The Cosmic Web

Observations and simulations of the intergalactic medium reveal the largest structures in the universe

Robert A. Simcoe

There is no such thing as empty space. The idea of absolute emptiness realizes its closest approximation in the barren expanses between the stars and the galaxies, but even the most remote corners of the universe are suffused with very low density gas—which becomes increasingly rarefied as one ventures farther away from the places where galaxies consort. Consider this fact: In the air we breathe, each cubic centimeter contains roughly 5×10^{19} atoms. In contrast, the intergalactic medium has a density of only 10^{-6} particles per cubic centimeter—each atom inhabits a private box a meter on each side. This would seem to suggest that there is not much matter in the intergalactic medium. But, given the enormous volume between the galaxies, it quickly adds up: The combined atomic mass of intergalactic gas exceeds the combined atomic mass of all the stars and galaxies in the universe—possibly by as much as 50 percent! There is indeed something in empty space.

As cosmologists construct new narratives of the universe's evolution from its beginning–the Big Bang–to the present day, it is becoming clear that we must understand the physics of intergalactic matter if we are to write the history of how the galaxies, stars and planets formed. In the past decade, rapid advances in both the design of telescopes and computing power have allowed us to study the remote corners of intergalactic space in unprecedented detail. These new results deepen our understanding of how the grandest structures in the universe formed and evolved.

Robert A. Simcoe is a Pappalardo Fellow in Physics at the Massachusetts Institute of Technology. His research on galaxies and intergalactic matter in the early universe often takes him to the twin 6.5-meter Magellan telescopes in the Chilean Atacama desert. He has also worked on the development of instrumentation for large, ground-based observatories, including the Palomar 200-inch Hale telescope. Address: MIT Center for Space Research, 77 Massachusetts Avenue #37-664B, Cambridge, MA 02139. Internet: simcoe@mit.edu

In the Red

Intergalactic gas is so tenuous and dark (producing no light of its own) that you might well ask how astronomers can hope to observe it. The trick is to detect it indirectly, by seeing how it influences light coming from faraway sources. The most common object for these observations is a quasar, a special type of galaxy containing a supermassive black hole at its center. Gas around the black hole emits intense radiation, which often outshines the average galaxy by 100 or more times. Because quasars are so bright, we can observe them at great distances and so measure the effects of intergalactic gas over substantial portions of the universe.

Using the world's most powerful telescopes, we can collect photons from these distant beacons and sort them by their wavelengths into spectra *(Figure 4)*. The strongest feature in such a record is an emission line that is produced by hydrogen atoms near the quasar's black hole. The electrons in these atoms are excited to a single quantum level above their ground state. When they settle back to ground, photons are emitted with the precise wavelength of 121.56701 nanometers—called the Lyman- transition. Yet we observe the emission line at a much longer wavelength, 560 nanometers. This is because the quasar is racing away from us, carried by the general expansion of the universe (see "The Hubble Constant and the Expanding Universe," January–February 2003). The expansion is such that objects far from us recede proportionally faster than those that are close. As an object moves away from us, the light that it emits is stretched to longer wavelengths in much the same way that the Doppler effect lowers the pitch of a receding train whistle. Astronomers use the term *redshift* to describe this phenomenon, since the colors of ever-more-distant objects become systematically redder.

Now consider what happens to the light of a quasar when it is transmitted through the intergalactic medium. As light from the quasar heads toward the Earth, some of its photons will intercept hydrogen atoms along the way. If one of these photons has a wavelength of 121.56701 nanometers, it will be absorbed by the atom, which then has one of its electrons kicked out of the ground state. When the electron loses energy and falls back to the ground state, the photon is re-emitted, but in an arbitrary direction, which is not likely to be toward Earth *(Figure 2)*. So a cloud of hydrogen atoms will absorb light at a very specific wavelength and scatter it away—we see this as a dark "hole" in the spectrum.

The intergalactic medium contains many hydrogen clouds at different distances from us. And because clouds at different distances have different redshifts, a quasar spectrum shows many absorption lines at different wavelengths. The wavelengths below the hydrogen emission line thus appear to be "eaten" away according to the location of each cloud between the quasar and us. In the past decade new instruments on large telescopes have allowed us to examine the spectra of quasars at very-fine-wavelength resolution and high signal-to-noise ratio. These "zoomed-in" views *(Figure 4 bottom)* resolve the intergalactic medium into individual clouds.

Spinning a Cosmic Web

When the absorption lines of quasars were first studied, it was not at all clear how to interpret them, particularly without the benefit of the high-quality data we have today. From the late 1970s to the early 1980s, Wallace Sargent's team at Palomar Observatory made a series of measurements that convinced most astronomers that these absorption lines represent intergalactic matter. However, a number of theoretical explanations were consistent with the available data, and most models explained the lines as clusters of discrete spherical clouds of gas.

