Church's thesis meets the N-body problem

Warren D. Smith[*]
Email: WDSmith@fastmail.fm

1 March 1999; revised for publication Jan. 2005

Abstract —
"Church's thesis" is at the foundation of computer science. We point out that with any particular set of physical laws, Church's thesis need not merely be postulated, in fact it may be decidable. Trying to do so is valuable. In Newton's laws of physics with point masses, we outline a proof that Church's thesis is false; phsyics is unsimulable. But with certain more realistic laws of motion, incorporating some relativistic effects, the Extended Church's thesis is true. Along the way we prove a useful theorem: a wide class of ordinary differential equations may be integrated with "polynomial slowdown." Warning: we cannot give careful definitions and caveats in this abstract – you must read the full text – and interpreting our results is not trivial.

Keywords — Newtonian N-body problem, Church's thesis, computability, numerical methods for ordinary differential equations.

Contents

1 Background

You should know what a Universal Turing Machine is and what the Church thesis is [mi67], and the basics of Newtonian Mechanics (especially the 2-body problem [co89]) and Special Relativity. The careful reader would want to know a fair amount about numerical methods for ordinary differential equations and their terminology [bu87] [ha76], and would want to have carefully examined Gerver's paper [ge91].

See the appendix for explanations of some of our notation.

[*] Work done while at the now-defunct NEC Research Institute, 4 Independence Way Princeton NJ 08540 USA. Non-electronic correspondence to: 21 Shore Oaks Drive, Stony Brook NY 11790.

2 Introduction. Our results and their interpretation.

"Church's thesis," or the "Church-Turing thesis" [tu36] [mi67], states that the *set of things commonly understood to be computation is identical with the set of tasks that can be carried out by a Turing machine.*

At first, Church's thesis seems merely to be a definition of the word "computation" and thus content-free. Indeed, it does have some of a character somewhere between that of a definition and an assertion, which is why it is always stated in an intentionally slightly vague way.

However, it can also be interpreted as a profound claim about the physical laws of our universe, i.e.: *any physical system that purports to be a "computer" is not capable of any computational task that a Turing machine is incapable of.*

Definition 1 *If computer A will always complete a task whose input is L bits long in time $T(L)$, and computer B always does the same task in time $\leq P(T(L), L)$ where P is a polynomial, then B is said to have "polynomial slowdown" relative to A.*

The "extended" Church thesis states that a Turing machine can do anything any other kind of (physically realizable) computer can do, with at most polynomial slowdown[1].

Church's thesis lies at the heart of theoretical computer science and physics; if it were false, much of the life's work of most computer scientists and theoretical physicists would become worthless, or at least, worth less.

So an important question is now to try to formulate certain sets of physical laws and to try to determine whether Church's thesis or the extended Church's thesis would be valid in a universe with those physical laws. A way to prove the (extended) Church's thesis is to construct an (efficient) algorithm for simulating any physical system. A way to disprove Church's thesis is to show how to use the laws of physics to construct a "computer" that can do something that Turing machines cannot do.

Let N be a fixed whole number. The "Newtonian N-body problem" is to describe the motion of N point masses whose initial locations and velocities are given, assuming that Newton's law of gravity $F = Gm_1m_2/r^2$ and acceleration $F = ma$ hold. We will sketch a proof (it is based on J. Gerver's proof of the "Painlevé conjecture" in the plane) of theorem 4 that an *uncountably* infinite number of *topologically distinct* trajectories are possible in 1 second, among the planar N-body problems with fixed masses and whose initial locations lie within certain disjoint balls and whose velocities are bounded. Meanwhile, of course, Turing machines can only experience a finite number of possible histories in a finite time. As a consequence, it is impossible for a Turing machine to compute a correct qualitative description of the motion that N bodies can make in 1 second. (Here, of course, "1 second" could be replaced by any finite interval of time.)

This result can be interpreted as "the (unextended) Church thesis would not be valid in a universe with point masses and Newton's laws of motion." This interpretation of our result is muddled by the fact that Newtonian physics involves real numbers specified infinitely precisely. However, the topological distinctness statement that we prove is discrete. Let us be clearer.

2.1 First way to interpret theorem 4: N-bodies are unsimulatable

A Turing machine could be given N real numbers as input by simply providing it with N infinite tapes each containing the binary representation of a number. This is in fact the realistic model of input if the task is N-body simulation. Such a Turing machine *could*, in infinite computational time, calculate the topology of the trajectories of N bodies up to the point in time (if any) where a singularity occurs. In fact, if the real numbers specifying the initial locations and momenta of the bodies were "usual," the Turing machine would in fact succeed in completing this calculation (for any particular desired amount of simulated time) in *finite* time. But, if the real numbers happened to be "unusual," in fact, if they happened to correspond to one of our examples whose trajectory topologies achieved singularity and infinite complexity in 1 second, then in finite computer time it could only partially describe the topology of the trajectory and would have to keep reading more input bits forever. (Probably our examples truly are unusual, in the sense that[2] they are measure zero in the space of real number tuples. However, we have not proven this.) Thus, this way of looking at it makes it apparent that the N-body system really can do something a Turing machine cannot do.

2.2 Second way to interpret theorem 4 in which N-bodies "solve" the general halting problem in 1.2 seconds.

[1] In the event that the physical system is not deterministic, then the Turing machine has to be given a random bit generator, and the criteria for completion of a computational task would have to become statistical. We will not concern ourselves with this in the present paper. Also, it is naturally essential, in order to give the physical universe any chance in the competition with a Turing machine having an infinite tape, that the laws of physics consider the universe to have infinite extent.

[2] In which case, it could be argued that these unusual examples correspond to a zero volume set in phase space. Compressing the phase space volume of a physical system exponentially small (e^{-n}) seems to require (by the second law of thermodynamics) increasing the entropy of the outside universe (outside this system, that is) by at least nK_B, which, assuming the outside universe was at temperature T, would require dissipating energy $\geq nk_BT$. Thus, setting up the initial conditions with infinite accuracy would require infinite energy expenditure, a conclusion which certainly must be taken into account when evaluating my claim to have "disproved Church's thesis." I still claim there is a sense in which N point masses with Newton's laws are more powerful than a Turing machine, but the present footnote exhibits a sense in which that is not the case.

Definition 2 *The "halting problem" is the following computational task. Input: a (finite length) description of a Turing machine T and the non-blank part of its tape. Output: would T eventually halt – yes or no?*

As is well known [mi67] the halting problem is "undecidable," that is, no Turing machine exists which can decide the halting problem in finite time.

Definition 3 *A "computable real number" is a real number R such that some Turing machine T with finite non-blank tape segment, and equipped with an "output," exists, that will output the digits of R in order. (Note: we allow the use of real number representations such as "5-digit binary" in which the digits a_i, $i \geq -M$, are in $\{-2, -1, 0, +1, +2\}$ and $R = \sum_{i \geq -M} a_i 2^{-i}$. This permits a certain amount of "correctability" of "numerical errors.")*

Our unsimulatable N-body examples are computable real numbers, in the sense that a Turing machine T_1 exists that, given sequential access to a description of the desired trajectory topology, will output the real numbers R corresponding to an initial configuration that would evolve according to that topology. (Strictly speaking, these real numbers are computable if and only if the trajectory topology is describable by a computable bit string. Otherwise, they are only computable in an "extended sense" where the input to T_1 in definition 3 is allowed to be infinite. T_1, of course, does not care, it will happily output the bits of R forever regardless of the size of its input.)

This view leads to a different way to interpret our results. (See figure 0.) Consider a Turing machine T_2 which, given initial real number data in such a form that it can access more bits on demand, by some ODE-timestepping scheme (suitable good ones will be described later in this paper) simulates the motion of the N bodies, to sufficient accuracy to be confident it knows the topology of the trajectory the bodies take in 1 second. (It keeps restarting the simulation using more precision, if necessary, to achieve such confidence.) Now, such a Turing machine, if asked to compute said topology and then halt, will halt iff and only if the N bodies do not reach singularity in 1 second (and in our examples, they will also physically reach infinite distance, in 1 second). Hence, the N-bodies are solving a Turing machine halting problem in 1 second. Of course (the reader may and should object) this is not *the* general halting problem, but rather a particular halting problem. But in fact, I claim, the general halting problem is not harder than this halting problem. Because, the system "$T_0|T_1|T_2$" in which the output of T_0 is piped into the input of T_1 and the output of T_1 is piped into the input of T_2 – the total combination is equivalent to another Turing machine we will call T_3[3] – halts if and only if, T_0 halts, and there is no restriction whatever on T_0.

Hence, we have two machines, the N-body system, and a Turing machine T_3, and we may make a face-to-face comparison. Given *finite* input, namely the computable real number described by $T_0|T_1$, T_2 either halts, or not, iff the unrestricted halting problem for T_0 halts, or not. Meanwhile, given the "same" input but expressed in *its* language (namely, actually as real numbers), the N-body system either explodes to infinity in 1 second, no body going within the unit circle, or it does not go infinitely far in 2 seconds and a large number of "asteroids" hit the unit circle (which contains, say, an unlucky cat) in 1.2 seconds[4]. This is about as fair a comparison as it is possible to make. Each machine swallows input in its own "language," namely, the Turing machine swallows a finite number of bits (the description of $T_0|T_1$) and the N-bodies swallow computable real numbers. There is no way to avoid feeding the N bodies infinitely long real numbers. Had they been fed reals with terminating expansions, for example, they would not even "know" it, since the infinite string of zeros at the end would matter. Each machine has a $\leq$ 2-state output, readily recognizable, the simplest we could ask for: the cat dies, or not. The N-body machine can solve all such problems in 2 seconds, but no Turing machine can solve all such problems in any finite time bound per problem.

Thus (as a typical consequence; one could write any number of statements of this kind) there exist initial configurations of N bodies in the plane in which several of the bodies will hit the unit circle within 1.2 seconds if and only if the Riemann hypothesis is false[5].

2.3 *Church's thesis as a statement about simulability versus as a statement about buildability of supercomputers*

The present paper demonstrates (I claim) that *unsimulable physical systems exist* in Newton's laws of gravity and motion for point masses. However, it does *not* appear to demonstrate, that, if we lived in a universe governed by those laws, we could actually *build* a device with super-Turing computational power. This is because (conjecturally – this is an important open problem) the set of initial conditions corresponding to unsimulable behavior, has measure zero. If so, then there is no way to achieve such behavior if our placement accuracy is non-infinitesimal. Indeed, there appear to be fundamental thermodynamic reasons why imprecision, in classical mechanics, cannot be avoided. If we wish to compress our initial conditions into a very small "ball of imprecision" in phase space, of volume V, that would require

[3] More details: T_3 runs the UNIX$^{\text{TM}}$ operating system and hence has no trouble with IO pipes... also, in the event that T_2 halts, it sends "kill" signals to T_1 and T_0 so that T_3 will halt too... also, T_0 is not completely unrestricted, since we will demand that it output a character every state transition. This is of course not a real restriction... finally, the set of "topological types" that we allow T_1 and T_2 to think about happen to be in a simple 1-1 correspondence with finite and infinite bit strings (see §3.6) and singularity occurs if and only if the bit string is infinite.

[4] In place of "1.2" one could have used any upper bound on $1 + 1/(2\pi)$.

[5] and the trajectories may be interpreted as a description of a counterexample.

increasing the entropy of something else (the rest of the universe) by an additive amount proportional to $\ln V$, which would require energy of order $k_B T \ln V$ where T is the temperature of the rest of the universe. As $V \to 0$, this would be infinite.

2.4 Einstein to the rescue

Next, we show (theorem 8) that if Newton's laws of gravity and motion are replaced by certain more Einsteinian laws (which are basically intermediate between Newton's laws and General Relativity), then the number of topologically distinct trajectories that can happen in finite time, is finite, and indeed, efficient and accurate simulation is possible by a Turing machine with only polynomial slowdown.

There are at least 3 ways to alter Newton's laws so that Church's thesis is saved. The simplest is simply to alter Newton's laws so that there are short-range repulsive forces, for example, the force between two bodies of masses m_1 and m_2 at separation r could be $Gm_1 m_2(r - [m_1^{1/3} + m_2^{1/3}]K)/r^3$.

Another is the "special relativistic theory of gravity" (SRTG).

$$\vec{v}_i = \frac{d\vec{x}_i}{dt}, \quad \vec{p}_i = m_i \vec{v}_i$$

$$m_i = (m_{\text{rest}})_i / \sqrt{1 - |\vec{v}_i|^2/c^2}$$

$$\frac{d\vec{p}_i}{dt} = G \sum_{j \neq i} m_i m_j' \frac{\vec{x}_j' - \vec{x}_i}{|\vec{x}_j' - \vec{x}_i|^3}$$

where primed quantities are understood to apply at a retarded moment t', i.e. $\frac{1}{c}|\vec{x}_i - \vec{x}_j'| = t - t'$. ($t'$ depends on i and j, a fact obscured by our notation.) SRTG is an appealing theory of gravity, until one sees that it must be incorrect by considering certain general relativistic gedanken experiments ([mi73] pp. 187-9). SRTG is mathematically more complicated than Newton plus repulsion, due to the "retarded potentials" (leading to a "delay differential equation" instead of an ordinary differential equation, and forcing the initial data to be provided by an "oracle" who can provide data about the state of the system *before* $t = 0$) and also to the not-time-reversible possibility of "coalescing black holes."

The third and most complicated set of laws we discuss (and the one we will concentrate on) is "modified linearized general relativity," which has all the mathematically annoying features of SRTG, and also has tensors.

Our result can be interpreted as "Church's extended thesis would be valid in a universe having only a finite number of point masses, all obeying these laws of motion." But this interpretation is again muddled by the question of what initial conditions one is "allowed" to specify. To prove our simulation result we must assume that the initial rest masses and kinetic energies must be given in *unary*, although their initial positions and directions are written in binary. This seems "fair;" it would be ridiculous to allow you to have a mass-energy of 10^{100} grams for only 100 dollars and still desire that your power (computational or any other kind!) be only polynomially greater than mine.

