

Infinitely Complex Machines

Prof. Eric Steinhart, Department of Philosophy, William Paterson University.

Email: esteinhart1@nyc.rr.com

Website: www.ericsteinhart.com

Published as: Steinhart, E. (2007) Infinitely complex machines. In A. Schuster (Ed.) *Intelligent Computing Everywhere*. New York: Springer, 25-43.

ABSTRACT: Infinite machines (IMs) can do supertasks. A supertask is an infinite series of operations done in some finite time. Whether or not our universe contains any IMs, they are worthy of study as upper bounds on finite machines. We introduce IMs and describe some of their physical and psychological aspects. An accelerating Turing machine (an ATM) is a Turing machine that performs every next operation twice as fast. It can carry out infinitely many operations in finite time. Many ATMs can be connected together to form networks of infinitely powerful agents. A network of ATMs can also be thought of as the control system for an infinitely complex robot. We describe a robot with a dense network of ATMs for its retinas, its brain, and its motor controllers. Such a robot can perform psychological supertasks – it can perceive infinitely detailed objects in all their detail; it can formulate infinite plans; it can make infinitely precise movements. An endless hierarchy of IMs might realize a deep notion of intelligent computing everywhere.

KEYWORDS: infinite computer, infinite mind, supertask, complexity.

1. Introduction

We discuss a variety of infinitely powerful machines and their infinitely complex operations. We won't argue for the existence of such machines.¹ Our only purpose here is to start to map out the logical space of all possible machines and minds (see Doyle, 1991). If we have a map of that logical space, we can locate ourselves there. We can then ask questions about our cognitive powers. For instance, if our brains are only finitely powerful machines, we can ask whether that finiteness is necessary or merely contingent (e.g. on the fact that our brains are made of certain stuff). And if we have a map of that logical space, we can study the upper bounds on the powers of our artifacts. Perhaps we can make artificial intellects far more powerful than our brains. They could use entirely different physics (e.g. quantum mechanical computers). If these artifacts are only finitely powerful, we can wonder why. If they are infinitely powerful, then having a map of logical space of all possible intellects helps us understand what they can and cannot do. So a map of the logical space of possible intellects is useful in several ways.

2. Infinite Progressions

2.1 The Progression of Ordinals

We all know that the series of whole numbers (natural numbers) is endless. We can write 0, 1, 2, 3, and so on. We all know that for any natural number n , there exists a next natural number $n+1$. Natural numbers are more precisely known as *ordinals*. So we'll talk about ordinals rather than natural numbers. Although it's well-known that the series of finite ordinals is endless, it's less well-known that the endless finite series is bounded above by the first *limit ordinal*. The first limit ordinal is the first transfinite number. The ordinal number line thus extends beyond the finite into the transfinite. A useful description of the ordinal number line is given by Cantor's three number generating rules. For a more complete description, you'll need to look at the full development of the ordinals in set theory (see Drake, 1974 or Hamilton, 1982). The rules look like this:

1. **Initial Rule.** The initial ordinal is 0. The initial ordinal 0 is the set of all ordinals less than 0. There are no whole numbers less than 0. So 0 is the empty set $\{\}$.
2. **Successor Rule.** For every ordinal n , there exists an ordinal $n+1$ that is greater than n . The ordinal $n+1$ is the successor of n . It is a successor ordinal. Each ordinal $n+1$ is the set of all ordinals less than itself. It is the set $\{0, \dots, n\}$. For instance, $1 = \{0\}$; $2 = \{0, 1\}$; $3 = \{0, 1, 2\}$; and so it goes.
3. **Limit Rule.** For any endlessly increasing series of ordinals, there exists a limit ordinal greater than every ordinal in the series. Since the series of finite ordinals 0, 1, 2, 3, \dots is endlessly increasing, there exists a limit ordinal ω greater than every finite ordinal. The ordinal ω is the set of all ordinals less than itself. It is the set of all finite ordinals. It is $\{0, 1, 2, \dots\}$. Since ω is an ordinal, it has a successor $\omega+1$. And so it goes.

2.2 The Progression of Zeno Points

We'll say that a *progression* is a series of objects defined by three rules: An *initial rule* that associates the initial ordinal 0 with some object. A *successor rule* that associates each successor ordinal $n+1$ with some object. And a *limit rule* that associates the limit ordinal ω with some object. For example, consider the following progression of fractions:

1. **Initial Rule.** The progression starts with the initial object 0. $Z_0 = 0$.
2. **Successor Rule.** For every object Z_n in the progression, there is a successor object Z_{n+1} in the progression. $Z_{n+1} = Z_n + 1/2^{(n+1)}$. Hence we get the series $Z_1 = 1/2$; $Z_2 = 3/4$; $Z_3 = 7/8$; $Z_4 = 15/16$; and so on.
3. **Limit Rule.** The progression ends with a limit object Z_ω . The limit object is the limit of the Z_n as n goes to ω . So $Z_\omega = 1$.

The progression $Z_0, Z_1, \dots, Z_n, Z_{n+1}, \dots, Z_\omega$ is the *Zeno progression*. If we think of the fractions in the Zeno progression as points on a spatial line, they are *Zeno points*. If we think of them as fractions of a time interval, then they are *Zeno moments*.

Zeno progressions are useful for defining supertasks. A *supertask* is an infinite series of operations done in some finite region of space-time. Although some supertasks are ill-defined (and thus seem to be paradoxical), many supertasks have consistent recursive definitions and converge to well-defined objects at transfinite limits (Earman & Norton, 1996; Koetsier & Allis, 1996). We can think of a supertask done in the temporal interval $[0, 1]$ as a progression of finitely complex (*finitary*) tasks done at Zeno moments during that interval. The first task is done by time $1/2$; the second by time $3/4$; and so on. At time 1, as many tasks have been done as there are natural numbers.

