Mathematical and Physical Perspectives on Gravitational Radiation

B F Schutz
Max Planck Institute for Gravitational Physics
(Albert Einstein Institute)
Golm, Germany

August 2, 2002

1 Purposes

1. Survey the (various) mathematical foundations of gravitational radiation theory and their relationship to what physicists plan to observe.

2. Identify places where further mathematical investigation might be fruitful.

2 Defining gravitational radiation

- No universal definition of gravitational radiation: cannot invariantly define local GW content of an arbitrary solution.

- Worse than similar situation in EM, because of *equivalence principle* and *nonlinearity*.

- Equivalence principle implies no *local* definition possible in any situation: must attempt "regional" definition in regions at least as large as a wavelength. Any definition must be either an *approximation* or a *limit*. Definitions will generally have to cope with gauge problems.

- Nonlinearity compounds the difficulty of formulating approximations.

- Generally, expect to be able to define radiation if the wavelength is shorter than other relevant lengths and if there is a stationary "background" metric. Helps if there is a preferred frame and/or gauge.

- Since sources of radiation move, there is a close link between the study of gravitational radiation and the study of equations of motion in GR.

- Physically, definitions should be useful: it should be possible to relate them to what real detectors expect to see, or to other observable physical effects (like "radiation reaction"). Detectors work in the weak-field limit, but sources can be strongly nonlinear.
- A theoretical goal is to fit gravitational radiation into the rest of physics where possible, e.g. by proving energy conservation in some form. This has no real practical importance for experiments, but could be important in attempts to unify GR with the rest of fundamental physics.

3 Mathematical preliminaries

Almost all approximations used for studying gravitational radiation are forms of *perturbation theory*, so we begin by defining that. Then we mention a couple of useful background facts.

- Perturbation theory starts with a full (possibly nonlinear) solution (possibly with sources), and examines small departures from the solution (uniformly over the whole manifold, where possible). It is used to prove stability results (both structural stability and dynamical stability). Perturbations can be taken to any order:

  $$g_{\mu\nu}(x^\alpha;\varepsilon) = g^{(0)}_{\mu\nu}(x^\alpha) + \varepsilon h^{(1)}_{\mu\nu}(x^\alpha) + \frac{1}{2}\varepsilon^2 h^{(2)}_{\mu\nu}(x^\alpha) + \ldots.$$

  We can regard this as a family of metrics on a single manifold, or a family of manifolds. We will return to this issue below. The zero-order metric $g^{(0)}_{\mu\nu}$ is called the "background" metric; $h^{(1)}_{\mu\nu}$ is the first-order perturbation; $h^{(2)}_{\mu\nu}$ the second-order; etc. An ε-dependent coordinate transformation which limits to the identity at $\varepsilon = 0$ is called a *gauge transformation*.
- There is an important difference between *equations* and *solutions* in perturbation theory. It is not enough to derive equations at various orders. Solutions are defined by auxiliary conditions, such as initial data and/or boundary/asymptotic data. This data must also fit the approximation scheme or the solution will not be consistent with the equations. In most cases, the auxiliary data are essential for the approximation scheme to be well defined.
- The above equation could be either an *asymptotic approximation* or the first few terms of a *convergent summation*. Here is a simple example:

  Physicists use asymptotic approximations in calculations most often, but sometimes convergence is important. The slow (or non-existent!) convergence of the post-Newtonian approximation is currently a serious handicap to generating reliable waveforms for inspiral of neutron stars.

  Below I shall treat perturbation theory as an asymptotic approximation, not least because its convergence is generally unknown. I will regard perturbation theory as an approximation describing a sequence of solutions.

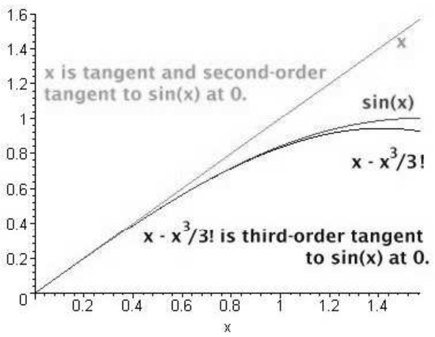

Figure 1: Illustration of meaning of asymptotic approximations to $sin(x)$.

