lecting the data, the GP-B scientists struggled to understand and model the distribution of the charge patches and the effects caused by those charges. Once they imposed the corrections implied by their complicated model, the data from the four gyros were seen to be mutually consistent, as shown in figure 2. The cost of all that complexity is the large error bars in the GP-B results.

Not as easy as it sounds

Everitt and company had hoped to measure the geodetic precession to an accuracy of 0.01% and the frame-dragging precession to 1%. Had they succeeded, they would have achieved by far the most precise tests of those general-relativistic effects. However, as early as 1996, lunar-ranging data had confirmed the geodetic effect to 0.7%, and improved lunar-ranging experiments are currently under way. In 2004, just after GP-B began taking data, Ignazio Ciufolini and Erricos Pavlis announced the results of their laser-ranging experiments on two artificial *LAGEOS* satellites. They reported confirming the frame-dragging effect to 10%, but they had to assume a priori the correctness of the general-relativistic geodetic effect. A third satellite, the Italian Space Agency's *LARES*, is scheduled for launch later this year. It may enable an order-of-magnitude improvement in the precision with which frame dragging can be tested.

Fairbank remarked early on, "No mission could be simpler than GP-B: It's just a star, a telescope, and a spinning sphere." But conceptual simplicity is no guarantee that a cutting-edge experiment will come off as anticipated. "Did I have the slightest idea," muses Everitt, "that it was going to take 50 years? No. If I had known it would be 50 years, would I have done it? Probably not. But one has never been bored."

Steven K. Blau

Reference

1. C. W. F. Everitt et al., *Phys. Rev. Lett.* **106**, 221101 (2011).

Ultracold Bose gases deviate from the textbook picture

Interactions among identical bosons remove the upper bound on the number of particles in excited states.

Statistical-physics textbooks typically present Bose–Einstein condensation in a way that's similar to Albert Einstein's original description. In a gas of identical bosons, they say, the statistical weighting of each state is such that the total occupation of the excited states is capped at a critical value N_c, which depends on temperature. If the temperature is decreased, or the particle number is increased, so that the number of identical particles exceeds N_c, the thermally excited portion of the system becomes saturated, and all additional particles must occupy the ground state.

The textbook picture doesn't include interparticle interactions, but it does assume thermal equilibrium, which cannot exist without interactions of some form. It's no surprise, therefore, that the picture doesn't describe real systems exactly. But how far wrong is it? Physics is full of simplified descriptions that neglect some of the complexities of real systems but are still close enough to be applicable in many typical cases. Is the saturation picture of Bose–Einstein condensation one of them?

Zoran Hadzibabic and colleagues at Cambridge University have now shown that it's not.[1] Even when the interparticle interactions are moderate, the thermal component of the gas continues to grow far beyond N_c, and the condensate's population is only about half what the saturation picture says it should be. By tuning the interaction strength, the researchers showed that as the system converged to the noninteracting limit, the saturation result was recovered. Since a noninteracting Bose–Einstein condensate (BEC) is itself not physically realizable, the demonstration of convergence constitutes the first experimental verification of Einstein's description.