In recent years, the advances in observing techniques have been joined by increasingly powerful computer models, which together deliver a more sophisticated picture of the intergalactic medium. This work involves several collaborations and requires months of supercomputer time. In these simulations, an imaginary box is designed to resemble a large representative volume of the universe. The box is divided up into a three-dimensional grid of cells, and matter is distributed throughout the grid in an initial state—according to conditions set by observations of the early universe. All of the physical processes that affect the evolution of the intergalactic medium are dialed into the model. Then the simulation is "turned on," allowing matter and energy to flow from cell to cell in the box, governed simply by the physics. The final product resembles a cosmic time-lapse movie with millions of years compressed into each frame. The computer code examines the distribution of matter in the box at each frame, or *time step*, and calculates the total force acting on each particle to determine where it should move in the next step. At regular intervals, the computer records the density of the gas throughout the intergalactic medium, and these results are

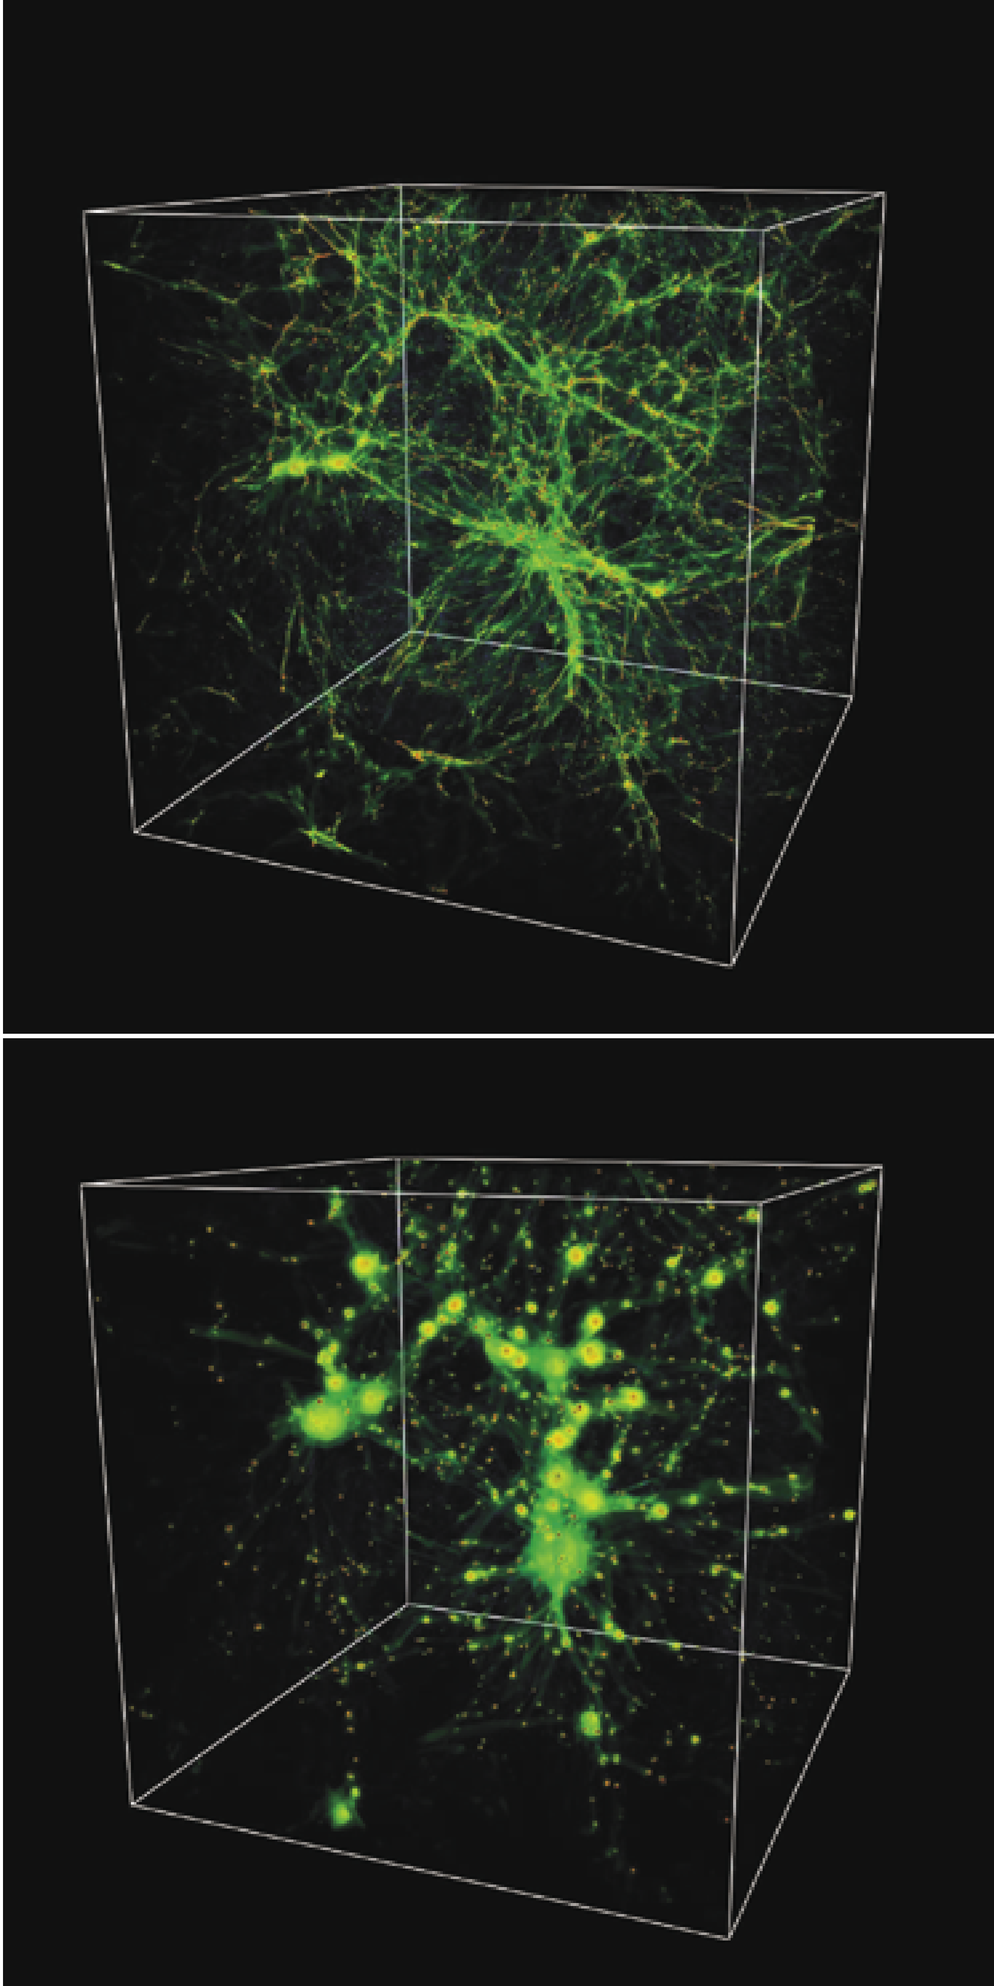

Figure 1. Numerical simulations show the segregation of intergalactic gas into filaments *(green)* and voids *(black)*—the cosmic web. Such computer modeling, paired with telescopic observations, reveals that the cosmic web had a gauzy texture when the universe was about 2 billion years old *(top)* but appears clumpy today *(bottom)*, nearly 12 billion years later. The change is a natural outcome of a "gravitational-runaway" effect, in which slight density enhancements grow progressively larger as they accumulate more and more mass. The author describes what astronomers have learned about the cosmic web and the subtle interactions between galaxies and the intergalactic medium. The top simulation box is about 30 million light-years on a side, whereas the bottom box is four times larger (owing to 12 billion years of cosmic expansion). At the scales depicted here, the Milky Way galaxy would be a tiny fraction of a millimeter across—effectively invisible. (Images courtesy of Renyue Cen, Princeton University.)

