the size of memristance effects increases as the inverse square of device size. Strukov *et al.*[2] use a simple model to show how memristance arises naturally in a nanoscale system when electronic and atomic transport are coupled under an external voltage. The authors realize this memristive system by fabricating a layered platinum–titanium-oxide–platinum nanocell device. Here, the hysteretic current–voltage characteristics relate to the drift back and forth of oxygen vacancies in the titanium oxide layer driven by an applied voltage[4].

This observation provides a wonderfully simple explanation for several puzzling phenomena in nanoscale electronics: current–voltage anomalies in switching; hysteretic conductance; multiple-state conductances (as opposed to the normal instance of just two conductance states, ON and OFF); the often mischaracterized 'negative differential resistance', in which current decreases as voltage increases in certain nanoscale two-terminal devices; and metal–oxide–semiconductor memory structures, in which switching is caused by the formation and breakdown of metal filaments owing to the movement of metal atoms under applied bias.

But what of Moore's Law? Established by Intel co-founder Gordon Moore in 1965, this empirical rule states that the density of transistors on a silicon-based integrated circuit, and so the attainable computing power, doubles about every 18 months. It has held for more than 40 years, but there is a sobering consensus in the industry that the miniaturization process can continue for only another decade or so.

The memristor might provide a new path onwards and downwards to ever-greater processor density. By fabricating a cross-bar latch, consisting of one signal line crossed by two control lines[5], using (two-terminal) memristors, the function of a (three-terminal) transistor can be achieved with different physics. The two-terminal device is likely to be smaller and more easily addressable than the three-terminal one, and more amenable to three-dimensional circuit architectures. That could make memristors useful for ultra-dense, non-volatile memory devices.

For memristor memory devices to become reality, and to be readily scaled downwards, the efficient and reliable design and fabrication of electrode contacts, interconnects and the active region of the memristor must be assured. In addition, because (unlike with transistors) signal gain is not possible with a memristor, work needs to be put into obtaining high resistance ratios between the ON and OFF states. In all these instances, a deeper understanding of the memristor's dynamic nature is necessary.

It is often the simple ideas that stand the test of time. But even to consider an alternative to the transistor is anathema to many device engineers, and the memristor concept will have a steep slope to climb towards acceptance. Some will undoubtedly trivialize the realization of this ubiquitous nanoscale concept, whereas others will embrace it only after the demonstration of a well-functioning, large-scale array of these densely packed devices. When that happens, the race towards smaller devices will proceed at full steam. ■

James M. Tour and Tao He are in the Departments of Chemistry, Computer Science, Mechanical Engineering and Materials Science, and the Smalley Institute for Nanoscale Science and Technology, Rice University, Houston, Texas 77005, USA.
e-mail: tour@rice.edu

1. Chua, L. O. *IEEE Trans. Circuit Theory* **18,** 507–519 (1971).
2. Strukov, D. B., Snider, G. S., Stewart, D. R. & Williams, R. S. *Nature* **453,** 80–83 (2008).
3. Chua, L. O. & Kang, S. M. *Proc. IEEE* **64,** 209–223 (1976).
4. Yang, J. J. *et al. Nature Nanotech.* (in the press).
5. Kuekes, P. J., Stewart, D. R. & Williams, R. S. *J. Appl. Phys.* **97,** 034301 (2005).

CLIMATE CHANGE

Natural ups and downs

Richard Wood

The effects of global warming over the coming decades will be modified by shorter-term climate variability. Finding ways to incorporate these variations will give us a better grip on what kind of climate change to expect.

Climate change is often viewed as a phenomenon that will develop in the coming century. But its effects are already being seen, and the Intergovernmental Panel on Climate Change recently projected that, even in the next 20 years, the global climate will warm by around 0.2 °C per decade for a range of plausible greenhouse-gas emission levels[1]. Many organizations charged with delivering water and energy resources or coastal management are starting to build that kind of warming into their planning for the coming decades. A confounding factor is that, on these timescales, and especially on the regional scales on which most planning decisions are made, warming will not be smooth; instead, it will be modulated by natural climate variations. In this issue, Keenlyside *et al.* (page 84)[2] take a step towards reliably quantifying what those ups and downs are likely to be.

Their starting point is the ocean. On a timescale of decades, this is where most of the 'memory' of the climate system for previous states resides. Anomalously warm or cool patches of ocean can be quite persistent, sometimes exchanging heat with the atmosphere only over several years. In addition, large ocean-current systems can move phenomenal amounts of heat around the world, and are believed to vary from decade to decade[3,4].