While proving this result, we also prove an intermediate theorem (theorem 7) of general interest: a wide class of ordinary differential equations is simulatable with polynomial slowdown[6]. Besides infinite precision, the two key facts about the Newtonian N body problem which make it unsimulatable are (1) an infinite amount of potential energy is available by moving two point masses arbitrarily close together, (2) a body can move at arbitrarily high velocity by providing it with enough energy. In the new more realistic laws of gravity and motion: nothing can move faster than the speed of light, and if any two masses get closer than half their Schwarzschild radius R, the simulator is allowed to assume that they instantly combine into one point mass.

A more precise description of one thing we can do about Schwarzschild radii is as follows:

1. If masses approach within $(R_1 + R_2)/2$, always combine.

2. If two masses never get closer than $R_1 + R_2$, never combine.

3. If the closest approach is between these two values, then the decision is at the discretion of the simulator.

This is not as silly as it sounds, because "Kerr black holes" have event horizons whose radius is in $(R/2, R)$ where R is the Schwarzschild radius, but these event horizons are a function *not* only of the mass of the hole (known to the simulator) but also of its angular momentum (not known to the simulator; angular momentum of a point mass not being a Newtonian concept). Also, there is no simple theory of what the event horizons do during near-collisions, especially nonbinary ones.

[6] This is the first polytime result for ODEs. Previous papers related to this one are Moore's paper [mo90] showing that a point particle bouncing between a finite number of perfect parabolic and planar mirrors can simulate a Turing machine, and hence has undecidable long term behavior. Moore then observed that the mirrors could be "softened" so that the long term behavior of a particle moving in a certain smooth and finitely describable potential in 3D, is undecidable. Brockett [br89] shows how to construct a finite system of ODEs which will emulate a finite state machine or (with the aid of the sine function) a counter. By putting these together (which, for some reason, Brockett did not actually mention) one obtains "counter machines" which [mi67] have the full power of a Turing machine. Thus, once again, one has a small system of ODEs with undecidable long term behavior. Brockett's systems have the advantage that they are "self correcting;" a bounded amount of error in the real number initial data that specifies the counter machine that the ODE is simulating, will not matter. Also, since Brockett's systems are defined mathematically and not required to have any connection to physics, one could also "speed up time" so that Brockett's machines will have undecidable behavior in finite "time." Ko [ko83] showed various results, mostly negative, about the abstract computational complexity of ODEs. For example: The solution $y(t)$ of $\dot{y} = F(y, t)$, even if F is polynomial-time computable, need not be computable. A different kind of undecidability result for systems of ODEs may be found in chapter 9 of [ma93].

2.5 Apology

I apologize to the reader for the excessive number of footnotes and the length of the discussion of the interpretation of the results. The latter was necessary because interpreting our results – as is often the case in the land of uncomputability – really is quite difficult! The former was felt to be necessary in view of the almost equally excessive amount of flak that the author has had to put up with. I feel that certainly any questioning of Church's thesis *should* be examined narrowly, but I feel that most of the criticisms that have been directed at this paper have been, while seemingly well motivated, actually unfounded. (Of course, an additional footnote was then required to refute each such criticism.) On the other hand, I feel that a few such criticisms actually *are* well founded, or at least worth discussing. For such discussion, plus some more on interpreting our results, and also a few open questions, go to the end of this paper.

3 The Newtonian N-body problem is unsimulatable (even in the plane)

3.1 A little background

The Newtonian N-body problem is governed by the laws of motion

$$\ddot{\vec{x}}_i = G \sum_{j \neq i} m_j \frac{\vec{x}_j - \vec{x}_i}{|\vec{x}_j - \vec{x}_i|^3}, \tag{1}$$

where $\vec{x}_i$ is the position of body i at time t, m_j is the mass of body j, G is Newton's gravitational constant, and dots denote time derivatives. Some elementary theorems are that the total momentum $\vec{p} = \sum_j m_j \dot{\vec{x}}_j$, total angular momentum $\vec{L} = \sum_j m_j \vec{x}_j \times \dot{\vec{x}}_j$, and total energy $K + P$, where $K = \frac{1}{2} \sum_j m_j |\dot{\vec{x}}_j|^2$ is the kinetic energy and $P = -G \sum_i \sum_{j<i} \frac{m_i m_j}{|\vec{x}_i - \vec{x}_j|}$ is the potential energy, are each conserved, and the virial theorem $\ddot{J} = 4K - 2P$ where $J = \sum_j m_j |\vec{x}_j|^2$ is the moment of inertia, holds. The flow in $6N$-dimensional momentum-position "phase space" induced by Newton's laws, is volume-preserving. Saari ([sa73] and references therein) has shown that with any fixed set of N masses, the subset of phase space which will evolve to a collision singularity, is of measure zero.

In the 2-body problem, the bodies follow trajectories which in polar coordinates with the origin at the center of mass are $r = P/(1 - E \cos(\theta - \theta_0))$, i.e. conic sections with focus at the origin, where E is the "eccentricity," namely $E = 0$ for a circle, $E = 1$ for a parabola (infinitely long ellipse), $0 < E < 1$ for an ellipse, and $E > 1$ for a hyperbola. It is convenient to regard the eccentricity E is a *vector* $\vec{E}$ so that $|\vec{E}| = E$ and $\arg(E) = \theta_0$. P is a constant, namely (for ellipses) P is the radius when $\theta - \theta_0 = \pi/2$.

The speeds of the bodies along these trajectories follow from Kepler's law of "equal areas (of the conic) are swept out in equal times" which is really just conservation of angular momentum, and the law of conservation of total energy. For a circular orbit ("Kepler's third law") the orbital year is proportional to the 3/2 power of the orbital radius. The total energy (which is negative) is proportional to the reciprocal of the orbital radius.

We note that in a hyperbolic orbit with asymptote angle bounded below π, the minimal separation r between the two bodies will be of order $r \approx Gm_2/v_\infty^2$ where v_∞ is the speed of body 1 at ∞, and "all the action" (an arbitrarily large fixed fraction of the curvature of trajectory 1, that is) will take place during a time interval of order $t \approx Gm_2/v_\infty^3$.

It is sometimes possible to treat the motion of two bodies in an N-body problem as a small perturbation of the solution of the 2-body problem. In particular, if the forces exerted on the 2 bodies by the other $N - 2$ bodies have magnitude bounded above by a small constant ϵ, then their final momenta will be perturbed by at most $O(\epsilon t)$ and their final positions by at most $O(\epsilon t^2)$ after a duration t, in the limit $\epsilon \to 0+$ with all other initial data remaining fixed. Thus, considering the previous paragraph, hyperbolic swingbys with bend angle bounded within $(0, \pi)$, executed by sufficiently fast-moving objects, or near-circular orbits with sufficiently small radius, will be essentially unaffected by external forces during the characteristic time scales of these swingbys/orbits.

3.2 Gerver's example of a noncollision singularity in the plane

In the Newtonian N-body problem, a "singularity" is a moment of nonanalyticity of the motion. A "collision singularity" is a singularity in which there is a point of space which is simultaneously approached, as $t \to t_{\text{singular}}$, by two or more bodies. A "noncollision singularity" is a moment of nonanalyticity of the motion, at which no such point exists. In 1991, J. Gerver [ge91] proved that $3N$ point masses in the plane, for N sufficiently large, but fixed, could evolve under Newton's laws in finite time to a noncollision singularity[7]. Gerver's 68-page paper, and even Gerver's key lemmas, whose statements alone can take 3 pages, are too long to repeat, so we will content ourselves with a sketchy description of Gerver's argument, enough to indicate why our modifications to it will work.

Gerver's example is illustrated in figure 1. Although there are $3N$ bodies, there is exact N-way rotational symmetry, so that there are "really" only 3 bodies, the rest are just "images." Of these three bodies, two are "stars" each having unit mass. These two stars orbit about their common center of mass roughly circularly (or, to a better approximation, elliptically), forming a "binary star," with small orbital radius, located near a vertex of a regular N-gon. (The images of this binary lie near the other vertices of the N-gon, of course.) The

[7] A noncollision singularity for the 5-body problem in 3-space had been shown slightly earlier by Xia [xi92].

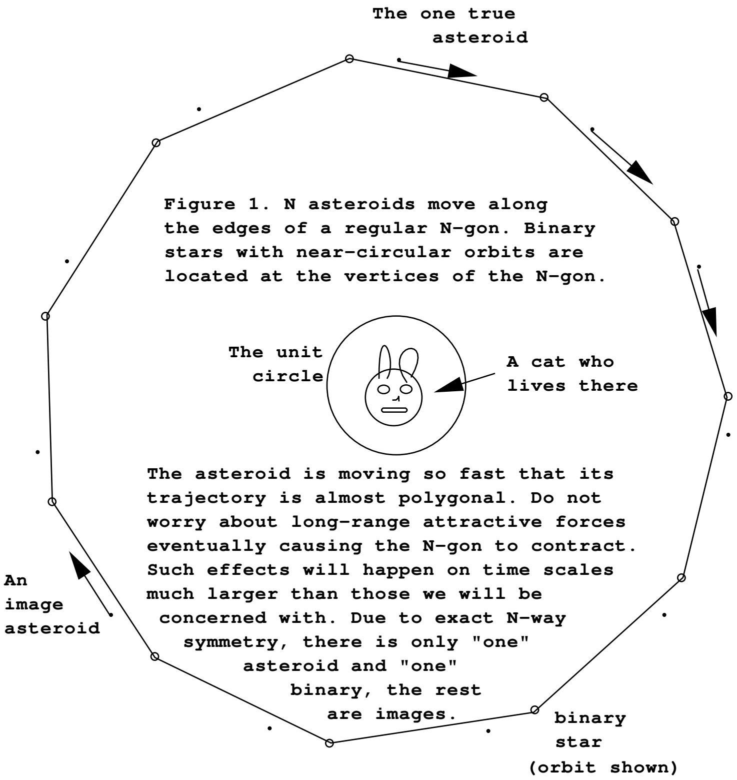

Figure 1:

third mass is an "asteroid" of mass μ^2 ($\mu^2 \approx 0.001$, say) which travels between the binary and its images in sequence, roughly moving along the perimeter of the N-gon. (Simultaneously, all the image asteroids are moving along each other edge of the N-gon, of course.) In the limit in which the radius of the binary went to zero while the N-gon remained of constant size and μ^2 went to zero, the binary would of course travel in an exact elliptical trajectory.

The binary may be said to have a "phase angle" in $[0, 2\pi]$ arising from Kepler's "equal areas in equal times" principle; that is, the area of the ellipse traversed so far, divided by the area of the whole ellipse, times 2π, is the "phase angle."

Each time the asteroid gets near a binary, it interacts with it in such a way that the following properties hold.

Property I. The asteroid winds up getting deflected at precisely the right angle to start moving along the next N-gon edge toward an eventual interaction with the next binary. As a consequence, the binary's center of mass is slightly accelerated away from the center of the N-gon, increasing the rate of expansion of the N-gon. The N-gon expands between interactions by a factor of $1 + 2\pi^2 N^{-2}\mu + o(N^{-2}\mu)$. (We explain the o and O symbols after property III.)

Property II. The asteroid extracts enough energy from the binary so that its speed is increased by a roughly constant factor $1 + \mu + o(\mu)$. Consequently, the radius of the binary contracts slightly to $1 - 2\mu + o(\mu)$ times its former value, and its orbital speed increases by a factor $1 + \mu + o(\mu)$, so that its orbital "year" is multiplied by a factor $1 - 3\mu + o(\mu)$. Considering property I, the transit time of the asteroid along an N-gon edge is multiplied by a factor of $1 - \mu + o(\mu)$, so that these transit times are decreasing, but measured in binary-star years, they are actually increasing by a factor of $1 + 2\mu + o(\mu)$ every interaction.

Property III. The asteroid's speed and deflection angle are in fact carefully adjusted so that it will intercept the next binary at precisely the right phase angle and position so that I, II, and III will also happen next time.

In the above, the uses of the symbols "o" and "O" are to be interpreted as pertaining to a hypothetical limiting process in which $N \to \infty$ and $N^{20}\mu \to 1$. Of course, N and μ are actually fixed, the O's and o's are merely a convenient labor saving device for proving the existence of suitable N and μ values.

As a consequence of properties I-III, the entire N-gon expands in roughly geometric progression to infinite size, but the asteroid traverses the N-gon edges in durations of time which shrink roughly geometrically toward zero. In consequence, the asteroid travels an infinite number of circuits around the N-gon (and this is an infinite distance), as the N-gon grows to infinite size, in a finite time t_{singular}, bounded by

$$t_{\text{singular}} < 2\mu^{-1}\tau_0 \tag{2}$$

where τ_0 is the initial time required for the asteroid to traverse an N-gon edge. But as $t \to t_{\text{singular}}$ there is no point of space approached by any body (much less by more than one of them simultaneously!) so this singularity is *not* a collision.

Furthermore, Gerver shows, during this process, we may require that the binary star's elliptical orbit never becomes very noncircular. Specifically, its "instantaneous eccentricity" E is always bounded by $0 \le E < 10\mu$. Also, the quotient of the speed of the asteroid by the binary star's orbital speed (i.e. with respect to its center of mass) remains roughly constant; it is $\sqrt{2}\mu^{-1}(1 + O(N^{-2}))$. Call these two statements **Property IV**.

3.3 Crude justification of Gerver's properties and proof

It seems to me that the two key ideas that Gerver required to make properties I-IV, and hence his proof, hold, are: (1) estimates, many of which in fact may be made by crude "dimensionality arguments," justifying properties I and II although not some of the unimportant specific constants in them, and an "expansion argument" justifying property III.

We now outline the main crude dimensionality arguments required. We will not work carefully enough to get Gerver's specific constants, such as $\sqrt{2}$, but this will not matter; it seems to me the results we will obtain are good enough for our purposes despite the fact that the analysis is much simpler than, and sloppier than, Gerver's.