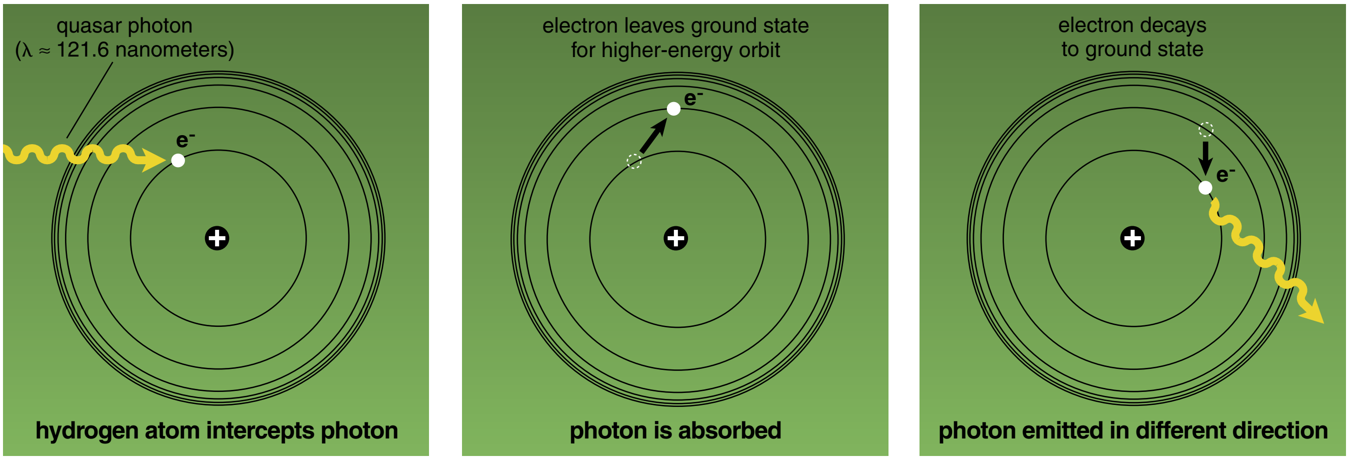

Figure 2. Hydrogen atoms absorb and emit photons at a wavelength (λ) of 121.56701 nanometers, which is characteristic of the energy difference between the first and second electron-orbital levels—the Lyman-α transition. When an atom intercepts a photon from a distant quasar *(left)*, the absorbed energy sends the lone electron into a higher-energy orbit *(middle)*. When the electron decays back into the ground state *(right)*, a photon is emitted, but in a random direction, which is unlikely to be toward the Earth. This explains why clouds of hydrogen gas in the intergalactic medium appear to dim the light of a distant quasar even though they emit as many photons as they absorb.

compared with actual observations of quasar spectra to test the accuracy of the physical models.

One such output, from a simulation run by Jeremiah Ostriker and Renyue Cen of Princeton University, is shown in the top panel of Figure 1. This particular view shows the universe when it was about 15 percent of its present age, or about 2 billion years old. The most striking feature seen is a tendency for gas to collapse into a network of filamentary tendrils that crisscross through vast, low-density voids. This pattern is a common feature of the new computational models and has been nicknamed "the cosmic web."

To test this depiction of the universe against concrete observations, large numbers of artificial quasar spectra are generated by drawing random lines through the simulation box. By evaluating the variations of gas density along any single line, astronomers can calculate the amount of absorption that would be observed in a spectrum measured along that line of sight *(Figure 3)*. It is as though an observer stood on one side of the box and measured the spectrum of a quasar on the other side.

Statistically, the "spectra" from these artificial universes are nearly indistinguishable from the spectra of real quasars. The models accurately predict the number of absorption lines, the distribution of their strengths and widths, and their evolution through time. At a basic level, these models have captured the physical processes that dominate the evolution of the universe on the largest scales.

