

Deposit Insurance

Chips are made inside sealed vacuum chambers, and controlling what goes on within is still to some degree a black art.

The silicon and gallium arsenide chips that carry integrated circuits or optical devices are really more like nachos. The basic slab—the “substrate”—is layered with multiple thin films of assorted semiconductors to make the finished circuit, just as tortilla chips are slathered with refried beans, sour cream, guacamole, cheese, and black olives to make nachos. The recipe for nachos is largely a matter of individual taste and what’s available in the fridge. Chip recipes, however, call for a very specific composition for each layer—five germanium atoms per every 95 silicon atoms, for example. But what actually winds up there merely approximates the recipe, and the chip won’t work if that approximation strays too far. Chips are made inside sealed vacuum chambers, and controlling what goes on within is still to some degree a black art. It’s easy to sample the finished product, but it’s far too late to rescue a bad batch by then. Tasting a layer as it is being deposited is a challenge. Most of the methods used on finished chips either won’t work in the chamber’s harsh conditions or get in the way of the deposition process. The methods that have been developed for minding the stove generally don’t give very detailed information, or don’t work fast enough, or work only with certain semiconductors.


A new method, called reflection electron energy-loss spectroscopy or REELS, reveals the composition of any kind of semiconductor layer as it grows, and can be added to a standard chip-making chamber by essentially bolting it on. The method was developed by Harry Atwater, assistant professor of applied physics; Shouleh

Nikzad, a postdoc in Atwater’s group; and Channing Ahn, senior research fellow in materials science and lecturer in materials science. “A lot of very sensitive tools for analyzing thin-film surfaces have been developed in the last 30 years,” says Atwater, “but this one can be done during growth.”


Chips are grown by a process called epitaxy. (Epitaxy means that the atoms in the growing layer align themselves with the crystal structure of the layer below, nesting like rows of Lego blocks. The layer-to-layer alignment is crucial—a typical ultra-large-scale integrated-circuit or optical-device chip has more than a dozen layers on it at any given spot, all of which must be seamlessly stacked for the device to function.) The substrate, a three-inch-diameter wafer, hangs upside down in the high-vacuum chamber. Below the wafer sit hemispheric lumps of pure silicon, gallium, arsenic, or whatever—the source of the layer material. Each lump is roughly the size of a golf ball cut in half. Layer deposition begins after the wafer has been heated up to about 400°C and the chamber pumped down to five-billionths of atmospheric pressure. A heating element on the flat underside of each hemisphere heats the lump until atoms evaporate off of its surface. As the vaporized atoms waft upward, they stick to the substrate positioned just overhead to collect them. If the layer being grown is an alloy of two or more elements, the proper mix of atoms in the vapor is set by adjusting each lump’s heater. But things start to get complicated when the atoms reach the substrate. Some atoms prefer to cling to their fel-

Looking like a mutant fire hydrant dressed up for Christmas, the vacuum deposition chamber stands in Atwater’s lab.

From left: Atwater, Nikzad, and Ahn admire their handiwork. The aluminum foil swaddling the vacuum chamber insulates it.



A schematic look inside the vacuum chamber. The wafer to be coated (A) hangs upside down at the top of the chamber. Stray gas molecules are swept up and stuck to the chamber walls by argon ions emitted by an ion gun (B). Thickness monitors (C), placed in the flight path of the vaporized atoms, vibrate like tuning forks. Some of the atoms deposit themselves on the monitors, increasing their mass, which lowers their resonant frequency by a measurable amount. Back-calculation gives the number of atoms deposited on the monitor, and, eventually, the thickness of the layer deposited on the wafer. A lump of the material to be deposited (D) sits in its heater (E) like a muffin in its pan.




lows instead of mixing uniformly with the other alloy components. Some components, as well as many contaminants, tend to float on the growing surface like the foam on a pitcher of beer as it fills. And some atoms just plain don't stick very well. "As in any sort of process that you're trying to control, there's often a deviation between the surface that you intend to make and the surface that's actually formed," says Atwater. "The surface composition can change at a rate equal to the material's growth rate." With real-time information about what's actually growing, the heaters could be adjusted to compensate. A thin film typically grows a fresh monolayer—a layer one atom thick—every second or so.