Since the asteroid is deflected an angle $\approx 2\pi/N$ during an interaction, and its mass is of order μ^2 times the mass of binary, it follows from conservation of momentum that the binary will be accelerated outward to obtain a velocity increment Δv corresponding to the expansion rate needed to cause the N-gon to expand by a factor 1 plus order $\mu^2 N^{-2}$ during the time it takes the asteroid to traverse the next N-gon edge. (The gravitational attraction between the binaries, causing N-gon shrinkage, is asymptotically negligible.) This leads to a linear difference equation among the N-gon radii and expansion velocities during the kth time interval whose solution grows exponentially as a function of k with growth rate 1 plus order μN^{-2} as claimed in property I. The fact that the asteroid's speed is increased by a factor of 1 plus order μ during an interaction, is apparent by conservation of momentum and the shape of the trajectories round each individual star – each time such an individual "slingshot" event occurs, the star is given a total impulse on the order of $\mu^2 v$ where v is the speed of the asteroid when it is out of the immediate neighborhood of the star, namely v is of order μ^{-1} times larger than the star's own orbital speed. A fraction of order 1 of this impulse is directed in opposition to the star's motion, and a fraction of order 1 is traverse to it. Since $|\vec{w}+\vec{\Delta}|^2 - |\vec{w}|^2 = 2\vec{\Delta}\cdot\vec{w} + |\vec{\Delta}|^2$ where $\vec{w}$ is the orbital velocity of the star and $\vec{\Delta}$ is the increment to it, where $|\vec{\Delta}|/|\vec{w}|$ is of (small) order μ, only the nontraverse part matters as far as the energy of the star's orbit is concerned, and this shows that the star indeed loses a fraction of its energy and orbital momentum of order μ. The fact that the eccentricity of the binary star's orbit stays small is one of the trickier arguments[8]. Because the stars's orbital velocities get changed, during an interaction, by a fraction of order μ, the change $\vec{\Delta}$ in $\vec{E}$ during an interaction is also of order μ (in fact, Gerver shows its magnitude is bounded by $(7 + \sqrt{8} + o(1))\mu$). It turns out that by choosing the orbital phase correctly (i.e. as the correct function of $\vec{E}$) at the beginning of the interaction, one may always assure that $\vec{\Delta}\cdot\vec{E} < 0$, in fact that the angle between $\vec{\Delta}$ and $\vec{E}$ is bounded inside $(100, 260)$ degrees. Choosing the initial phase angle anywhere within an interval of width almost $\pi/2$ will always suffice to assure this, in fact, that is, approximately half of the possible reasonable initial phase angles are eccentricity-accentuating, and approximately half are eccentricity-opposing[9] (and if the initial eccentricity is large enough, in fact eccentricity-decreasing). If a damping-type choice is always chosen, $|\vec{E}|$ will never get larger than order μ (in fact, Gerver shows the bound 10μ). In our alternative scenario in figure 2, the same sort of eccentricity behavior must happen, since the total momentum-transfer properties of the *two* swingbys of each star are the same, except for slight changes in the constants which do not affect the key sign – in other words, again, approximately half the possible initial phase angles have the good eccentricity-limiting behavior. Finally, the fact that the ratio of the asteroid's speed to the binary's orbital speed, stays roughly constant, arises from conservation of energy over the long term, since the decrease in the star's orbital energy has to balance the increase in the asteroid's kinetic energy (all other energy terms being negligible in comparison) their speeds must stay in constant ratio.

The parts of properties I, II, and IV that are needed, have now been justified.

We now outline the "expansion argument" that proves that initial conditions exist, that will force property III to keep holding an infinite number of times. The idea is that small changes in the phase angle of the binary and the

---

[8]Actually, there is probably a simpler proof than the one Gerver gave. With $G = 4$, the following initial data $\vec{x}_{\text{star1}} = -\vec{x}_{\text{star2}} = (0.18268833855, 0.98317087577)$, $\vec{v}_{\text{star1}} = -\vec{v}_{\text{star2}} = (-0.99682983499, 0.07956305716)$, $m_{\text{star1}} = m_{\text{star2}} = 1$; $\vec{x}_{\text{asteroid}} = (-90.9402646668, 0.28622819645)$, $\vec{v}_{\text{asteroid}} = (6.32455532034, 0)$, $m_{\text{star2}} = 1/20$; causes the asteroid to swing past the binary with an asymptotic deflection angle of 20.000006° (note, $360/18 = 20$). The eccentricity of the binary changes from 0.10388522 to 0.10388549 during this process. The asteroid speeds up from 6.32 to 6.97. The binary's ellipse, whose long axis was oriented at 175.437738° above the x-axis, ends up at 175.437766°. In other words, up to errors in the 6th significant figure, this process causes the shape and orientation of the binary's orbit to be *preserved* and only its size to change, the resulting energy being imparted to the asteroid. This, together with arguments concerning the locally linear behavior of small perturbations, is a numerical "proof" that Gerver's construction works when $N = 18$, and it also simplifies his proof, since we now do not need to keep correcting the eccentricity to keep the binary near-circular; instead the binary stays elliptical, but the ellipse is exactly preserved.

[9]This "half one way, half the other way" behavior of course is plausible abstractly from the time-reversibility of Newton's laws. Eccentricity increase tends to happen if the swingby, which siphons off some of the orbital angular momentum of the stars, happens during the "high point" of their elliptical orbits, but decrease tends to happen if the swingby occurs at a "low point."

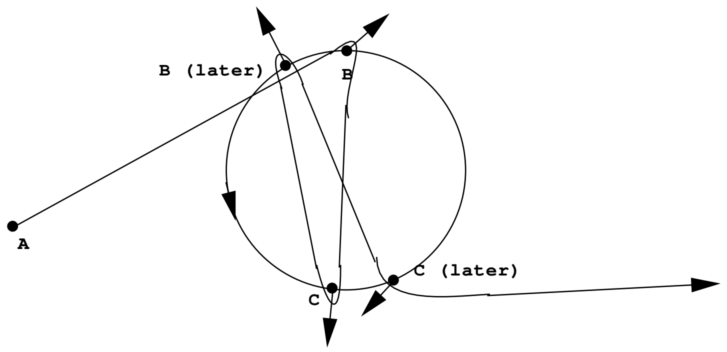

Figure 2. Detailed behavior during "type 1" interaction of asteroid with binary. Stars B and C are rotating happily in the roughly circular orbit shown. Asteroid A swings past star B then past star C, then B again (later in its orbit) and then finally C (again) and then begins the long trek along the N-gon edge toward the next binary. During the 4 hyperbolic swingbys, 4 momentum transfers from the asteroid to the star occur, roughly as shown by the vectors. The total effect of these momentum transfers has a component in opposition to the orbital velocity of the stars. Note: some distances have been exaggerated for clarity. In fact, the angular separation between B and B(later) or between C and C(later), is O(μ) (that is, small), and the separation between the asteroid and any star at the point of closest approach, is extremely small, so that the trajectory should in fact look much more like a chain of 5 line segments, than it does.

Figure 2:

deflection angle of the asteroid just prior to interaction number k, will cause large changes in these quantities just prior to interaction number $k+1$. This implies that a single fine tuning before interaction number 1 will suffice to adjust the infinite number of parameters controlling future interactions.

Firstly, since the behavior during a "slingshot event" is locally approximately hyperbolic, and since the asymptotes of the hyperbola may be adjusted to any angle in $(0, \pi)$ by fine adjustments in the initial trajectory of the asteroid, it is intuitively clear that any sequence of slingshot events you want can be forced by correct choice of conditions at the first one – so long as all the required hyperbola angles are bounded within $(0, \pi)$ and the asteroid has sufficiently tiny mass compared to the stellar mass and sufficiently large speed compared to the stellar speeds (both are assured by making μ small). Therefore, making trajectories, during a single asteroid-binary interaction, of the right qualitative type is no problem. Secondly, we observe that we have a very long "lever arm" (namely, an N-gon edge) so that it is obvious that a tiny change in initial data will cause a huge change at the next asteroid binary interaction. The only thing which is not tremendously obvious is that a infinitesimal change in the phase angle will cause a much larger infinitesimal change in the next phase angle. This is because changes in the phase angle affect the angles of momentum-transfers during slingshot events, relative to the direction of the star's orbital speed, in the first order, so that the number of star-years during the asteroid's trek to the next binary, which, recall, is a large number which grows geometrically, is affected fractionally in the first order. Thus, the expansion in the (phase angle, incoming trajectory location) 2-space between interactions, can be assured to be large in all directions. This is still the case even if the phase angles need to be chosen so that the change in the eccentricity vector $\vec{E}$ contains a decent component that is opposed to $\vec{E}$, since the size of the range of permissible phase angles is bounded above zero (indeed is $\approx \pi/2$).

Again, the preceding argument has not been nearly as careful or detailed as Gerver's (Gerver analyzed all 64 elements of the 8×8 Jacobian matrix of the recurrence map), but it seems to me, is adequate.

3.4 Side remark: a problem about baseball pitchers

Thor Johnson and Jade Vinson (students in lectures I gave at Princeton) suggested to me that there is a simpler problem whose analysis will give you some understanding of how it can be that Gerver's scenario will explode to infinite speed and radius in finite time, and also how it can be that the exponential growth rate for radius is not the same as for velocity.

Their problem, and a sketch of its solution, is as follows. A baseball pitcher on a frictionless plane surface flings a ball against a wall repeatedly (and catches it on the rebound). Each throw, he employs K times larger arm speed than on the previous throw.

Let M be the pitcher's mass, m be the ball's mass, and let v_1 be the velocity of the first fling. After F fling-catch

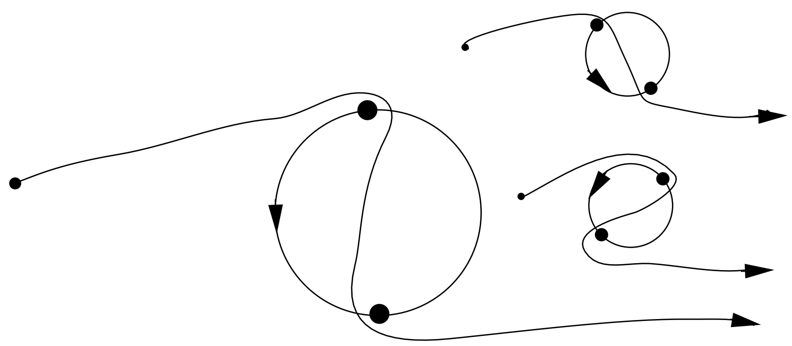

`Figure 3. Detailed behavior of "type 2" interaction of asteroid with binary. this simple trajectory was the one originally proposed by Gerver and it has the same sort of energy transfer properties as our alternate proposal in figure 2, in the limit when μ is small. Incidentally, in figures 2 and 3, the range of initial phase angles is quite flexible (there is a range of size nearly π), as is shown in the two smaller alternative versions of the present figure that are off to the right. This flexibility is essential to keep control of the eccentricity of the (slightly elliptical) orbit.`

Figure 3:

cycles, I find the pitcher is receding at speed

$$V_F = 2v_1 \frac{m}{M} \frac{K^F - 1}{K - 1}. \tag{3}$$

This grows at the same growth ratio as the ball speed. But the pitcher's distance D_F from the wall after F cycles does *not* grow at the same exponential growth ratio, due to the fact that the fling-catch cycles do not take constant time (which would have implied this) but in fact happen faster and faster (indeed exponentially so, explaining the different exponential growth rates), so that an infinite number of fling-catch cycles happen in finite time[10].

3.5 Modifications of Gerver's example

We propose the alternative detailed behavior during an asteroid-binary interaction in figure 2. This contrasts with Gerver's original proposal in figure 3. All the arguments of the preceding simplified version of Gerver's proof, still hold with this alternate type of interaction, only the constants possibly differ. (Note: since the asteroid now slingshots past each star *two* times, one must figure the *total* momentum transfer in the crude momentum-conservation arguments. The essential fact is that the angle this asteroid-to-star impulse transfer makes with the star's orbital velocity, is bounded above 90°. The quadruple swingby is essentially the same, from the standpoint of all momentum transfers, as a double swingby at a different initial orbital phase angle. The fact that the quadruple swingby is much more sensitive to initial conditions actually makes all the expansion arguments in the proof work *better*.)

[10]Can the difference equation for D_F be solved in closed form? We are not sure.

Therefore, we claim, initial conditions must exist in which each successive interaction will be[11] of type 1 or type 2, for *any* desired infinite sequence of 1's and 2's.

Furthermore, should we so desire, we can make one interaction be of a new "type 3" – causing the asteroid to be deflected *not* along the next N-gon edge, as usual, but instead toward the center of the N-gon. Naturally, if ever an interaction of type-3 occurred, the geometric expansion process would then cease and no further asteroid-binary interactions, nor any singularity, need occur. Furthermore, the N asteroids would in this case pass near the N-gon's center (intersect the unit circle, say, killing an N-way symmetric 2-dimensional cat who lives there) after $< \frac{1}{2\pi} t_{\text{singular}}$ more seconds, although in the case where interactions of types 1 and 2 only occurred, this would not happen.

Next, we observe that by walking backward through simulated time, say using rigorous bounds and interval arithmetic, we can actually perform the "fine adjustments" in the parameter values occurring just prior to previous interactions, computationally, restricting them to smaller and smaller sets every time we backstep one more collision. Thus our existence argument is not merely nonconstructive; arbitrarily good approximations to the real numbers involved are in fact computable.

[11]A different way to distinguish 1's and 2's would be, to make the asteroids, instead of going to the cyclically *next* binary star, instead *skip* one binary star and go to the one cyclically 2 ahead mod N. This would allow just using Gerver's original interaction, from figure 3, alone, making the validity of our argument especially clearly equivalent to the validity of his.

3.6 Conclusion

In an appropriate moving and rotating coordinate system (a, b) the two stars may be regarded as fixed at $(0, -1)$ and $(0, 1)$. The asteroid's trajectory than winds around these two fixed points and the "topological type" just means the homology of this path with respect to these points[12]. By simply recording whether each asteroid-binary interaction was of type 1 or type 2, we see that the topological types of the trajectories which arise in our examples are in 1-to-1 correspondence with the finite strings of '1's and '2's, ending with a '3', unioned with the set of infinite strings of '1's and '2's.