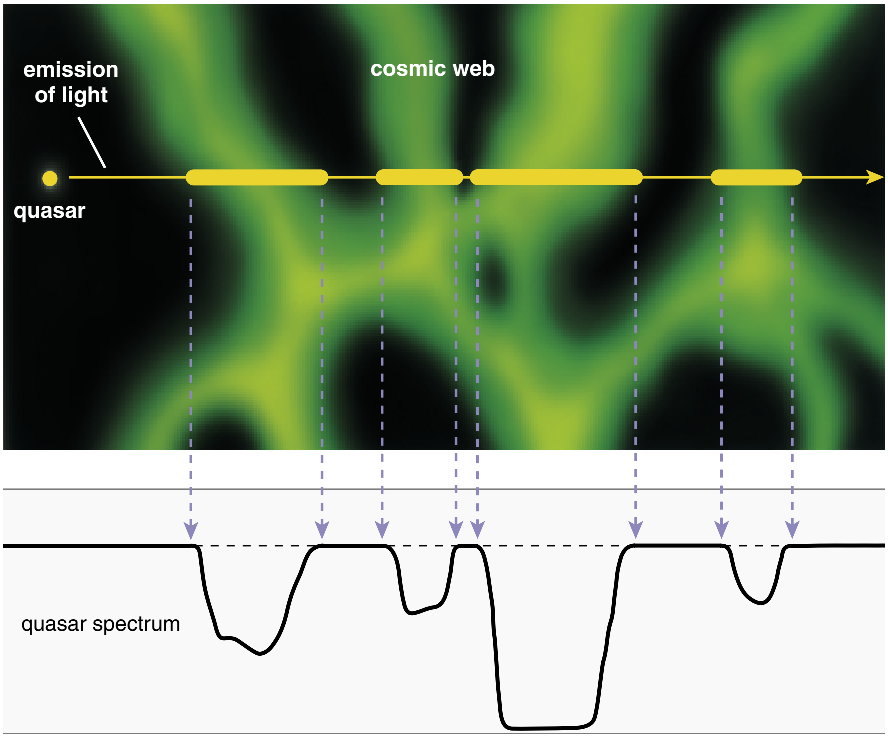

Figure 3. Cosmic web of gas filaments *(top, green)* in the intergalactic medium interrupts the light of a distant quasar before it reaches the Earth, producing the absorption lines in a quasar's spectrum *(bottom)*.

Lumps and All

Technology has thus given us the tools to observe the remote corners of the intergalactic medium and to interpret these observations in the context of a cosmological model. Having described the methods, now let us step back to examine the model itself by offering a narrative that explains the formation of galaxies and intergalactic structure.

The story begins more than 13 billion years ago, roughly 380,000 years after the Big Bang, when the universe was very different from today. There were no stars, galaxies or webs yet, just a uniform soup of free-floating protons and electrons. In fact, the gas was so evenly distributed that its peak densities differed by only 1 part in 100,000 from the cosmic average. But some-time between then and now it evolved into a very lumpy place, where vast stretches of nearly empty space are interrupted by "dense" strands of galaxies and gas. Today, the range of densities is much greater: The difference between the atomic density of the Sun's interior and intergalactic space spans about 32 orders of magnitude!

Astronomers believe that this transition from smooth to lumpy was driven by gravity. Imagine a box containing a perfectly uniform distribution of matter, so that the density of the particles is constant. Suppose that at one location in the box the particles are somehow stirred, leading to a slight density enhancement at this particular spot. This tiny new concentration of mass will create a gravitational force, which tugs on the surrounding particles and causes them to fall inward. The infalling matter increases the clump's mass, which in turn increases its gravitational pull, allowing it to assemble even more material, and so on. Given enough time, this "gravitational runaway" transforms what was originally a tiny density enhancement into a dense clump, containing most of the mass that was distributed throughout the volume.

This simple phenomenon is the basis for theories of how the large-scale structure of the universe was formed. Yet in order for it to work, the universe must have been "imprinted" at some earlier time with a network of primordial density perturbations that would later collapse into the structures we see today. As it happens, the signature of these ripples has been observed—as tiny variations in the temperature distribution of microwave photons coming from different parts of the sky. Characterization of this microwave background is currently a major focus of astronomical research, as the ripples represent the ancient gravitational seeds of cosmic structure.

It would seem that we have all the elements needed to explain the origin of the cosmic web. We have observed density variations in the early universe, and we have a powerful model that explains how they could evolve into larger structures. However, there is one problem: The primordial variations were so small that 13.7 billion years is still not enough time to grow them into the assemblages we observe today! This puzzle received a great deal of attention during the 1970s, perhaps fueled by Cold War politics.