Many promising ideas for monitoring the process had been tried before. One method collects and counts the atoms that bounce off the substrate, in order to back-calculate how many atoms stuck. This method works well and can be used on gallium arsenide chips, but not for, for instance, silicon and germanium, whose atoms stick so well that there's not much to collect. The most obvious way to look at something you can't touch is to shine a light on it—spectroscopy. Unfortunately, those portions of the spectrum so beloved of chemists—the infrared, visible, and ultraviolet—are absorbed and emitted by the bonds between atoms rather than the atoms themselves; and photons—particles of light—tend to penetrate too deeply into the crystal, muddling data from the growing surface with "noise" from the interior. Another method uses polarized light: under some circumstances, the plane of polarization changes when the grow-

Like a brick wall leaving its imprint on the face of a cartoon character who runs full tilt into it, an atom reflecting an electron sometimes leaves an identifying mark.

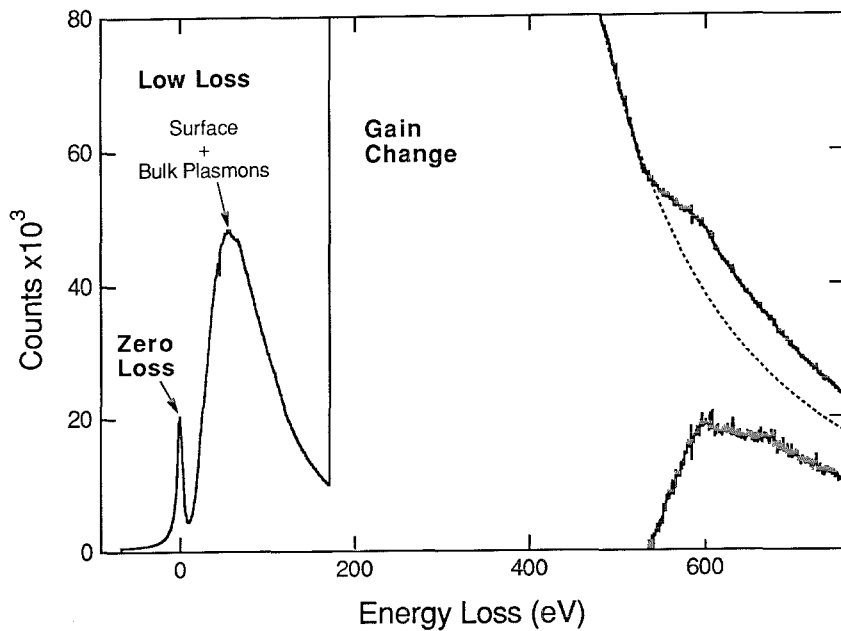
An electron (bold arrow) entering a crystal at a shallow angle can travel some distance before being reflected, yet not penetrate very deeply. Like ranchers in Montana, the valence electrons that define an atom's boundary are spread thinly over a vast region. The core electrons, on the other hand, cluster around the atom's nucleus like suburbanites around the metropolis.



ing crystal reflects the light. Crystals of gallium arsenide, for example, consist of alternating monolayers of gallium atoms and arsenic atoms, and arsenic atoms have a slightly greater affinity for electrons than do gallium atoms. Thus each arsenic atom acquires a very slight negative charge at the expense of a neighboring gallium atom, which becomes equally positively charged. These partial charges balance each other out inside the crystal, but don't always at the growing surface, what with fresh atoms coming in all the time. These locally unbalanced charges, or dipoles, rotate the plane containing the electric component of the electromagnetic wave that we call light. The amount of rotation oscillates through a regular cycle as the crystal fills itself in. One complete cycle corresponds to the deposition of one monolayer each of gallium and arsenide atoms. The way the rotation varies with time shows how fast the crystal is growing, and, indirectly, what its composition is. Unfortunately, the method only works for compound semiconductors such as gallium arsenide. It doesn't work for silicon semiconductors, which are the industry's mainstay. And, even in the best cases, "it's a very difficult measurement, because the polarization change is so small. It's only been shown to be possible at the hands of experts. The casual user probably won't have much success."

