distinguish between "good singularities," with which we can peacefully coexist, and "bad singularities," with which even detente is impossible. He thought that all singularities must be excluded from a complete scientific theory (see chapter 1).

33. Let $\psi: \Sigma \to M$ denote the imbedding of Σ as a Cauchy surface of M, g_{ab} and let $\tilde{\Psi}: M \to \tilde{M}$ and $\tilde{\tilde{\Psi}}: M \to \tilde{\tilde{M}}$ denote the isometric imbeddings of M, g_{ab} into $\tilde{M}, \tilde{g}_{ab}$ and $\tilde{\tilde{M}}, \tilde{\tilde{g}}_{ab}$ respectively. Chruściel and Isenberg (1993) note that if the Cauchy surface $\psi(\Sigma)$ has a privileged status, then one might not want to count $\tilde{M}, \tilde{g}_{ab}$ and $\tilde{\tilde{M}}, \tilde{\tilde{g}}_{ab}$ as equivalent under an isometry $\varphi: \tilde{M} \to \tilde{\tilde{M}}$ if φ *moves* Σ in the sense that $\varphi \circ \tilde{\Psi} \circ \psi \neq \tilde{\tilde{\Psi}} \circ \psi$.

34. See Wald (1984a) for a discussion of the positive mass proof. Negative mass Schwarzschild spacetime (see note 14) has negative ADM mass and a naked singularity. It is not a counterexample to the positive mass theorem, which requires that the initial hypersurface be singularity free. Nor, as noted above, would this example be regarded as a violation of the form of cosmic censorship that excludes only those naked singularities that develop from regular initial data, since the singularity has been present for all times.

35. See chapter 6 for a discussion of the nature of physical laws.

36. A determinism maximal spacetime is also known in the literature as a *hole-free* spacetime. Being hole free does not entail being globally hyperbolic (think of Gödel spacetime). Nor does the implication go in the opposite direction (think of a truncated version of Minkowski spacetime with all of the points such that $t \geq 1997$ removed). However, a globally hyperbolic and inextendible spacetime is necessarily hole free.

37. See chapter 2 for a discussion of this matter.

4

Supertasks

4.1 Introduction

Is it possible to perform a supertask, that is, to carry out an infinite number of operations in a finite span of time? In one sense the answer is obviously yes since, for example, an ordinary walk from point a to point b involves crossing an infinite number of finite (but rapidly shrinking) spatial intervals in a finite time. Providing a criterion to separate such uninteresting supertasks from the more interesting but controversial forms is in itself no easy task,[1] but there is no difficulty in providing exemplars of what philosophers have in mind by the latter. There is, for instance, the Thomson lamp (Thomson 1954–55). At $t = 0$ the lamp is on. Between $t = 0$ and $t = 1/2$ the switch at the base of the lamp is pressed, turning the lamp off. Between $t = 1/2$ and $t = 3/4$ the switch is pressed again, turning the lamp on. Etcetera. The upshot is that an infinite number of presses are completed by $t = 1$. Then there is the super π machine. Between $t = 0$ and $t = 1/2$ it prints the first digit of the decimal expansion of π. Between $t = 1/2$ and $t = 3/4$ it prints the second digit. Etcetera. The result is that the complete expansion has been printed at $t = 1$. More interestingly from the point of view of mathematical knowledge there is the Plato machine which checks some unresolved existential conjecture of number theory for '1' during the first half-second, for '2' during the next quarter-second. Etcetera. The result is that the truth value of the conjecture is determined at the end of one second.

Thomson thought that such devices are logically or conceptually impossible. The operation of the Thomson lamp (a misnomer if Thomson were correct) entails that (i) for any t such that $0 < t < 1$, if the lamp is off at t, then there is a t' such that $t < t' < 1$ and the lamp is on at t', and (ii) for any t such that $0 < t < 1$, if the lamp is on at t, then there is a t' such that $t < t' < 1$ and the lamp is off at t'. Thomson thought that it followed from (i) that the lamp cannot be off at $t = 1$ and from (ii) that the lamp cannot be on at $t = 1$, a contradiction since it is assumed that the lamp must be in one or the other of these states at any instant. The fallaciousness of the argument was pointed out by Benacerraf (1962).

Others have held that though conceptually possible such devices are physically impossible. Benacerraf and Putnam (1983, p. 20), for example, seem to have thought that these devices are kinematically impossible due to the fact that relativity theory sets c (the velocity of light) as the limit with which the parts of the device can move. Again, however, the impossibility is not as obvious as claimed. A demonstration is needed to rule out as a kinematic possibility that the operation of the device is arranged so that with each successive step the distance the parts have to move (as in an ordinary stroll from a to b) shrinks sufficiently fast that the bound c is never violated. Of course, even if the device can be shown to pass muster at the kinematic level, it may still fail to satisfy necessary conditions for a dynamically possible process (see Grünbaum 1968, 1969 for a discussion).

I have nothing new to add to this discussion here.[2] Rather, my focus will be on the ways that the relativistic nature of spacetime can be exploited so as to finesse the accomplishment of a supertask. Very crudely, the strategy is to use a division of labor. One observer has available to her an infinite amount of proper time, thus allowing her to carry out an infinite task in an unremarkable way. For example, she may check an unresolved conjecture in number theory by checking it for '1' on day one, for '2' on day two, etc., ad infinitum. (Or if, as the numerals increase, she needs increasing amounts of time to complete the check, she can allow herself $f(n)$ days to check the conjecture for 'n', where $f(n)$ is any increasing function of n as long as $f(n) < \infty$ for all n.) A second observer, who uses up only a finite amount of his proper time, is so situated that his past light cone contains the entire world line of the first observer. The second observer thus has direct causal access to the infinite computation of the first observer, and in this way he obtains knowledge of the truth value of the conjecture in a finite amount of time. If this is genuinely possible in relativity theory, there is an irony involved. Prima facie relativity might have been thought to make supertasks more difficult if not impossible by imposing kinematic limitations on the workings of Thomson lamps, Plato machines, and the like. But on further analysis relativity theory seems to open up a royal road that leads to the functional equivalent of the accomplishment of a supertask. The rough sketch just given contains an unjustified optimism. We will see that relativistic spacetimes do provide opportunities for carrying out the functional equivalents of supertasks, but we will also see that they do so at a price. One approach is to set the supertask in a well-behaved spacetime (see section 4.2). Here a double price has to be paid; for the second observer who tries to take advantage of the infinite labor of the first observer must submit himself to unbounded forces that end his existence, and in any case he never observes the completion of the infinite labor at any definite time in his existence.

Alternatively, both of these difficulties can be overcome by exploiting spacetimes with unusual structures which I will dub *Malament–Hogarth spacetimes*. A large part of this chapter will be devoted to articulating the senses in which these spacetimes are physically problematic. As Hogarth has already shown, they are not globally hyperbolic (Lemma 4.1, section 4.3), so that they violate strong cosmic censorship. And they may also violate other requirements one would expect a physically realistic spacetime to fulfill (section 4.6). It will turn out that the failure of global hyperbolicity occurs in a way which necessarily defeats attempts to control disturbances to the signaling between the first and second observer from singularities and other sources. This signaling will prove to be problematic in other ways. It may demand that the second observer pursue his own mini-supertask in his neighborhood of spacetime, forfeiting the advantage that a Malament–Hogarth spacetime was supposed to offer (section 4.7). Again, the signaling will be associated with indefinite blueshifts (Lemma 4.2, section 4.5), so that the energy of the signals can be indefinitely amplified, threatening to destroy the second observer who receives them.