We then have proved[13]

Theorem 4 *For some finite fixed value of N, N point masses in the plane, whose initial positions and velocities lie inside a cube in $\mathbf{R}^{4N}$, can describe an uncountably infinite number of topologically distinct trajectories in 1 second. The initial locations and velocities of the bodies required to force a future trajectory of desired topological type, are computable (or extended sense computable) real numbers.*

4 With more realistic physical laws, the extended Church thesis is saved

The plan of this section is as follows. First, we will prove a general result (theorem 7) saying that a wide class (cf. "Assumptions II") of systems of ODEs (ordinary differential equations) are "simulatable on a Turing machine, with polynomial slowdown." Next, we will apply this result to the N-body problem with suitably modified laws of motion. We will also provide auxiliary discussions of which numerical schemes have the right behavior to make some version of theorem 7 hold (§4.3), and also of which Newton-like laws of motion are reasonably compatible with general relativity (§4.4-4.8). These auxiliary discussions hardly need to be read if you don't want to know.

Consider any N-dimensional system of ordinary differential equations of the form

$$\dot{\vec{x}} = \vec{F}(\vec{x}). \tag{4}$$

Apparently more general equations, such as allowing higher order than 1, allowing F to depend explicitly on time t or on time derivatives of elements of $\vec{x}$, and so on, in fact are easily put into the form above by adding a linear number of extra variables[14].

4.1 Euler's numerical method isn't good enough

Assumptions I: Suppose $\vec{F}$ is differentiable and computable to B bits of precision in time polynomial in B and N. Suppose also that the initial data $\vec{x}$ at $t = 0$ is specified as binary fixed point numbers. Also, suppose it is known that in the time interval $0 < t < T$, the absolute value of each component of $\vec{F}$, $\vec{x}$, and $\dot{\vec{x}}$, are *bounded* and that the absolute value of each partial derivative of $\vec{F}$ with respect to any one of its N arguments is similarly bounded (and consequently, the same is true for $\ddot{\vec{x}}$), by bounds that are polynomial functions of T, B, and N. Then

Theorem 5 *Under "assumptions I" above, the system (4) may be solved numerically on a Turing machine, so that the values of $\vec{x}$ at any desired time t with $0 \le t \le T$, may be calculated, accurate to $\pm\epsilon$, for any desired $\epsilon > 0$, by using "Euler" time stepping. The computation required during any time step is only polynomially large, but an exponentially large number of time steps (but not more) may be required.*

Proof Sketch. "Euler" time stepping, the simplest numerical scheme for solving ODEs, will suffice to show this rather weak theorem. If enough bits B of precision are used so that $2^{-B} \ll \delta$ and if $\vec{x}$ were exactly correct at the beginning of a time step, then the error made by a single iteration of simulated time δ, will be of order δ^2 times some bound P_1 polynomial in B, N. However, if the initial data were itself in error by an amount ϵ_{prev}, then the error at the end of the time step will be

$$\epsilon_{\text{new}} = \delta^2 P_1 + \delta P_2 \epsilon_{\text{prev}}. \tag{5}$$

The difficulty is the second term, which causes errors to accumulate potentially exponentially from time step to time step. The growth ratio of the exponential is bounded by $1 + \delta P_3$. Hence, after T/δ time steps, the total numerical error is bounded by $(1+\delta P_3)^{T/\delta}\delta^2 P_4$, which is less than (but of the same order as) $\exp(P_3 T)\delta^2 P_4$.

---

[12] It will be necessary, if the path is finite, to adjoin an infinite ray to its end, in order to clarify the homology. Should anyone object that our notion of "topology" is inherently 2-dimensional and thus not relevant to the 3D behavior of the bodies (restricted to a plane though they may be), we point out that the bodies could be emitting "laser beams" in the directions normal to the plane, which impinge on a sheet of photographic paper parallel to it...

[13] Sci-fi fans may enjoy the following scheme, related to our proof, for making a spaceship approach the speed c of light. Find a tight-binary B, one of whose members is a black hole (the other is some other compact object), orbiting at mild relativistic speeds, e.g. $0.01c$. Nearby (but much further away than the orbital radius of B) should be another black hole A. (Probably such a configuration exists somewhere in the universe; mildly relativistic neutron star binaries are known to exist [kl93].) Repeatedly make round trips between A and B, and arrange that each slingshot through B increases your speed by about $0.02c$. After 50 round trips one should near lightspeed and, with a final slingshot, one may fly off into the universe in any desired direction. The only energy input needed is tiny midcourse corrections performed while approximately midway between A and B. Note that although slingshot trajectories around A in Newtonian mechanics are hyperbolas with asymptote opening angle θ always obeying $0 < \theta < \pi$, in general relativistic gravity, "self-crossing" trajectories with $\theta < 0$ also are possible. That gives us enough control to make this work. There will be enormous accelerations during swingbys (e.g. from $+.9c$ to $-.9c$ in only 10 kilometers) but this by itself presents no difficulty because we are always in free fall. The difficulty arises from *tidal* forces during the swingbys, whose effects should be roughly equivalent to a large bomb exploding nearby. By minimizing the physical size and maximizing the strength of the spaceship ("brilliant pebble?") perhaps the tides could be survived.

[14] For example $\ddot{x} = F(x)$ would be rewritten as $\dot{y} = F(x)$, $\dot{x} = y$.

Thus, *if* we choose δ to be exponentially small, the simulation is performed to the desired accuracy and each iteration only requires polynomial compute time – but an exponentially large number of iterations are required. $\Box$

The theorem above is of course a rather weak result since the simulation involves exponential slowdown. (On the other hand, it does suffice to place the problem of simulating ODEs satisfying assumptions I in the complexity case PSPACE, and this is true even if we make ϵ exponentially small.)

If we widen "assumptions I" to state that F is $(k+1)$-time differentiable and that any partial derivative of F and any time derivative of $\vec{x}$, having degree $\le k+1$, is bounded, then we could use some other timestepping method (instead of the Euler method) of fixed degree k, where $k > 1$. (The Euler method has degree $k = 1$.) Any such method would also require an exponentially large number of iterations to make our error analysis assure a fixed accuracy ϵ (such as $\epsilon = 0.1$) after time T. However the growth constant of the exponential may be decreased. Using a kth degree method would tend to replace the expression $\exp(P_3T)\delta^2P_4$ for the total numerical error, by $\exp(P_3T)\delta^{k+1}P_4$ so that you could choose $\delta = (\epsilon/P_4)^{1/(k+1)}\exp(-P_3T/(k+1))$, thus cutting down the growth factor of the exponential to its $(k+1)/2$th root (if P_3 and P_4 remained the same – which they probably would not).

So it should now be clear that this exponential error buildup, and consequent need for exponential time slowdown, is in fact *unavoidable* with bounded degree time stepping methods with fixed accuracy goal, since it will tend to happen for virtually any system of ODEs having positive Liapunov exponent[15].

Note: This subsection has been essentially the same as Henrici's error analysis of the Euler method [he62]. The observation that integrating (EQ 4) with Lipshitzian polytime F is in PSPACE, was apparently first stated by Ko [ko83], and in fact Ko went beyond this by showing PSPACE-completeness, i.e. his PSPACE result was best possible. The present subsection has been included merely to set the stage for the next subsection.

4.2 But a scheme involving Runge-Kutta methods of unboundedly large degree, is good enough

Since bounded degree timestepping schemes are not good enough, we will use schemes of unboundedly large degree. Specifically, we will use the Implicit "Gaussian" Runge-Kutta methods devised by J.C. Butcher [bu64], and choose their degree to grow linearly with T.

Butcher's schemes have a number of pleasant properties. One of the ones most touted in the literature is "A-stability," described in [bu87]. In fact, Butcher's book defines a large variety of possible stability properties. The implications among them are pictured in figure 4. Note that "algebraic stability" and "L-stability" together imply all the other kinds of stability defined by Butcher. The "Gaussian" methods we use here are algebraically stable but are not weakened-L stable[16]. However, we will not need these stability properties.

In the below, let $k = 2v$ be even. (The degree of the scheme will be $k - 1$.) We *will* use the following properties:

Property I. The $v^2 + v$ coefficients in Butcher's kth degree scheme may be computed numerically to B bits of precision, since they arise from zeros of Legendre polynomials P_v of degree v as described on page 58 of [bu64], in computational time polynomial in k and B. Specifically, if $k = 2v$, let $c_1, c_2, ..., c_v$ be the v roots of $P_v(2c-1) = 0$ in increasing order. Then find a_{ij}, $i, j \in \{1, 2, ..., v\}$ as solutions to the v^2 linear equations (but only the v equations with i fixed need be considered at a time)

$$\sum_{j=1}^{v} a_{ij} c_j^{k-1} = \frac{1}{k} c_i^k \qquad (6)$$

for $i, k \in \{1, 2, ..., v\}$. Similarly find b_j, $j \in \{1, 2, ..., v\}$, as solutions to the v linear equations

$$\sum_{j=1}^{v} b_j c_j^{k-1} = \frac{1}{k}, \quad k \in \{1, 2, ..., v\}. \qquad (7)$$

Then Butcher's Runge-Kutta scheme for numerically solving (4) is

$$\vec{g_i} = \vec{F}\left(\vec{x}_{\text{old}} + \delta \sum_{j=1}^{v} a_{ij}\vec{g_j}\right), \quad i \in \{1, 2, ..., v\}, \qquad (8)$$

$$\vec{x}_{\text{new}} = \vec{x}_{\text{old}} + \delta \sum_{i=1}^{v} b_i \vec{g_i} \qquad (9)$$

where $\vec{x}_{\text{new}}$ is $\vec{x}$ at a time t which is δ larger than the value of t yielding $\vec{x}_{\text{old}}$, and $\vec{g_1}, ..., \vec{g_v}$ are defined implicitly by (8). (One has to know how to compute zeros of Legendre polynomials P_v efficiently. Just using their zero-dovetailing property and interval bisection is good enough for our purposes.)

Property II. The fact that the coefficients a_{ij} do not grow very large, is assured by the following new

Lemma 6 *The coefficients a_{ij} in Butcher's degree $2v$ Runge-Kutta scheme, $v \in \{1, 2, 3, \ldots\}$, always obey $a_{ij} \le$*

[15] And in particular, virtually any Hamiltonian system, in particular the Newtonian N-body problem with $N \ge 3$, will almost always have a positive Liapunov exponent (since phase space volume is preserved).

[16] The very similar "Lobatto" and "Radau" Runge-Kutta methods are both algebraically stable and L-stable, but it is not clear to me that this makes them preferable to the Gaussian methods. It could be argued that it is a great virtue for the boundary between the stable and unstable region to be precisely the imaginary line, and this "symmetry" property is incompatible with weakened-L stability and apparently is *uniquely* enjoyed by the Gaussian RK methods. In any case, the theorems of this paper still hold if you use Butcher's Radau or Lobatto RK methods in place of his Gaussian RK methods throughout.

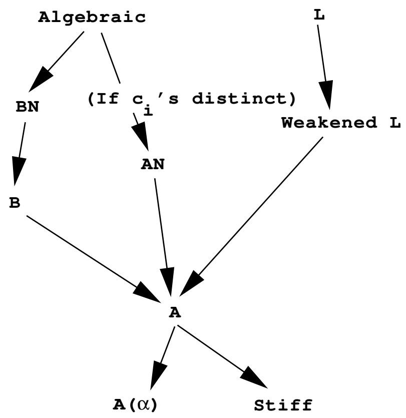

Figure 4.

Various kinds of "stability" properties that numerical methods for ODEs can have (as defined in Butcher 1987). "A→B" is to be read "A implies B."

Figure 4:

1. *Also, the matrix* A *of these coefficients has maximum eigenvalue* 1*, minimum eigenvalue* $1/v$*, and the Euclidean length, or the sum, of any row (or column) of* A *is* <1.

Proof.[17] Actually the statement about row sums is immediate from equation (EQ 6) with $k=1$, which shows that the ith row sum is c_i and clearly $0 < c_1 < c_2 < \ldots < c_v < 1$.

The other statements of the lemma are slightly less elementary. Equation (6) defining the a_{ij} may be written as a matrix equation among $v \times v$ matrices as follows:

$$A = DCRC^{-1} \tag{11}$$

where the i-down, j-across entry of A is a_{ij}, and of C is c_i^{j-1}, and D and R are diagonal matrices whose ith diagonal entries are c_i and $1/i$ respectively. Observe that the minimum eigenvalue of D is c_1, $c_1 > 1/v^2$, and the minimum eigenvalue of CRC^{-1} (which of course has the same eigenvalues as does R) is $1/v$, and the maximum eigenvalue of D is c_v, $c_v < 1$, and the maximum eigenvalue of R is 1. Since $D^{-1/2}AD^{1/2}$ is symmetric, A has only real eigenvalues, and in fact, its eigenvalues are identical to those of R. As a consequence, the maximum eigenvalue of A is 1 and the minimum eigenvalue of A is $1/v$. Using the maximum principle for eigenvalues now allows one to deduce that $\max_{i,j} |a_{ij}| < 1$ ($=1$ would only be possible if $A_{ii}=1$ and the other elements of row i were 0, but since the row sum is c_i, this is impossible, so 1 is a strict upper bound.) and that the Euclidean length, or the sum, of any row (or column) of A is <1. □

[17] Numerical evidence suggests some conjectures which would be even stronger than this lemma. Namely, it would appear $\min_{i,j} a_{ij}$ is a monotone increasing function of v, and $\max_{i,j} a_{ij}$ is a monotone decreasing function of v, when $v \geq 2$. These would imply that

$$-0.0387 \approx \frac{1}{4} - \frac{\sqrt{3}}{6} \leq a_{ij} \leq \frac{1}{4} + \frac{\sqrt{3}}{6} \approx 0.5387, \tag{10}$$

for $1 \leq i,j \leq v$, with strict inequality when $v > 2$. It also appears $\min a_{ij}$ behaves like $-c/v$ where $0.13 < c < 0.14$, and $\max a_{ij}$ like c/v where $1.6 < c < 1.7$, for v large. It appears $\max_j b_j$ and $\min_j b_j$ are both monotone decreasing functions of v and $\max_j b_j \leq 1$, with equality only when $v=1$, and $\min_j b_j > 0$. Min b_j seems to behave like c/v^2 where $3.4 < c < 3.7$, and $\max b_j$ like c/v where $1.4 < c < 1.7$, for v large.