First-order perturbations are the tangent space to the sequence at its origin; second-order perturbations are in the second-tangent space.

- The Sachs-Stewart-Walker theorem on gauge invariance asserts that *any geometric quantity in perturbation theory is gauge-invariant at its lowest non-zero order.* The most commonly encountered example is the Riemann tensor of linearized theory. The metric tensor is never gauge-invariant because its unperturbed value is normally non-zero. [R. Sachs, *Relativity, Groups and Topology*, (eds B. deWitt and C. DeWitt) New York: Gordon and Breach (1964); J.M. Stewart and M. Walker, *Proc. Roy. Soc. (London)* **A341**, 49 (1974).]

- The treatment of radiation in GR often parallels the treatment in EM where possible. We inherit some nomenclature and ideas from EM. For example, the *near zone* is the region containing the source and extending out to a distance of order one wavelength (if such a region exists). Similarly the *far zone* or *wave zone* is the region extending from a distance of a few wavelengths to infinity, provided it is smooth on the scale of a wavelength.

- *Radiation reaction* is a central concept and the main aim of many calculations. Actually, *it is a misnomer.* It represents the self-field forces on a dynamical system. Because of energy conservation, any energy lost by the system turns up as radiated energy far away, but there is no sense in which the system is reacting to this radiation: the self-interaction occurs immediately, there is no delay to discover what is happening in the wave zone. Importantly, there could in principle be conservative but non-trivial self-interaction forces that would not put energy into radiation but would affect the details of the radiated waveform. Here is an example from EM to illustrate the mechanics of how radiation reaction forces arise.

  – First consider a single charge q following an arc of a circular orbit of radius r with angular velocity ω, accelerated by a non-EM force.

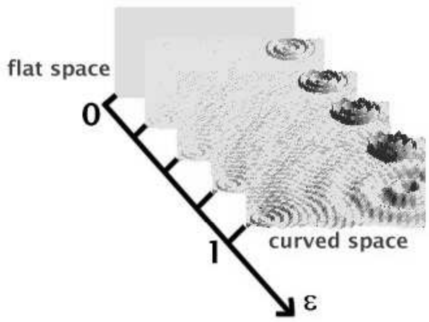

Figure 2: The sequence of metrics for perturbation theory seen as a sequence of manifolds parametrized by ε. This has the structure of a fiber bundle whose fibers are the manifolds. Shown here are 2D spatial sections from each of the manifolds, which are of course 4-dimensional.

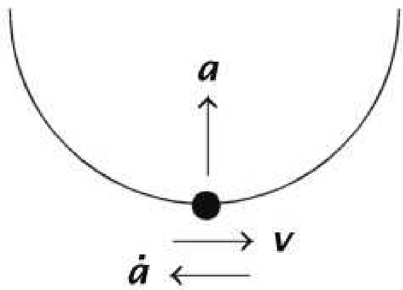

Figure 3: An accelerated charge moving on the arc of a circle. This will radiate dipole radiation.

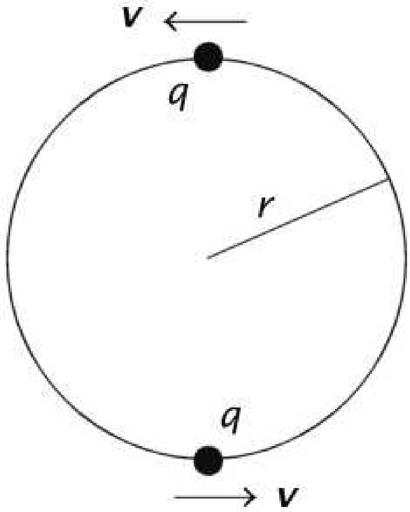

Figure 4: Two identical charges moving on the same circle at diametrically opposite points. This system has zero dipole moment and will radiate only quadrupole radiation.

We know it will radiate dipole radiation to lowest order in its velocity, and it will experience the dipolar radiation reaction force

$$F^j_{\rm react} = -\frac{2}{3} q^2 \dot{a}^j / c^3.$$

This force comes from the self-field of the accelerated charge, and it can be derived in a local manner, taking a limit of a small charged sphere as its size shrinks to zero. One sees clearly that the force comes from the retarded interactions of parts of the sphere with one another.