3. Some Supertasks

3.1 Drawing the Royce Map

A system or agent that performs a supertask is said to *accelerate* (see Weyl, 1963: 42; Grunbaum, 1969; Boolos & Jeffrey, 1989: 14 - 16). We'll follow convention and say that *Zeus* is an agent able to perform supertasks. We'll informally describe some supertasks for Zeus. Our first supertask is drawing the *Royce Map*. Royce describes a perfectly accurate map of England drawn on the surface of England (1899: 506 - 507). Since the map has to exactly describe the part of England on which it is drawn, it contains an exact copy of itself (which contains an exact copy of itself . . .). The Royce Map contains an endlessly nested series of exact copies of itself. It is endlessly recursive.

For simplicity, say England is just a square crossed by a north-south road and an east-west road. Zeus is going to draw the perfect map of England. He writes with a special pencil that can always write twice as thin and that never runs out of graphite. He can always write twice as fast and twice as precisely. He writes on a sheet of paper that is continuously divided. For any two real numbers x and y varying between 0 and 1, there is a point (x, y) on the paper that can be either marked by the pencil or left unmarked. Drawing the Royce Map is a supertask. It consists of ω many tasks. Each task is finitely complex and all tasks have the same complexity. Figure 1 shows the first four tasks in the construction of the Royce Map. Zeus draws the complete Royce Map by following these three rules:

1. Initial Rule. The initial map M_0 = a square with a cross drawn in it.
2. Successor Rule. The successor map M_{n+1} is the previous map M_n + a cross drawn in the lower right square of M_n . Zeus can draw the map M_{n+1} in half the time he takes to draw the map M_n . So Zeus has drawn M_{n+1} by the Zeno moment Z_{n+1} .
3. Limit Rule. The limit map M_ω is the superimposition of all the finite maps. Zeus has drawn the Royce Map M_ω at the limit Zeno moment $Z_\omega = 1$. Following Dedekind, we

say a structure is infinitely complex (*infinitary*) iff it has a proper part that can be put into a 1-1 correspondence with the whole (that is isomorphic to the whole). Any lower right square of the Royce Map is isomorphic with the whole.




Figure 1. The first four iterations of a Roycean self-nested map.

3.2 Drawing the Hilbert Paper

Our second supertask is the construction of the *Hilbert Paper*. The Hilbert Paper is a finitely sized square piece of paper on which every natural number is written in base 1 notation (as a stroke series). Writing all the natural numbers down on the Hilbert Paper is a supertask. It consists of ω -many tasks. Although each task is only finitely complex, each successive task is more complex than its predecessor (the complexity goes up linearly). The n -th finite Hilbert Paper has the numbers 1 to n written on it. The Hilbert Paper itself is the limit of the progression of finitary Hilbert Papers. The construction of all the Hilbert Papers is given by the three rules below and is illustrated in Figure 2. The rules are:

1. Initial Rule. The initial Hilbert Paper H_0 = a piece of paper divided in half vertically and horizontally with a single stroke | in the upper left quarter.
2. Successor Rule. The successor Hilbert Paper $H_{n+1} = H_n + \text{Zeus}$ divides the right column in half vertically and divides the bottom row in half horizontally; Zeus copies the bottom row of strokes into the next lower row and adds one stroke on the right. Zeus can draw the successor paper H_{n+1} in half the time he takes to draw the paper H_n . So Zeus has drawn H_{n+1} by the Zeno moment Z_{n+1} .
3. Limit Rule. The limit Hilbert Paper H_ω = the super-imposition of all the H_n with n finite. Zeus has drawn the limit paper H_ω at the limit Zeno moment 1. The Hilbert Paper is the limit paper. The Hilbert Paper is infinitely complex: any square in the lower right hand corner has exactly the same structure as the whole Hilbert Paper.

The rows and columns of the Hilbert Paper form a *Zeno Matrix*. A *Zeno Matrix* is a piece of (continuous) paper that is divided up into ω many columns and rows. The first row takes up 1/2 of the paper; the next row takes up the next 1/4 of the paper; and so on. The first column takes up 1/2 of the paper; the next row takes up the next 1/4 of the paper; and so on. For any n and m , there is a cell on the paper at row n and column m .

The Hilbert Paper is a Zeno matrix with all cells in the lower left triangle filled in with 1s.




Figure 2. A few iterations towards the Hilbert Paper.

3.3 Computing on Hilbert Papers

Zeus loves to compute. His writing tablet is made of special sheets of paper. Each sheet is divided into infinitely many rows. The top row is $1/2$ the sheet; the next row takes up the next $1/4$ of the sheet; and so it goes. Each sheet is divided into two columns. The left column is a copy of the Hilbert Paper. So the natural numbers are inscribed on the left column, starting with I, then II, then III, and so on. Each number in any row in the left column is associated with an empty cell in the same row in the right column.

Zeus wants to determine the locations of the primes in the natural numbers. For any n , it is a finitely complex task to determine whether n is prime or not. Of course, the complexity of these tasks increases without bound as n increases without bound. But that doesn't bother Zeus. Since Zeus is a super-agent, he can do any finitely complex task in any finite interval of time. He can do any task of complexity less than ω in any time interval greater than $1/\omega$. For any n , and for any finite time, Zeus can determine if n is prime.

Zeus is ready to compute. At time 0, every right column in every row is blank. At time $t = 0$, Zeus puts his pen on the square in the initial row and right column of his tablet. Within $1/2$ second, he does two tasks. First, he looks at the number written in the left column of that row (the number 1) and records the fact that it is not prime by writing a 0 in the right column of that row. Second, he moves his pen down one row. His pen is now over the square in the right column of the second row. Within $1/4$ second, he does two tasks. First, he looks at the number in the left column of that row (the number 2) and records the fact that it is prime by writing a 1 in the right column of that row. Second, he moves his pen down one row. His pen is now over the square in the right column of the third row. By the n -th Zeno moment in the unit time interval, Zeus has either written 0 or 1 in the right column of the n -th row. At the ω -th Zeno moment 1, Zeus has written 0 or 1 in the right column of every row. Hence at time 1, for every row n , the right column of n is 0 if n is not prime and is 1 if n is prime. Zeus has thus located every prime number. For any even number p , it is a finitary task to determine whether or not p is the sum of two primes. Goldbach's Conjecture says that every even number is the sum of two primes. By performing another supertask in the next unit of time, Zeus can test Goldbach's Conjecture.