REELS uses the next best thing to photons: electrons. A beam of electrons, shining at a shallow angle to the crystal, is both reflected and diffracted, like a beam of light. The diffracted electrons are useful in their own way, but the reflected electrons carry the critical data. Like a brick wall leaving its imprint on the face of a cartoon character who runs full tilt into it, an atom reflecting an electron sometimes leaves an identifying mark. The mark isn't a physical imprint—instead, the electron loses a characteristic amount of energy to the atom.

If all the reflected electrons bore the mark, things would be very easy—a beam of known energy fired into the sample would produce a reflected beam of lower energy, and the energy difference would show what the beam hit. In fact, the reflected electrons cover a whole spectrum of energies, from the input energy—typically 30,000 electron volts (eV)—on down. (One electron volt is the amount of energy imparted to an electron passing through a one-volt electric field.) Plotting the number of electrons versus the amount of energy lost reveals that several things are happening. A sharp spike at zero electron volts shows that a goodly number of the electrons collide "elastically" with the




A typical energy-loss spectrum, in this case from oxygen atoms in a silicon dioxide film on a silicon substrate. The vertical line at "gain change" shows where the recorder's sensitivity was increased by a factor of about 1,000. Even so, the core-loss peak is quite small—it's the difference between the solid line and the dashed line showing the continuation of the plasmon peak's tail. Subtracting the plasmon peak out and doubling the remaining core-loss signal gives the peak in the lower right-hand corner.

atoms, losing no energy at all. The main peak comes a little further down the energy-loss spectrum—somewhere around 20 eV—but, unfortunately, it also contains no useful information. It represents what are called "plasmon modes," in which the outermost, or valence, electrons of many atoms in the crystal get excited simultaneously to shared energy levels. These shared levels don't reveal anything about the individual atoms, but they're where the vast majority of the reflected electrons leave the atoms they hit. Moving further down the spectrum, the marked electrons—the ones with the characteristic energy loss—show up as a tiny bump on the tail of the plasmon peaks. These electrons have lost their energy to the "core electrons" that make up an atom's inner shells of electrons. Core electrons are very tightly bound to their atoms. It takes a lot more energy—hundreds or thousands of electron volts—to excite these electrons. "The excitation energy of a core electron is strongly dependent on the atomic number—the number of protons in the atom's nucleus—which also determines the atom's elemental identity," notes Atwater. "And there is often more than one observable core-level excitation energy for a given element, which allows you to determine its signature even more unambiguously. So there's a built-in cross-check."


The REELS system uses an electron gun mounted inside the chamber to fire a beam of electrons toward the growing crystal at a very shallow angle—about 2 degrees. The detector is 90 degrees' worth of curved, tubular magnet positioned so that the reflected electrons will

Below: A RHEED pattern. Electrons diffracted by atoms within the crystal cause the bright set of vertical lines, while the dim set is caused by electrons diffracted by the exposed atoms on the crystal's surface. The surface atoms do not continue the interior crystal structure exactly, since they don't have neighbors on all sides to hold them in position.





The marked electrons show up as a tiny bump on the tail of the plasmon peaks.

Schematic of the hybrid RHEED-REELS setup. PMT stands for photomultiplier tube. E_0 is the path of the elastically colliding electrons. ΔE is the path of electrons that have lost energy. Si and Ge are silicon and germanium, respectively.




enter it. Once inside the tube, the magnetic field bends each electron's flight path, with the degree of curvature inversely proportional to the electron's energy. An electron having just the right energy for the field strength will coast around the bend, passing through a series of focusing slits to hit a scintillator plate that emits a flash of light on impact. A photomultiplier tube picks up and amplifies this signal, and sends it on to the computer. By starting with a very weak field and then gradually cranking up the power, the REELS system records the number of electrons at all energies.

