NEWS & VIEWS

PARTICLE PHYSICS

Mass by numbers

Frank Wilczek

A highly precise calculation of the masses of strongly interacting particles, based on fundamental theory, is testament to the age-old verity that physical reality embodies simple mathematical laws.

In a milestone paper, Dürr *et al.*[1] report a first-principles calculation of the masses of strongly interacting particles (hadrons, such as the proton), starting from the basic equations for their constituent particles (quarks and gluons), and including carefully documented estimates of all sources of error. Their results, published in *Science*, highlight a remarkable correspondence between the ideal mathematics of symmetry and the observed reality of the physical world.

Quantum chromodynamics (QCD), the theory of the so-called strong force or strong interaction, postulates elegant equations for quarks and gluons. Those equations embody enormous symmetry, which largely dictates their form. A dramatic reflection of this conceptual rigour is that the equations contain very few freely disposable parameters — just a mass for each 'flavour' of quark (u, d, s, c, b, t) and an overall coupling constant. This makes QCD, in principle, an extremely powerful predictive framework. In fact, it's even tighter than this accounting suggests: for many purposes one can ignore the heavy quarks (c, b, t) and absorb the coupling constant into an overall scaling factor.

QCD predicts, however, that quarks and gluons are not observable particles. Rather, they occur only as building blocks inside more complex objects, collectively dubbed hadrons. The most familiar hadrons are protons and neutrons, from which ordinary atomic nuclei are assembled. Over decades of investigation, dozens of additional hadrons have been discovered. Most of these relatives of protons and neutrons are highly unstable, but their properties — notably their mass, charge and spin — can be measured[2].

If QCD is valid, its equations should account for all the properties of hadrons. But it has proved extremely challenging to solve those equations with enough precision to enable a sharp, quantitative comparison between theoretical predictions for, and experimental measurements of, hadron properties. The full power of modern, massively parallel computing has been brought to bear on this problem. Several impressive partial results have been announced in recent years[3]. Now Dürr *et al.*[1] have assembled all the pieces systematically, and added some new refinements, to achieve a fully convincing, successful comparison at a level of precision of 1–2% (Fig. 1).

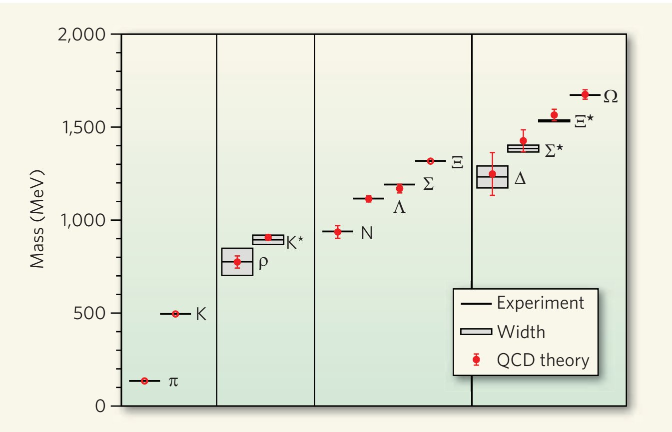

Figure 1 | Theory meets experiment. The masses of light hadrons (expressed in energy units; 1 MeV corresponds to 10^6 electronvolts) computed by Dürr *et al.*[1] using quantum chromodynamics (QCD) calculations (filled circles) are in remarkable agreement with the experimental values (horizontal lines). Each symbol (π, K and so on) refers to a different type of hadron, according to quark complement, spin and so on. The widths of the bands indicate the experimental decay widths, which are related to the finite lifetime of the particles. The vertical error bars denote the theoretical error estimates. Three of the hadrons (π, K and Ξ) have no error bars because they are used to fix the theory's parameters. The vertical lines divide the particles into four groups according to spin: 0, 1, ½, 3/2, reading from left to right. (Modified from ref. 1.)

A key aspect of their calculations is the estimation of errors. We know the equations of QCD precisely, but practical calculations require several approximations. To appreciate these approximations, we must first briefly review some specific, unusual features of the equations and their solution.