4.2 Pitowsky spacetimes

The first published attempt to make precise the vague ideas sketched in section 4.1 for using relativistic effects to finesse supertasks was that of Pitowsky (1990). His approach is encapsulated in the following definition.

DEFINITION 4.1
M, g_{ab} is a *Pitowsky spacetime* just in case there are future-directed timelike half-curves $\gamma_1, \gamma_2 \in M$ such that $\int_{\gamma_1} d\tau = \infty$, $\int_{\gamma_2} d\tau < \infty$, and $\gamma_1 \subset I^-(\gamma_2)$.

The blandest relativistic spacetime of all, Minkowski spacetime, is Pitowskian, as shown by Pitowsky's own example. (It seems a safe conjecture that this example can be generalized to show that any relativistic spacetime that possesses a timelike half-curve of infinite proper length is Pitowskian.) Choose an inertial coordinate system $(\mathbf{x}, t)$. Let γ_1 be the timelike half-geodesic $\mathbf{x}(t) = \text{constant}$, $0 \leq t < +\infty$. Choose γ_2 to be a timelike half-curve that spirals around γ_1 in such a way that it keeps γ_1 in its causal shadow and that its tangential speed is $u(t) = [1 - \exp(-2t)]^{1/2}$, $c = 1$. The proper time for γ_2 is $d\tau = \exp(-t)\, dt$, so that $\int_{\gamma_2} d\tau = 1$. Those familiar with the "twin paradox" may wish to take this example as the extreme case of the paradox with γ_2 as the ultimate traveling twin who ages biologically only a finite amount while his stay-behind twin ages an infinite amount. But admittedly this example does not conform to the usual twin paradox scenario where the twins hold a final reunion.

Pitowsky tells the following story about this example.

> While [the mathematician] M [γ_2] peacefully cruises in orbit, his graduate students examine Fermat's conjecture one case after the other. . . . When they grow old, or become professors, they transmit the holy task to their own disciples, and so on. If a counterexample to Fermat's conjecture is ever encountered, a message is sent to [M]. In this case M has a fraction of a second to hit the brakes and return home. If no message arrives, M

disintegrates with a smile, knowing that Fermat was right after all. (Pitowsky 1990, p. 83)

(The example is now somewhat dated since a proof of Fermat's last theorem has been offered. However, some lingering doubts may remain since the purported proof is over 200 pages. In any case, the punch of Pitowsky's story can be preserved by substituting for Fermat's theorem any unresolved conjecture of number theory with a prenex normal form consisting either of all universal quantifiers or else all existential quantifiers. Or the logician may wish to contemplate the problem of deciding for a formal system strong enough for arithmetic whether or not a given well-formed formula is a theorem.)

There are two things wrong with this story. The first concerns the notion that "M [γ_2] cruises peacefully in orbit." For ease of computation, assume that the mathematician γ_2 undergoes linear acceleration with $u(t)$ as before. The magnitude of acceleration $a(t) =: (A_b(t)A^b(t))^{1/2}$, where A^b is the four-vector acceleration, is $\exp(t)/[1 - \exp(-2t)]^{1/2}$, which blows up rapidly. (To stay within a linearly accelerating γ_2's causal shadow, γ_1 would also need to accelerate. But γ_1's acceleration can remain bounded. Indeed, γ_1 can undergo constant ("Born") acceleration, which guarantees that γ_1's velocity approaches the speed of light sufficiently slowly that its proper length is infinite.) Thus, any physically realistic embodiment of the mathematician will be quickly crushed by g-forces. The mathematician disintegrates with a grimace, perhaps before learning the truth about Fermat's conjecture. What is true in this example is true in general since any ultimate traveling twin in Minkowski spacetime must have unbounded acceleration. If the ultimate traveling twin moves rectilinearly and has an upper bound to his acceleration, then another traveler, Born accelerated at this upper bound, would achieve equal or greater velocity at each instant and therefore age less. But this Born accelerated traveler's world line has infinite proper length. Therefore the rectilinearly accelerated traveller must have no upper bound to his acceleration if he is to have finite total proper time. This result holds a fortiori for the general case of a traveler in curvilinear motion, for part of his acceleration will be transverse to the direction of motion, thus generating no velocity over time and no resultant clock slowing.

The second and conceptually more important difficulty with Pitowsky's story concerns the claim that the mathematician γ_2 can use the described procedure to gain sure knowledge of the truth value of Fermat's conjecture. If Fermat was wrong, γ_2 will eventually receive a signal from γ_1 announcing that a counterexample has been found, and at that moment γ_2 knows that Fermat was wrong. On the other hand, if Fermat was right, γ_2 never receives a signal from γ_1. But at no instant does γ_2 know whether the absence of a signal is because Fermat was right or because γ_1 has not yet arrived at a counterexample. Thus, at no definite moment in his existence does γ_2 know that Fermat was right. The fictitious mathematical sum of all of γ_2's stages

knows the truth of the matter. But this is cold comfort to the actual, non-mathematical γ_2. By way of analogy, if your world line γ is a timelike geodesic in Minkowski spacetime and you have drunk so deep from the fountain of youth that you live forever, then $I^-(\gamma)$ is the entirety of Minkowski spacetime. So the fictitious sum of every stage of you can have direct causal knowledge of every event in spacetime. But at no definite time does the actual you possess such global knowledge.

4.3 Malament–Hogarth spacetimes

Malament (1988) and Hogarth (1992) sought to solve the conceptual problem with Pitowsky's example by utilizing a different spacetime structure.

DEFINITION 4.2
M, g_{ab} is a *Malament–Hogarth spacetime* just in case there is a timelike half-curve $\gamma_1 \subset M$ and a point $p \in M$ such that $\int_{\gamma_1} d\tau = \infty$ and $\gamma_1 \subset I^-(p)$.

This definition contains no reference to a receiver γ_2. But if M, g_{ab} is a Malament–Hogarth (hereafter, *M–H*) spacetime, then there will be a future-directed timelike curve γ_2 from a point $q \in I^-(p)$ to p such that $\int_{\gamma_2(q,p)} d\tau < \infty$, where q can be chosen to lie in the causal future of the past endpoint of γ_1. Thus, if γ_1 proceeds as before to check Fermat's conjecture, γ_2 can know for sure at event p that if he has received no signal from γ_1 announcing a counterexample, then Fermat was right.

Such arrangements can also be used to "effectively decide" membership in a recursively enumerable but non-recursive set of integers.[3] To decide whether or not a given n is a member, γ_1 proceeds to enumerate the members of the set. By assumption, this can be done effectively. As each new member is generated, γ_1 checks to see whether it is equal to n. This too can be done effectively. γ_1 sends a signal to γ_2 just in case she gets a match. Consequently, γ_2 knows that n is a member precisely if he has received a signal by the time of the M–H event.