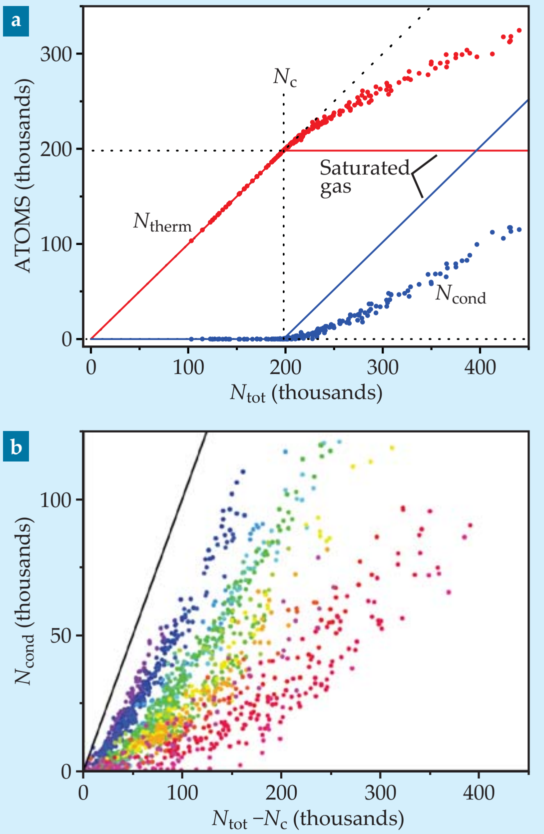

Figure 1. (a) The textbook picture (red and blue lines) of a Bose–Einstein condensate at constant temperature says that the number N_{therm} of thermally excited atoms cannot increase beyond the critical number N_c. So when the total atom number N_{tot} exceeds N_c, all the excess atoms would enter the condensate. But in typical experimental data (red and blue dots) for a potassium-39 system, the number N_{cond} of condensed (ground-state) atoms increases more slowly than predicted, and N_{therm} increases far beyond N_c. **(b)** Experimental data for 18 combinations of temperature (T ranging from 115 nK to 284 nK) and interaction strength (scattering length a ranging from 40 to 356 Bohr radii). Bluer colors represent lower values of a and T, and show the approach to the saturation picture (black line). (Adapted from ref. 1.)

Inspiration from flatland

Studies of interparticle interactions in BECs are not new. Ever since the first experimental demonstration of Bose–Einstein condensation in 1995, some researchers have been exploring the properties of the condensed component, which are entirely controlled by interactions. Others have shown that, thanks to interactions, the critical number N_c (or the critical temperature T_c for constant particle number) isn't quite what the textbook picture would predict.[2]

But the question of whether the thermal component actually becomes saturated for particle numbers above N_c has received much less attention.

Hadzibabic became interested in the saturation of Bose gases while, as a postdoc with Jean Dalibard at the École Normale Supérieure in Paris, he was studying two-dimensional systems of ultracold atoms (see PHYSICS TODAY, August 2006, page 17). Einstein's statistical argument for saturation doesn't work in two dimensions, but there is still a phase transition (called a Berezinskii-Kosterlitz-Thouless transition) to a superfluid state. The Paris researchers' plan was to look side by side at the 2D system (which, as expected, did not show saturation) and a 3D gas (which they expected would saturate). But to their surprise, the 3D system was not saturated either.

"At the time, we could not make sense of it theoretically," recalls Hadzibabic. "And we also could not prove anything experimentally because we were working with rubidium-87, so we had no access to a Feshbach resonance," which would have allowed them to tune the interatomic interaction strength with an applied magnetic field. (For more on Feshbach resonances, see the Reference Frame by Daniel Kleppner in PHYSICS TODAY, August 2004, page 12.) Although ^{87}Rb does have Feshbach resonances, they're all either too narrow or at too high a magnetic field to allow practical control over the interactions.

Another alkali atom isotope, potassium-39, does have a convenient Feshbach resonance. But ^{39}K is much harder to cool to BEC temperatures, in part because of its tiny scattering length a, a quantity related to the square root of the scattering cross section extrapolated to zero collision energy. Applying a field near the Feshbach resonance can increase a, but that works only if the atoms are confined to an optical trap, for which they must already be quite cold. But the Feshbach resonance was so appealing that Hadzibabic, who by that time had joined the faculty at Cambridge, worked with his group to create one of the only ^{39}K-condensing machines in the world. (The first such apparatus was created by Massimo Inguscio and colleagues to study Anderson localization; see the article by Alain Aspect and Inguscio in PHYSICS TODAY, August 2009, page 30.) They achieved the initial cooling by using ^{87}Rb to sympathetically cool the ^{39}K atoms.

