to solve, but for which a solution, once found, can be efficiently recognized by a classical computer. QMA can be seen as the quantum generalization of NP. All QMA-hard problems are also NP-hard, but there are conjectured to be NP-hard problems that are not QMA-hard (see Fig. 1). Now, it's reasonable to ask: once we know a problem is NP-hard, isn't that hard enough? NP-hardness already establishes a problem as intractable in the worst case, at least under the famous 'P ≠ NP' conjecture. So why go to the additional step, as Schuch and Verstraete do, of proving the problem QMA-hard?

Here is where things get interesting. Suppose we consider a slight variant of the electron-ground-state problem, where we want to minimize the energy over all pure states, but are not interested in mixed states (which are thermal mixtures of pure states). In that case, minimizing single-electron energies could already be a difficult NP problem. If we found a fast algorithm to compute the universal functional, the consequence would be, not to solve QMA problems, but 'merely' to make the class QMA equal to the class NP — which is again considered unlikely. Thus, here we can get evidence that a practical problem is hard, but only by reasoning about a hypothetical collapse of 'higher level' computational classes. The conclusion really does depend on the fine-toothed distinction between QMA and NP, between quantum proofs and classical proofs.

In important respects, the result of Schuch and Verstraete is illustrative of quantum information science as a whole. This field does nothing to challenge the laws of quantum mechanics, the framework for almost all of physics since the 1920s. But it does ask a new set of questions about those laws. (In this case, what is the complexity of computing the DFT universal functional?) Because such questions straddle disciplines, they can look strange at first both to physicists and to computer scientists. But often enough they've turned out to have illuminating and non-trivial answers. ❒

Scott Aaronson is in the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA.
e-mail: aaronson@csail.mit.edu

RELATIVISTIC THERMODYNAMICS

Always look back

Formulating a consistent framework for relativistic thermodynamics has been the subject of intense debate over the past century. Defining quantities with respect to the observer's past lightcone could open new vistas.

Fabrice Debbasch

Traditional models of continuous media, such as fluids or elastic systems[1,2], were developed in the nineteenth century. They conform, therefore, to the Galilean world view, where time and space are absolute. Generalizing these models to Einstein's relativistic framework has been the subject of active research since the beginning of the last century (for a review, see for example ref. 3). Considerable progress has been made, but some still find a certain conceptual unease with present relativistic continuous-media theories. Writing on page 741 of this issue[4], Jörn Dunkel, Peter Hänggi and Stefan Hilbert look at the problem from a fresh perspective and put forward a bold suggestion, one that should provide new insight, but — if taken to its logical conclusion — will also generate new questions.

Let us first consider traditional macroscopic continuous-media theories, such as non-relativistic hydrodynamics (one example is the Navier–Stokes model[5]). A non-relativistic continuous medium is a physical system which, at any given time t, occupies a certain region $R(t)$ in space. The macroscopic attributes of the medium are represented by time-dependent fields defined over $R(t)$. Some of these fields represent extensive quantities, such as the number of particles, the energy, the momentum or the entropy. Furthermore, we can define densities of these extensive quantities (particle density, energy density, and so on), with respect to some reference volume measure (or element) $\mathrm{d}V$. Other fields represent intensive quantities, such as temperature or pressure, and these are not densities. But continuous-media theory never can dispense with extensive quantities, and all of these are represented, at any given time, by spatial densities, that is, by densities with respect to a certain (typically) three-dimensional (3D) spatial volume element. And this is where the potential conceptual conflict with the relativistic world view crops up. Indeed, the concept of spatial density inherently breaks the relativistic symmetry between time and space[6] and it is, *prima facie*, not obvious how it should be generalized to the relativistic world.

Standard relativistic continuous-media theories address and solve this issue in the following way: consider an arbitrary observer in a relativistic continuous medium (which is now defined as a physical system that occupies a certain domain $\mathcal{R}$ of the four-dimensional (4D) spacetime). This observer is represented by a choice of 4D coordinates $(ct, \mathbf{r})$, where c is the speed of light and $\mathbf{r}$ represents the spatial coordinates used by the observer. This observer thus slices the 4D region $\mathcal{R}$ into a collection of 3D subregions $R(t)$, where the time parameter t typically belongs to a certain real-valued interval; the subregion $R(t)$ is, from the point of view of this observer, the 3D spatial region occupied by the continuous medium at time t. The observer can now proceed and build continuous-media theories using, at each time t, spatial densities with respect to some reference volume element $\mathrm{d}V(t)$ defined over $R(t)$.

Of course, the choice of slicing depends, by definition, on the observer as does, *de facto*, the reference volume element. It would thus seem that when relativistic theories are built in this way, they should depend on the observer, and should therefore be ruled out as objective models of the intrinsic (that is, observer-independent) behaviour of a system. This is not so, however, because it turns out that all fields involved in such theories can nevertheless be collected into tensor fields[3,7–9]. The resulting theories are therefore actually observer independent, even though the precise interpretation — as opposed to predictions — of a certain theory does depend on the observer using that theory. Note that the same situation is encountered in Maxwellian electromagnetism[10], where the particle and current densities are spatial densities. They can, however, be represented as a tensor field, the so-called 4-current, and this makes the Maxwell theory compatible with Einstein's theory of relativity. The same remark applies to all Yang–Mills gauge theories[11,12], both quantum and non-quantum.