"We're standing on the shoulders of giants, in a sense, with apologies to Isaac Newton," says Atwater. "There's another method, called RHEED, that is almost a universal tool in this field." RHEED stands for reflection high-energy electron diffraction. As mentioned earlier, some electrons are diffracted by the exposed atoms on the crystal's surface, the way light is diffracted by a grating. The resulting diffraction pattern contains all the information needed to calculate how the atoms that generated the pattern are arranged. "RHEED tells you where atoms are on the surface, but it doesn't give you their chemical identity. Sometimes this is enough information for you to be able to infer the composition, but very often it's not. Suffice it to say that if it were possible to accurately measure surface composition under wide-ranging conditions with RHEED, it would have been done." RHEED also uses an electron gun at a shallow angle to the substrate, but the detector is a phosphorescent screen—like a TV screen—positioned within the chamber in the reflected electrons' path. Each point on the screen lights up with an intensity proportional to the number of electrons hitting that point, making the diffraction pattern visible. A video camera equipped with a frame-grabber converts the pattern to a digital image and sends it to a computer that calculates the atoms' positions. RHEED only records the




Above: Layers deposited on a chip can be seen directly—if belatedly—by doing transmission electron microscopy (TEM) on a slice, a few hundred angstroms thick, cut from the wafer. The two black lumps (above, right) are germanium islands on a silicon substrate. A nanometer (nm) is a billionth of a meter. Below: REELS picks up a noticeable signal from a layer less than a monolayer thick. A germanium monolayer is about 0.3 nm thick.



cumulative number of reflected electrons; their tell-tale energy spectrum is lost.

Atwater and Ahn realized that a hybrid system could run both methods at once. A half-centimeter-diameter hole drilled in the RHEED screen passes all the electrons hitting the screen at that point on to the REELS detector. Moving the RHEED screen—and thus the hole—around allows the energy spectrum for the electrons from any point in the diffraction pattern to be found. “Various features of the RHEED pattern correspond to different structures on the crystal. Some literally come from the first layer of atoms, others come from a very short distance below the surface. So if you focus a particular part of the pattern into the REELS spectrometer, you can infer where those electrons came from.”

The electrons return useful information from a depth of less than a monolayer’s thickness to at least 10 monolayers down—roughly 20 to 30 angstroms (ten-billionths of a meter)—and possibly deeper. REELS can tell how deep the electron got before being reflected, allowing the experimenters to discriminate between the monolayers. “If the chemical makeup of the third monolayer is very different from the fourth one—say, if monolayer three is germanium atoms and monolayer four is silicon—we can identify each monolayer unambiguously. If the monolayers vary gradually, then it’s somewhat harder, but we can still tell approximately how deep we’re looking. It surprised us that the physics was as simple as it turned out to be. We assumed—correctly—that most electrons reflect once, lose energy, and come back out.



The REELS detector. The cylinder at the upper left is the photomultiplier tube, with the curved magnet visible below it.

The mean free path that an electron at our incident energy will travel before it hits something is well known from other experiments. We estimated what the odds of our electron reflecting should be as a function of the film's thickness, and it's a beautiful fit to what we observed.

"This is essentially a material-independent technique. We've looked at gallium arsenide layers as well as silicon. Typically, of course, we're using this tool in an environment where we have a pretty good idea of what we're depositing; we just don't know exactly how much. And one always has to be wary of unintended contaminant species. This method, in fact, has shown a depressing level of sensitivity. It's been a real eye-opener."

Their current setup uses a borrowed magnet-sector spectrometer—the curved magnet—as its detector, and a conventional RHEED gun as the beam source. "Nothing has been optimized yet. We're working at a much lower beam energy than the magnet was designed for—we had to modify the power supply to keep the beam from hitting the magnet's side walls, and that wasn't trivial. Nor is this the world's fastest spectrometer—it takes us a few seconds to collect a complete spectrum. This idea is a very modest extension of existing hardware, so it should be very easy to adopt, but it's also a very powerful method. From what little we've seen so far, it's the most promising technique for watching surface growth." Adds Ahn, "It's so simple that everybody we've talked to is a little bit amazed that no one had tried it before." □—DS

"This idea is a very modest extension of existing hardware, so it should be very easy to adopt, but it's also a very powerful method. From what little we've seen so far, it's the most promising technique for watching surface growth."

"We will demonstrate how to expand your mind, using electrons."