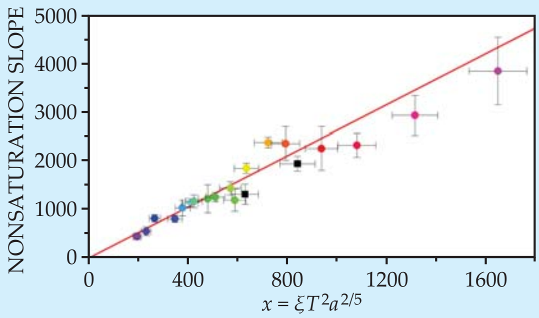

Figure 2. Convergence to the noninteracting limit. Hartree–Fock theory predicts that the nonsaturation slope—a measure of the growth of the thermal atom number—should be proportional to the dimensionless parameter x, which depends on temperature T and scattering length a. That theory is borne out by experiment: The 18 colored points correspond to the 18 potassium-39 data series in figure 1b, and the two black points represent rubidium-87 data at two different temperatures. The red line, the best linear fit to the data, passes through the origin within experimental uncertainty and thus confirms Einstein's picture of a saturated gas in the noninteracting limit. (Adapted from ref. 1.)

Unsaturated gas

To vary the number of atoms in their ^{39}K gas, Hadzibabic and colleagues hold the gas in an optical trap for up to two minutes. During that time, some atoms escape (through collisions with the background gas and other processes) and the rest re-equilibrate. The researchers then release the gas from the trap and allow it to freely expand; the cloud of thermal atoms expands faster than the condensate, so the researchers can segregate the two components and determine their respective populations. Repeating that process some 100 times, varying the holding time while keeping the temperature and magnetic field constant, they could obtain a plot like the one in figure 1a. When the temperature is kept stable enough (within a few nanokelvin) over the course of the experiment, the data clearly show that the thermal population increases beyond the expected point of saturation.

Why would an ultracold Bose gas in an optical trap have more thermal atoms than the saturation picture predicts? Qualitatively, the researchers attribute the effect to repulsion between the thermal and condensed components, which turns the effective potential for the thermal atoms from a harmonic well into a Mexican hat. (The Mexican hat shape stems from the statistics of identical bosons, which cause the condensed atoms to interact twice as strongly with the thermal atoms as they do with one another. Were that not the case, the effective potential in the region of the condensate would be flat.) As a result of the repulsion, the thermal atoms spread out over a larger volume, and additional excited states become thermally accessible. The larger the population of the condensate, the more the thermal atoms spread out.

Figure 1b shows similar data series for 18 sets of experimental conditions, with bluer colors representing lower temperatures and shorter scattering lengths. As either the temperature or the scattering length decreases, the population of the condensate appears to approach the saturated-gas prediction shown by the solid black line, but that qualitative observation doesn't prove that the system actually converges to the saturation picture at $T = 0$ or $a = 0$. For temperature, though, the convergence is trivial: At zero temperature, by definition, there is no thermal component, so $N_c = 0$ and the population of the condensate equals the total number of atoms in the gas.

To the limit

To figure out which numbers they should look at to quantify the convergence to the noninteracting limit, the researchers simulated the gas using Hartree–Fock theory. Hartree–Fock is a mean-field theory in that it treats each particle as moving in the potential energy landscape created by all the others. It's inexact because it neglects interparticle correlations, but it provides a pretty good approximation of many quantities for many systems (including the electrons in molecules; see the article by Martin Head-Gordon and Emilio Artacho in PHYSICS TODAY, April 2008, page 58).

According to Hartree–Fock theory, for atom numbers greater than N_c, the thermal population N_{therm} should be linearly related to the condensed population N_{cond} raised to the ⅖ power. Furthermore, the slope of that line (which the researchers called the nonsaturation slope, since the slope for a saturated gas is zero) should be directly proportional to $x = \xi T^2 a^{2/5}$, where ξ is a combination of system parameters that makes x dimensionless. As a goes to zero, therefore, the nonsaturation slope should also go to zero, and the saturation picture should be recovered.

For small condensates (with atom numbers just above N_c), the data matched the Hartree–Fock predictions well. For larger condensates, the plot of N_{therm} versus $N_{\text{cond}}^{2/5}$ strayed from the Hartree–Fock linear relationship, but the researchers were able to redefine the nonsaturation slope, and they found, as shown in figure 2, that it was still proportional to x. Even two data points taken with ^{87}Rb condensates at different temperatures (remember, the scattering length for ^{87}Rb can't be changed) lie on the line.