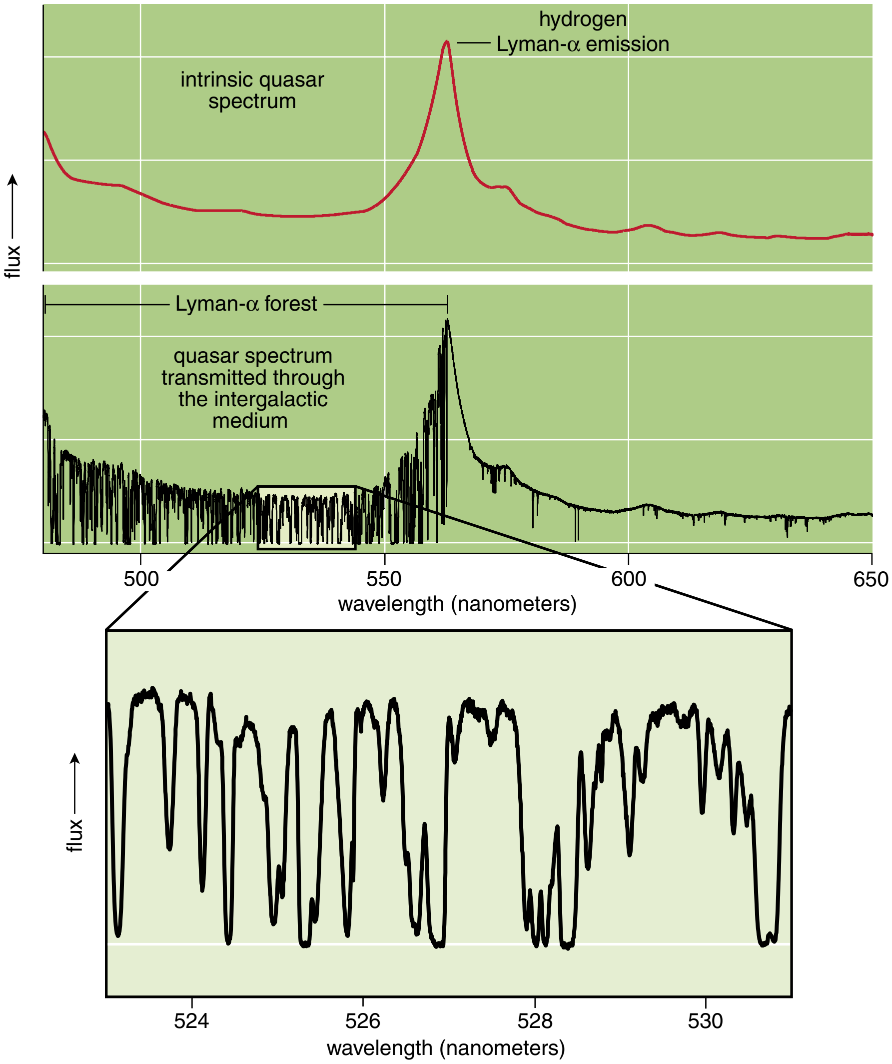

Figure 4. Intrinsic spectrum of a quasar *(top)* reveals the amount of light it emits at different wavelengths. The central peak is produced by hydrogen atoms, which give off photons at a wavelength of 121.56701 nanometers—known as the Lyman-α emission line. Here this line appears at about 560 nanometers because the expansion of the universe stretches light to longer wavelengths, from ultraviolet toward red—the so-called redshift. As the quasar's light travels through the intergalactic medium its spectrum is eaten away *(middle, hash to left of peak)*, because intervening clouds of hydrogen gas absorb light that locally has a wavelength of 121.56701 nanometers, but these absorptions take place at different redshift-distances from the Earth, producing a "Lyman-α forest." Some heavy elements absorb photons with longer wavelengths, producing the few absorption lines to the right of the peak. In recent years, high-resolution spectra *(bottom)* have allowed astronomers to study the individual clouds in the intergalactic medium.

Two competing theories of structure formation emerged, one devised by Yakov Zel'dovich at the School of Russian Astrophysics in Moscow, and the other by James Peebles and his collaborators at Princeton University. The ensuing debate exposed significant weaknesses in both theories. The solution required the introduction of an entirely new ingredient—ominously named *dark matter*—in the cosmological models. This proved to be one of the most important discoveries in modern cosmology.

This dark stuff is quite different from the ordinary matter that makes up stars, planets and people. Not only does dark matter not shine, it interacts with "our" kind of matter only through the force of gravity. It is largely believed to consist of exotic particles that have no other effects on ordinary atoms and molecules. Furthermore, dark matter appears to outweigh normal matter throughout the universe by a factor of four to one. This notion is indeed odd, and it has met with resistance since it was first suggested by the eccentric astronomer Fritz Zwicky in