Remark. Lemma 6 implies a "generalized interior point property" of Butcher's degree-k Runge-Kutta schemes. Specifically, an "interior point property" would assert that all the "intermediate points" $x_{\text{old}} + \delta \sum_j a_{ij} g_j$ lie "inside the time step interval" $[x_{\text{old}}, x_{\text{new}}]$. Of course, when we are in more than 1 dimension so that these x and g values are vectors, the word "interval" no longer makes sense (hence the word "generalized" in the preceding sentence) but if we are in 1 dimension and $F(\delta + x_{\text{old}}) = F_0 + F_0'\delta + O(\delta^2)$, then the lemma above shows that for all sufficiently small δ, $|g_i - F_0| < c\delta$ where the constant c does not depend on $k = 2v$, the degree of the scheme, but only upon F_0'. Hence all the intermediate points lie inside the interval $[x_{\text{old}}, x_{\text{old}} + Q\delta F_0 + O(\delta^2)]$, where Q is the maximum row sum of the A-matrix. Since by the lemma, Q is bounded below 1 for each v, and since $x_{\text{new}} = x_{\text{old}} + \delta F_0 + O(\delta^2)$, we conclude that: *If* F *is locally linear at* x_{old}*, then each of Butcher's schemes obey the interior point property for all sufficiently small* δ*.*

The interior point property is obviously a very desirable property for numerical ODE-solving schemes to possess, but, surprisingly, it has not been previously proven for any nontrivial class of Runge-Kutta methods, nor is it even mentioned in Butcher's book [bu87] or similar books. Of course, for any *particular* RK-scheme, the interior point property (or its falsity) is generally readily apparent.

Property III. With these coefficients known, and assuming that F obeys "assumptions II" below, one can actually use the Runge-Kutta scheme (8) (9) to perform a time step, in time polynomial in N, k, and B. This is actually not immediately clear, since (8) involves solving some nonlinear equations and the solution is only defined implicitly. However, as Butcher [bu64] shows in his appendix, and also as is shown in his later book ([bu87] section 341) in more generality, so long as F is Lipshitzian: $|F(\vec{x}) - F(\vec{y})| < |\vec{x} - \vec{y}| \cdot K$, where K is a constant such that

$$|\delta| \cdot |A| \cdot K < 1, \tag{12}$$

(here the symbol $|x|$ denotes absolute value if x is a scalar, Euclidean length if x is a vector, and Euclidean operator norm [largest eigenvalue] if x is a matrix) then the solution of equations (8) exists and is unique and the explicit iterative process (where the superscript denotes the iteration number s)

$$g_i^{(s)} = F\left(\vec{x}_{\text{old}} + \delta \sum_{j=1}^{i-1} a_{ij} g_j^{(s)} + \delta \sum_{j=i}^{v} a_{ij} g_j^{(s-1)}\right) \tag{13}$$

will converge to it geometrically in the Euclidean norm. In fact, due to lemma 6, we see that $|A| < 1$, so that $|K\delta| < 1$ suffices for convergence.

Property IV. The error ϵ_{step} incurred by using Butcher's degree-$(k-1)$ process, $k = 2v$, $v \in \{1, 2, 3, ...\}$, to perform a time step with simulated duration δ, may be bounded by using the bounds on the "principal error term" given on pages 58-59 of [bu64], the bound 3^k on the number of "elementary differentials" [bu63] of degree k, and Taylor's theorem with remainder. The resulting bound[18] is $\epsilon_{\text{step}} < \delta^{k+1} B_{k+1} 3^{k+1} (v!)^4$ where B_k is some bound on the maximum magnitude of the kth derivative of $\vec{x}$.

Assumptions II: Suppose $\vec{F}$ is infinitely differentiable and computable to B bits of precision in time polynomial in B and N. Suppose also that the initial data $\vec{x}$ at $t = 0$ is specified as binary fixed point numbers. Also, suppose it is known that in the time interval $0 < t < T$, the absolute value of each component of $\vec{F}$, $\vec{x}^{(k)}$, and the absolute value of each partial derivative of $\vec{F}$ with respect to any of its N arguments, having total differentiation-degree k, is similarly bounded, by bounds that are $(NkTB)^{O(k)}$. Then

Theorem 7 *Under "assumptions II" above, the system (4) may be solved numerically on a Turing machine, so that the values of $\vec{x}$ at any desired time t with $0 \le t \le T$, may be calculated, accurate to $\pm\epsilon$, for any desired $\epsilon > 0$, by using Butcher's implicit Runge-Kutta schemes for time stepping. The computation required during any time step is only polynomially large and the number of time steps that will be required depends only polynomially on T,N,B, and $\min(\epsilon, 1)^{-1/\max(1,TP_3)}$, where P_3 is a polynomial function of T,N,B.*

Proof sketch. Similarly to the previous proof, but now using Butcher's Runge-Kutta scheme of degree k instead of Euler's method, we have, using property IV,

$$\epsilon_{\text{new}} = \delta^{k+1} (kP_1)^{O(k)} + \delta P_2 \epsilon_{\text{prev}} \tag{14}$$

where P_1 and P_2 are polynomial functions of N, B, T. We use this to see that after T/δ time steps, the total numerical error is bounded by $(1+P_3\delta)^{T/\delta} \delta^{k+1} (kP_4)^{O(k)}$, which is less than but of the same order as $\exp(P_3 T) \delta^{k+1} (kP_4)^{O(k)}$. Thus, if we choose δ so that

$$\delta < \epsilon^{1/(k+1)} (kP_4)^{-O(1)} \exp(-P_3 T/k + 1) \tag{15}$$

the desired simulation is performed with final error bounded by ϵ, and each iteration only requires polynomial compute time by properties I, III. But if k is chosen, *not* to be constant, but instead to be a linear function of TP_3, then in fact $1/\delta$ will be only as large, at most, as some *polynomial* in T,B,N,$\min(\epsilon, 1)^{-1/\max(1,TP_3)}$. (Also, by lemma 6 and property III, a polynomially small δ suffices to ensure the quick geometric convergence of the iteration performed at each Runge Kutta step.) Incidentally, we need the result that B (the number of bits of precision we need to carry, see assumptions II) may also be chosen only polynomially large, which requires such standard observations [pa98] as the fact that the exact inversion of a matrix with integer entries, may be performed with precision only polynomially large in the total number of bits in all the (integer ratio) entries of the matrix. Conclusion: The total "slowdown" is only polynomial. □

In other words, the numerical integration of an N-dimensional ODE system (4) obeying assumptions II, over a time interval of duration T, is performable to accuracy ϵ, by a Turing machine, in compute time polynomial in N,T,B and $\min(\epsilon, 1)^{-1/\max(1,TP_3)}$.

This result would seem to be of considerable general interest, aside from the specific use we will put it to in this paper.

4.3 Read the Numerical Analysis literature, and be surprised

Most previous workers, for example we mention C.W. Gear's high quality DIFSUB algorithm [ge71] as a canonical example, advocated time stepping schemes of bounded degree to solve ODEs. DIFSUB uses "implicit multistep" schemes of degrees 1-6 having a certain "stiff stability" property and attempts to choose the degree and the step length sensibly at each step. But, like all schemes of bounded degree, DIFSUB cannot integrate the vast bulk of ODEs of the sort we care about, to fixed accuracy, without incurring exponential slowdown.

It is possible that somebody has advocated schemes of unbounded degree before, but certainly that has never been the view of the mainstream numerical ODE literature.

Since the gap between exponential and polynomial is rather severe (!), the question is, why?

Probable answers: (A1) the numerical analysis community was never particularly concerned about bit complexity and polynomial time. (A2) there is little point to using schemes of unbounded degree if your machine represents numbers with bounded word length.

[18] This bound is very conservative. In fact, the term $3^{k+1}(v!)^4$ almost certainly could be replaced by a rapidly *decreasing* function.

Re (A2), one of the things that we now understand is that to achieve fixed accuracy at fixed time, the *first* time step has to be made extremely (exponentially) accurately, since its error will later be amplified exponentially. Of course, this requires a polynomially large number of bits and is incompatible with bounded word length.

Finally, we point out that our choice of Butcher's Implicit Runge-Kutta methods may not be the only way, or the best way, to solve ODEs in polynomial time. We now mention some other numerical methods of unbounded degree for ODEs, the "bad" ones first.

(1) Taylor series methods are conceptually simpler than, and might well be more efficient than, the methods of the present paper. However, the resulting theorem would be of less wide utility, since Taylor series methods can't be used if you don't know how to take derivatives.

(2) Richardsonian methods (involving "extrapolation to the limit" of zero step size) will only achieve degree k with $\approx 2^k$ substeps, if step size halving is used; generally this is not good enough for polynomial time.

(3) At the present time, all known *explicit* RK schemes of unbounded degree k involve $\gtrsim 3k^2/8$ stages, and are thus much less efficient than Butcher's implicit RK schemes of comparable degrees, at least as far as function evaluation count is concerned. (Also, no explicit scheme can be A-stable.) On the other hand, with explicit methods the linear algebra work is smaller. This area needs more investigation; especially one would like to know what number of stages truly are required as a function of k.

While the preceding three ideas look more or less unpromising, the next two may be quite promising. The reader should recall that only the crudest stability statements were needed in our proof of polynomial slowdown, and exponentially growing (as a function of the degree k of the ODE timestepping scheme) error constants turn out not to be a problem – they cannot kill the proof. When I started this research I chose to investigate Butcher's Runge-Kutta schemes because they had all sorts of pleasant theoretical properties, but most of these properties were not used in my analysis, and perhaps it would be better if I had chosen a timestepping scheme less oriented toward good stability and error properties, and more oriented toward computational efficiency. But be careful: good stability properties and slowly growing error constants may matter in a more careful analysis in which one attempts to say something precise about the polynomial governing the slowdown. The question of what the best such polynomial is and how to get it, is open.

(4) Multistep methods have the attractive feature that you need only ≈ 1 more function evaluation per timestep (except for the annoying problem of getting them started) regardless of the degree of the method. For this reason they may dominate all Runge-Kutta methods. Multistep methods of high degree do have unattractive stability properties. The only multistep method that is A-stable is the implicit trapezoidal rule of degree 2. The "Dahlquist barrier" [da56] [da59], which forces stable multistep methods to have only about half the accuracy-degree one might have hoped for, does not bother us, since that is only a constant factor. The multistep methods of "Adams type," which are highly touted in many numerical analysis books, both implicit and explicit, have finite stability regions, and indeed the width of these regions, as a function of the degree k, is exponentially decreasing toward zero. (For example, the widths w of the stability interval $[-w, 0]$ of the implicit Adams methods of degree k with $k = 1, 2, 10, 20$, are $w = \infty, 6, 0.115, 0.00034$ respectively [ha74].) Thus the Adams methods of high degree are actually quite bad. Still, results of R. Jeltsch [je81] assure the existence of certain infinite families of multistep formulae with non-shrinking stability regions and only exponentially growing error constants. These, or something like them, are probably suitable.

(5) "General linear" methods [bu87] have the potential to combine the best features of both Runge-Kutta methods and multistep methods, but so far are little explored.

*4.4 General Relativity**

For the reader's conventience, we present an essentially complete description of general relativity condensed down to two pages.

In general relativity, spacetime is a (3+1)-dimensional manifold. (The 4th dimension is customarily ct where t is time[19].) "Position" on this manifold is specified by 4 "coordinates," that is, a unique real 4-tuple is associated to each point on the manifold by a diffeomorphism to $\mathbf{R}^{3+1}$. Because the manifold and $\mathbf{R}^{3+1}$ may not have the same topology, it may not be possible to handle the entire manifold at once with a single coordinate system, in which case one needs a finite or countable number of overlapping "coordinate patches." The "metric tensor" is a function of position denoted by $g_{\alpha\beta}$. This is a 4×4 symmetric matrix whose indices (taking values in {1,2,3,4}) are α and β. The element ds of infinitesimal "length" on the manifold is $g_{\alpha\beta}\mathrm{d}x^\alpha \mathrm{d}x^\beta$, where we are using the Einstein summation convention in which repeated indices are summed over ("$Q_\alpha Z^\alpha$" means $\sum_{\alpha=1}^4 Q_\alpha Z^\alpha$). This is not really a "metric," as the word is commonly understood, since the distance between two points can be negative. Positive distances are called "spacelike" and negative ones are "timelike."

In a coordinate system which is locally Euclidean at a particular point (and one always exists) $g_{\alpha\beta} = \eta_{\alpha\beta} \equiv \mathrm{diag}(1, 1, 1, -1)$ there. We use $g^{\alpha\beta}$ to denote the inverse matrix to $g_{\alpha\beta}$, and more generally, a quantity Q with an index that is a *subscript* is related to a quantity Q with a corresponding index that is a *superscript*

[19] Beware! A vast multitude of differing notational conventions are used by different authors in general relativity, including the use of "imaginary time," different overall signs for the metric and the Riemann tensor, and different summation conventions for Greek and Latin indices.

by $g_{\alpha\beta}Q^\alpha = Q_\beta$, $g^{\alpha\beta}Q_\alpha = Q^\beta$. This somewhat strange-looking, and at first confusing, notational convention is actually convenient, especially for "tensors." Tensors are (usually indexed) quantities which depend only on position in the manifold and not on coordinate system, in the sense that they transform in the obvious ways

$$Q^\alpha_{\text{new}} = c^\alpha_\beta Q^\beta_{\text{old}}, \quad Q^{\text{new}}_\alpha = c^\beta_\alpha Q^{\text{old}}_\beta, \tag{16}$$

where $c^\beta_\alpha = dx^\beta_{\text{old}}/\partial x^\alpha_{\text{new}}$, when one changes coordinate systems from x^μ_{old} to x^μ_{new}.

The "Christoffel symbols," which are 3-indexed non-tensor functions of position x, are defined by

$$\Gamma^\alpha_{\gamma\beta} = \Gamma^\alpha_{\beta\gamma} = \frac{1}{2} g^{\alpha\mu} \left(\frac{\partial g_{\mu\beta}}{\partial x^\gamma} + \frac{\partial g_{\mu\gamma}}{\partial x^\beta} - \frac{\partial g_{\beta\gamma}}{\partial x^\mu} \right). \tag{17}$$

Interpretation: $\Gamma^\mu_{\alpha\beta}$ is the μ-component of change in e_α (an infinitesimal unit vector in the x^α direction) caused by parallel transport along e_β.