4. Accelerating Machines

4.1 Accelerating Turing Machines

A conventional Turing machine (a CTM) is a digital computer with an unbounded memory. Many good descriptions of CTMs are available (e.g. Weizenbaum, 1976; Hopcroft, 1984). We won't describe CTMs here. We show how to extend a CTM so that it can accelerate. We thus describe an *accelerating Turing machine* (an ATM). An ATM can do supertasks (Copeland, 1998a; Davies, 2000; Hamkins & Lewis, 2000).

An ATM has two main parts: a head and a tape. The tape of an ATM is only finitely long. It has some unit length. It is a line segment running from 0 to 1. The tape is divided into ω many cells. Each cell can hold either a 0 or a 1. The initial cell 0 occupies the interval 0 to $1/2$ without including the endpoint at $1/2$. More precisely, cell 0 occupies the interval $[0, 1/2)$. The next cell 1 occupies the interval $[1/2, 3/4)$. Each cell $n+1$ occupies an interval half as long as its predecessor. The boundaries of the cells are thus the Zeno points in the interval $[0, 1]$. Such a tape is therefore a *Zeno Tape*.

The head of an ATM moves across the tape. The head can read the content of any cell and can write a character on a cell. The head has a finite set S of possible states. It is always actually in some state from S . The head operates according to a fixed list of transition rules. Each transition rule is an instruction to the head of this form: if you're in state A and you read character B , then write character C , make a motion D , change to state E . The states in a rule are from S . The characters are from the character set $\{0, 1\}$. The set of motions is $\{\text{left, stay put, right}\}$. The head has a rule for each (state, character) pair. The size of the head varies as it moves. As the head moves right, it shrinks to half

its size. As it moves left, it expands to twice its previous size. If the head moves to either boundary of the tape (to 0 or 1), it shrinks to an infinitesimal size. It shrinks to the size of a point.

An ATM runs on an accelerating clock. An ATM performs its first operation in $1/2$ second. It performs its next operation in $1/4$ second. It performs each next operation in half the time it took to perform its previous operation. Its operations thus fill the Zeno moments in the interval $[0, 1]$. At time 1, the ATM has performed ω many operations. We can precisely describe the behavior of an ATM by three rules:

1. Initial Rule. The initial tape T_0 is inscribed with the input to the ATM. The clock is at time 0. The head is over cell 0. It is in state 0.
2. Successor Rule. For every tape T_n , there is a successor tape T_{n+1} . The successor tape T_{n+1} is defined by applying some transition rule. The head reads the value of the cell beneath it and looks at its state. It finds a rule whose antecedent matches its current (state, character) pair. It then applies that rule to write a character into the cell, make a motion, and change into a new state. The head accelerates. It makes each next transition twice as fast. Each successor tape is computed twice as fast. An ATM has computed tape T_{n+1} by the Zeno moment Z_{n+1} .
3. Limit Rule. Either the progression of tapes T_n for n finite converges to a limit tape T_ω or else it does not.² If it converges, then the limit configuration of the ATM is the limit tape with the head at its limit position. Thus if the progression of tapes converges, the head is either located at some finitely indexed cell or at the endpoint 1. If the progression of tapes does not converge, then the limit configuration of the ATM is the completely unmarked tape with the head positioned exactly over the point 0. Thus if the progression of tapes does not converge, the head is not on any cell at all.

4.2 Examples and Powers of Accelerating Turing Machines

We describe an ATM that converges and an ATM that does not. The convergent ATM starts with a blank tape (all cells unmarked). It starts in state 0. It always acts according to this rule: it writes a 1, it moves right, it stays in state 0. At the n -th clock tick, it has filled the first n cells with 1s. At the limit clock tick (at time 1), the head is on cell ω and the tape is entirely filled with 1s. The ATM has thus converged. The ATM that does not converge (that diverges) also starts with a blank tape and in state 0. It always acts according to two rules. The first rule says: if you are over a cell with value 0 and you are in state 0, then change the value of the cell to 1, stay put, and go into state 1. The second rule says: if you are over a cell with value 1 and you are in state 1, then change the value of the cell to 0, stay put, and go into state 0. This ATM *oscillates*. It always stays over cell 0. It marks the cell; it unmarks the cell; it marks the cell; it unmarks the cell; and so it goes. This ATM is the equivalent of the *Thompson lamp* (Thomson, 1954). This ATM does not converge to any state at the limit time 1. Hence at time 1, its head is on point 0 and its tape is blank.

An ATM can solve problems that cannot be solved by any CTM. Copeland (1998b) has shown that a universal ATM (a universal CTM that accelerates) can solve the halting problem for CTMs. It can thus fill in the *halting table* for CTMs. The halting table is recorded on a piece of paper divided into a Zeno Matrix. For any n and m , there is a cell on the paper at row n and column m . There is an ATM that can fill in each cell at row n and column n with 0 if the n -th CTM does not halt on input m and with 1 if the n -th CTM does halt on input m . A universal ATM can use the halting table to compute all the Rado numbers (Boolos & Jeffrey, 1989: ch. 4).

5. Intellects and Games

5.1 Infinite State Machines

Although Turing machines (both conventional and accelerating) are interesting in many ways, they are also rather dull. They don't interact. If an ATM realizes a mind, it is a solipsistic mind. But it is plausible that intelligence requires interaction with another agent either directly or through an environment (Russell & Norvig, 1995; Maes, 1995). For if perception and action are genuine, then another agent is necessary. We therefore move beyond Turing machines to consider infinitary machines that interact.

An *infinite state machine* (ISM) has a set I of possible inputs, an set S of possible states, and a set O of possible outputs. Its input and output sets may be finite or infinite. Its state set must be infinite. We think of the input, output, and state sets as sets of ordinals. An ISM has a function F that maps its current (input, state) pair onto its next state and a function G that maps its current (input, state) pair onto its next output. So the tuple (I, S, O, F, G) specifies a species of ISM. Of course, ISMs can accelerate.