These scenarios cannot be carried out in Minkowski spacetime, as follows from

LEMMA 4.1 *An M–H spacetime is not globally hyperbolic.*

A formal proof of Lemma 4.1 was given by Hogarth (1992). A simple informal proof follows from the facts that a globally hyperbolic spacetime M, g_{ab} contains a Cauchy surface and that a spacetime with a Cauchy surface can be partitioned by a family of Cauchy surfaces. Suppose for purposes of contradiction that M, g_{ab} is both globally hyperbolic and contains an M–H point $p \in M$, i.e., that there is a future-directed timelike half-curve γ such that $\gamma \subset I^-(p)$ and $\int_\gamma d\tau = \infty$. Choose a Cauchy surface Σ through p, and extend

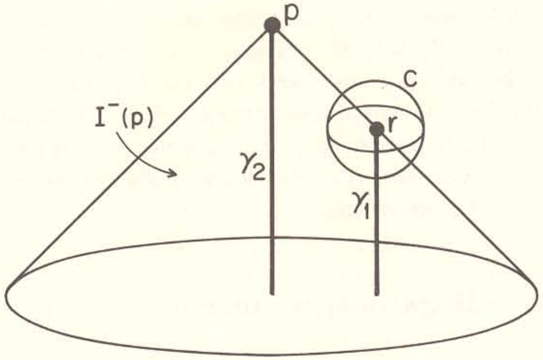

Fig. 4.1 A toy Malament–Hogarth spacetime

γ maximally in the past. This extended γ' is also contained in $I^-(p)$. Since γ' has no past or future endpoint, it must intersect Σ. But then since there is a timelike curve from the intersection point to p, Σ is not achronal and cannot, contrary to assumption, be a Cauchy surface.[4]

What of the problem in Pitowsky's original example that the receiver γ_2 has to undergo unbounded acceleration? In principle, both γ_1 and γ_2 can be timelike geodesics in at least some M–H spacetimes. The following toy example illustrates the point and also serves as a useful concrete example of an M–H spacetime. Start with Minkowski spacetime $\mathbb{R}^4$, η_{ab} and choose a scalar field Ω which has value 1 outside of a compact set C (see Fig. 4.1) and which goes rapidly to $+\infty$ as the point r is approached. The M–H spacetime is then M, g_{ab} where $M = \mathbb{R}^4 - r$ and $g_{ab} = \Omega^2\eta_{ab}$. Timelike geodesics of η_{ab} in general do not remain geodesics in g_{ab}, but Ω can be chosen so that γ_1 is a geodesic of g_{ab} (e.g., if γ_1 is a geodesic of η_{ab}, choose an Ω with γ_1 as an axis of symmetry).

4.4 Paradoxes regained?

Consider again the super π machine which is supposed to print all the digits in the decimal expansion of π within a finite time span. Even leaving aside worries about whether the movement of the parts of the machine can be made to satisfy obvious kinematic and dynamic requirements, Chihara (1965) averred that there is something unintelligible about this hypothetical machine.

> The difficulty, as I see it, is not insufficiency of time, tape, ink, speed, strength or material power, and the like, but rather the inconceivability of how the machine could actually finish its supertask. The machine would supposedly print the digits on tape, one after another, while the tape flows through the machine, say from right to left. Hence, at each stage in the calculation, the sequence of digits will extend to the left with the last digit printed being "at center." Now when the machine completes its task and shuts itself off, we should be able to look at the tape to see what digit was printed last. But if the machine finishes printing all the digits which constitute the decimal expansion pi, no digit can be the last digit printed. And how are we to understand this situation? (Chihara 1965, p. 80)

Note first that the baldest form of Chihara's worry does not apply to the setup that has been imagined for M–H spacetimes; for the tape will not be available for γ_2's inspection since γ_1 goes crashing into a singularity or disappears to infinity. However, it might seem that a more sophisticated version of Chihara's conundrum can be mapped onto the M–H set up as follows. γ_1, who has available to her an infinite amount of proper time, prints the digits of π, say, one per second. And at the end of each step she sends a light signal to γ_2 announcing the result. γ_2 has a receiver equipped with an indicator which displays 'even' or 'odd' according as the case may be. By construction there is a $p \in \gamma_2$ at which γ_2 has received all of the signals from γ_1. One can then ask: What does the indicator read at that moment?

Any attempt to consistently answer this query fails. How the failure is reflected in any attempted physical instantiation will depend on the details of the physics—in one instantiation the indicator device will burn out before the crucial moment, in another the indicator will continue to display but the display will not faithfully mirror the information sent from γ_1, etc. But independently of the details of the physics, we know in advance that the functional description of the device is not self-consistent. Does this knowledge constitute a general reductio of the possibility of using M–H spacetimes to create the functional equivalents of Plato machines? No, for the inconsistency here can be traced to the conditions imposed on one component of the π machine—the receiver–indicator—and such conditions are not imposed in mimicking Plato machines.

If the M–H analogue of the super π machine is to operate as intended, then the receiver–indicator must satisfy three demands: (a) the indicator has a definite state for all relevant values of its proper time τ, (b) the indicator is faithful in the sense that, if it receives an odd/even signal at τ, then it instantly adopts the corresponding odd/even indicator state, and (c) the indicator does not change its state except in response to a received signal in the sense that if τ_{ns} is a time at which no signal is received, then the indicator state at τ_{ns} is the limit of indicator states as τ approaches τ_{ns} from below. These demands are supposed to guarantee that at the crucial moment the indicator displays the parity of the "last digit" of π. That such a component is possible by itself leads to contradictions if it is assumed that the receiver–indicator device is subject to infinitely many alternating signals in a finite time. The limit required by (c) does not always exist, contradicting (a). I take the impossibility of such a component to be the lesson of forlorn attempts to construct an M–H analogue of the super π machines.

Denying the use of such functionally inconsistent devices will not affect attempts to construct M–H analogues of Plato machines and to use them to gain new mathematical knowledge. The computer γ_1 is an infinity machine in the innocuous sense that it performs an infinite number of operations in an infinite amount of proper time. I see no grounds for thinking that such machines involve any conceptual difficulties unless they are required to compute a non-existent quantity. The uses to which I will put them make no such demand. Similarly, a conceptually non-problematic receiver–indicator device can be coupled to the computer through M–H spacetime relations in order to determine the truth values of mathematical conjectures. To flesh out the suggestion already made above, imagine, as in Pitowsky's example, that γ_1 is the world line of a computer which successively checks a conjecture of number theory for '1', for '2', etc. Since it has available to it an infinite amount of proper time the computer will in the fullness of time check the conjecture for all the integers. It is arranged that γ_1 sends a signal to γ_2 if and only if a counterexample is found. γ_2 is equipped with a receiver and an indicator device that is initially set to 'true' and which retains that state unless the receiver detects a signal, when the indicator shifts to 'false' and the receiver shuts off. By reading the display at the M–H point, γ_2 can learn whether or not the conjecture is true. Although I can give no formal proof of the consistency of this functional description, I see no basis for doubt. However, I will show below that attempts to physically instantiate this functional description run into various difficulties. But the difficulties have nothing to do with the paradoxes and conundrums of Thomson lamps and the like.