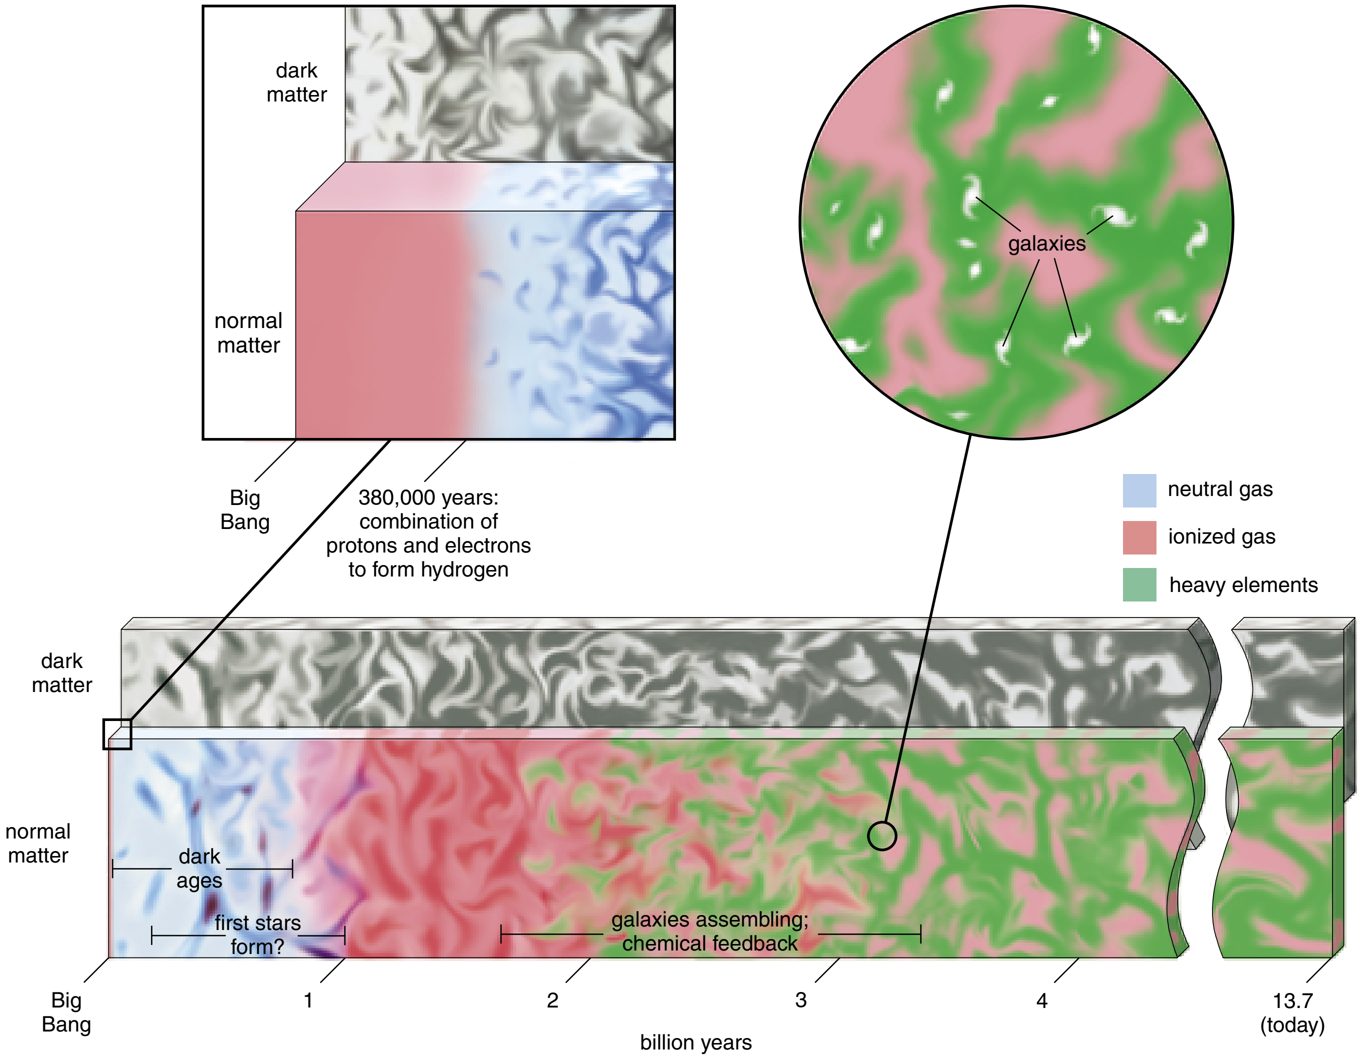

Figure 5. Evolution of the intergalactic medium begins shortly after the Big Bang, when the density of normal matter—protons, electrons and photons—was oscillating under the competing influences of gravity and pressure. At the same time, dark matter was quietly assembling into a primitive cosmic web. About 380,000 years after the Big Bang, the protons and electrons combined to form neutral hydrogen gas *(blue)*, which was quickly drawn into the existing lattice of dark matter by gravity. The gas and the dark matter continued to assemble in the cold darkness until the first stars were born in the highest-density regions of the intergalactic medium. The light from these stars eventually reionized and reheated the universe, probably within the first one billion years *(pink)*. As the galaxies began to take shape, the cumulative production of heavy elements by the stars began to "pollute" the intergalactic medium *(green)*.

the 1930s. However, cosmologists have now grown to accept its existence as nearly certain in the face of overwhelming evidence from a variety of observations. Although we may not understand exactly what dark matter *is*, we do understand what it *does*—it holds galaxies together, bends light, slows down the universe's expansion and drives the formation of intergalactic structure.

To understand this last point, we need to return to the early history of the universe. During the first 380,000 years, the relic heat from the Big Bang kept the universe so hot (greater than 3,000 kelvins) that electrons and protons in the primordial soup could not combine to form neutral hydrogen atoms. Such ionized gas, in this case consisting of dissociated electrons and protons, is known as a *plasma*. When plasma particles are in their free-floating state, they can interact with light, exchanging energy and momentum. In the early universe, this scattering increased the gas pressure within the cosmic soup. So, when gravity tried to collapse the first density perturbations, the gas pressure pushed back—much as a balloon does when it is squeezed. As long as the electrons and protons were separated, the gas could not form larger structures. Instead, the potential structures churned and oscillated as the inward pull of gravity fought the outward push of gas pressure.

Then, when the universe was 380,000 years old, a major event took place. As the universe was expanding, it was also cooling, and at this point it became cold enough for electrons and protons to combine, forming hydrogen atoms. Suddenly, these new atoms became decoupled from the photons—they no longer interacted so strongly with light—which drastically reduced the pressure that had kept gravity at bay. With gravity free to work on all the newly formed hydrogen atoms, structures could form in earnest.

How did dark matter fit into the picture? While the protons, electrons and photons were oscillating under the competing influences of gravity and pressure, the dark matter followed a different storyline. Because dark matter interacts with normal matter only through gravity, the pressure that kept the normal gas from collapsing couldn't act on it. Particles of dark matter enjoyed an unimpeded assembly into large structures long before the normal gas could begin to get organized. By the time normal matter decoupled from the photons, the dark matter had already grown into a primitive web-like network. As soon as the normal matter lost its support from the photon pressure, the gravity from the pre-existing dark-matter structures quickly pulled normal gas into the web. In this way, normal matter was given a gravitational "head-start" by the dark matter.