We can think of the ordinals in the input set I as corresponding to the possible configurations of an input device. If we think of these input configurations in cognitive terms, they are the possible *perceptions* of the ISM. The ordinals in the output set O correspond to the configurations of an output device. These are the possible *actions* of the ISM. The states in S are the possible internal mental states (*ideas*) of the ISM. An ISM has an operation cycle. It consumes an input; it changes its state by applying F ; it produces an output by applying G . We can interpret the operation cycle psychologically as involving a perception, a calculation, and an action. We have discussed the psychological aspects of infinite minds elsewhere (Steinhart, 2003), and so we do not discuss them here.

Since ISMs have inputs and outputs, they can be coupled together. One way to couple two (or more) ISMs together is to link them to a common structure. This common structure is their shared environment. And a nice way to model the interactions of two (or more) agents is to have them play a game. Many ethical, social, and political concepts can be analyzed in game-theoretic terms.³ Of course, ISMs will only be interested in playing infinitely complex games. There is a large literature on infinitary games.⁴

5.2 Infinitary Board Games

An m,n,k -game is a game in which two players (black and white) take turns placing marks of their own colors on the points of an m by n grid. You win iff you get k marks of your color in a line. For example, tic-tac-toe is a $3,3,3$ -game. Freestyle gomoku is a $19,19,5$ -game. Other m,n,k -games include Pente and Connect6. Finitary m,n,k -games are of no interest to superminds. We can extend m,n,k -games to the infinite by allowing the grid to be infinite (with grid points indexed by pairs of integers). We can easily compress a grid whose points are indexed by pairs of integers into a finite area. The compression is by Zeno compression on each direction of the x and y axes. Each next step away from the origin is twice as small as the previous step. The resulting *Zeno Grid* has its origin at $(0, 0)$ and its limit edges at points with x or y coordinates $+1$ or -1 .

We mechanize an infinitary m,n,k -game by extending the mechanization of a finitary m,n,k -game (e.g. tic tac toe). For any infinitary m,n,k -game, each player has at least an eye, a brain, and a hand. The eye surveys the board. The eye is an infinitary array of sensors with the same structure as the Zeno Grid. It is an infinitary retina positioned over and looking at the game board. Each configuration of the board corresponds to a number in I . The eye sends this number to the brain. The brain computes the next move according to the player's strategy. This strategy is encoded in the functions F and G . The brain uses F to change its mental state. It uses G to generate its output. Its output is sent to its hand. The hand is a device that can mark a point on the Zeno Grid. The output instruction tells the hand which point to mark. For instance, an output number can encode a pair of numbers (i, j) . If the hand gets output instruction (i, j) , then it moves to the i -th row and j -th column and puts its mark on the grid point with those coordinates. Since the hand must be able to move from any point to any other point in finite time, it has to accelerate.

Suppose that Zeus and Hera are ISMs playing some infinitary m,n,k -game. Each player has to move in one clock tick and the clock accelerates. Hera goes first. She moves in the first $1/2$ second. Zeus moves in the next $1/4$ second. And so it goes. Of course, this simplistic scheme doesn't enable us to select a winner. We need some way to determine when a player wins. We could let each player decide when he or she wins. But that might introduce conflict. So we add a *referee* who checks the board after each move. We can call this referee *Apollo*. Apollo is also an ISM. According to this scheme, each move is divided into two parts: a player operates; the referee checks for a winner. So we can think of Zeus and Hera playing various infinitary m,n,k -games, carefully watched by Apollo. Other board games (like chess) can also be extended to the infinite (see Pickover, 1995).

5.3 Infinitary Number Selection Games

An infinitary number selection game involves two players (Zeus and Hera) who construct an infinitely long sequence of numbers. They take turns selecting numbers from a fixed set and by adding their choices to the sequence. Suppose the set of numbers is just the set of decimal digits (see Hamilton, 1982: 189). It is the set $\{0, \dots 9\}$. The game is played on a Zeno Tape. Hera and Zeus take turns reading from and writing on the tape. The first few plays in an infinitary number selection game might go like this: Hera writes 3 in tape cell 0 in the $1/2$ second; Zeus writes 4 in cell 1 in the next $1/4$ second; Hera writes 9 in cell 2 in the next $1/8$ second; and so on. This game makes a sequence that starts with 349.

Each player is an ISM. Each player consists of a controller, an input device, and an output device. The input and output devices always rest over some position on the tape. We can think of these devices very simply as combined into a single read/write head that rests over some position of the tape just like the read/write head of an ATM. The controller is an infinitely complex CPU. It can have an infinitary memory and can perform any of its basic operations in any finite non-zero time. As an ISM, the logic of the controller is defined by the tuple (I, S, O, F, G) . Its input set I is $\{\text{blank}, 0, \dots 9\}$. The blank is used to start the game. Its state set S is the set of all finitely long digit sequences. Its output set O is $\{0, \dots 9\}$. The functions F and G are stored in a Zeno Matrix in the controller. The matrix that stores F and G is the long-term memory of the player. The short-term memory of each player includes the variables used to track the player's current configuration. The program that regulates the behavior of the players is given in this footnote.⁵

As the clock accelerates to the limit moment 1, the players complete an infinitely long sequence $\langle d_0, d_1, d_2, d_3 \dots \rangle$ on the tape. Since each d_i is from the set $\{0, \dots 9\}$, we may think of the tape as the decimal expansion of the real number $0.d_0d_1d_2d_3 \dots$ that lies between 0 and 1. To make this into a game, we need to add some definition of winning. A player wins by hitting a real number. Before the game begins, the referee (*Uranos*) partitions the set of real numbers in $[0,1]$ into two sets H and Z . Every real in $[0, 1]$ is either in H or in Z but not in both. At the end of the game, *Uranos* checks whether the tape defines a real in H or Z . If it is in H , Hera wins. If it is in Z , Zeus wins.