4.5 Characterization of Malament–Hogarth spacetimes

It was seen in section 4.3 that M–H spacetimes are not globally hyperbolic and thus violate Penrose's version of strong cosmic censorship. The converse is generally not true: some spacetimes that are not globally hyperbolic can fail to be M–H spacetimes. (Trivial example: Minkowski spacetime with a closed set of points removed does not contain a Cauchy surface but is not an M–H spacetime.) Some M–H spacetimes are acausal. Gödel spacetime is causally vicious in that for every point $p \in M$ $(=\mathbb{R}^4)$ there is a closed future-directed timelike curve through p (see chapter 6). In fact, for any $p \in M$, $I^-(p) = M$. Since Gödel spacetime contains timelike half-curves of infinite proper length, every point is an M–H point. I will not discuss such acausal spacetimes here. The reason is not because I think that the so-called paradoxes of time travel show that such spacetimes are physically impossible; indeed, I will argue just the opposite in chapter 6. But such paradoxes do raise a host of difficulties which, though interesting in themselves, only serve to obscure the issues about supertasks I wish to emphasize.

Therefore, in what follows I will restrict attention to causally well-behaved spacetimes. In particular, all of the spacetimes I will discuss are

stably causal, which entails the existence of a global time function (see chapter 6). I claim that among such spacetimes satisfying some subsidiary conditions to be announced, the M–H spacetimes are physically characterized by divergent blueshifts. The intuitive argument for this assertion is straightforward. During her lifetime, γ_1 measures an infinite number of vibrations of her source, each vibration taking the same amount of her proper time. γ_2 must agree that an infinite number of vibrations take place. But within a finite amount of his proper time, γ_2 receives an infinite number of light signals from γ_1, each announcing the completion of a vibration. For this to happen γ_2 must receive the signals in ever decreasing intervals of his proper time. Thus, γ_2 will perceive the frequency of γ_1's source to increase without bound. (This argument does not apply to acausal M–H spacetimes. The simplest example to think about is the cylindrical spacetime formed from two-dimensional Minkowski spacetime by identifying two points (x_1, t_1) and (x_2, t_2) just in case $x_1 = x_2$ and $t_1 = t_2$ modulo π. γ_2 can be chosen to be some finite timelike geodesic segment and γ_1 can be a timelike half-geodesic that spirals endlessly around the cylinder. The light signals from γ_1 may arrive at γ_2 all mixed up and not blueshifted.)

The main difficulty with this informal argument, as with all of the early literature on the redshift/blueshift effect (see Earman and Glymour 1980) is that the concept of frequency it employs refers to the rate of vibration of the source at γ_1 and to the rates at which γ_1 sends and γ_2 receives signals. But the effect actually measured by γ_2 depends on the frequency of the light signal (photon) itself. Thus, we need to calculate the blueshift using the definition of the emission frequency of a photon from a point $p_1 \in \gamma_1$ as $\omega_1 =: -(k_a V_1{}^a)|_{p_1}$ and the measured frequency of the photon as received at the point $p_2 \in \gamma_2$ as $\omega_2 =: -(k_a V_2{}^a)|_{p_2}$, where the timelike vectors $V_1{}^a$ and $V_2{}^a$ are respectively the normed tangent vectors to the world lines γ_1 and γ_2, and the null vector k^a is the tangent to the world line of the photon moving from the first to the second observer (see Fig. 4.2). Then the redshift/blueshift effect is given by the ratio

$$\frac{\omega_2}{\omega_1} = \frac{(k_a V_2{}^a)|_{p_2}}{(k_a V_1{}^a)|_{p_1}} \tag{4.1}$$

The following key fact is established in the appendix at the end of this chapter.

Lemma 4.2. *Let M, g_{ab} be a Malament–Hogarth spacetime containing a timelike half-curve γ_1 and another timelike curve γ_2 from point q to point p such that $\int_{\gamma_1} d\tau = \infty$, $\int_{\gamma_2} d\tau < \infty$, and $\gamma_1 \subset I^-(p)$. Suppose that the family of null geodesics from γ_1 to γ_2 forms a two-dimensional integral submanifold in which the order of emission from γ_1 matches the order of reception at γ_2. If the photon frequency ω_1 as measured by the sender γ_1 is constant, then the time-integrated photon frequency $\int^{p_2} \omega_2\, d\tau$ as measured by the receiver γ_2 diverges as p_2 approaches p.*

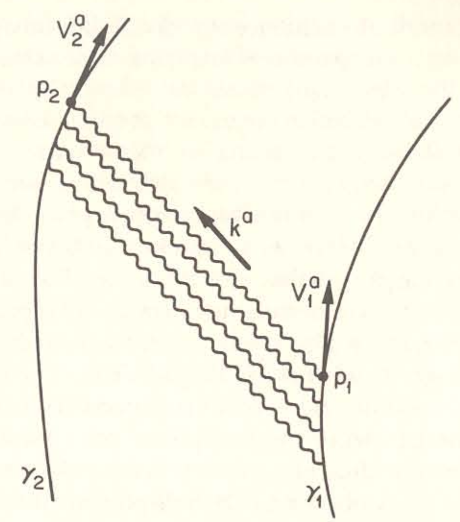

Fig. 4.2 The redshift/blueshift effect

Parameterize γ_2 by a t such that γ_2's past endpoint q corresponds to $t = 0$ and p corresponds to $t = 1$. Then it follows from Lemma 4.2 that $\lim_{t \to 1} \omega_2(t) = \infty$ if the limit exists. If not, then $\lim_{t \to 1} \omega^{\text{lub}}(t) = \infty$, where $\omega^{\text{lub}}(t) =: \text{lub}\{\omega_2(t'): 0 < t' \le t\}$. Thus, one can choose on γ_2 a countable sequence of points approaching p such that the blueshift as measured by γ_2 at those points diverges. Typically this behavior will hold for any such sequence of points on γ_2, but there are some mathematically possible M–H spacetimes where γ_2 measures no red or blueshift at some sequence of points approaching p.

The following example (due to R. Geroch and D. Malament) illustrates this counterintuitive feature. As in the toy model in Fig. 4.1, start with parallel timelike geodesics of Minkowski spacetime. Parameterize γ_1 by the proper time τ of the Minkowski metric and adjust the curve so that the past endpoint corresponds to $\tau = 0$ and r corresponds to $\tau = 1$. At the points on γ_1 corresponding to $\tau = \tau_n = 1 - (3/4)(1/2^n)$ draw a sphere of radius $r_n = 1/2^{n+3}$ (as measured in the natural Euclidean metric). On the nth sphere put a conformal factor Ω_n which goes smoothly to 1 on the surface of the sphere and which has its maximum value at the point on γ_1 corresponding to τ_n. Construct the Ω_n such that the proper time along γ_1 in the conformal metric $\Omega^2\eta_{ab}$ is infinite. For instance, if γ_1^n is the part of γ_1 within the nth sphere, set Ω_n so that $\int_{\gamma_1^n} \Omega_n \, d\tau = 1$. The result is an M–H spacetime. But at the points on γ_2 that receive photons from the points on γ_1 corresponding to $\tau = 1/2, 3/4, 7/8$, etc., there is no blue- or redshift.