The "Riemann curvature tensor" is a 4-indexed function of position defined by

$$R^\mu{}_{\alpha\beta\eta} = \Gamma^\mu_{\alpha\eta,\beta} - \Gamma^\mu_{\alpha\beta,\eta} + \Gamma^\mu_{\nu\beta}\Gamma^\nu_{\eta\alpha} - \Gamma^\mu_{\nu\eta}\Gamma^\nu_{\beta\alpha}. \tag{18}$$

Interpretation: parallel transport of a vector A^μ around an infinitesimal quadrilateral, close to being a "parallelogram," with legs u^γ and w^δ, will cause a change in A^μ of $\Delta A^\alpha = R^\alpha{}_{\beta\gamma\delta} A^\beta u^\gamma w^\delta$. Space is flat iff all elements of the Riemann curvature tensor are zero.

The "Ricci curvature tensor" is a 2-indexed function of position got by "contracting" the Riemann tensor:

$$R_{\nu\mu} = R_{\mu\nu} = R^\alpha{}_{\mu\alpha\nu} \tag{19}$$

and the "curvature scalar" is

$$R = g^{\alpha\beta} R_{\alpha\beta} \tag{20}$$

(The "Gaussian curvature" is $-R/2$.) Now we are ready to state "Einstein's field equations"

$$R_{\alpha\gamma} = \frac{8\pi G}{c^4} (T_{\alpha\gamma} - \frac{1}{2} g_{\alpha\gamma} T^\beta_\beta) \tag{21}$$

where G is Newton's gravitational constant. Here the 2-indexed function of position "$T_{\alpha\beta}$" describes the mass-energy density and momentum density at that point in spacetime, specifically

$$T^{\alpha\beta} = \frac{p^\alpha p^\beta}{E/c^2} \delta(\vec{x} - \vec{x}(t)) \tag{22}$$

for a point particle located at $\vec{x}(t)$ whose momentum-energy 4-vector is p^α. (The spatial coordinates $\alpha = 1, 2, 3$ of p^α are momentum and its time coordinate $\alpha = 4$ is mass-energy E/c^2 times c. To be precise, the 4-momentum p^μ of a particle with velocity $\vec{v}$ and rest mass m_0 is

$$p^\mu = \frac{m_0}{\sqrt{1 - |\vec{v}|^2/c^2}} (\vec{v}, c) \tag{23}$$

where the common factor outside of the 4-vector in parentheses is the "(non-rest) mass." This expression makes it clear that in Lorentz-invariant laws of motion no body can ever exceed the speed c of light. Note that $T_{\alpha\beta}$ is zero in empty space free of electromagnetic fields. We are assuming the point particle has no multipole moments, nor any angular momentum. δ is an appropriate kind of spatial Dirac delta function.

These field equations, which are a second degree nonlinear system of 10 partial differential equations (although the derivatives are hidden in the notation!) describe *how matter (mass-energy) affects the metric of spacetime.*

The second equation that one needs describes the *laws of motion of matter* in such spacetime; specifically (assuming the matter is unaffected by nongravitational forces) matter moves along "geodesics" of the manifold, i.e. those curves $x^\mu(\lambda)$ which obey the geodesic[20] equation

$$\frac{d^2 x^\alpha}{d\lambda^2} + \Gamma^\alpha_{\mu\nu} \frac{dx^\mu}{d\lambda} \frac{dx^\nu}{d\lambda} = 0. \tag{24}$$

Strictly speaking, these laws of motion may not be needed, since Einstein and Infeld [ei49] [ke59] have shown that the vacuum field equations alone (plus some smoothness assumptions) suffice to force singularities of the metric (such as point masses) to move along geodesics.

Finally, gauge freedom allows us to specify 4 of the 10 entries (due to symmetry of $g_{\alpha\beta}$, there were 10 and not 16 free parameters in it) of the metric. Four gauge conditions which are often imposed are the "harmonic gauge"

$$g^{\mu\nu} \Gamma^\alpha_{\mu\nu} = 0. \tag{25}$$

4.5 Why we haven't tried to simulate full GR

There are many difficulties involved in the numerical simulation of general relativity. First, we no longer have ODEs to simulate, as with Newton's laws, but PDEs. Since these PDEs are inherently nonlinear, there is presumably no way to remove the fields, similarly to the use of Lienard-Wiechert retarded potentials in Maxwell's linear electrodynamics, and only consider the particles. (The fact that, according to Einstein-Infeld [ei49], in GR we may remove the *particles,* is little consolation.) So it seems that one requires $10\infty^3$ numbers to specify the metric at any "time." Singularities of the metric seem to be virtually generic and so the simulator cannot generally hope to avoid them. Penrose's "cosmic censorship hypothesis" which would force any singularities and possible violations of "causality" to be forever hidden behind "event horizons," and thus invisible to any outside observer, but this hypothesis – this would seem to be rather important for simulators of GR! – has never been proven. Even if Penrose's hypothesis is true, the fact that light from an external source can wind unboundedly many times around a Schwarzschild black hole be-

[20] Note: these are *not* the same thing as "shortest paths," due to the fact that the "metric" is not a metric. Rather, they are curves which locally are like lines in flat $(3+1)$-space, and "keep going in the same direction."

fore escaping ([mi73] page 674), for example, leads one to suspect that very complicated behavior is possible[21].

A method of splitting Einstein's field equations (EQ 21) into timelike and spacelike parts was devised by Arnowitt, Deser, and Misner [ar62]. In this formulation of GR, the manifold is sliced into hypersurfaces containing only spacelike distances, and the time-evolution from one such surface to the next is determined by the field equations. In the event that "causality" holds, data on one such surface will suffice to determine the metric at all future times. In the event that one could then prove a theorem that a sufficiently wide class of PDEs are simulatable with polynomial slowdown (the initial data would have to be provided in the form of an oracle) – where simulation has to continue despite the appearance of singularities and despite possibly being forced to change coordinate systems on the manifold! – then one could prove the Extended Church Thesis in general relativity. We are not prepared to undertake this task in the present paper[22] and do not know whether the Extended Church Thesis is true under GR.

[21] Still, it doesn't seem possible to kill the Extended Church Thesis, at least by naive use of relativistic effects. For example, suppose you program your computer to solve a problem and report its answer to you by laser, then you jump into a black hole. You argue that from the viewpoint of the computer, you will take infinite time to fall into the black hole, whereas from your own viewpoint, you will take finite time. Thus you hope to get the answer to the halting problem by laser in a finite amount of your time. (Admittedly, you will die.) But actually, if the laser signal is emitted too late, it can never catch up to you before you have been crushed into a point ([mi73] page 835-6). You could also go into near orbit exponentially close to a black hole and emerge much later to hear your computer's answer. Unfortunately, this would only work if your physical dimensions were exponentially small, and also the "photon orbit" is located at $r = 3M$ which is well above the event horizon at $r = 2M$ so that no circular orbit this close to the hole actually exists without artificial forcing. You could get exponentially large computational speedup by waiting out the computation while traveling in a very fast rocket, thus taking advantage of special relativistic time dilation; but this seems to require an exponentially large expenditure of energy. Generally if two observers's energy supplies and the reciprocal of their linear dimensions are polynomially bounded, then superpolynomial time dilation between them is impossible.

[22] In general relativity with a "cosmical constant," it is in principle possible for the global structure of space to be essentially time invariant and of constant negative curvature, i.e. "hyperbolic nonEuclidean geometry." There is a natural embedding of an infinite tree, each of whose nodes has valence 3 and each of whose edges had fixed length, in the hyperbolic plane, because, e.g. the area and perimeter of a circle of radius r grows exponentially, not polynomially, with r in such a geometry. (No nice embedding of such a tree is possible in Euclidean geometry.) In the event that our universe actually was of this sort and *static* (which it isn't), and infinite, and if there were an infinite number of alien civilizations scattered throughout the cosmos one in each ball of radius 1000 light years, then we could send a radio message telling any civilization within 1000 light years to build a computer and send its own similar radio message. The result would be an exponential number of computers being built, and communicating by radio, after time t, resulting in a fairly clear violation of the Extended Church Thesis. However, our universe is expanding, not static. For an analysis of whether GR permits an infinite Turing machine, see [sm03].

4.6 *Linearized general relativity*

These difficulties are almost all avoided by a theory of gravity intermediate between GR and Newton, namely "linearized GR," also called "weak field GR." It involves a 4×4 matrix-valued gravitational field. This theory is what happens to GR when one formally pretends that the gravitational constant G is infinitesimal, so that all terms of order G^2 may be dropped[23].

In linearized GR, the metric is $g_{\alpha\beta} = \eta_{\alpha\beta} + h_{\alpha\beta}$ where h is a small (order G) perturbation to the flat-spacetime metric $\eta_{\alpha\beta}$. The field equations are then linear in $h_{\alpha\beta}$ since all nonlinear terms are dropped, and under the adoption of harmonic gauge, it turns out that they reduce to a matrix D'Alembertian equation with solution

$$h_{\alpha\beta}(\vec{x}, t) = 4G \sum_{i=1}^{N} \frac{S_{\alpha\beta}\left(\vec{x}'_i, t - \mathrm{dist}(\vec{x}, \vec{x}'_i)/c\right)}{\mathrm{dist}(\vec{x}, \vec{x}'_i)}. \tag{26}$$

Here the quantity in parenthesis to the right of $S_{\alpha\beta}$ is its argument, i.e. the 4-vector of coordinates at which we evaluate $S_{\alpha\beta}$, and $\vec{x}'_i$ refers to the spatial position of particle i at a *retarded* moment, namely the moment t'_i when

$$\frac{\mathrm{dist}(\vec{x}, \vec{x}'_i)}{c} = t - t'_i, \text{ i.e. } (x^\alpha - x_i^{\alpha\prime})\eta_{\alpha\beta}(x^\beta - x_i^{\beta\prime}) = 0. \tag{27}$$

There is a unique such retarded moment since no particle can exceed c. Here

$$\begin{aligned} S_{\alpha\beta} &= \int (T_{\alpha\beta} - \frac{1}{2}\eta_{\alpha\beta}T^\nu_\nu)\mathrm{dspatial} \\ &= \frac{(m_0)_i}{\sqrt{-v_\delta v^\delta}}(v^\alpha v^\beta - \frac{1}{2}\eta_{\alpha\beta}v_\delta v^\delta) \end{aligned} \tag{28}$$

(where $v_\delta v^\delta = |\vec{v}|^2 - c^2$ and v^δ is the velocity 4-vector with components $(\vec{v}, c)$); the spatial integration in a neighborhood of $\vec{x}'_i$ merely serves to convert it into an energy-momentum tensor (but with a slightly altered definition) for a point particle instead of a energy-momentum *density* tensor. Since the particles are points and the T-tensor is a delta function, this integration is trivial. The $\mathrm{dist}(\vec{x}, \vec{y})$ function in (EQ 26) refers to spatial distance and might at first be thought to be a difficult thing to get ahold of, considering we are living on a curved manifold. However, using the *Euclidean* distance function $|\vec{x} - \vec{y}|$ is entirely satisfactory since it differs by

[23] Keeping all terms up to some power G^k, $k \geq 2$, is also possible; the resulting theories have been called "Post-Minkowskian expansions" [be56] [be60] [sc86] [ke59] [da82] [de97]. An alternative expansion in which the speed of light c is assumed to be near-infinite, so that only terms of order c^{-k} or larger are retained, is also possible; it is called "Post-Newtonian expansion" [ch65] [ed38] [ei38] [lo17] [da85] [sc85] [wi72]. The post-Minkowskian expansions have the advantage over post-Newtonian expansions that they are Lorentz-invariant. One may also linearize the Einstein-Maxwell equations of gravity and electromagnetism. It is not clear to me whether these expansion processes can be continued indefinitely (larger and larger k) or whether they must eventually break down (nor whether, even if they can be so continued, that they will converge). We will not be concerned with these issues in this paper.

terms of order G from the true distance, and thus introduces negligible errors of order G^2 into (EQ 26).

This observation also legitimizes our use of an absolute time coordinate t. The coordinates $\vec{x}, t$ really refer to spatial position and time in the flat space metric $\eta_{\alpha\beta}$ *before* the perturbation by $h_{\alpha\beta}$, which may be thought of as a matrix-valued gravitational field at each point $x^\mu = (\vec{x}, ct)$ of flat spacetime.

The solution (EQ 26) combined with the geodesic equations of motion (EQ 24) (and in which it is understood, when one is computing the motion of particle j, that the sum in (EQ 26) should not include the term $i = j$, since we are only interested in the field produced by the other $N-1$ particles) are a Newton-like system of almost-ODEs for solving the N-body problem in linearized general relativity – we say "almost" since they involve retarded times, i.e. the evolution of the system now depends not only upon the present state of the system, but also upon its state in the past.

The equations of motion (EQ 24) involve first derivatives of the h-field (inside the Christoffel symbol) namely, we may take (accurate to order G)

$$\Gamma^\mu_{\alpha\beta} = \frac{1}{2}\eta^{\nu\mu}\left(\frac{\partial h_{\beta\nu}}{\partial x^\alpha} + \frac{\partial h_{\alpha\nu}}{\partial x^\beta} - \frac{\partial h_{\beta\alpha}}{\partial x^\nu}\right). \qquad (29)$$

These derivatives are easily got in closed form from (EQ 26) and (EQ 27) and will depend on the retarded velocities of the bodies, but *not* on their accelerations, since these accelerations are $O(G)$ so that their effect is $O(G^2)$.