– Now put another identical charge on the same circle, diametrically opposite the first, and moving with the same angular velocity. This system has zero dipole moment, so it cannot emit dipole radiation.

– Yet the motion of each individual charge is the same as before, so each charge acts on itself with the dipole "reaction" force: this force is still there despite the absence of dipole radiation! But clearly this force must be cancelled by something else, because otherwise energy would not balance. *The cancellation comes from the retarded field of the other charge:* each charge exerts a force $+\frac{2}{3} q^2 \dot{a}^j / c^3$ on the other at dipole order, which can be found by expanding the expression for the retarded electric field of one point charge at the location of the other in powers of the velocity $\omega r/c$. What is left, at the next nontrivial order (c^{-5}), is a residual quadrupolar "reaction" force that is identical on both charges and which accounts for the energy radiated in electric quadrupole radiation.

– This shows that the "reaction" forces arise by self-interactions and take place on the time-scale of the light-crossing time of the system. There is no reason to suppose, especially in GR, that all such self-forces are linked to radiation: they may make orbits precess or do other conservative things.

4 A catalog of approaches to defining gravitational waves

Here is a list of six variations of perturbation theory and one non-perturbative method that are used in different circumstances to define gravitational waves and, in some cases, to study their sources. Most of them are normally valid in restricted regions of spacetime and under other conditions as well.

1. Simplest case is *linearized gravity.* This is a perturbation theory away from flat spacetime, taken to just first order. Here one finds the basic properties of waves: speed of light, polarisation. The choice of gauge is crucial to simplifying it: TT gauge is allowed. Note that Riemann tensor is gauge-invariant (Sachs-Stewart-Walker) so that one can use the Riemann tensor to describe the radiation invariantly in any region of spacetime where linearised gravity is valid. The approach is valid if $h^{(1)}$ is small and $h^{(2)}$ negligible. The source of the field must obey special-relativistic equations of motion: there is no self-gravity, no self-interaction. Therefore linearized theory cannot handle energy. It was first investigated by Einstein, who derived the quadrupole formula (error of factor 2 corrected by Eddington).

2. *Test matter* in a given curved background metric. The matter does not create a field at lowest order. This is perturbation theory where the perturbation is a new matter field rather than an alteration of the one that creates the zero-order metric. This is normally used in at least two cases, for detectors and for particle orbits in external fields (geodesics). It is actually quite good for detectors, since the self-field of a bar detector, for example, is of order 10^{-25}. For particle orbits, it is often iterated to the next level, where the self-field of the particle is taken into account. (See post-test approximation, below)

3. The post-linear or *post-Minkowskian* (PM) approximation simply iterates the field equations in small amplitude beyond linearised theory. So it is a general perturbation theory with special relativity as the zero-order metric. It is valid for weak fields, but it does treat energy. Leads to pseudo-tensors. However, the gauge is free and energy is gauge-dependent because of equivalence principle. Isaacson stress-energy tensor is derived

in this approximation, and it shows how energy can be localizedd to within a region comparable to a wavelength:

$$\text{Isaacson:} \qquad T^{\rm gw}_{\mu\nu} = \frac{1}{32\pi}\langle h^{(1)}_{\alpha\beta,\mu} h^{(1)\alpha\beta}{}_{,\nu}\rangle,$$

where indices are raised with the Minkowski metric and angle brackets $\langle\,\rangle$ denote an average over the localization region. This stress-energy is the effective source of the smooth non-zero curvature that appears at second order:

$$G^{(2)}_{\mu\nu} = 8\pi T^{\rm gw}_{\mu\nu}.$$

The PM approximation can also be used to investigate the exterior field of a source, far enough away that the field is weak. This has been done in a systematic way by Thorne. [Thorne, K.S., "Multipole expansions of gravitational radiation" *Rev. Mod. Phys.* **52** 299-340 (1980).] Damour has used the PM approximation as the starting point for his post-Newtonian work.