While mathematically well defined, such examples are physically pathological. In particular, I do not know of any examples of M–H spacetimes which are solutions to Einstein's field equations for sources satisfying standard

energy conditions (see section 4.6) and which have the curious feature that the blueshift as measured by γ_2 diverges along some but not all sequences of points approaching the M–H point. Thus, although the slogan that M–H spacetimes involve divergent blueshifts is literally incorrect, it is essentially correct in spirit.

It may help to fix intuitions by computing the blueshift in some concrete examples. For the toy model pictured in Fig. 4.1 the result is

$$\frac{\omega_2}{\omega_1} = \frac{\Omega_{p_1}}{\Omega_{p_2}} = \Omega_{p_1} \tag{4.2}$$

This ratio diverges as γ_1 approaches the (missing point) r and γ_2 approaches the M–H point p.

Another stably causal M–H spacetime is obtained by taking the universal covering of anti-de Sitter spacetime (Hawking and Ellis 1973, pp. 131–134). Suppressing two spatial dimensions, the line element can be written as $ds^2 = dr^2 - (\cosh^2 r)\, dt^2$. Following Hogarth (1992) we can take γ_2 to be given by $r = r_2 = \text{constant}$ and γ_1 to be given by a solution to $dr/dt = \cosh r/\sqrt{2}$ (see Fig. 4.3). The blueshift is

$$\frac{\omega_2}{\omega_1} = \frac{\cosh r_1}{\cosh r_2(\sqrt{2} - 1)} \tag{4.3}$$

which diverges as $r_1 \to \infty$ and p_2 approaches the M–H point p.

We can also pose the converse question as to whether a divergent blueshift behavior indicates that the spacetime has the M–H property. The answer is positive in the sense that the proof of Lemma 4.2 can be inverted.

The fact that an M–H spacetime gives an indefinitely large blueshift for the photon frequency implies that the spacetime structure acts as an arbitrarily powerful energy amplifier. This might seem to guarantee

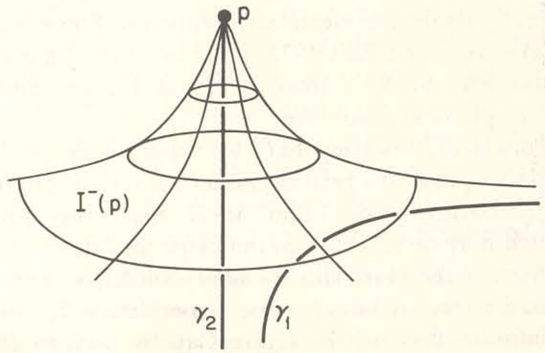

Fig. 4.3 Anti-de Sitter spacetime is a Malament–Hogarth spacetime

unambiguous communication from γ_1 to γ_2. But this first impression neglects the fact that a realistic instantiation of γ_1 will have thermal properties. The slightest amount of thermal radiation will be amplified indefinitely, which will tend to make communication impossible. In order not to destroy the receiver at γ_2, γ_1 will have to progressively reduce the energy of the photons she sends out. This means that there will be a point at which the energy of the signal photons will be reduced below that of the thermal noise photons. The indefinite amplification of the thermal noise will in any case destroy the receiver. Perhaps this difficulty can be met by cooling down γ_1 so as to eliminate thermal noise or by devising a scheme for draining off the energy of the signal photon while in transit. But even assuming a resolution of this difficulty, still further problems dog the attempt to use M–H spacetimes to accomplish supertasks.

4.6 Supertasks in Malament–Hogarth spacetimes

Are supertasks in Malament–Hogarth spacetimes to be taken seriously? The question involves three aspects. The first concerns whether M–H spacetimes are physically possible and physically realistic. As a necessary condition for physical possibility, general relativists will want to demand that the spacetime be part of a solution to Einstein's field equations for a stress–energy tensor T^{ab} satisfying some form of energy condition, weak, strong, or dominant (see chapter 3). The toy model of Fig. 4.1 can be regarded as a solution to Einstein's field equations with vanishing cosmological constant Λ by computing the Einstein tensor $G_{ab}(g)$ and then defining $T_{ab} =: (1/8\pi)G_{ab}$. But as conjectured in chapter 3, such models may be ruled out by the energy conditions. Anti-de Sitter spacetime, another M–H spacetime, can be regarded as a vacuum solution to Einstein's field equations with $\Lambda = R/4$, R (<0) being the curvature scalar; then the energy conditions are trivially satisfied. However, if it is required that $\Lambda = 0$, anti-de Sitter spacetime is ruled out by the strong energy condition if a perfect fluid source is assumed (see chapter 3).

None of these concerns touch Reissner–Nordström spacetime which is the unique spherically symmetric electrovac solution of Einstein's field equations with $\Lambda = 0$ (Hawking and Ellis 1973, pp. 156–161). Since this spacetime is an M–H spacetime, at least some M–H spacetimes meet the minimal requirements for physical possibility.

It is far from clear, however, that M–H spacetimes meet the (necessarily vaguer) criteria for physically realistic spacetime arenas. For one thing, it was seen in the preceding section that M–H spacetimes involve divergent blueshifts, which may be taken as an indicator that these spacetimes involve instabilities. Such is the case with Reissner–Nordström spacetime, where a small perturbation on an initial value hypersurface Σ (see Fig. 4.4) can produce an infinite effect on the future Cauchy horizon $H^+(\Sigma)$ of Σ (see Chandrasekhar and Hartle 1982).

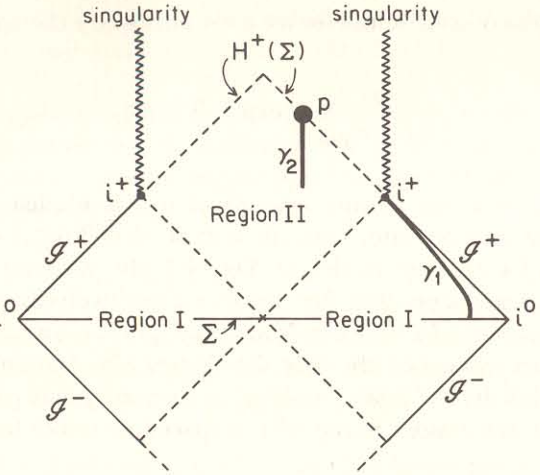

Fig. 4.4 Reissner–Nordström spacetime is a Malament–Hogarth spacetime

For another thing, various M–H spacetimes run afoul of one or other versions of Penrose's cosmic censorship hypothesis. By Lemma 4.1 all M–H spacetimes violate strong cosmic censorship, and many examples of M–H spacetimes violate weaker versions as well. Thus, evidence that cosmic censorship holds for physically reasonable spacetimes is ipso facto evidence against the physical reasonableness of M–H spacetimes. Conversely, one might take the rather bizarre scenarios that can be concocted in M–H spacetimes as grounds for thinking that a cosmic censor should be at work. But then again, those with a taste for the bizarre may hope that cosmic censorship fails just so that they can own the functional equivalent of a Plato machine.