Carrying out the differentiation explicitly is made easier with the aid of the following formula for the differential of the retarded time

$$\mathrm{d}t' = \frac{|\vec{x'}_i - \vec{x}|}{\vec{v'}_i \cdot (\vec{x'}_i - \vec{x})}\left(\frac{(\vec{x'}_i - \vec{x})\cdot \mathrm{d}\vec{x}}{|\vec{x'}_i - \vec{x}|} - c\mathrm{d}t\right)/\left(1 - \frac{c|\vec{x'}_i - \vec{x}|}{\vec{v'}_i \cdot (\vec{x'}_i - \vec{x})}\right). \qquad (30)$$

Then, within errors of order $O(G^2)$, the components of $\partial h_{\alpha\beta}$ are given by

$$-4G\sum_{i=1}^{N}\frac{S_{\alpha\beta}(\vec{x}'_i, t')}{(t-t')^2 c}\mathrm{d}t'. \qquad (31)$$

To summarize: linearized GR obeys Newton-like, but retarded, equations of motion (EQ 24), where (EQ 29), (EQ 26), (EQ 27), (EQ 28) hold. The explicit differentiations called for in (EQ 29) may be carried out as in (EQ 30),(EQ 31).

4.7 Non-collision singularities are impossible in linearized GR

This is because the speeds of all bodies are bounded above by a constant, namely c. Hence, the position of any body as a function of time is Lipshitzian, which makes it obvious that as $t \to t*$, each body must approach a particular point of space, so that any singularity must be a collision.

Now, considering the results of Saari and Hulkower [sa81] that in the *Newtonian* N-body problem, any collision singularity, as it is approached, has velocity vectors for each particle which *have limiting orientations*, i.e. "infinite spins" on the approach to a collision are impossible[24], one might therefore be led to suspect that the same is true in linearized GR, so that *no* singularity of the motion, in linearized GR, could exhibit complex behavior.

Whether or not this is true is not relevant to the real world, since, while linearized GR may be an acceptable approximation to GR in some regimes (weak fields, small masses) it is certainly not a good approximation to GR when a singularity is approached. In particular, GR exhibits "black holes." If ever a body gets within a distance Gm/c^2 (half the "Schwarzschild radius") of a body of mass m, it is within an "event horizon" about that body[25]. Once within the event horizon, paths toward the body become timelike and, since nothing can exceed c, the two bodies will inevitably combine into a single black hole, indeed in a short finite time. The linearized theory is unaware of such "never get out again" behavior which would appear to it to violate conservation of energy. (In full GR, one realizes that the necessary energy loss is provided by gravitational radiation.) We wish to *modify* linearized GR to make its treatment of close encounters a little more realistic, or at least, to confess its inadequacy.

4.8 Modified linearized GR

Is the same as linearized GR, except that if ever two bodies approach each other more closely than the sum of their Schwarzschild radii, then the simulator is allowed, at his discretion, to pretend that they instantly combine into a single body with summed momentum and energy and located at the center of mass of the original bodies[26].

[24] This remark proves that noncollision singularities are the *only* way in which the Newtonian N-body problem can exhibit infinitely complicated computer-like behavior in finite time.

[25] The event horizon of a Kerr-Newman black hole, essentially the most general possible stationary solution of the Einstein-Maxwell equations, has a radius which can range from 1/2 to 1 times the Schwarzschild radius, depending on the angular momentum and charge of the hole. Nonstatic solutions can have event horizons which behave in a complex manner, but it is a theorem of Hawking [ha72] that once two event horizons have merged into a connected component, they can never un-merge.

[26] The "at his discretion" was required to prevent a discontinuity and consequent problems for the simulator, when a body is very close to the critical radius, with deciding whether it ever got inside it. If the reader prefers, we can change the rules so that (1) if closest separation $> R$, they escape, (2) if $< R\cdot(1-\epsilon)$ they coalesce, and (3) if in between, it is at the simulator's discretion, where ϵ is polynomially small.

This is of course a crude approximation to what is actually going on, since (for example) the resulting black holes would probably have angular momentum and thus generate a Kerr and not a Schwarzschild metric, some mass-energy would have been lost to gravitational waves, and so on. Nevertheless, it is better than nothing.

4.9 The N-body problem in modified linearized GR is simulatable with polynomial slowdown by a Turing machine

The fact that the laws of motion in modified linearized GR obey delay-ODEs is no great obstacle since by binary search and interpolation in the stored past-time solution, we may make a efficient subroutine that accurately evaluates the retarded fields at any time, and with the aid of this subroutine, the simulator feels exactly as though he/she[27] is simulating a system of genuine ODEs.

There is a slight difficulty with the initial data, however; it must be specified not just at one particular time ("$t = 0$") but in fact also at all times during the past, far enough back into the past so that the outgoing h-waves, propagating at c, which originated before history began at time $t = -\tau$, will never be able to affect the future motion. We will assume that our Turing machine is provided with an oracle who will answer any such question about the state of the N bodies at time t where $-\tau \le t \le 0$, to as many decimals of accuracy as it desires.

Finally, to make our general-purpose ODE-simulation theorem 7 apply, we need bounds on the magnitudes of the kth time derivatives of the solution which grow at most like polynomial$(L, k)^k$ where L is the length of the input in bits. Because the bodies are always *separated* by at least their Schwarzschild radii, and because (consequently) mass-energies are always bounded above by a constant at all times, and because (consequently) these Schwarzschild radii never can get very small, and (also consequently) all body's speeds are bounded below c so that retarded distances are within a not-large factor of unretarded distances, and finally because we demand that all the initial masses and energies have to be described in *unary*... one easily verifies that we in fact *have* such bounds so long as no two bodies get closer than their Schwarzschild radius at any time during the simulation[28].

In the event that such a close approach does happen, and is detected (and since we may take $1/\epsilon$ at least polynomially large, for any desired polynomial, we will detect any such close approach except possibly for those that approach the Schwarzschild radius very closely – we can't decide whether inside or outside it – which doesn't matter since the combining-holes decision has been left to the simulator's discretion) we instantly combine the two colliding holes and restart the simulation from there. We conclude:

Theorem 8 *The N-body problem in modified linearized general relativity, as described above (in particular, essential assumptions include*

1. *initial rest masses and kinetic energies specified in unary, although directions and positions may be binary fixed point numbers*
2. *combining bodies within Schwarzschild radius separation into a single black hole at the simulator's discretion*
3. *the input includes a record of the past, i.e. at negative times,*

is simulatable by a Turing machine with at most polynomial slowdown.

The "oracle" who tells us about the past on demand, could in fact be replaced by a polynomially large amount of data stored in tables with high degree interpolation (of degree of order τ) between tabulated points, by an analysis exactly similar to the analysis in the proof of theorem 7, with the provision that the needed duration τ of the historical record of the past, before the onset of simulation, would then be an additional *argument* of this polynomial. Since the needed τ might be rather large, one could argue that perhaps the oracle really cannot be replaced by a small lookup table and is inherently a large lookup table. In any case, the "polynomial slowdown" claim in theorem 8 does not require any knowledge of the size of the oracle; it refers to the runtime of the simulator (a Turing machine) being bounded by a polynomial in T, L, N, where L is the number of bits in a description of the state of the N bodies at time $t = 0$ *only*, and T is the duration of simulated time that is selected.

5 Final remarks

5.1 Where I stand

It is not the intent of this paper to undermine Church's thesis and thus to render much of computer science (in particular, most work on "polynomial time") and theoretical physics irrelevant.

I am in fact rooting *for* Church's thesis to be *true* in whatever laws of physics actually hold in our universe.

I am, however, pointing out that

1. Church's thesis need not merely be stated and then taken on "faith" or justified merely by heuristic arguments such as those found in [tu36][29] – An attempt can and should be made to prove or disprove it. Not only is this task of great philosophical importance, it will also shed light on physics by elucidating which parts of physics are simulable, and which parts are not, and how to go about simulating the simulable parts.
2. Care is necessary! There may be more "gotchas" such as theorem 4 lurking.

[27] from now on: "he!"

[28] Basically, the essential fact is that the kth derivitives of Newton-like potentials $1/x$, are $\pm k!/x^{k+1}$, so that if $|x|$ is bounded below (by a Schwarzchild radius) then these values grow like $k^{O(k)}$ as required.

[29] Turing's two main commonsense arguments for the universal power of a Turing machine involved (1) the idea that any computor must have only a *finite* number of possible "states of mind" or possible fundamental "symbols" he can write, since otherwise he might become "confused," and (2) the idea that one computational step (reading and writing a symbol, and changing mental state) must take one time unit. Both of these ideas are undermined, especially the second, in theorem 4.

3. There seems to be a distinction between *unsimulability* and *buildability of super-Turing computers*, see §2.3.

5.2 Movie

With the aid of Henry Cejtin, computer graphics, and a numerical ODE solver, I produced a short movie showing some of the essential elements of Gerver's proof, and my modification of it needed for theorem 4, in action.

5.3 Infinite energy source

One criticism of this paper is that an essential ingredient of theorem 4 was the use of an infinite source of energy, namely, the possibility of moving two point masses arbitrarily close together.

If that is your trick (the criticizers continue) why not just build a mechanically driven Turing machine and power it with an infinite energy source so that it will go faster and faster and do an infinite amount of computation in finite time? One suggested energy source was a pair of orbiting classical point charges, which will radiate infinite energy in finite time.

To reply to this, I agree with the first paragraph, but the suggestion of the second paragraph I find aesthetically displeasing, since it requires postulating laws of physics in which infinitely strong rigid objects (such as gears and driveshafts) which could withstand infinitely high energy fluxes without melting, were available. I prefer to make as few assumptions about the laws of physics as possible...

Naturally, it is easy to construct made-up laws of physics in which one can easily obtain immense computational power, but such claims are only of interest if the laws of physics seem realistic. One might argue that the fact the Newton's laws can provide an infinite energy source, makes them unrealistic. True – and this illustrates why it is important and valuable to investigate Church's thesis: attempts to refute it tend to focus on, and throw into sharp relief, the unrealistic or inconsistent features of a physical theory.

5.4 Water world

We will now mention a different set of made-up physical laws in which Church's thesis is false, namely "water world." This world consists of three ingredients, namely (1) "water," an infinitely subdividable continuous fluid which supports pressure waves obeying the wave equation, (2) "steel" which is infinitely strong, rigid, continuous stuff which perfectly reflects pressure waves impinging on it from outside, and (3) you can cut up the steel with plane cuts by use of a "laser" and you can weld the steel using "glue."

By making N plane cuts and welds, we can create a set of (nonconvex) polyhedral steel obstacles in a pond of water. Then to determine the length of the shortest obstacle-avoiding path from A to B, "flick your finger" (create a delta function impulse) at A and measure the time before B "hears anything." This calculation took only $O(N)$ "operations" plus $O(1)$ time. However, Canny and Reif [ca87] have shown that deciding whether the shortest obstacle-avoiding A-to-B path is shorter than any given rational number "L" is NP-complete. Thus the extended Church's thesis is false in water world, or else P=NP.

The unrealistic aspect of water world's laws of physics, which permitted us to derive this, was perhaps not so much the huge computational parallelism implied by the fact that each "water molecule" is doing some "data processing," but rather the fact that Canny and Reif's NP-completeness proof required the use of polyhedra with some exponentially small edges. In the real world, atoms are *not* exponentially small and the shortest path would only be determined to an accuracy no better than the size of an atom. The determination of L to within a fixed error ϵ, in compute time polynomial in (L/ϵ), is *not* an NP-hard problem, indeed it is easy.

5.5 Real numbers

My critics also argue that the key feature of Newton's laws which made them unsimulable was that they involved real numbers, and it is somehow unfair to compare Turing machines, which can only input a finite number of bits in unit time, with laws of physics with real numbers. Real numbers make these modern-day Kroneckers very uncomfortable.

Well, I disagree! Face it, physics *does* involve real numbers, and real numbers perhaps *can* be used to do things Turing machines cannot! But if you think my entire paper is resting on the preceding sentence, you are very mistaken; the truth is deeper. For example, the laws of physics in §4 also involve real numbers, but Church's thesis, indeed even its extension, is, I argue, true in those laws of physics. (Better work this one out before complaining to me!)

What matters is *not* that real numbers are involved. As we've shown, this by itself need not prevent a Turing machine from simulating physics to superb accuracy (e.g. exponentially more accurate than any physical length scale) with only polynomial slowdown. There is nothing preventing a Turing machine from having real number inputs. What matters is whether the laws of physics permit the *abuse* of real numbers. Or, stated differently, whether the computational power of real numbers is physically *accessible*. Thus a very important question is whether the laws of physics permit one to build an "infinite amplifier with finite delay[30]." The fact that this is possible in Newton's laws and impossible in the "repaired" laws is certainly an important reason, and, some might argue, possibly the only reason, for our results.

[30] An important related question is whether the geometric time-mean of the Liapunov exponent (ℓ, where close phase-space trajectories can diverge proportionally to $\exp(\ell t)$) of the physical system is necessarily polynomially bounded. If it is not, then certainly Church's extended thesis cannot hold.

Nevertheless, some people have also demanded that in any comparison between physics and a Turing machine, that my initial data for physics had to be integers, or perhaps rationals. I regard it as unlikely (although I have not proven this) that the Newtonian N-body problem will ever do anything uncomputable in finite time (or even in infinite time) if all initial data is rational. Hence, my critics would argue, theorem 4 shows nothing.

I reply that this criticism is in fact ludicrous. The physical system *does not know* that you have granted it the "boon" that all initial data were integers. As far as the physical system is concerned, you had to input an infinite number of zeros after the decimal point. The "boon" exists only in the finite mind of my critic. Also, why are rationals OK for my critics? Surely even my honorable critic cannot do an infinite long division in finite time? This point makes it apparent that what my critic really wants is not so much real numbers with terminating decimal expansions, but in fact he wants real numbers that are *finitely describable* – real numbers that he feels he can understand. Perhaps the critic feels he can understand numbers like $3+2^{1/4}$ which are constructible with ruler and compass. Or perhaps he has advanced beyond the ancient Greeks and will accept numbers such as $2^{1/3}$, or even γ and e. And once we have advanced to this point, we see that what is really wanted is the field of "computable real numbers" – the real numbers with finite algorithmic descriptions – and that is exactly what was accomplished by the argument of §2.2!