The primary objects in the theory of QCD are quantum fields — the quark and gluon fields. Quantum fields are entities that fill all space and exhibit spontaneous activity. That spontaneous activity is often referred to as quantum fluctuations, or 'virtual particles'. In the mathematical formulation, there is a master wavefunction for the quantum fields. The wavefunction is a superposition of different possible patterns of excitation in the fields, each occurring with some definite amplitude. The central problem involved in solving the equations of QCD, to predict the census of hadrons and their properties, is to compute this wavefunction: that is, to determine the numerical value of the amplitudes. Having constructed the wavefunction of 'empty space', we can inject different combinations of quarks and gluons and study the equilibria they settle into. Those equilibria correspond to observable particles — the hadrons.

The possible patterns of excitation in continuous quark and gluon fields map out a space of infinite dimensionality — roughly speaking, we should have to specify 84 $((3 \times 3 \times 4) + (8 \times 6))$ numbers at each point in space. For the quark fields there are three flavours, three colours, and four components accounting for spin and antiparticles; for the gluon fields there are eight directions in the space of its symmetry group, and for each direction there are six fields: three electric and three magnetic. In principle, we should calculate the amplitude of each such pattern. But no computer can handle an infinite number of variables, so two types of approximation seem unavoidable: the spatial continuum must be replaced by a discrete lattice of points; and the calculations must focus on a finite volume.

The process of discretization, which might seem to be a drastic mutilation of the theory, is actually well controlled theoretically, owing to QCD's central property of asymptotic freedom[4]. In this context, asymptotic freedom implies that the short-wavelength fluctuations of the fields, which are the ones we lose track of when we discretize space, are of a very simple form. In technical jargon they approach Gaussian random fields, or what physicists call 'free fields'. Thus the effects of the missing fluctuations can be computed analytically and added back in.

The approximation of finite volume is mitigated by the fact that in QCD the fundamental interactions occur among field variables at neighbouring points (we say they are local interactions). The patterns of equilibrium that define hadrons are, however, generally spatially extended, and so it is important to take a large enough volume so they fit comfortably. It is possible to control finite-volume errors by varying the simulated volume and making theoretically informed extrapolations.

Two additional approximations have also proved unavoidable, and troublesome, in practice. One is that as the *u* and *d* quark masses are taken down to their (very small) physical values, the equations get harder to solve. (For experts: this is because there are long correlation lengths, and the equations become numerically 'stiff'.) Like the compromise of assuming a finite volume, this is handled by sophisticated, theoretically informed extrapolation from simulations using larger mass values. Finally, even after acceptable levels of discretization and restriction to finite volume, the space that should be surveyed by the wavefunction is far too large for even the most powerful modern computer banks to handle. So in place of a complete survey, we must content ourselves with a statistical sample of the wavefunction. This introduces errors that can be estimated by the standard techniques of statistics.

For optimal use of resources, one should bring all the important sources of error to the same level. This involves a delicate balancing act. For example, using larger volumes or smaller quark masses requires lengthier calculations, which degrade the sampling rate of the wavefunction. The technical feat of Dürr *et al.*[1] is to achieve such a balance, keeping all the errors demonstrably small.

Of course, overwhelming evidence for the validity of QCD has been accumulating for decades, from very different sorts of calculations and experiments. Although quarks and gluons do not exist as isolated particles, they can be reconstructed from the patterns of energy–momentum flow they imprint on hadrons. In high-energy collisions, the emerging hadrons are found to be organized into jets of particles moving in approximately the same direction as each other. According to QCD, if we replace the jets by fictitious single particles with the same total energy and momentum as the jets, those fictitious particles will obey the equations of elementary quarks and gluons. This is another aspect of asymptotic freedom. Through the study of jets, the basic equations of QCD have been verified in exquisite detail.

So what value is added by using already-validated equations to compute already-measured hadron masses? One answer is practical. The same techniques that are used to compute known hadron masses can also be used to compute other interesting quantities that are very difficult to measure experimentally. For example, some key reactions involving small nuclei and unstable particles (hyperons) are very important in stellar nucleosynthesis and supernova dynamics, but are impracticable to measure. Having numerical techniques that reliably reproduce what is known, we can address the unknown confidently.