4. The *post-Newtonian* (PN) approximation is the best-studied and most-used of all the ones on our list. It is valid for slow motion and weak fields, linked in such a way that

$$\Phi_{\rm Newtonian} \sim v^2_{\rm typical}.$$

This can be seen as a refinement of the PM approximation, but it creates a re-ordering of terms because some parts of the stress-energy tensor are linked to the nonlinearity by the above equation. [Recent review: Luc Blanchet, "Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries", Living Rev. Relativity 5, (2002), 3. [Online article]: cited on 01 August 2002, http://www.livingreviews.org/Articles/-Volume5/2002-3blanchet/] A number of remarks are in order:

   - In the PN approximation, the wave operator degenerates to the Laplacian, so strictly speaking this approximation *cannot represent gravitational radiation.* It is only a description of the *near zone* of a gravitational-wave source.
   - The PN approximation is *incomplete* without a boundary condition or matching or extension to wave zone, and when these are applied correctly then the near-zone equations contain "radiation reaction".
   - PN orders are labeled 1PN, 2PN, 2.5PN, etc, marking powers of Φ or v^2 beyond the Newtonian order. But how can an approximation method have *half-order steps*!? The answer goes back to the difference between equations and solutions mentioned above: the dynamical equations only have integer steps, but the solutions have half-steps because they incorporate radiation conditions that are not natural to the PN approximation.

- The zero-order limiting solution is often thought of as Newtonian theory expressed geometrically, rather than SR, but this is not the only way to construct the perturbation theory. (See later.)
- The PN approximation is our main tool for calculating the generation of gravitational waves in realistic sources. It does not offer a definition of energy in radiation (since it does not represent radiation), but it does treat energy of self-field as source of field at higher orders, and does give radiation-reaction forces when handled carefully. Agreement between PN calculations and radio observations of the orbital decay of the Hulse-Taylor binary pulsar system provide our best evidence for the correctness of the Einstein model of gravitational radiation, and underpinned the Nobel Prize won by Hulse and Taylor in 1993. [Will, C.M., "The Confrontation between General Relativity and Experiment", Living Rev. Relativity 4, (2001), 4. [Online article]: cited on 01 August 2002 http://www.livingreviews.org/Articles/Volume4/2001-4will/]
- Importantly, the PN approximation shows that the Newtonian gravitational potential energy contributes to the effective mass-energy of radiation source, so the nonlinearity of GR respects the equivalence principle.
- There are several PN varieties, differing in two respects. One is the way the matching or extension to the wave zone is done. The other is in how compact sources are handled. Because often one is not interested in the interior details of the stars in a binary system, the PN approximation is often formulated using delta-function point particles. These are difficult in GR, especially due to nonlinearity. Different approaches "regularize" these singularities in different ways. The only real test of their validity is the consistency among the answers obtained by different methods. This is another *incompleteness* of many PN schemes: the representation of the source interior in a regular way. This is an unnecessary incompleteness, as I will explain below.

5. The post-test (PT) approximation allows the *self-field of a test particle* to be taken into account at higher orders to get the corrected motion of the particle. This differs from PN and PM because the zero-order solution is a full nonlinear solution of GR. It is basically a restricted version of nonlinear perturbation theory. For the particle-orbit problem, this self-interaction leads to "radiation reaction". This approach is not well-developed yet: we do not yet have any way of calculating the self-forces on a particle orbiting in the Kerr metric, for example. A small number of people are working very hard on it, because it is regarded as essential if LISA is to detect signals from gravitational captures of stellar-mass black holes by supermassive black holes. The LISA International Science Team (LIST) last year identified this as a priority problem for scientists (and funding agencies!) wishing to support the LISA project.

6. *Asymptotic null infinity* is the first place we think of when we want a good definition of radiation, but I have saved it for last because it makes an interesting comparison with the other schemes. Bondi was the first to give a good definition of radiation in the far field, following the outgoing radiation to what we now call asymptotic null infinity. But the concept of null infinity came later, introduced by Penrose in order to make Bondi's ideas more tangible. Some remarks are appropriate here:

   - The *historical importance* of this cannot be overestimated. It finally put gravitational radiation on firm ground, with an invariant local energy conservation theorem, that the apparent mass of an isolated system decreased with the energy lost to gravitational radiation.
   - However, null infninity does not tell us anything directly about sources. Basically it is a clever and rigorous limit of the exterior post-Minkowski approximation. It is therefore *incomplete* in the complementary way to that of the PN approach.
   - The reason that an invariant energy is possible at infinity is that there is a *preferred observer*, who in fact defines infinity, and who is flat (asymptotically). Moreover, this observer is infinitely large, so that grad is short wavelength and can be localized. This is the one place where we can define a universally agreed energy flux carried by gravitational waves.
   - Null infinity is unfortunately *less connected* than the other schemes to practical computations of sources. While it provides a reference for ideas, it has proved hard to link it, for example, to the PN calculation to set a no-incoming-radiation condition.
   - From a physical point of view, null infinity is *very* far away. A measure of how far one has to get from a source to be "near" infinity is to consider the divergence of the true curved-space light-cones from their flat-space approximations, which wind up at spatial infinity. Martin Walker first pointed out the enormous distance required to separate these cones by just one wavelength or period of the gravitational wave, a reasonable length scale for a radiation problem. The separation is something like $2M\ln(r/M)$. Setting this equal to λ for the Hulse-Taylor pulsar, we solve for r and find that it is a bit more than 10^{10^9} km! This is unimaginably bigger than the observable Universe, whose radius is a mere 10^{23} km.
   - This highlights the true physical problem with null infinity: radiation leaving a source encounters all kinds of problems — black holes, lenses, caustics — when it travels only a short fraction of the size of the Galaxy; by the time real outgoing light cones reach anywhere near asymptotic null infinity, they do not resemble their mathematical idealization at all. What physicists should have as an outer wave solution region, to complete the PN scheme, is something much closer, something that could reasonably fit between the Hulse-Taylor

binary and its nearest neighboring star in the Galaxy. This is where the physicists' "far zone" lives. It might start at, say, 10 times the wavelength, which would be about 10^{10} km, or 0.001 pc. I will propose below a mathematical construction for this asymptotic zone.

7. *Numerical approximations* are the only non-perturbative approach on my list. They are typically based on approximation by use of basis functions. Here the small parameter is the resolution scale. Orders of approximation are simply refinements of resolution. Usual finite-difference schemes are approximations using piecewise continuous polynomials. Other schemes that use global basis functions, e.g. orthogonal polynomials, have been used in relativity by the Paris group to good effect. The simplest applications of numerical methods to gravitational wave problems have been local, solving a bounded region of spacetime that includes the source and some approximation to the far zone. A more elegant and potentially powerful scheme is to incorporate conformal techniques to bring null infinity to a finite point on the grid, then can incorporate infinity into the computational domain (Friedrich, Husa, Lechner, Frauendiener all attending this meeting). Interestingly, when the system is as highly relativistic as a two-black-hole collision, our measure of how far away null infinity is becomes much more reasonable, perhaps as small at 10^{10} km. The far zone in this case starts much closer, perhaps at 10^4 km, but the divergence of the light cones may already be significant before the cones reach any other stars.

5 Stitching together local approaches in different regions

Given the number of different approaches on my list, it is perhaps remarkable that we all seem to agree that they all deal with the same thing, gravitational radiation, in one way or another. Clearly some are derived from others, but some are also incomplete. Since PN and the external PM are incomplete, each needing to be matched to the other, there have been many attempts to unite them. There are simple (and overly simple!) outgoing wave boundary conditions in numerical relativity, matched asymptotic expansions and other matching techniques in PN and PM approaches, and attempts to match to the Bondi system. What is needed is a more unified view of the whole spacetime and how the

different parts fit together. I propose one here. [See B.F. Schutz, "Motion and Radiation in General Relativity", in Bressan, O., Castagnino, M., Hamity, V., eds., *Relativity, Supersymmetry, and Cosmology*, (World Scientific, Singapore, 1985), pp. 3–80.] Consider a sequence of spacetimes, as before, only this time drawn with their time-axes, and showing a sketch of the binary orbit of two stars. The systems are constructed to have a Newtonian limit in the following way: let the density of the stars decrease uniformly toward zero (off the left of our picture) while the stars remain at the same orbital radius. This requires that the velocities decrease, and it is easy to see that the relation between the density-dependent Φ and v^2 is correct. Therefore, as one proceeds along the sequence from right to left, in the direction of decreasing density, the orbital time takes longer. I have drawn a map between the spacetimes connecting events at similar spatial positions and proper times.