Now, one of my most fire-breathing critics was undeterred by this logical setback and demanded that in any comparison between physics and a Turing machine in which one purports to refute Church's thesis, the initial data for physics ought to come with error bounds (which are also part of the "input") and the physical system would then proceed to do the *same* uncomputable thing *no matter what* the actual initial data was, so long as it was within the error bounds. In other words, he wanted the physical machine to be "self correcting." Imposing these demands certainly kills my theorem 4, or anything like it, provided it is the case (which I suspect, but have not proven) that the set of initial data for the N-body problem which will evolve to a singularity, is of measure zero.

But in fact, I now argue (heuristically) that, if those are the demands, then Church's thesis is presumably true... but, the argument should also make it clear that such demands were inadmissible in the first place. Because: in any laws of physics (such as Newton's laws) which are time-reversible, or, more generally, which preserve phase-space volume, self-correcting behavior is *impossible*! This is because any such behavior has to involve "attraction" and thus "dissipation" or "friction" – phenomena incompatible with the assumptions. (Incidentally, Brockett's ODEs, cf. footnote 6, which simulate a Turing machine, *are* self-correcting and *do* involve dissipative behavior – but they do not correspond to fundamental-sounding physics.)

Now, to dramatically illustrate a final facet of the real number issue, consider the following satirical comparison of a "physical" computer with a Turing machine, intended to make the present author look silly. The Turing machine T is given a tape on which is pre-written an infinite number of '1's and '0's. It is clearly impossible for T to always determine, in finite time, whether there is a '1' on the tape. Meanwhile, the physical system P is two unit masses at (0,0) and (0,1), both with velocity vectors having $v_x = 0$, but body #1 has $v_y = 0$ while body #2 has $v_y = R$, where R is the real number .0010100 (or whatever) corresponding to what is written on the tape. A collision of the bodies occurs in finite time if and only if $R = 0$. This "comparison" of two "computers" demonstrates Smith's essential silliness in the starkest possible manner... right?

Well, no. First, realize that, binary collisions are "regularizable" by elastic bounces, that is, if the motion is continued after the collision in this way, then the {final state} will be a continuous function of the {initial data}. In other words, binary collisions, are, in some sense, not "really" singularities at all![31] Thus as the impact parameter is decreased, the bodies will swing by each other in an ellipse or hyperbola that gets more and more like a 180° bounce-back. Indeed, if the motion of the bodies were observed through "eyes" which were not infinitely acute, there would be no way to tell whether a collision had occurred. Meanwhile, the physical computer in the scenario of §2.2 makes two easily distinguished choices involving a cat living or dying, and does so with the aid of a complicated process which certainly seems much more like genuine "computation" and in which each successive swingby is known to go either to the left or right after examining only a finite number more bits of the initial data – there isn't any nonsense about having to wait for a huge number more bits before you know what to do at any particular swingby.

5.6 *Open questions*

What if quantum mechanics is put in? The recent advent of "quantum computers [sh97]," "quantum error correcting codes," and "quantum fault tolerance techniques [pr96]" suggests that Church's extended thesis (appropriately redefined) is *false*, at least with a naive interpretation of quantum mechanics and some simplistic models of "decoherence [zu91]." The un-extended Church thesis seems open. (Later note: I have established the validity of Church's thesis in quantum mechanics in a different paper [sm99]. Amazingly, the same n-body problem which, under Newton's laws, we have seen falsifies Church's thesis, is in fact simulable in quantum mechanics, so that, surprisingly, nonrelativistic quantum mechanics is actually *easier* than nonrelativistic classical mechanics.)

[31] Triple collisions, however, are not regularizable, although due to the "limiting orientations theorem" of [sa81], the term "elastic bounce" can still be given an unambiguous meaning for k-way collisions.

At the other end of the spectrum, we have the pinnacle of classical (meaning deterministic and continuous) field theories: General Relativity, as embodied by the Einstein-Maxwell equations.

There are plausibility arguments that this is simulatable, although certainly I have no proof. (Some discussion is in §4.5.) For example, in GR, point masses and point charges *cannot* serve as infinite energy sources, nor can information propagate faster than c. I also point out, although this is not commonly appreciated, that such phenomena as radiation reaction, which were thought to reveal inconsistencies or limitations in classical field theories, in fact seem to be handled entirely self-consistently by GR [sc85].

On the other hand, GR exhibits phenomena which cast doubt on its simulability, such as the fact that black holes can form out of nothing (e.g. from colliding gravitational waves in pure vacuum) and the fact that light can spiral an unboundedly large number of times around a black hole before finally escaping (or being swallowed).

6 Appendix on notation

We use the acronyms ODE=ordinary differential equation, PDE=partial differential equation, RK=Runge-Kutta (see EQ 8, 9), GR=general relativity, iff="if and only if." When we say x is "bounded above" y we mean that a constant $c > 0$ exists so that $x \geq c + y$, uniformly within some set of parameters that x and y are functions of and that c is independent of – this set should be clear from the context. Similarly for "bounded below" and "bounded within (an interval);" such phrases are of course stronger than the usual "greater than." We use O and o and "order" asymptotic notation: $f(x) = O(g(x))$ means there exists a constant $c > 0$ so that $0 \leq f(x) < cg(x)$ for all sufficiently large (or small; which limit is intended should be apparent from the context) x. $f(x) = o(g(x))$, where $f(x), g(x) > 0$, means that in the appropriate x-limit, $f(x)/g(x) \to 0$. If $f(x)$ and $g(x)$ are "of the same order" that means $f(x) = O(g(x))$ and $g(x) = O(f(x))$. In order to avoid confusion regarding the word "order" (two uses already in this sentence!) we have used the word "degree" when speaking either of the number of differentiations in an equation or of the degree of accuracy of a timestepping scheme ("degree k" means having error of order δ^{k+1}). A quantity is "polynomially large" if it grows more slowly than some polynomial as certain quantities (which should be clear from the context...) tend to ∞. On the other hand, if a quantity grows like exp of a polynomially large quantity, it is "exponentially large." You are "polynomially small" if your reciprocal is polynomially large, and similarly for "exponentially small." One must be careful when using these terms. For example, "to represent a fixed-point number to exponential accuracy requires polynomially many bits" is a true statement.

References

[ar62] R. Arnowitt, S. Deser, and C.W. Misner. The dynamics of general relativity. In L. Witten, editor, *General relativity, an introduction to current research*, pages 227–265. Wiley, 1962.

[be56] B. Bertotti. On gravitational motion. *Il Nuovo Cimento*, 4(4):898–906, October 1956.

[be60] B.Bertotti & J.Plebanskii: Theory of gravitational perturbations in the fast motion approximation. em Annals Physics NY 11 (1960) 169-200.

[br89] R.W. Brockett. Smooth dynamical systems which realize arithmetical and logical operations. In H. Nijmeijer and J.M. Schumacher, editors, *Three decades of mathematical systems theory*, number 135 in Springer-Verlag lecture notes in control and information science, pages 19–30. Springer-Verlag, 1989.

[bu63] J.C. Butcher. Coefficients for the study of Runge-Kutta integration processes. *J. Australian Math. Soc.*, 3:185–201, 1963.

[bu64] J.C. Butcher. Implicit Runge-Kutta processes. *Math. of Computation*, 18:50–64, 1964.

[bu87] J.C. Butcher. *The numerical analysis of ordinary differential equations, Runge Kutta and general linear methods.* J. Wiley, 1987.

[ca87] J. Canny and J.H. Reif. New lower bounds for robot motion planning problems. In *28th IEEE Symposium on Foundations of Computer Science*, pages 49–60. IEEE, 1987.

[ch65] S. Chandrasekhar, V. Nutku, and Paul F. Esposito. Post-newtonian equations of hydrodynamics in general relativity. *Astrophysical Journal.* I: AJ 142,4 (November 1965) 1488-1512, II: AJ 158,1 (October 1969) 55-79, III: AJ 160,1 (April 1970) 153-179.

[co89] George W. Collins. *The foundations of celestial mechanics.* Pachart publishing, Tuscon AZ, 1989.

[da56] G. Dahlquist. Convergence and stability in the numerical integration of ordinary differential equations. *Math. Scand.*, 4:33–53, 1956.

[da59] G. Dahlquist. Stability and error bounds in the integration of ordinary differential equations. *Trans. Royal Inst. Technol. Stockholm*, (130), 1956.

[da82] Thibault Damour. Gravitational radiation and the motion of compact bodies. In Nathalie Deruelle and Tsvi Piran, editors, *Gravitational Radiation*, pages 59–144. North-Holland, 1982.

[da85] T. Damour and Gerhard Schäfer. Lagrangians for n point masses and the second postnewtonian approximation of general relativity. *General Relativity and Gravitation*, 17(9):879–905, 1985.

[de97] Steven Detweiler & Lee H., Brown Jr.: Post-Minkowski expansion of general relativity, Phys. Rev. D (3) 56,2 (1997) 826-841.

[ed38] A. Eddington and G. Clark. The problem of n bodies in general relativity theory. *Proc. Royal Soc.*, 166:465–475, 1938.

[ei49] A. Einstein and L. Infeld. On the motion of particles in general relativity. *Canad. J. Math.*, 1:209–241, 1949.

[ei38] A. Einstein, L. Infeld, and B. Hoffman. The gravitational equations and the problem of motion. I: Annals Math. (Princeton) 39 (1938) 65–100, II: Annals Math. (Princeton) 41,2 (1940) 455-464.

[fl16] L. Flamm. Beitrage zur Einsteinschen Gravitationstheorie. *Zeitschrift f. Physik*, 17:448–454, 1916.

[ge71] C.W. Gear. Algorithm 407, difsub, for solution of ordinary differential equations. *Communications Assoc. Comput. Machinery*, 14:185–190, 1971.

[ge91] Joseph L. Gerver. The existence of pseudocollisions in the plane. *J. Differential Equations*, 89:1–68, 1991.

[ha76] G. Hall, J.M. Watt, and (many contributors). *Modern Numerical methods for Ordinary Differential Equations*. Clarendon Press Oxford, 1976.

[ha74] George Hall. Stability analysis of predictor-corrector algorithms of Adams type. *SIAM J. Numer. Anal.*, 11(3):494–505, June 1974.

[ha72] Stephen R. Hawking. Black holes in general relativity. *Commun. Math. Phys.*, 25:152–176, 1972.

[he62] P. Henrici. *Discrete Variable methods in ordinary differential equations*. Wiley, 1962.

[in51] L. Infeld and A. Scheidegger. Radiation and gravitational equations of motion. *Canad. J. Math*, 3:195–207, 1951.

[in49] L. Infeld and A. Schild. On the motion of test particles in general relativity. *Reviews in Modern Physics*, 21:408–413, 1949.

[je81] Rolf Jeltsch and Olavi Nevanlinna. Stability and accuracy of time discretizations for the initial value problem. *Numerische Math.* I: 37 (1981) 61-91, II: 40 (1982) 245-296.

[ke59] R. P. Kerr. The Lorentz-covariant approximation method in general relativity. *Il Nuovo Cimento.* I: 13,3 (August 1959) 469-491; II (second order): 13,3 (August 1959) 492-502; III (Einstein-Maxwell field): 13,4 (August 1959) 673-689; IV (quasi-static approx.): 16,1 (April 1960) 26-62.

[kl93] D.Kleppner: The gem of general relativity, Physics Today 46 (April 1993) 9-11.

[ko83] Ker-I Ko. On the computational complexity of ordinary differential equations. *Information and Control*, 58(1-3):157–194, 1983.

[kr60] Martin D. Kruskal. Maximal extension of Schwarzschild metric. *Physical review*, 119:1743–1745, 1960.

[lo17] H.A. Lorentz and J. Droste. (1917 paper). In *Collected papers of H.A.Lorentz (volume 5)*, pages 330–, The Hague, 1937. Nijhoff.

[ma93] Yuri V. Matiyasevich. *Hilbert's Tenth Problem.* MIT Press, Cambridge MA, 1993.

[mi67] Marvin L. Minsky. *Computation: finite and infinite machines.* Prentice Hall, 1967.

[mi73] C.W. Misner, K.S. Thorne, and J.A. Wheeler. *Gravitation.* Freeman, 1973.

[mo90] Christopher Moore. Unpredictability and undecidability in dynamical systems. *Physical Review Letters*, 64(20):2354–2357, May 1990.

[pa98] Christos H. Papadimitriou, Kenneth Steiglitz: Combinatorial optimization: algorithms and complexity Dover Publications 1998.

[pr96] John Preskill. Reliable quantum computers. *Talk at ITP Conf. on Quantum Coherence and Decoherence*, 15–18 Dec. 1996; Quant-ph/9705031

[sa73] D.G. Saari. Collisions are of first category. *Proc. Amer. Math. Soc.*, 47:442–445, 1975.

[sa81] D.G. Saari and N.D. Hulkower. On the manifolds of total collapse... for the n-body problem. *J. Differential Equations*, 41:27–43, 1981.

[sc85] Gerhard Schäfer. The gravitational quadropole radiation-reaction force and the canonical formalism of ADM. *Annals of Physics (NY)*, 161:81–, 1985.

[sc86] Gerhard Schäfer. The ADM hamiltonian at the postlinear approximation. *General Relativity and Gravitation*, 18(3):255–270, 1986.

[sh97] Peter W. Shor. Polynomial time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM J. Computing*, 26: 1484–1509, 1997.

[sm99] Warren D. Smith. Church's thesis meets quantum mechanics, `http://math.temple.edu/ wds/homepage/works.html` #49.

[sm03] Warren D. Smith. Turing machine engineering and Immortality. `http://math.temple.edu/ wds/homepage/works.html` #74.

[tu36] Alan M. Turing. On computational numbers with an application to the entscheidungsproblem. *Proc London Math Soc.*, 42:230–265, 1936. correction: PLMS 43 (1937) 544-546.

[we72] Sheldon Weinberg. *Gravitation and Cosmology*. Wiley, 1972.

[wi72] C.M. Will and K. Nordtvedt Jr. Conservation laws and preferred frames in general relativity. *Astrophys. J.*, 177:757–774, 1972.

[xi92] Z. Xia. The existence of noncollision singularities in the n-body problem. *Annals Math.*, 135(3):411–468, 1992.

[zu91] Wojciech H. Zurek. Decoherence and the transition from quantum to classical. *Physics Today*, pages 36–44, October 1991.