I now turn to the second aspect of the question of how seriously to take the possibility of completing supertasks in M–H spacetimes. This aspect concerns whether γ_1 can be implemented by a physically possible/physically realistic device which, over the infinite proper time available to it, carries out the assigned infinite task. Once again the question is made difficult by the fact that there is no agreed upon list of criteria that identify physically realistic devices. I will make the task of tackling this question tractable by confining attention to dynamical constraints that physically realistic γ_1 should satisfy. (One doesn't have to worry about dynamical constraints on γ_2 since typically γ_2 can be chosen to be a geodesic.) Minimally, the magnitude of acceleration of γ_1 must remain bounded, otherwise any device that one could hope to build would be crushed by g-forces. This condition is satisfied in the anti-de Sitter case (Fig. 4.3) where $a(r_1) = \sqrt{2}\,[\exp(2r_1) - 1]/[\exp(2r_1) + 1]$, which approaches $\sqrt{2}$ as $r_1 \rightarrow \infty$. However, we must also demand a finite bound on the total acceleration of γ_1: $\mathrm{TA}(\gamma_1) =: \int_{\gamma_1} a\, d\tau$. For even with perfectly efficient rocket engines, the final mass m_f of the rocket and the mass m_{fuel} of the fuel needed to accelerate the rocket must satisfy (see the appendix to this chapter)

$$\frac{m_f}{m_f + m_{fuel}} \leq \exp(-\mathrm{TA}(\gamma_1)) \qquad (4.4)$$

Thus if $\mathrm{TA}(\gamma_1) = \infty$ an infinite amount of fuel is needed for any finite payload. In the anti-de Sitter case, $dr_1 = d\tau$ so that $\mathrm{TA}(\gamma_1) = \infty$, and the demand fails. In the toy model of Fig. 4.1 the demand is met since $\mathrm{TA}(\gamma_1) = 0$, γ_1 being a geodesic; but the spacetime involved was ruled out as not physically possible. In Reissner–Nordström spacetime a timelike geodesic γ_1 can be chosen to start on the time slice Σ (see Fig. 4.4) and to go out to future timelike infinity i^+.[5] And $\gamma_1 \subset I^+(p)$ for an appropriate point $p \in H^+(\Sigma)$. But again there are reasons to regard this spacetime as not being physically realistic.

Finally, since a physically realistic device must have some finite spatial extent, we are really concerned not with a single world line γ_1 but with a congruence Γ_1 of world lines. Even if Γ_1 is a geodesic congruence it cannot be instantiated by a physically realistic computer (say) unless the tidal forces it experiences remain bounded. Since the tidal forces are proportional to the Riemann curvature tensor,[6] one can satisfy this demand in Reissner–Nordström spacetime, which is asymptotically flat in the relevant region. One simply starts the geodesic congruence sufficiently far out towards spatial infinity and has it terminate on future timelike infinity i^+.

To summarize the discussion up to this point, it is not clear that any M–H spacetime qualifies as physically possible and physically realistic. But to the extent that M–H spacetimes do clear this first hurdle, it seems that the role of γ_1 can be played by a world line or world tube satisfying realistic dynamical constraints. However, Pitowsky (1990) feels that, for other reasons, γ_1 cannot be instantiated by a computer that will carry out the assigned infinite task. I will take up his worry in section 4.8 below. Before doing so I turn to the third aspect of whether M–H can be taken seriously. It concerns discriminations that the receiver γ_2 must make.

4.7 Malament–Hogarth spacetimes and unresolved mathematical conjectures

Can Malament–Hogarth spacetimes be used to gain knowledge of the truth values of unresolved mathematical conjectures? Suppose now for sake of discussion that some M–H spacetimes are regarded as physically possible and physically realistic and that in these arenas there are no barriers to a physically possible and physically realistic instantiation of γ_1 by a computer which carries out the task of checking Fermat's last theorem or some other unresolved conjecture of number theory. Nevertheless there are reasons to doubt that γ_2 can use γ_1's work to gain genuine knowledge of the truth value of the conjecture. The pessimism is based on a strengthening of Lemma 4.1.

LEMMA 4.3. *Suppose that $p \in M$ is a M–H point of the spacetime M, g_{ab} (that is, there is a future-directed timelike half-curve $\gamma_1 \subset M$ such that $\int_{\gamma_1} d\tau = \infty$ and $\gamma_1 \subset I^-(p)$). Choose any connected spacelike hypersurface $\Sigma \subset M$ such that $\gamma_1 \subset I^+(\Sigma)$. Then p is on or beyond $H^+(\Sigma)$.*

PROOF: If $p \in \mathrm{int}[D^+(\Sigma)]$ then there is a $q \in D^+(\Sigma)$ which is chronologically preceded by p. $M' =: (I^-(q) \cap I^+(\Sigma)) \subset D^+(\Sigma)$, and the smaller spacetime $M', g_{ab}|_{M'}$ is globally hyperbolic. Choose a Cauchy surface Σ' for this smaller spacetime which passes through p. Since $\gamma_1 \subset M'$ we can proceed as in the proof of Lemma 1 to obtain a contradiction.

Lemma 4.3 is illustrated by the Reissner–Nordström spacetime (Fig. 4.4). Any M–H point involved with a γ_1 starting in region I must lie on or beyond $H^+(\Sigma)$.

Think of Σ as an initial value hypersurface on which one specifies initial data that, along with the laws of physics, prescribes how the computer γ_1 is to calculate and how it is to signal its results to γ_2. Since by Lemma 4.3 any M–H point $p \in \gamma_2$ must lie on or beyond $H^+(\Sigma)$ for any appropriate Σ, events at p or at points arbitrarily close to p are subject to non-deterministic influences. In typical cases such as the Reissner–Nordström spacetime illustrated in Fig. 4.4 there are null rays which pass arbitrarily close to any $p \in H^+(\Sigma)$ and which terminate in the past direction on the singularity. There is nothing in the known laws of physics to prevent a false signal from emerging from the singularity and conveying the misinformation to γ_2 that a counter-example to Fermat's conjecture has been found.[7] (γ_2 need not measure an infinite blueshift for photons emerging from the singularity; at least there is nothing in Lemma 4.2 or the known laws of physics that entails such a divergent blueshift.) Of course, the receiver γ_2 can ignore the signal if he knows that it comes from the singularity rather than from γ_1. But to be able to discriminate such a false signal from every possible true signal that might come from γ_1, γ_2 must be able to make arbitrarily precise discriminations. In the original situation it was the Plato machine that had to perform a supertask by compressing an infinite computation into a finite time span. The trick adopted here was to finesse the problems associated with such a supertask by utilizing two observers in relativistic spacetime. But we have found that the finesse also involves a kind of supertask—not on the part of the computer but on the part of the receiver who tries to use the work of the computer to gain new mathematical knowledge.