Although the researchers observed that the scaling behavior is universal with respect to temperature and atomic species, it's not universal with respect to trap geometry: The amount by which the thermal atoms can spread out—and whether they can spread out at all—is expected to depend critically on the size, shape, and harmonic nature of the optical trap. One of the next steps the researchers hope to take is to look at the effect of different trapping potentials, especially disordered potentials that might change the effective dimensionality of the system.

Johanna Miller

References

1. N. Tammuz et al., *Phys. Rev. Lett.* **106**, 230401 (2011).
2. J. R. Ensher et al., *Phys. Rev. Lett.* **77**, 4984 (1996); F. Gerbier et al., *Phys. Rev. Lett.* **92**, 030405 (2004); R. Meppelink et al., *Phys. Rev. A* **81**, 053632 (2010).

Current-driven magnetic domain walls gather speed

The key, according to new experiments, is to house the walls in a sandwich of platinum, cobalt, and aluminum oxide.

Just because a magnet has a net moment—a north pole and a south pole—that doesn't mean that all of its unpaired spins point in the same direction. More likely, several magnetic domains having different orientations will coexist, each separated from its neighbors by thin transition regions called domain walls.

And those domain walls don't always stay put. They shift, for example, when an external magnetic field causes some domains to expand and others to shrink. Domain walls can also be displaced by spin-polarized current. As an

physics update

These items, with supplementary material, first appeared at http://www.physicstoday.org.

High-spin early stars. Stellar nucleosynthesis by exothermic fusion ends with iron, the most tightly bound of all nuclei. Stars produce heavier species by two neutron-capture mechanisms—the rapid "r-process" and the slow "s-process"—and expel them in supernovae or stellar winds. The r-process avails itself of the enormous neutron flux during the supernova to convey nuclei up the atomic-number scale by way of short-lived stages faster than those stages can decay back down. The s-process, making do with weaker neutron fluxes inside quiescent stars, requires long-lived intermediate stages. It was long thought that the s-process occurs only in stars of modest mass. But now Cristina Chiappini (Leibniz Institute for Astrophysics, Potsdam, Germany) and coworkers have reported anomalously high abundances of two s-process products—yttrium and strontium—in spectra of a cluster of Milky Way stars so old that their Y and Sr must have been created in an early generation of very massive stars that exploded less than a billion years after the Big Bang. Yet model calculations of the s-process in early massive stars yield far too little Y and Sr to account for the overabundances—unless those stars were spinning very fast. Chiappini and company suggest that first-generation stars were not only very massive, as is generally thought, but also very rapidly spinning, with surface speeds as high as 800 km/s. They calculate that the resultant centrifugal mixing of interior layers boosted the s-process's efficiency by the requisite four orders of magnitude. The spectra were taken with the European Southern Observatory's Very Large Telescope, shown here. (C. Chiappini et al., *Nature* **472**, 454, 2011.)

—BMS

Optical pump-probe diagnosis for melanoma? Whether rosy or rich, the color of human skin comes from a pigment called melanin, which occurs in two forms: Pheomelanin acts as a photosensitizer for UV radiation, while eumelanin has a protective role against light and seems to be overabundant in melanoma, a dangerous form of skin cancer. Melanoma is typically diagnosed by removing a lesion and examining it microscopically, layer by layer. Several noninvasive imaging techniques have been explored to detect the cancer, but they lack the molecular specificity needed to distinguish the two melanins. Enter researchers from Duke University who address that shortcoming with a nonlinear optical pump–probe technique. They excite the molecules with an ultrafast laser pulse and, after a varying time delay, measure the absorption of a second pulse—the probe. The contrast between the two melanins arises from their different time-delay

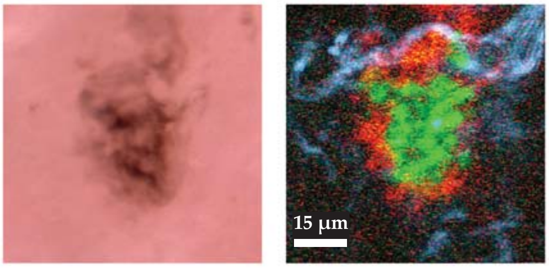

absorption profiles. The white-light image on the left shows a melanoma lesion on human skin grafted onto a live mouse. The false-color image on the right, taken at a depth of 45 μm with the