But perhaps a more profound answer is philosophical. A great vision of science — stretching from Pythagoras' credo "All things are number", to Kepler's ordering of the planets based on Platonic solids, to Wheeler's slogan "Its from bits" — has been that physical reality embodies ideally simple mathematical laws. As physics developed before the quantum revolutions of the twentieth century, the basic equations emphasized dynamics (how given systems evolve in time) as opposed to ontology (the science of what exists). Kepler's system was stillborn, but in the world of QCD and hadrons, the great vision lives and thrives.

Finally, let me add a note of critical perspective. The accurate, controlled calculation of hadron masses is a notable milestone. But the fact that it has taken decades to reach this milestone, and that even today it marks the frontier of ingenuity and computer power, emphasizes the limitations of existing methodology and challenges us to develop more powerful techniques. QCD is far from being the only area in which the challenge of solving known quantum equations accurately is crucial. Large parts of chemistry and materials science pose similar mathematical challenges. There have been some remarkable recent developments in the simulation of quantum many-body systems, using essentially new techniques[5]. Can the new methods be brought to bear on QCD? In any case, it seems likely that future progress on these various fronts will benefit from cross-fertilization. The consequences could be enormous. To quote Richard Feynman[6]: "Today we cannot see whether Schrödinger's equation contains frogs, musical composers, or morality — or whether it does not. We cannot say whether something beyond it like God is needed, or not. And so we can all hold strong opinions either way." ■

Frank Wilczek is in the Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142, USA.
e-mail: wilczek@mit.edu

1. Dürr, S. *et al. Science* **322,** 1224–1227 (2008).
2. http://pdg.lbl.gov
3. Aubin, C. *et al. Phys. Rev. D* **70,** 094505 (2004).
4. Wilczek, F. in *Les Prix Nobel* 100–124 (Almqvist & Wiesell Int., 2004).
5. Verstraete, F. & Cirac, J. I. Preprint at http://arxiv.org/abs/cond-mat/0407066 (2004).
6. Feynman, R., Leighton, R. & Sands, M. in *The Feynman Lectures on Physics* Vol. 2, Ch. 41, 12 (Addison-Wesley, 1964).

PALAEONTOLOGY

Turtle origins out to sea

Robert R. Reisz and Jason J. Head

Various aspects of turtle evolution are the subject of vigorous debate among vertebrate palaeontologists. A newly described fossil species, the oldest yet discovered, adds grist to the mill.

During the Late Triassic, some 220 million years ago, primitive turtles about 40 centimetres in length were preserved in sedimentary deposits in what is now southwestern China. These fossils are examples of a new species of a very early turtle, named *Odontochelys semitestacea*, which is described by Li *et al.* on page 497 of this issue[1] and which will change ideas about turtle origins and the evolution of their striking body plan.

Turtles are remarkable animals[2]. They have a horny beak rather than teeth, and a shell like that of no other animal, one that is composed of an upper carapace and a lower plastron, jointed together by a bony bridge. The shell is a composite structure derived from ribs, parts of the shoulder girdle and specialized dermal bones. This precludes the typical costal respiration of tetrapods, in which movable ribs allow the chest cavity to expand and contract. Turtles have overcome this obstacle by having the muscles that control breathing use the limb pockets at the borders of the shell. This shell has become modified as turtles diversified and adapted to terrestrial, amphibious and aquatic environments (Fig. 1). The evolutionary relationships and ecology of turtles through time, and the developmental and evolutionary origins of the shell, are major controversies in studies of vertebrate evolution.

Previously, the fossil evidence for turtle origins came largely from *Proganochelys quenstedti* from Germany, which lived between 204 million and 206 million years ago, and other less-well-known early turtles. *Proganochelys* is known from several skeletons. It has a massive shell and spiked armour on the neck and tail, but also retains teeth on the roof of the mouth and has other primitive features in the skull and skeleton. Its osteology has been used to propose[3] that turtles are related to pareiasaurs, a group of extinct parareptiles that includes species with extensive dermal armour. And on the basis of evidence from *Proganochelys*