To know and predict the state of the ocean requires an approach similar to weather forecasting: one sets up (initializes) a mathematical model of the climate system using observations of the current state, and runs it forwards in time for the desired forecast period. With a given climate model, enough observations to set the ball rolling and a large-enough computer to move it onwards, the exercise is conceptually straightforward.

But does it actually produce anything useful? We don't expect to be able to predict the details of the weather at a particular time several years in the future: that kind of predictability runs out after a week or two. But even predicting, say, that summers are likely to be unusually wet during the coming decade would be useful to many decision-makers. Only recently, with the study from Keenlyside *et al.*[2] and another

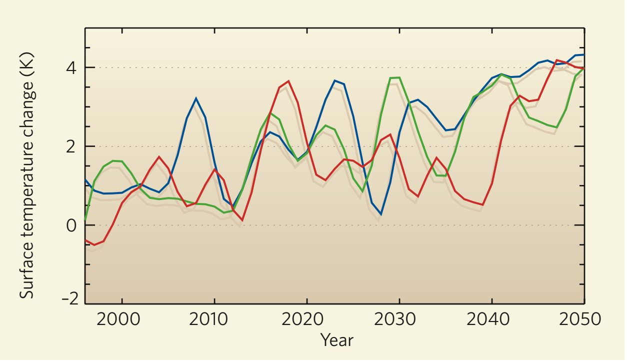

Figure 1 | Heat up? These three possible trends of winter temperature in northern Europe from 1996 to 2050 were simulated by a climate model using three different (but plausible) initial states[6]. The choice of initial state crucially affects how natural climate variations evolve on a timescale of decades. But as we zoom out to longer timescales, the warming trend from greenhouse gases begins to dominate, and the initial state becomes less important. Keenlyside and colleagues[2] use observations of the sea surface temperature to set the initial state of their model. Their results indicate that, over the coming decade, natural climate variability may counteract the underlying warming trend in some regions around the North Atlantic. (Figure courtesy of A. Pardaens, Met Office Hadley Centre).

from researchers at my own institution[5], have climate modellers begun to explore whether such predictions are possible.

Keenlyside and colleagues' model[2] uses a very simple ocean initialization method in which they add heat to or remove it from the ocean surface until sea surface temperatures across the globe are close to observed values. They use their model to produce a set of retrospective 'forecasts' starting from earlier states, which they test against what actually happened. Their system produces refined temperature predictions a decade ahead for large parts of Europe and North America.

The enhanced accuracy of the model seems to stem from a greater ability to simulate natural variations in the meridional overturning circulation (MOC). This is a giant conveyor belt that brings warm water northwards into the North Atlantic, releases its heat to the atmosphere, and returns the cooled water to the south. There is evidence that the strength of this circulation can fluctuate naturally over periods of decades[3]; when it is strong, the climate in the North Atlantic region passes through a warm phase.

The authors use their model to predict that the MOC will weaken over the next decade, with a resultant cooling effect on climate around the North Atlantic. Such a cooling could temporarily offset the longer-term warming trend from increasing levels of greenhouse gases in the atmosphere. That emphasizes once again the need to consider climate variability and climate change together when making predictions over timescales of decades (Fig. 1).

These results provide encouragement that such predictions may be possible, but substantial points require clarification. Chief among these is whether the authors' initialization, which takes into account only sea surface temperatures, is in fact suitable to characterize the state of the MOC. The MOC extends to a depth of several kilometres and depends not just on temperature, but on the salinity of the ocean water. The answer will be some time in coming, as regular monitoring of the MOC has only just begun[4].

In addition, there is the fact that, although it seems to improve predictions around the North Atlantic compared with models that have no initialization, the model's accuracy is less over other regions such as the tropical North Atlantic and central Africa. This might be because of deficiencies in the climate model or in the initialization procedure. The two studies that have attempted decadal-scale modelling to date[2,5] differ in the regions for which their predictions are most accurate, so clearly there is much still to be understood.

Keenlyside *et al.*[2] have so far used their model to predict decadal average temperature only. If it can be shown to accurately forecast other variables such as precipitation, as well as their variation during specific seasons, its range of applications would be greatly increased.

But climate predictions for a decade ahead will always be to some extent uncertain. There are inherent limits to the predictability of climate variability, and unpredictable perturbations to the climate, such as volcanic eruptions, cannot be factored in; predictions must inevitably be probabilistic in nature.

