and realist mathematics require rejection of inductive formalism.

Hilbert: C->E + inductive formalism; inconsistency with realist mathematics was proved by Godel
Godel: realist (Platonic) mathematics + inductive formalism; explains Godel's rejection of C->E
Cantor, Finsler: C->E + realist mathematics; implies rejection of inductive formalism

Godel's acceptance of the principle that inductivly defined objects exist while coinductively defined objects do not led to his incompleteness result. Given his realst principles, Godel had to choose between the general principle of C->E and the more specialized inductive formalism of Hilbert, and chose inductive formalism. Had Godel instead accepted C->E and recognized along with Finsler that this implied accpetance of coinductive reasoning, the evolution of logic might have been very different.

6. Interactive Software Technology

The evolution of computer architecture from mainframes to personal computers and networks, of software engineering from procedure-oriented to object-oriented and component-based systems, and of AI from logic-based to agent-oriented and distributed systems has followed parallel paths [We1]. According

to the capsule history below, the 1950s through the 1970s were concerned with the development and refinement of algorithm technology for mainframes, sequential interaction became the dominant technology of the 1980s, while distributed interaction became the dominant technology in the 1990s:

1950s: machine language, assemblers, hardware-defined action sequences
1960s: procedure-oriented languages, compilers, programmer-defined action sequences
1970s: structured programming, composition of action sequences, algorithm composition
1980s: object-based languages, personal computers, sequential interaction architecture
1990s: structured object-based programming, networks, distributed interaction architecture

Whereas the shift from machine to procedure-oriented languages involves merely a change in the granularity of actions, the shift from procedures to objects is more fundamental, involving a qualitative extension from algorithms to interaction. The extension from sequential to distributed interaction requires a further fundamental paradigm shift in models of computation. SIMs express the shift from algorithms to sequential interaction architecture, while MIMs express the further shift to distributed interaction.

Figure 1 illustrates the extension from algorithms to interaction along a number of dimensions. Each algorithmic concept in the left-hand column is paralleled by a more expressive interactive concept in the right-hand column. Moreover, each right-hand concept has both a single-agent (sequential) and a multi-agent (distributed) form whose expressiveness is specified by SIMs and MIMs. The domain-independent generality of SIMs and MIMs is both an advantage in providing a uniform modeling framework and a drawback in that general models provide little guidance for domain-specific applications.

The transition from the view that computing is primarily concerned with input-output transformations to the view that computing systems provide services over time arises in many different contexts. Services over time are specified by models of interaction that cannot be reduced to or expressed by algorithms or TMs. Algorithms are time-independent (instantaneous) episodes in the life-cycle of an interactive systems. The one-dimensional quantitative performance metric of algorithmic complexity becomes the multidimensional qualitative performance metric of quality of service (QoS), which is an increasingly central focus for research in the database and human-computer interaction communities.

Algorithmic Concepts	Interactive Concepts
input-output transformation	services over time (QoS)
procedure-oriented programming	object-oriented programming
structured programming	structured object-oriented prog.
compositional behavior	emergent behavior
programming in the small	programming in the large
logic and search in AI	agent-oriented (distributed) AI
closed systems	open systems
algorithmic computer science	empirical computer science

Figure 1: Parallel Extensions from Algorithms to Interaction

Procedures and objects both determine a *contract* between providers and clients of a resource, but objects provide richer services to clients that cannot be expressed by algorithmically specified procedures. Algorithms are like sales contracts, guaranteeing an output for every input, while objects are like marriage contracts, describing ongoing contracts for services over time. An object's contract with its clients specifies its behavior for all contingencies of interaction (in sickness and in health) over the lifetime of the object (till death us do part) [We1]. The folk wisdom that marriage contracts cannot be reduced to sales contracts is computationally expressed by interaction not being reducible to algorithms.

Though object-based programming has become a dominant technology, its foundations are still shaky. Everyone talks about it but no one knows what it is. "Knowing what it is" has proved elusive because of the implicit belief that "what it is" must be defined in terms of algorithms. Interactive models have the liberating effect of providing a broader framework than algorithms for defining "what it is". Component-based

software technology is even less mature than object-based technology: it is the technology underlying interoperability, coordination models, pattern theory, and the World Wide Web. Knowing what it is in turn requires liberation from sequential object-based models.

Structured programming [Di] proved too weak as a model for program structure because the transition to object-oriented programming made procedural structured programming based on composing algorithms and while statements obsolete. Objects have behavior that cannot be compositionally expressed in terms of the behavior of their components. Structured programming for actions (verbs) can be formally defined by function composition, while structured programming for objects (nouns) is modeled by design patterns that have no compositional formal specifications [GHJV]. As a consequence, the study of design pattern methods of component composition is an art rather than a science.

Compositionality is a desirable property for formal tractability of programs that has led to advocacy of functional and logic programming as a basis for computation. But it limits expressiveness by requiring the whole to be expressible as the sum of its parts. Actual object-oriented programs and computer networks exhibit noncompositional emergent behavior. There are inherent trade-offs between formalizability and expressiveness that are clearly brought out by the expressive limitations of compositionality. Arguments in the 1960s that go-tos are considered harmful for formalizability can be paralleled by arguments in the 1990s that compositionality is considered harmful for expressiveness.