Once this process was set in motion, the gravitational building blocks of the intergalactic medium were in place. Normal and dark matter continued to free-fall toward concentrations of mass until the rising gas pressure slowed the infall. The web-like lattice was taking shape, but stars had not yet begun to form and all of the gas in the universe was neutral. The universe had entered an age where matter drifted about in the darkness, quietly assembling under gravity's influence. So it continued until at some point—probably somewhere between 200 million years and one billion years after the Big Bang—a process began that would fundamentally alter the nature of the intergalactic medium and the universe as a whole: The first stars were born.

Fiat Lux

It seems preposterous that something as small as a star could affect the universe on intergalactic scales. After all, a star is only a few light seconds across, whereas the filaments of the cosmic web may extend for billions of light-years. How can a relatively tiny object impact such a large volume? The answer lies in how stars work, where they live and what happens when they die.

Before there were stars, the normal matter in the universe was composed almost entirely of hydrogen and helium. Astronomers refer to this mixture as a chemically pristine gas because it reflects the chemical composition of the cosmos just after the Big Bang. Since then, nearly every atom of every other element—from argon to zinc—was forged inside a star. Stars are effectively nuclear fusion reactors: They gravitationally compress gas to such high densities that light atomic nuclei smash together to form heavier elements. Such stellar nucleosynthesis releases enormous amounts of energy, and that's what makes the stars shine.

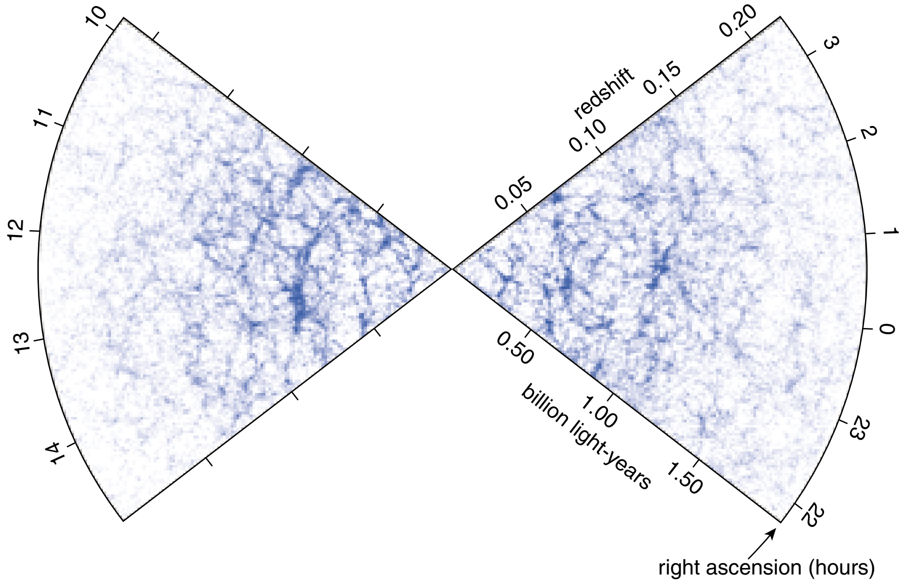

Figure 6. Survey of the nearby universe maps the distribution of about 75,000 galaxies *(small blue dots)*. The placement of each galaxy in the radial direction is proportional to its distance from the Earth (which is located at the intersection of the two wedges), and its angular position (or right ascension in hours of arc) corresponds to its location along a thin strip in the sky. The galaxies clearly trace a network of filamentary structures that are analogous to the cosmic web seen in numerical simulations. (Image courtesy of the 2dF Galaxy Redshift Survey team.)

Nucleosynthesis had several important effects on the intergalactic medium. First, it generated starlight, which escaped into intergalactic space and interacted with the neutral atoms. Later, the newly minted heavy elements were ejected into the intergalactic medium by strong *galactic winds*—powerful expulsions of hot gas—that stirred up and "polluted" vast regions of the universe.

Let's consider these processes in more detail by returning to the cosmic web *(Figure 1, top)*. Because galaxies are more than 10,000 times denser than the cosmic average, we would expect to find systems like the Milky Way within dense regions of the web itself, which contain the raw materials (gas reservoirs) needed to build the stars and galaxies.

In simulations, the densest regions are found within the web's filaments, especially where several intersect. Therefore, on cosmic scales, galaxies should behave like tiny particles trapped in the strands of the web, actually tracing the much larger structures outlined by intergalactic gas. Recent three-dimensional galaxy surveys, such as the Sloan Digital Sky Survey and the 2dF Galaxy Redshift Survey, have indeed revealed a filamentary pattern in the way that galaxies cluster *(Figure 6)*. Research groups, led by Max Tegmark at the University of Pennsylvania and Rupert Croft at Carnegie Mellon University, are currently investigating the clustering statistics of galaxies relative to those of the intergalactic gas as seen in quasar spectra. Their early results suggest that the same physics underlies the assembly of the intergalactic gas network and large-scale galaxy structures.

As the galaxies coalesced out of the web and began to shine, the universe was filled with the first new light since the Big Bang—the dark era had ended. And the stars dutifully began to churn out heavy elements. When enough stars had formed, the cumulative production of light and chemicals began to alter the nature of the intergalactic medium itself. Astronomers refer to these collective effects as "galaxy feedback," because the galaxies act on the surroundings from which they formed. Here I'll only consider two types of feedback, radiation and chemical pollution.