Now, because the velocity of light is finite, a given observer at each point $P(t)$ on his or her world line — the path on which the observer travels trough spacetime — will never have access to the whole 3D region $R(t)$, but only to the interior of their past lightcone; this is a 4D subdomain of the 4D spacetime, and its intersection with $R(t)$ is reduced to $P(t)$. As a consequence, considering fields defined over $R(t)$ and densities with respect to a 3D volume element defined over $R(t)$ may not seem really physical. Dunkel, Hänggi and Hilbert[4] therefore suggest that $R(t)$ should be replaced by the 3D past lightcone of the observer at point $P(t)$. (This past lightcone reduces to $R(t)$ when c tends to infinity, as is the case in the Galilean regime.)

This idea seems indeed reasonable and it has the advantage of being arguably more physically sound than the conventional procedure. But still, it remains to be seen where this suggestion will lead us. Among the open issues are the following: first, from a purely mathematical or physical perspective, there is no problem whatsoever with integrating on a lightcone. However, it is impossible to average on a lightcone in an intrinsic, observer-independent manner (this is because lightcones are so-called null surfaces[13], on which the normal vectors are also tangent vectors — remember that the relativistic line-element is not necessarily positive). All lightcone averages therefore involve an extra structure, typically the choice of an observer, and it is not clear if this constitutes a severe limitation or not. Second, when following in the footsteps of Dunkel, Hänggi and Hilbert[4], it is tempting to revisit all Yang–Mills theories and connect them with lightcone densities. Will this be possible? And will it have any influence on how we view quantization? The future will tell. ❐

Fabrice Debbasch is in the Equipe de Relativité, Gravitation et Astrophysique (ERGA) of the Laboratoire d'Etude du Rayonnement et de la Matière en Astrophysique (LERMA), Université Paris 6, 3 rue Galilée, 94200 Ivry, France.
e-mail: fabrice.debbasch@gmail.com

References

1. Landau, L. D. & Lifshitz, E. M. *Fluid Mechanics* (Pergamon, 1987).
2. Landau, L. D. & Lifshitz, E. M. *Theory of Elasticity* (Pergamon, 1986).
3. Israel, W. in *Relativistic Fluid Dynamics* (eds Anile, A. & Choquet-Bruhat, Y.) 152–210 (Lecture Notes in Mathematics, Vol. 1385, Springer, 1987).
4. Dunkel, J., Hänggi, P. & Hilbert, S. *Nature Phys.* **5,** 741–747 (2009).
5. Batchelor, G. K. *An Introduction to Fluid Dynamics* (Cambridge Univ. Press, 1967).
6. Wald, R. M. *Space, Time and Gravity* (Univ. Chicago Press, 1992).
7. Debbasch, F. *Physica A* **387,** 2443–2454 (2008).
8. Debbasch, F. & van Leeuwen, W. *Physica A* **388,** 1079–1104 (2009).
9. Debbasch, F. & van Leeuwen, W. *Physica A* **388,** 1818–1834 (2009).
10. Jackson, J. D. *Classical Electrodynamics* (Wiley, 1975).
11. Dubrovin, B. A., Novikov, S. P. & Fomenko, A. T. *Modern Geometry: Methods and Applications* (Springer, 1984).
12. Itzykson, C. & Zuber, J. B. *Quantum Field Theory* (Dover, 2006).
13. Wald, R. M. *General Relativity* (Univ. Chicago Press, 1984).

SAND SWIMMERS

In silica *in silico*

Swimming and flying are complicated processes to model, but at least the laws of fluid dynamics are known. In contrast, sand is a trickier medium to understand than water or air, because it can behave as a solid or as a fluid. Moreover, the presence of a 'swimmer' — such as the sand skink *Plestiodon reynoldsi* (pictured), seeking refuge from the heat of the Sun — changes the local properties of the sand, creating pockets of air and affecting the force chains between the granules. Consequently, there are no analytical equations of motion. To better understand the mechanism of swimming through a solid yet shifting medium, Takashi Shimada and colleagues have simulated the locomotion of a sand swimmer (*Phys. Rev. E* **80,** 020301; 2009), using a 'push-me-pull-you' model (pictured moving to the right) introduced by Joseph Avron and colleagues (*New J. Phys.* **7,** 234; 2005).

In essence, the push-me-pull-you model describes two disks connected by a spring. The disks inflate and shrink. To move forwards in fluid-like sand, the smaller anterior disk inflates as the spring lengthens. The initially fully inflated posterior disk acts as an anchor in solid-like sand. Once the anterior disk is fully inflated, it then acts as the anchor while the posterior disk shrinks and moves forwards as the spring contracts. To complete the move, the posterior disk inflates again, ready for the next stroke. Thus, a sand swimmer must deal with solidification near the anchor and fluidization near the moving disk at the same time.

The simulation's surprising result is that the optimal swimming frequency for maximum velocity is different from that for maximum efficiency. For example, if the swimmer moves too fast, the large voids created cause the swimmer to lose traction and slip. Hence the most efficient swimmer swims slowly. But move too slowly and the sand re-solidifies before any forward motion can be completed. Unexpectedly, the simulation also provides information on the fundamental time scales associated with granular packing.

MAY CHIAO

© ALESSANDRO CATENAZZI