Even so, the first attempts at decadal prediction[2,5] suggest that reasonably accurate forecasts of the combined effects of increasing greenhouse-gas concentrations and natural climate variations can be made. Climate scientists are gearing up to test and extend these ideas over the coming years, in the hope that we can then plan more confidently for the future. ■

Richard Wood is at the Met Office Hadley Centre, FitzRoy Road, Exeter, Devon EX1 3PB, UK.
e-mail: richard.wood@metoffice.gov.uk

1. Solomon, S. D. *et al.* (eds) in *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change* (Cambridge Univ. Press, 2007).
2. Keenlyside, N. S., Latif, M., Jungclaus, J., Kornblueh, L. & Roeckner, E. *Nature* **453,** 84–88 (2008).
3. Knight, J. R., Allan, R. J., Folland, C. K., Vellinga, M. & Mann, M. E. *Geophys. Res. Lett.* **32,** doi:10.1029/2005GL024233 (2005).
4. Cunningham, S. A. *et al. Science* **317,** 935–938 (2007).
5. Smith, D. M. *et al. Science* **317,** 796–799 (2007).
6. Stott, P. A. *et al. Science* **290,** 2133–2137 (2000).

QUANTUM PHYSICS

The squeeze goes on

Eugene S. Polzik

After 20 years of hard labour, squeezed states — light and matter whose quantum fluctuations have been arduously suppressed below standard levels of quantum noise — are coming of age and are ripe for application.

Light is a quantum-mechanical object. One of the manifestations of that fact is that its amplitude (roughly, the number of photons present) and phase cannot be known simultaneously with arbitrary accuracy. Take, for example, an ideal laser pulse containing a mean number of photons N. The actual number of photons present in a pulse will obey poissonian statistics, with an uncertainty $\sqrt{N}$. Obeying the diktat of Heisenberg's uncertainty principle, the variance in the pulse's phase will accordingly vary as $1/\sqrt{N}$. These fluctuations form the characteristic 'shot noise' of light, and are the consequence of the fact that photons are statistically uncorrelated objects, fully independent of one another.

Because the Heisenberg principle puts the limit only on the product of the variances of amplitude and phase in a pulse of light, it is possible to reduce one or other of them beneath the shot-noise values. The catch is that the photons must be made correlated through some kind of interaction with matter. Three recent papers[1–3] present the latest advances in this technique of 'squeezing' light, and a fourth[4] tells us why it matters. Following a much earlier proposal[5], this fourth paper describes a sensitivity-enhanced detector for gravitational waves — as-yet hypothetical tiny disturbances in space-time — that exploits squeezed-light states. This application is an illustration of how squeezed states have become central to

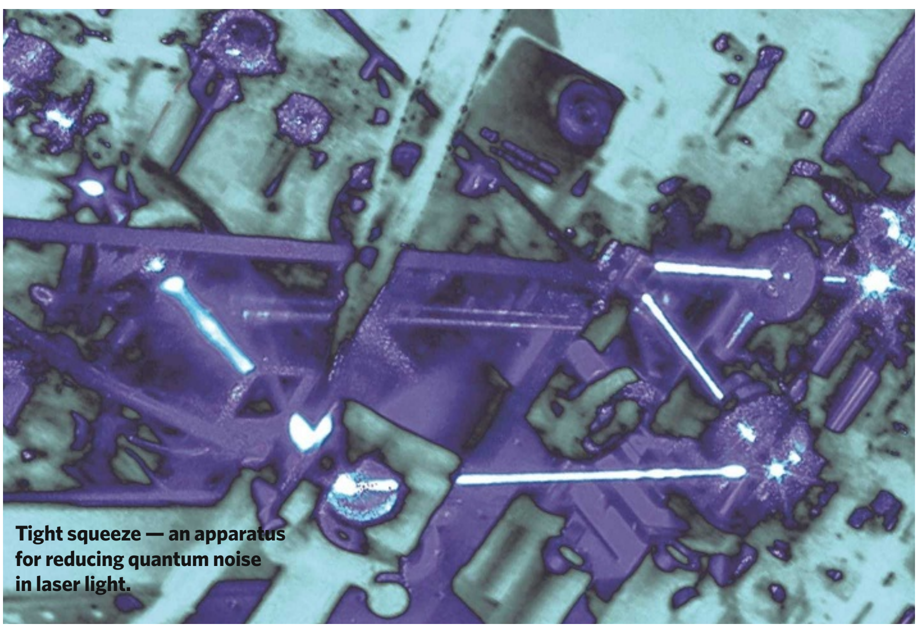
Tight squeeze — an apparatus for reducing quantum noise in laser light.

M. HØST, NIELS BOHR INST.