Programming in the large (PIL) is not determined by size, since a program consisting of a sequence of a million addition instructions is not PIL. PIL is synonymous with interactive programming, differing qualitatively from programming in the small in the same way that interactive programs differ from algorithms. Embedded and reactive systems that provide services over time are PIL, while noninteractive problem solving is not PIL even when the algorithm is complex and the program is large.

The evolution of artificial intelligence from logic and search to agent-oriented programming is remarkably similar to the evolution of software engineering. This paradigm shift is evident in research on agents [Ag], on interactive planning and control [DW], and in textbooks that systematically reformulate AI in terms of intelligent agents [RN]. AI illustrates more clearly than software engineering that reasoning is an inadequate basis for modeling [We1]. Though logic is a form of computation, computation cannot be entirely reduced to logic. The goals of the Fifth-Generation Computing Project of the 1980s, which aimed to provide a logic-based framework for universal computation, were in principle unrealizable.

Open systems can be precisely defined as interactive systems: interactive models provide a tool for classifying forms of openness and for analyzing open-system behavior. Empirical computer science can likewise be precisely defined as the study of interactive systems [We3].

7. References

[Ac] Peter Aczel, Non Well-Founded Sets, *CSLI Lecture Notes #14*, Stanford, 1988.

[Ag] Phil Agre, Computational Research on Interaction and Agency, *Artificial Intelligence*, January 1995.

[BM] Jon Barwise and Lawrence Moss, *Vicious Circles*, CSLI Lecture Notes #60, Cambridge University Press, 1996.

[CW] Luca Cardelli and Peter Wegner, On Understanding Types, Data Abstraction, and Polymorphism, *Computing Surveys*, December 1985.

[De] Michael Detlefsen, Philosophy of Mathematics in the Twentieth Century, in *Philosophy of Science, Logic, and Mathematics in the Twentieth Century*, Ed. S. Shanker, Routledge, 1996.

[Di] Edsger Dijkstra, *The Discipline of Programming*, Prentice-Hall, 1976.

[DW] Thomas Dean and Michael Wellman, *Planning and Control*, Morgan Kaufman 1991.

[EB] Encyclopaedia Britannica, Article on Realism

[Fe} Solomon Feferman, Kurt Godel, Conviction and Caution, in *Godel's Theorem in Focus*, Ed Shanker, Methuen, 1988.

[Fi] Paul Finsler, *Finsler Set Theory: Platonism and Circularity*, Eds. David Booth and Renatus Ziegler, Birkhauser, 1996.

[GHJV] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides, *Design Patterns: Elements of*

Reusable Object-Oriented Software, Addison-Wesley, 1994.

[GW] Dina Goldin and Peter Wegner, *Persistent Turing Machines*, Brown Technical Report, 1998.

[JR] Bart Jacobs and Jan Rutten, A Tutorial on Coalgebras and Coinduction, EATCS Bulletin 62, 1997.

[Kr] Saul Kripke, A Completeness Theorem in Modal Logic, *Journal of Symbolic Logic*, 1959.

[Mi] Robin Milner, Operational and Algebraic Semantics of Concurrent Processes, *Handbook of Theoretical Computer Science*, J. van Leeuwen, editor, Elsevier, 1990.

[Po] Karl Popper, *The Logic of Scientific Discovery*, Harper, 1965.

[Pr] Vaughan Pratt, Chu Spaces and Their Interpretation as Concurrent Objects, in *Computer Science Today: Recent Trends and Developments*, Ed. Jan van Leeuwen, LNCS #1000, 1995.

[RN] Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach,* Addison-Wesley, 1994.

[Sc] Dana Scott, The Lattice of Flow Diagrams, Lecture Notes in Mathemaitcs, Springer 1971.

[Tu1] Alan Turing, On Computable Numbers with an Application to the Entscheidungsproblem, *Proc. London Math Soc.*, 2:42, pp. 230-265, 1936.

[Tu2] Alan Turing, Computing Machinery and Intelligence, *Mind*, 1950.

[We1] Peter Wegner, Why Interaction is More Powerful Than Algorithms, *CACM*, May 1997.

[We2] Peter Wegner, Interactive Foundations of Computing, *Theoretical Computer Science*, Feb. 1998.

[We3] Peter Wegner, Towards Empirical Computer Science, *The Monist*, Issue on the Philosophy of Computation, January 1999, available at www.cs.brown.edu/people/pw.

[We4] Peter Wegner, Interactive Software Technology, *Handbook of Computer Science and Engineering*, Ed A. Tucker, CRC Press, 1996.

[We5] Peter Wegner, Tradeoffs Between Reasoning and Modeling, in *Research Directions in Concurrent Object-Oriented Programming*, Eds. Agha, Wegner, Yonezawa, MIT Press 1993.

[WG] Peter Wegner and Dina Goldin, Mathematical Models of Interactive Computing, Brown Technical Report, www.cs.brown.edu/people/pw