This verdict may seem unduly harsh. If γ_2 is to be sure beforehand that, whatever γ_1's search procedure turns up, he will obtain knowledge of the truth value of Fermat's conjecture, then γ_2 must be capable of arbitrarily precise discriminations. But, it may be urged, if γ_2 is capable of only a finite degree of precision in his signal discriminations, he may yet learn that Fermat's conjecture is false (if indeed it is) if he receives a signal long enough before the M–H point so that it lies within his discrimination range. This, however, would be a matter of good fortune. One can pick at random a quadruple of numbers (x, y, z, n), $n \geq 3$, and check whether $x^n + y^n = z^n$. If one is lucky, a counterexample to Fermat's conjecture will have been found. But the interest in Platonist computers and their M–H analogues lay in the notion that they do not rely on luck.

Of course any observer faces the problem of filtering out spurious background signals from those genuinely sent from the system observed. But it is usually assumed that sufficiently thorough attention to the experimental setup could at least in principle control all such signals. What Lemma 4.3 shows, however, is that no such efforts can succeed even in principle in our case. No matter how carefully and expansively we set up our experiment—hat is, no matter how large we choose our initial value hypersurface—we cannot prevent spurious signals from reaching p or coming arbitrarily close to p.

The problem can be met by means of a somewhat more complicated arrangement between γ_1 and γ_2 by which γ_1 not only sends a signal to γ_2 to announce the finding of a counterexample but also encodes the quadruple of numbers that constitutes the counterexample. A false signal may emerge from the singularity, but γ_2 can discover the falsity by a mechanical check. With the new arrangement γ_2 no longer has to discriminate where the signal came from since a counterexample is a counterexample whatever its origin. Unfortunately, γ_2 may still have to make arbitrarily fine discriminations since the quadruple sent will be of arbitrarily great size (=number of bits) and must be compressed into a correspondingly small time interval at γ_2.

The worry about whether γ_2 can gain knowledge of Fermat's conjecture by using γ_1's efforts also involves the concern about γ_2's right to move from 'γ_1 has not sent me a signal' to 'Fermat's conjecture is true'. The correctness of the inference is not secured by the agreement γ_1 and γ_2 have worked out, for even with the best will in the world γ_1 cannot carry out her part of the agreement if events conspire against her. As suggested above, the most straightforward way to underwrite the correctness of the inference is for there to be a spacelike Σ such that $\gamma_1 \subset D^+(\Sigma)$ and such that initial conditions on Σ together with the relevant laws of physics guarantee that γ_1 carries out her search task. And if, as is compatible with at least some M–H spacetimes (e.g., Reissner–Nordström spacetime), the M–H point p can be chosen so that $\Sigma \subset I^-(p)$, it would seem that γ_2 could in principle come to know that the conditions which underwrite the inference do in fact obtain. But the rub is that p or points arbitrarily close to p may receive a false signal from the singularity, indicating that conditions are not conducive to γ_1's carrying out her task. If so, γ_2 is not justified in making the inference unless he can discriminate false signals as such. This, of course, is just another version of the difficulty already discussed. But the present form does not seem to have an easy resolution.

4.8 Can γ_1 carry out the assigned task?

γ_1 is supposed to check an unresolved conjecture of number theory for each of the integers. By construction, γ_1 has time enough. But Pitowsky feels that γ_1 never has world enough.

> The real reason why Platonist computers are physically impossible *even in theory* has to do with the computation space. According to general relativity the material universe is finite. Even if we use the state of every single elementary particle in the universe, to code a digit of a natural number, we shall very soon run out of hardware. (Pitowsky 1990, p. 84)

In response, I note that general relativity theory does not by itself imply a spatially or materially finite universe. Further, it was seen that there are spatially infinite M–H spacetimes, such as Reissner–Nordström spacetime, that are live physical possibilities in the minimal sense that they satisfy Einstein's field equations and the energy conditions. A γ_1 who wanders off into the asymptotically flat region of this spacetime certainly has space enough for any amount of hardware she needs to use. But she cannot avail herself of an unlimited amount of hardware without violating the implicit assumption of all of the foregoing; namely, that γ_1 and γ_2 have masses so small that they do not significantly perturb the background metric. Here Pitowsky's objection has some bite.

Perhaps there are solutions to Einstein's field equations where the spacetime has the M–H property and there is both space enough and material enough for a physically embodied computer with an unlimited amount of computation space.[8] Pending the exhibition of such models, however, one must confine oneself to tasks that can be accomplished in an infinite amount of time but with a finite amount of computation space. Whether there are such tasks that deserve the appellation 'super' remains to be seen.

4.9 Conclusion

Thomson lamps, super π machines, and Platonist computers are playthings of philosophers; they are able to survive only in the hothouse atmosphere of philosophy journals. In the end, M–H spacetimes and the supertasks they underwrite may similarly prove to be recreational fictions for general relativists with nothing better to do. But in order to arrive at this latter position requires that one first resolve some of the deepest foundation problems in classical general relativity, including the nature of singularities and the fate of cosmic censorship. It is this connection to real problems in physics that makes them worthy of discussion.

There are also connections to the philosophy of mathematics and to the theory of computability. Because of finitist scruples, some philosophers have doubted that it is meaningful to assign a truth value to a formula of arithmetic of the form $(\exists x_1)(\exists x_2)\ldots(\exists x_n)F(x_1, x_2, \ldots, x_n)$. It seems to me unattractive to make the truth of mathematical statements depend on the contingencies of spacetime structure. The sorts of arrangements considered above can be used to decide the truth value of assertions of arithmetic with a prenex normal form that is purely existential or purely universal.[9] (Fermat's last theorem, for example, has a purely universal form.) For such an assertion γ_1 is set to work to check through the (countably infinite) list of n-tuples of numbers in search of a falsifier or a verifier according as the assertion to be tested is universal or existential, and γ_1 reaps from these labors a knowledge of the truth value of the assertion. But as soon as mixed quantifiers are involved, the method fails.[10] However, Hogarth (1994) has shown how more complicated arrangements in general relativistic spacetimes can in principle be used to check the truth value of any arithmetic assertion of arbitrary quantificational complexity. Within such a spacetime it is hard to see how to maintain the attitude that we do not have a clear notion of truth in arithmetic.[11]

The computational arrangements between γ_1 and γ_2 envisioned might also seem to bring into doubt Church's proposal that effective/mechanical computability is to be equated with Turing computability or recursiveness, for apparently γ_1 and γ_2 can in concert obtain a resolution to recursively unsolvable problems by means that certainly seem to merit the appellations of 'effective' and 'mechanical'. But putting the matter this way is a little unfair to Church since any account of effective/mechanical computability that implies that there are subsets of numbers which can be effectively/mechanically enumerated, but whose complements cannot be, will be subject to the one-upmanship of bifurcated supertasks. Perhaps the most illuminating way to state the moral to be drawn from bifurcated supertasks is that two levels of computation need to be distinguished: the first corresponding to what the slave computer γ_1 can do, the second to what γ_2 can infer by having causal access to all of γ_1's labors. Church's proposal is best construed as aimed at the first level and as asserting that Turing computability is an upper bound on what any physical instantiation of γ_1 can accomplish. Read in this way, there is nothing in present concerns to raise doubts about Church's proposal.[12]