The limit of this sequence of manifolds is Minkowski spacetime, since the density goes to zero.

But the Newtonian manifold is there. To see where, construct a different map, which maps times of fixed number of orbital periods, and spatial positions that preserve proper distance. Then the orbits are mapped onto one another, and each fiber looks the same. Rescaling the time-coordinate rescales the density so that it becomes constant on this map. The metric, however, is distorted in these coordinates because of the time-rescaling, and essentially the speed of light is sent to infinity. The limiting manifold (technically, the *tangent space* to the map at the far left end, i.e. at the Minkowski spacetime element of the sequence) is the Newtonian manifold.

We can construct even more interesting maps between fibers. Consider the next figure, where the map is drawn connecting times that are at similar periods of the orbit and thus of the gravitational waves, just as in the Newtonian map. But in this case, we re-adjust the spatial map to preserve the speed of light, so the unit of spatial distance is the wavelength of the gravitational wave. In these units, the orbits shrink. The end point of this map (again, the tangent space to the flow of the map) is a Minkowski spacetime with linearized gravitational waves and a line cut out at the center where the binary system is.

This is my proposal for the *wave zone* associated with this sequence of solutions. It is a 4-dimensional manifold, not just a boundary. Recall that all these maps thread through the same sequence, just looking at it in different ways. Using the waves as our distance standard, the Newtonian limit leaves behind just the radiation far from the system.

This picture has computational power. It is possible to show that the Isaacson energy flux, integrated over a sphere around the point-like source region in this limiting manifold, matches the expected quadrupole radiation. More interesting, the same is true for the integrated angular momentum flux. In this limiting Minkowskian wave zone, *angular momentum is perfectly well defined*. The supertranslation ambiguity arises at a higher order. [E. Nahmad-Achar, "A new derivation of the quadrupole formula for angular momentum", *J. Math. Phys.*, **30**, 1009–1012 (1989).]

This idea of maps along the sequence can also solve the other incompleteness problem of the PN approximations, the problem of point particles. Construct a sequence in which the source stars decrease in mass as in the original sequence, but also decrease in size so that M/R remains constant and is bounded below $1/2$. They therefore approach point-like size but remain regular regions of spacetime. By introducing a further map that scales as M and therefore examines the details of the interior, one can show that the interior solutions obey the Einstein equations for an isolated body in the limit. But they also continue to