At the end of an infinite series of moves (at the limit time 1), the referee *Uranos* checks whether the sequence on the tape is in H or Z . Of course, he only needs to check whether it is in H . If it is not, then it is in Z by default. The set H may be finite, countable, or uncountable. Let C be the cardinality of the real number continuum. The set H is stored in an array with C slots. The addresses of this array thus correspond to a well-ordering of an uncountably infinite set.⁶ This array is compressed into a tape of finite unit length. Each cell on this tape holds a series of ω many digits. It is a supertask to determine whether the sequence defined by Zeus and Hera is located in slot n of the array that stores H . For full iteration over H , *Uranos* must be able to perform uncountably many supertasks in finite time. *Uranos* is a machine that is far more powerful than either Hera or Zeus. *Uranos* works with spatio-temporal continua that are far richer than the continua

used by Hera and Zeus. It is not entirely clear how to define these continua. Since they are richer than the real number line, they are non-standard. Perhaps Uranos is working with continua that approximate Peirce's inexhaustible continua (Peirce, 1965: 3.563 – 570) or that approximate Conway's surreal number line (Conway, 2001).

A somewhat more complex version of this game allows the players to select any natural number at each turn (Neeman, 2004). We can think of these natural numbers as encoding various kinds of information. A natural number (expressed in base 2) is a bit string. It can be thought of as a file that describes an image (e.g. a JPG file). We can think of the read head of each player as eye that can recognize any of ω many images. We can think of the write head of each player as a hand that can paint any of ω many images. Thus when a player writes a number on the tape, he or she is painting a picture. And when a player reads a number from the tape, he or she is perceiving a picture. We might thus think of the whole sequence written on the tape as a movie. The sets H and Z define two kinds or styles of movies. Hera prefers those in H while Zeus prefers those in Z.

5.4 Infinitary Athletic Contests

Although so far we've only talked about games for super-nerds (infinitary board games and number games), we can define games for super-athletes. Consider infinitary tennis. We define it as a generalization of ordinary finitary tennis. Finitary tennis is played in a finite volume of space-time (a tennis court). A court has two sides. We can think of each side as a unit cube. Finitary tennis is a two-player game. We can refer to them as Hera and Zeus. Each player has a finitely sized racket. Each player occupies one side of the tennis court and can move freely in his or her unit cube. Tennis is played with a finitely sized ball. The players use their rackets to hit a finitely sized ball back and forth. We assume some familiarity with finitary tennis. We won't go into it in any more detail.

We start our definition of infinitary tennis by defining it at level 1. It is much like finitary tennis. In finitary tennis, the size of the ball is fixed. Since the size of the ball is fixed, we can divide each side of the tennis court into cubical cells. Each cell is exactly large enough to contain the ball. This division also divides the floor of the court into finitely sized square cells. For tennis at level 1, the ball always passes through (and only through) cubical cells. It thus lands on any bounce in some square cell on the floor of the court. In tennis at level 1, each player has a racket of normal finitary size. This racket is a disk whose diameter is larger than that of the ball. Each player has to move through his or her unit cube to hit the ball. Each player has more or less skill in getting to the ball and in hitting it back in such a way that it will land in the opposing player's court. Tennis at level 2 is much like tennis at level 1. The difference is just this: each cube in the court is divided in half on each axis to make 8 smaller sub-cubes. The ball shrinks to fit into these smaller cubes. The rackets shrink proportionally. The players remain the same size. They are trying to hit a smaller ball with a smaller racket. Tennis at level 3 is defined analogously. Each cube is again cut in half on each axis to make 8 even smaller cubes. The ball and rackets shrink again. We can iterate through tennis at level 4, level

5, and so on to level ω . At level ω , the players are playing with a ball that is the size of a point and rackets also the size of a point.

An easy version of infinitary tennis is played by moving from level n to level $n+1$ each time the ball passes over the net. A harder version allows the ball to either shrink or expand monotonically from any size to any other size as it moves across the court. Each volley in infinitary tennis (whether easy or hard) takes place in a unit time. Hera serves the ball in $1/2$ second; if Zeus returns it, the return takes only $1/4$ second; if Hera returns it, the return takes only $1/8$ second. And so it goes until someone fails to return the ball or returns it in an illegal way (so it lands out of bounds). If the volley goes to the limit, then the volley is counted for neither player. Playing infinitary tennis requires a body with infinitely precise perception, cognition, and volition. An infinitary tennis player has to have eyes with infinitary resolution. He or she has to have a brain that can compute the trajectories of arbitrarily small and arbitrarily fast objects. Its brain has to be able to accelerate. It has to have a motor system that can move in infinitely precise ways. Its muscles have to be able to accelerate. It has to have infinitely fine hand-eye coordination. An infinitary tennis player has to have an infinitary body. We can easily use analogous techniques to define infinitely precise versions of other games: infinitary handball, soccer, pool, and so on.

6. Infinitely Complex Bodies

A mind is an agent that interacts with an environment. The interaction is a loop involving perception, cognition, and action. A mind is more than merely a calculator or a brain. It has parts that enable it to perceive and to act. A mind is thus a whole body. An infinitary mind has to have an infinitary body. It has to have a body able to perform perceptual, cognitive, and volitional supertasks. It has to have a body able to at least perform operations with arbitrarily high finite precision (e.g. if it can hit a ball of some size moving at some speed, then it can hit a ball that is twice as small moving twice as fast). We describe an infinitary humanoid body here. Non-humanoid infinitary bodies are also possible.

As with any body, the infinitary body starts with a single cell (its zygote). The zygote encodes a growth plan that builds the infinitary body by means of an endless series of iterations. An infinitary body can only grow in a universe in which space, time, and matter are continuously divisible. The initial iteration of the growth plan takes $1/2$ unit of time. Each successor iteration happens twice as fast as its predecessor. So the body grows through infinitely many iterations in one time unit. The growth plan of the infinitary body is based on the generation of body-forms (phenotypes) by L-systems (Prusinkiewicz & Lindenmayer, 1990). It would be fun to present the infinite growth plan with greater biological realism, but to save space our presentation is merely schematic.⁷

The growth plan defines an initial iteration. On this first iteration, the body grows by ordinary cellular division into a central trunk with a head. Each offspring cell has the

same size as its parent. Hence the initial iteration fills out a finite volume. For simplicity, let the trunk just be a cylinder and the head a sphere. The trunk is the initial limb of the body. It has two *growth sites* at the top (for its arms) and two growth sites at the bottom (for its legs). After the initial iteration, the growth plan drives the body through endlessly many successor iterations. We discuss successor iterations for the limbs, eyes, and brain.⁸ The successor iterations converge to an infinite form in the limit. The limit form is infinitely complex. Infinite organs support infinite forms of perception, cognition, and action.