Appendix: Proofs of Lemma 4.2 and Equation 4.4

PROOF OF LEMMA 4.2. The null geodesics from γ_1 to γ_2 form a two-dimensional submanifold. For each of the null geodesics select an affine parameter λ which varies from 0 at γ_1 to 1 at γ_2. (This will always be possible since an affine parameter can be rescaled by an arbitrary linear transformation.) The null propagation vector $k^a = (\partial/\partial\lambda)^a$ satisfies the geodesic equation

$$k^a \nabla_a k^b = 0 \qquad \text{(A4.1)}$$

By supposition, these null geodesics form a submanifold. By connecting points of equal λ values, form a family of curves indexed by λ that covers the submanifold and interpolates between γ_1 and γ_2. Select any parameterization t of γ_1 and propagate this parameterization along the null geodesics to all the interpolating curves so that each null geodesic passes through points of equal t value. The indices λ and t form a coordinate system for the two-manifold. k^a and $\zeta^a = (\partial/\partial t)^a$ are its coordinate basis vector fields, which entails that they satisfy the condition $[\zeta, k]^a = 0$ so that

$$\zeta^a \nabla_a k_b - k^a \nabla_a \zeta_b = 0 \qquad \text{(A4.2)}$$

It follows that $(\zeta_a k^a)$ is a constant along the photon world lines. To show this it needs to be demonstrated that

$$\frac{d}{d\lambda}(\zeta_a k^a) = k^a \nabla_a(\zeta_b k^b) = 0 \qquad \text{(A4.3)}$$

This is done by computing

$$k^a \nabla_a(\zeta_b k^b) = k^b(k^a \nabla_a \zeta_b) + \zeta_b k^a \nabla_a k^b \qquad \text{(A4.4)}$$

The second term on the right-hand side of (A4.4) vanishes in virtue of (A4.1). Equation (A4.2) can then be used to rewrite the first term on the right-hand side as $\zeta^a k^b \nabla_a k_b = \frac{1}{2}\zeta^a \nabla_a(k_b k^b) = 0$ since k^a is a null vector.

Thus, for a photon sent from γ_1 to γ_2, we have $k_a \zeta^a_1 = k_a \zeta^a_2$, or $k_a V^a |\zeta^a_2| = k_a V^a |\zeta^a_2|$, where $V^a =: \zeta^a/|\zeta^a|$ is the normed tangent vector to the timelike world line. So from the definition (4.1) of photon frequency ratios one can conclude that $\omega_1|\zeta^a_1| = \omega_2|\zeta^a_2|$ which implies that

$$\int_{\gamma_1} \omega_1 |\zeta^a_1|\, dt = \int_{\gamma_2} \omega_2 |\zeta^a_2|\, dt \qquad \text{(A4.5)}$$

or

$$\int_{\gamma_1} \omega_1\, d\tau = \int_{\gamma_2} \omega_2\, d\tau \qquad \text{(A4.6)}$$

But $\int_{\gamma_1} d\tau = \infty$ and $\int_{\gamma_2} d\tau < \infty$. So if ω_1 is constant along γ_1, (A4.6) can hold only if $\int \omega_2\, d\tau = \infty$.

PROOF OF (4.4) (from Malament 1985). If V^n is the (normalized) four-velocity of the rocket and m its mass, the rate at which its energy–momentum changes is $V^p \nabla_p(mV^n)$, which must balance the energy–momentum J^n of its exhaust (it being assumed that the rocket's motor is the only source of propulsion). Thus,

$$J^n = V^n(V^p\nabla_p m) + mA^n \tag{A4.7}$$

where $A^n = V^m\nabla_m V^n$ is the four-acceleration. Since J^n is not spacelike, $J^n J_n \leq 0$. Consequently,

$$-(V^p\nabla_p m)^2 + m^2a^2 \leq 0 \tag{A4.8}$$

which uses $V^nV_n = -1$, $V^nA_n = 0$, and $a =: (A^nA_n)^{1/2}$. Furthermore, because the rocket is using up fuel, $V^p\nabla_p m \leq 0$. Thus,

$$a \leq -V^p\nabla_p(\ln(m)) = -\frac{d}{d\tau}(\ln(m)) \tag{A4.9}$$

So if m_i and m_f are the initial and final masses of the rocket, integration of (A4.9) yields

$$\mathrm{TA}(\gamma) \leq \ln(m_i/m_f) \tag{A4.10}$$

And since $m_i = m_f + m_{fuel}$,

$$\frac{m_f}{m_f + m_{fuel}} \leq \exp(-\mathrm{TA}(\gamma)) \tag{A4.11}$$

Notes

1. This task is taken up in Earman and Norton (1994).

2. A few new wrinkles are added in Earman and Norton (1994) concerning Ross's paradox (see Allis and Koetsier 1991; van Bendegem 1994) and some other paradoxes of the infinite.

3. A subset of $S \subset \mathbb{N}$ is said to be *recursively enumerable* (r.e.) just in case it is the range of a (partial) recursive function $f: \mathbb{N} \to \mathbb{N}$; informally, this means that there is an effective procedure for generating the members of S. S is said to be *recursive* if both S and the complement of S are r.e.; informally, there is an effective procedure for deciding membership in S. Key results on undecidability of formal systems hinge on there being sets that are r.e. but not recursive.

4. David Malament has pointed out that a quick proof of Lemma 4.1 can be obtained by using Prop. 6.7.1 of Hawking and Ellis (1973): For a globally hyperbolic spacetime, if $p \in J^+(q)$, then there is a non-spacelike geodesic from q to p whose length is greater than or equal to that of any other non-spacelike curve from q to p. Suppose that $\gamma \in I^-(p)$ is a timelike half-curve with endpoint q and that $\int_\gamma d\tau = \infty$. Since the endpoint q of γ belongs to $I^-(p)$, we could apply the proposition to p and q if the spacetime were globally hyperbolic. But then a contradiction results, since whatever the bound on the length of the timelike geodesic from q to p, we could exceed it by going along γ sufficiently far and then over to p. Robert Wald noted that an even quicker proof is obtained from the compactness of $J^-(p) \cap J^+(q)$ together with strong causality, which are consequences of global hyperbolicity.

5. In Fig. 4.4 i^0 labels spatial infinity and $\mathscr{I}^+$ and $\mathscr{I}^-$ respectively label future and past null infinity.

6. See Wald (1984, pp. 46–47) for a derivation of the formula for geodesic deviation.

7. One might also worry that a burst of noise from the singularity could swamp an authentic signal. But since any real signal arrives at γ_2 prior to the singularity noise, the former is not masked by the latter as long as the receiver can discriminate between a signal and noise.

8. The considerations raised here are similar to those discussed by Barrow and Tipler (1986) under the heading of "omega points."

9. Assuming that the relation quantified over is effectively decidable.

10. Showing this requires a more careful specification of how bifurcated infinity machines operate; see Earman and Norton (1994).

11. See Earman and Norton (1994) for more discussion of this and related matters.

12. But there are independent reasons to doubt Church's proposal; see Earman (1986) and Pitowsky (1990).