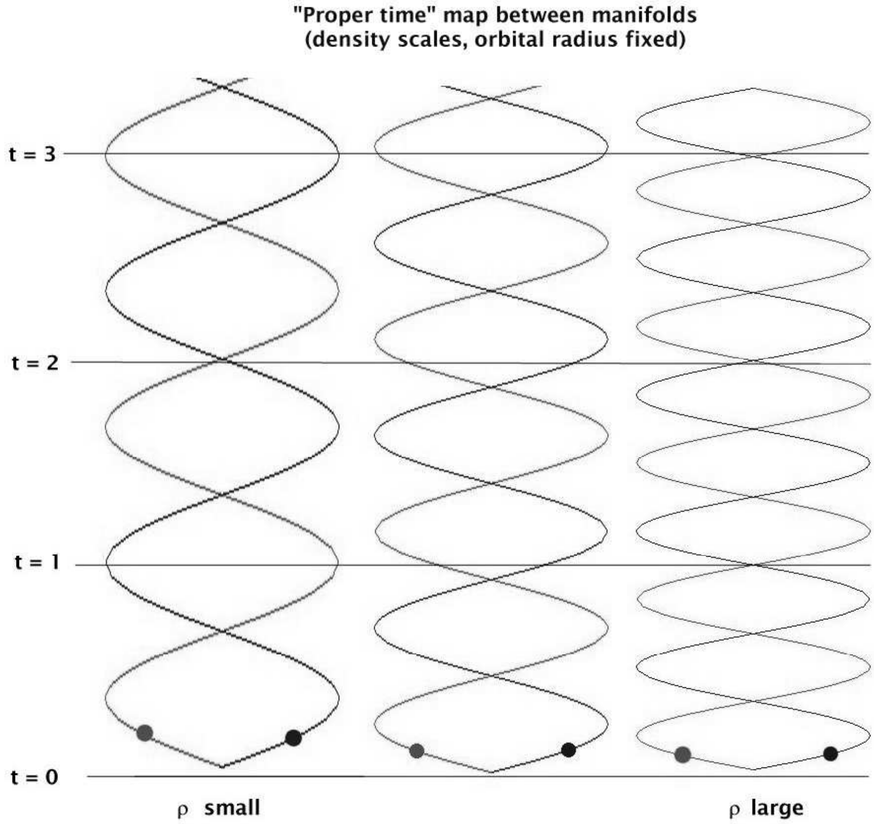

Figure 5: Three spacetimes from a sequence of spacetimes that has a Newtonian limit. The limiting spacetime is off the left of the illustration. Shown are sections including the time and one spatial direction, with the orbits of the two stars projected onto the sections. The sequence is arranged so that the densities of the stars decreases to the left, while their orbits remain in the same locations. The orbital period lengthens to the left. The horizontal lines are maps through the fiber bundle that identify points in the different manifolds with one another. These are drawn to identify points at the same proper time and distance. At the left end of the sequence is flat spacetime.

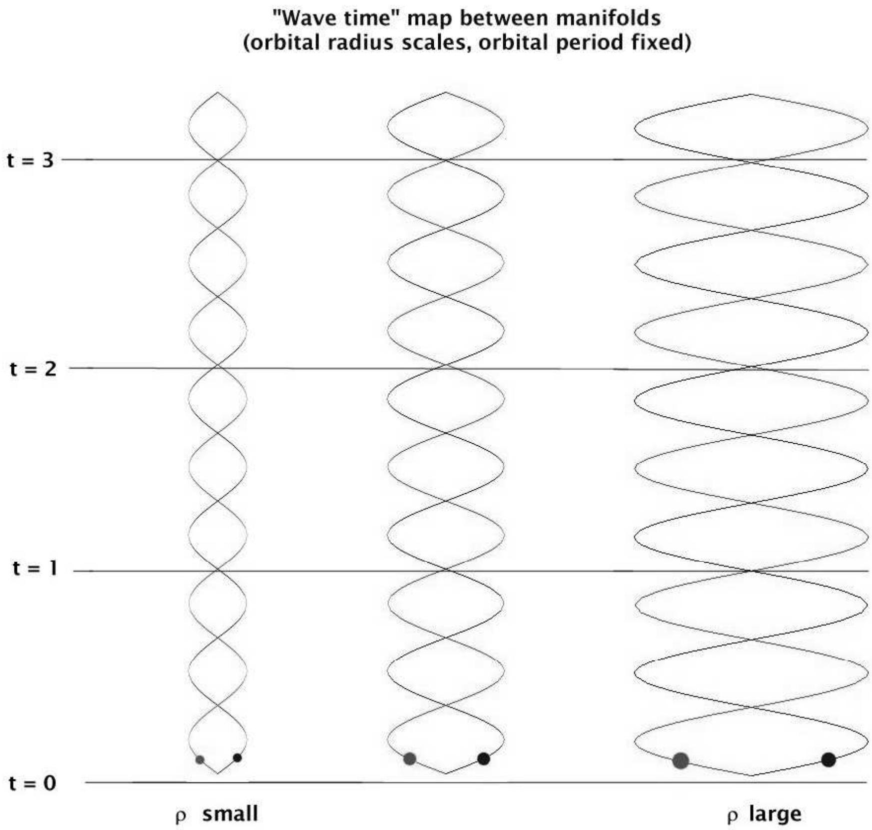

Figure 6: Three spacetimes from the same sequence of spacetimes as in the previous illustration, but with a different map identifying points in the different manifolds. The horizontal lines are maps that identify points with the same number of orbital periods in time and with the same number of gravitational wavelengths in space. This shrinks the spatial orbits, so that at the left end of the sequence, the tangent to the map is a flat spacetime with one line removed (the limit of the orbits) that contains the gravitational waves emitted by the system.

orbit one another in a Newtonian way. In this way, Futamase showed that two neutron stars will follow Newtonian orbits despite their strong internal gravity. [T. Futamase, *Phys. Rev. D*, **32**, 2566 (1985).]