The growth plan defines an endless series of successor iterations for the limbs. The limbs that exist at the start of the n -th iteration are the parent limbs. Parent limbs have growth sites. Each growth site generates an offspring limb. Each offspring limb has a bottom and a top. The bottom is attached to its parent by a joint. The top has twice as many growth sites as its parent. Each offspring limb is half as long and half as thick as its parent limb. The trunk grows two arms (each half as long and half as thick as the trunk). Each arm spouts four fingers (each half as long and half as thick as the arm). Each finger sprouts eight sub-fingers. The trunk also grows two legs (each half as long and half as thick as the trunk). Each leg grows four toes. Each toe grows eight sub-toes. Each n -th limb has the form of a branching binary tree with n levels and 2^n limb-tips (see Moravec, 1988: 102 - 108; 2000: 150 - 154). The series of finitely complex limbs converges in the limit to an infinitely complex limb with as many levels as natural numbers and as many tips as real numbers. These limbs can do infinitely complex motor tasks.

The growth plan defines an endless series of successor iterations for the eyes. The initial eye closely resembles a normal mature human eye. On each successor iteration, each photo-cell in the retina divides in an ideal way into two daughter photo-cell. Each daughter photo-cell is twice as small, fast, reliable, and efficient as its mother. So each next retina is twice as dense and twice as computationally powerful as the previous retina. The lens of each eye is progressively more perfect. Normal human eyes are sensitive to a limited part of the optical spectrum. But the idealization of the function of the eye implies a universalization of its function. It implies a widening of its sensitivity until it covers the entire spectrum of radiant energy in its universe. The series of finitely complex retinas converges in the limit to an infinitely complex retina. The infinitary eye has an infinitely dense retina (between any two photo-cells, there is another photo-cell). The infinitary eye can see an infinitely detailed object in a single glance. For instance, it can see Royce's perfect map of England in England in a single glance (Royce, 1899: 506 - 507). The infinitary eye can perform infinitely complex perceptual tasks.

The growth plan defines an endless series of successor iterations for the brain and nervous system. The initial brain closely resembles a normal mature human brain. On each successor iteration, the next brain is derived from the previous brain by dividing every neuron and doubling its connections. Each offspring neuron is twice as small, fast, reliable, and efficient as its parent. As the brain grows by ideal division, so the nervous system grows. A branching pattern of nerves follows the branching pattern of limbs. The nerves that control the next level of limbs are twice as small, fast, reliable, and efficient. Hence each next level of limbs is more precisely controlled. The series of finitely

complex brains converges in the limit to an infinitary brain. An infinitary brain is dense neural network (between any two neurons, there is another neuron wired to both of them). An infinitary brain has the computational power of an accelerating universal Turing machine (Copeland, 1998a; Steinhart, 2003). It can perform infinitely complex cognitive tasks.

7. Conclusion

We have informally discussed a variety of infinitely complex machines and minds. On the one hand, we might wonder whether any infinitary minds exist in our universe. On the other hand, we might wonder whether our universe exists in an infinitary mind. An infinitary mind can simulate any finitary process. For instance, it can exactly simulate or realize a finitary physical universe. A finitary universe has finitary space, finitary time, and finitary matter (there is some fixed finite upper bound on the number of bits of information that can be stored in any finite volume of space-time). Such universes may contain finitary living thinking things (e.g. finitary persons). An infinitary mind can simulate all possible finitary universes in exact detail in any finite time.⁹

Some writers have argued that our universe is only finitely complex (see Fredkin, 1991; Finkelstein, 1995; Steinhart, 1998). So, if infinitary minds exist, they can easily run simulations of our universe. Bostrom (2003) argues that this possibility must be taken seriously. Of course, theists have long argued that our universe is the product of an intelligent creator (an infinitary mind). Some theists have argued (roughly) that everything must have an explanation; since our universe is something, it must have an explanation; the best explanation for the existence of our universe is an intelligent creator. Skeptics have objected that if everything must have an explanation, then the intelligent creator must have an explanation. And if the best explanation for our universe is an intelligent creator C_1 , then by analogy the best explanation for the existence of an intelligent creator is an even more intelligent super-creator C_2 . But then C_2 is best explained by C_3 , and so it goes. The result is an endless regress of creators. Such a regress is possible.¹⁰ Since one kind of creation is simulation (as when an author imagines a character – see Moravec, 1988: 122 – 124), the endless series of creators can be thought of as an endless hierarchy of simulators. It is easy to describe an endless hierarchy of simulators:

1. **Initial Rule.** An initial universe is a computation that is not simulating any other universe. It is like a novel in which no character is writing another novel (and by *novel* we mean a novel that completely and exactly describes an entire universe). There is at least one initial computation (namely, our universe C_0). There may be many other initial computations (other initial universes).
2. **Successor Rule.** Every computation is a virtual machine running on a more powerful computation. For every computation C_n , there exists a more powerful computer C_{n+1} that runs C_n as a virtual machine. C_{n+1} simulates C_n . This is like saying that every novel is written by a character (a virtual author) in another novel. We thus have an endless series of virtual authors writing novels within novels.

3. Limit Rule. For every endless series of computers simulating computers, there exists a more powerful computer that simulates that entire series. The limit computer is more powerful than every computer in the series of which it is a limit. The first limit computer is C_ω . Following our literary analogy, C_ω writes a novel in which (1) some character writes a novel and (2) every novel written by a character contains a character who writes a novel. And C_ω is a character in yet another novel $C_{\omega+1}$.

We can gain even more precision by defining a chain of simulators as a function from the whole ZFC ordinal number line to simulations. One might conclude that every chain that spans the whole ZFC ordinal number line is the result of a simulation by an Absolute Mind whose power is proportional to the proper class of ordinals.¹¹ Obviously, we have not provided any argument for this endless series of nested simulations. And it is far from clear that any serious argument can be given for such a series. But if such a series does exist, then it is truly a case of intelligent computing everywhere. Our universe would be a thought in an infinitary mind; and that infinitary mind would itself be a thought in an even greater infinitary mind; and so on. If a scheme like this were true, it would be a computational version of the classical idealism of Berkeley and Royce. And if every possible universe is a member of a chain of nested simulations that converges to an Absolute Mind, then we would really have a case of intelligent computing *Everywhere*.

Notes

¹It has long been traditional in Western thought to argue that the theistic God is an infinite mind. Hence the old arguments for the existence of God would be arguments for the existence of an infinite mind. For good introductions to divine infinity, see Leblanc (1993) and Achtner (2005). However, their discussions are primarily historical. We do not appeal to the old arguments for the existence of God, and we do not see any easy way to link the modern Cantorian theory of the infinite with traditional theism.

²We think of convergence in terms of increasing resemblance. The idea is that as $n \rightarrow \omega$, the *difference* between the tape T_n and the limit tape becomes arbitrarily small. We think of each tape as encoding a real number and the difference between tapes as the difference between real numbers. Hence convergence for a sequence of tapes is defined using the familiar Weierstrass theory of limits. Another kind of convergence is set-theoretic. An infinite series of set-theoretic structures can converge to a set-theoretic limit. Set-theoretic limits are defined in terms of infinite unions: $\bigcup_{n < \omega} S_n$. For example: the set-theoretic limit of the von Neumann ordinals $\{\}, \{\{\}\}, \{\{\{\}\}\}$, and so on is the von Neumann ω .

³There is a large literature associated with the analysis of ethical, social, and political concepts in game-theoretic terms. Axelrod (1984) and Danielson (1992) develop ethical theories in the context of games played by machines. For a preliminary discussion of infinitary value theory, see Sorensen (1994).

⁴For some examples, see Gale & Stewart (1953); Galvin et al. (1978); Freiling (1984); Jech (1984); Ciesielski & Laver (1990); Scheepers (1993); Huddleston & Shurman (2002).

⁵Each player in the infinite number selection game has an data structure M (its memory). M is a Zeno Matrix with 4 columns labeled (input, current state, next state F , output G). For the sake of starting the game, M has an initial row whose input column is blank. Every row in M has the form (n, k, kno, o) . That is, the next state is the current state plus the input plus the output. Each player has a variable p that points at a row in M . The players are the First Player (player number 1) and the Second Player (player number 2). The tape is initially filled entirely with blanks. Player 1 starts with his or her head over cell 0 and player 2 starts with his or her head over cell 1. Each player runs the following algorithm: state = blank; the head is positioned over cell with address (player number – 1); the index p into M is initially 0; for i accelerating from 0 to ω do { in = read the content of the cell under your head; advance p until $M[p, \text{input}]$ is in and $M[p, \text{current state}] = \text{state}$; state = $M[p, F]$; out = $M[p, G]$; advance head; write out on the cell under your head; wait one clock tick for the other player; advance your tape head to catch up with the other player; }.

⁶We are working within the standard set theory known as ZFC. The Axiom of Choice in ZFC entails that every set can be well-ordered. Hence there is a well-ordering of every uncountable set. For example, there is a well-ordering of the real numbers. And yet no such well-ordering of the reals is known. We thus enter *terra incognita*.

⁷The growth plan of the human body is recursive: the rules in the growth plan are repeatedly applied to their own outputs. Since the growth plan is recursive, it generates a

body with a fractal structure (Goldberger et al., 1990; Bassingthwaighte et al., 1994). Our human bodies are generated by iterating the growth plan only a few times. The result is a fractal structure with only a few levels of depth (a fractal bush). But the growth plan can be iterated any number of times. More iterations increases the fractal depth of the body. The growth plan can be iterated to infinity, thus making an infinitely deep fractal body.

⁸We have defined idealizations only for the eyes, limbs, and brain. Idealizations of other organs are easily defined. These idealizations both universalize and infinitely amplify the powers of the organs. For example, the respiratory, digestive, and circulatory systems have the function of supplying the body with energy. These organs are universalized so the ideal body can obtain energy from any source in its universe.

⁹Since there are only ω many finitary universes, an infinitary mind can build a database in its memory of all finitary universes. This database is the finitary part of Leibniz's Palace of the Fates (Leibniz, 1996: secs. 414-417).

¹⁰Bostrom (2003: 253 – 254) writes: "Virtual machines can be stacked: it's possible to simulate a machine simulating another machine, and so on, in arbitrarily many steps of iteration. . . . Reality may thus contain many levels. . . . Although all the elements of such a system can be naturalistic, even physical, it is possible to draw some loose analogies with religious conceptions of the world. . . . Further rumination on these themes could climax in a *naturalistic theogony* that would study the structure of this hierarchy."

¹¹An absolutely infinite mind (with power proportional a proper class) would be the mind of God (see Taliaferro, 1985). But how would such a mind be defined? We cannot use any logical or mathematical techniques, since proper classes cannot enter into more complex structures. It does not seem likely that there is any machine or mind with complexity proportional to a proper class. On the contrary, set theoretic reflection principles suggest that anything that looks like an absolutely infinite machine or mind is really only an object at some non-absolute level of the set theoretic hierarchy.

References

- Achtner, W. (2005) Infinity in science and religion. *Neue Zeitschrift für Systematische Theologie und Religionsphilosophie* 47, 392 – 411.
- Axelrod, R. (1984) *The Evolution of Cooperation*. New York: Basic Books.
- Bassingthwaighte, J., Liebovitch, L., & West, B. (1994) *Fractal Physiology*. New York: Oxford University Press.
- Boolos, G. & Jeffrey, R. (1989) *Computability and Logic*. New York: Cambridge University Press.
- Bostrom, N. (2003) Are you living in a computer simulation? *Philosophical Quarterly* 53 (211): 243-255.
- Ciesielski, K. & Laver, R. (1990) A game of D. Gale in which one of the players has limited memory. *Periodica mathematica Hungarica* 22, 153 - 158.
- Conway, J. (2001) *On Numbers and Games* (2nd Edition). Natick, MA: A. K. Peters.
- Copeland, B. J. (1998a) Super Turing-machines. *Complexity* 4 (1), 30 - 32.
- Copeland, B. J. (1998b) Even Turing machines can compute uncomputable functions. In C. Calude, J. Casti, and M. Dinneen (Eds.), *Unconventional models of Computation*. New York: Springer-Verlag, 150 – 164.
- Danielson, P. (1992) *Artificial Morality: Virtuous Robots for Virtual Games*. New York: Routledge.
- Davies, E. B. (2000) Building infinite machines. *British Journal for the Philosophy of Science* 52 (4), 671 - 682.
- Doyle, J. (1991) The foundations of psychology: A logico-computational inquiry into the concept of mind. In R. Cummins & J. Pollock (Eds.) *Philosophy and AI: Essays at the Interface*. Cambridge, MA: MIT Press, pp. 39 - 78.
- Drake, F. (1974) *Set Theory: An Introduction to Large Cardinals*. New York: American Elsevier.
- Earman, J. & Norton, J. (1993) Forever is a day: Supertasks in Pitowsky and Malament-Hogarth spacetimes. *Philosophy of Science* 60, 22 - 42.
- Finkelstein, D. (1995) Finite physics. In R. Herken (ed.) *The Universal Turing Machine: A Half-Century Survey*. New York: Springer-Verlag, 323 - 347.

- Fredkin, E. (1991) Digital mechanics: An informational process based on reversible universal cellular automata. In Gutowitz, H. (1991) (Ed.) *Cellular Automata: Theory and Experiment*. Cambridge, MA: MIT Press, 254-270.
- Freiling, C. (1984) Banach games. *Journal of Symbolic Logic* 49 (2), 343 - 375.
- Gale, D. & Stewart, F. M. (1953) Infinite games with perfect information. In H. Kuhn & A. Tucker (Eds.), *Contributions to the Theory of Games Vol. 2*. Annals of Mathematical Studies 28. Princeton, NJ: Princeton University Press, 245 - 266.
- Galvin, F., Jech, T., & Magidor, M. (1978) An ideal game. *Journal of Symbolic Logic* 43 (2), 284 - 292.
- Goldberger, A., Rigney, D., West, B. (1990) Chaos and fractals in human physiology. *Scientific American* 262 (2), 42 -49.
- Grunbaum, A. (1969) Can an infinitude of operations be performed in a finite time? *British Journal of the Philosophy of Science* 20, 203 - 218.
- Hamilton, A. (1982) *Numbers, Sets, and Axioms: The Apparatus of Mathematics*. New York: Cambridge University Press.
- Hamkins, J. D. & Lewis, A. (2000) Infinite time Turing machines. *Journal of Symbolic Logic* 65 (2), 567 - 604.
- Hopcroft, J. (1984) Turing machines. *Scientific American* 250 (5), 86 - 107.
- Huddleston, S. & Shurman, J. (2002) Transfinite chomp. In R. Nowakowski (Ed.) (2002) *More Games of No Chance*. New York: Cambridge University Press.
- Jech, T. (1984) More game-theoretic properties of boolean algebras. *Annals of Pure and Applied Logic* 26, 11-29.
- Koetsier, T. & Allis, V. (1997) Assaying supertasks. *Logique et Analyse* 159, 291 - 313.
- Leblanc, J. (1993) Infinity in theology and mathematics. *Religious Studies* 29 (1), 51 - 62.
- Leibniz, G. W. (1996/1710) *Theodicy*. Trans. E. M. Huggard. Peru, IL: Open Court.
- Maes, P. (1995) Modeling adaptive autonomous agents. In C. Langton (Ed.), *Artificial Life: An Overview*. Cambridge, MA: MIT Press, 135 - 162.
- Moravec, H. (1988) *Mind Children: The Future of Robot and Human Intelligence*. Cambridge MA: Harvard University Press.

- Neeman, I. (2004) *The Determinacy of Long Games*. New York: Walter de Gruyter.
- Peirce, C. S. (1965) *Collected Papers of Charles Sanders Peirce*. Edited by C. Hartshorne & P. Weiss. Cambridge, MA: Harvard University Press.
- Pickover, C. (1995) *Keys to Infinity*. New York: John Wiley & Sons.
- Prusinkiewicz, P. & Lindenmayer, A. (1990) *The Algorithmic Beauty of Plants*. New York: Springer-Verlag.
- Royce, J. (1899) *The World and the Individual*. First Series. Supplementary Essay. New York: The Macmillan Company.
- Russell, S. & Norvig, P. (1995) *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, NJ: Prentice Hall.
- Scheepers, M. (1993) Variations on a game of Gale (I): Coding strategies. *Journal of Symbolic Logic* 58 (3), 1035 - 1043.
- Sorensen, R. (1994) Infinite decision theory. In J. Jordan (Ed.) *Gambling on God*. London: Rowman & Littlefield, 139 - 159.
- Steinhart, E. (1998) Digital metaphysics. In T. Bynum & J. Moor (Eds.), *The Digital Phoenix: How Computers are Changing Philosophy*. New York: Basil Blackwell, 117-134.
- Steinhart, E. (2003) Supermachines and superminds. *Minds and Machines* 13, 155 - 186.
- Taliaferro, C. (1985) Divine cognitive power. *International Journal for Philosophy of Religion* 18, 133 - 140.
- Thomson, J. (1954) Tasks and supertasks. *Analysis* 15, 1 - 13.
- Weizenbaum, J. (1976) *Computer Power and Human Reason: From Judgment to Calculation*. New York: W. H. Freeman and Co.
- Weyl, H. (1963) *Philosophy of Mathematics and Natural Science*. New York: Atheneum. (Original work in German 1927).