The first agent of galaxy feedback was starlight, which reionized the intergalactic medium. Recall that normal matter began to form large structures during the era of *recombination*, when the protons and electrons teamed up to form hydrogen atoms, so the gas in the universe was, for a time, entirely neutral. It was also very cold, reaching gas temperatures only a few tens of degrees above absolute zero. When the first stellar photons leaked out

Figure 7 "Galactic wind" streaming from galaxy M82 ejects enormous amounts of heavy-element-rich material *(red)* into the intergalactic medium. The center of M82 is experiencing a period of unusually vigorous star formation (a starburst). The most massive stars expire shortly after their birth in a concert of supernova explosions, blasting debris out of the galaxy. (Image courtesy of the Subaru Telescope, National Astronomical Observatory of Japan.)

from galaxies, they interacted with the hydrogen atoms, stripping away the electrons that had been in place since the era of recombination and reheating the resulting plasma up to temperatures near 10,000 kelvins. Reionization was initially confined within bubbles centered on the fledgling galaxies, because the starlight had not yet traveled far out into intergalactic space. As more galaxies began to shine, the ionized bubbles grew outward until those from adjacent galaxies began to overlap. Soon the entire volume of the universe was once again ionized *(Figure 5)*.

We now believe that the universe finally emerged from its "dark ages" and was reionized when it was less than 1 billion years old, or about 10 percent of its present age. Today, only about 1 hydrogen atom in 10,000 is in a neutral state and the average temperature of intergalactic gas is still very near 10,000 kelvins.

A Mighty Wind

It had long been assumed that the intergalactic medium was chemically pristine and that the production and distribution of new elements took place only within galaxies themselves. But astronomers also noticed that a few weak absorption features in quasar spectra appear redward of the hydrogen emission line. These other lines arise from different elements—in the case of Figure 2, carbon and silicon—whose characteristic wavelengths are redder (longer) than hydrogen's 121.56701 nanometers.

The absorption lines of these heavy elements are observed within regions that also contain a considerable amount of hydrogen. These zones correspond to gaseous halos around the first galaxies, whose stars were thought to supply the chemicals. However, in the early 1990s, quasar spectra taken by Lennox Cowie and Antoinette Songaila on the newly commissioned Keck telescopes revealed heavy elements far removed from any galaxy. Their discovery suggested that the chemical pollution of intergalactic space was much more efficient than originally believed.

The concentration of heavy elements in the intergalactic medium is very low: For example, only about one carbon atom can be found for every million (mostly hydrogen) atoms. So a box of intergalactic space that is 100 meters on a side would contain just a single carbon atom! Yet even this tiny amount reveals that some heavy elements were mixed throughout the cosmic web early in the history of the universe. How did they get out there—so far from the stars and galaxies in which they were made?

The evidence suggests that they were blown out into intergalactic space by violent *galactic winds*. These streams of matter flow out of galaxies where stars are actively forming. In all galaxies, the most massive stars burn brightly, and rapidly produce new elements. These stars burn so fast that they quickly exhaust their nuclear fuel and can no longer continue fusing light elements into heavier ones. When the reactor in a massive star turns off, the star ends its life in a tremendous explosion known as a supernova. The blast energy of a typical supernova rivals the simultaneous detonation of 10^{31} atom bombs, and the remnants of the dead star—including its newly fused heavy elements—are launched into surrounding space.

Despite its explosive power, a single supernova cannot pollute the intergalactic medium because the gravitational force from the star's galaxy traps the expanding debris before it can escape. However, galaxies occasionally experience bursts of unusually vigorous star formation where stars are born and die 10 to 50 times faster than usual. During these *starbursts,* multiple supernovae can be triggered in near succession. Their collective energy drives debris outward, like a rocket boosted by several stages, breaching the gravitational barrier and expelling heavy elements into the intergalactic medium. This phenomenon has been observed in a number of nearby galaxies *(Figure 7).*

Although we can study nearby starbursts and the resulting outflows in exquisite detail, these galaxies are the rare exception in the local universe. Most galaxies quietly go about forming stars and manage to retain the heavy elements they produce. But in the early universe, the situation was quite different. New observations of distant galaxies by Max Pettini at the University of Cambridge and his colleagues have revealed that outflows were extremely common when the universe was about 15 percent of its present age. This has two important implications. Nearly every galaxy we see today underwent some period of intense star formation in its past. And large quantities of heavy elements were launched into the intergalactic medium very early in the life of the universe. There was thus plenty of time for this material to coast out to large distances and mix with the chemically pristine intergalactic gas.

Studies of early galaxies and their feedback on the intergalactic medium define an important frontier of our knowledge about the first stars and cosmic structures. Several important questions remain open. For example, exactly when and where did the first stars form? Do heavy elements pervade the entire universe, or is there still chemically pristine gas left over from the Big Bang? Were the stars that triggered reionization the same stars that produced the observed intergalactic heavy elements?

For the past few years I have been investigating some of these questions with Wallace Sargent at the California Institute of Technology and Michael Rauch at the Carnegie Observatories. We have been measuring heavy-element concentrations in the early cosmic web to learn whether there are pristine corners of the universe that have not yet been reached by the galactic winds. So far we have detected heavy elements throughout all of the strands of the cosmic web, but it is still not clear whether the winds' sphere of influence extends beyond the filaments and into the intergalactic voids. In these remote regions the expected heavy-element densities are so low that even our most sensitive observations cannot reveal their absorption lines directly. Nevertheless, our results show that debris from galactic winds must have dispersed into most of the mass in the universe before the cosmos was a mere 20 percent of its present age.

We have also compared our observations with different models of star formation and chemical production to determine whether the stars that reionized the universe were the same ones that polluted the intergalactic medium. Our results suggest that the earliest stars did not produce most of the heavy elements, most likely because their heyday was too short *(Figure 8)*. Instead, we believe that galaxy feedback occurred in a series of waves. The first generation of stars reionized the universe, and later generations progressively enriched the intergalactic medium with chemicals.

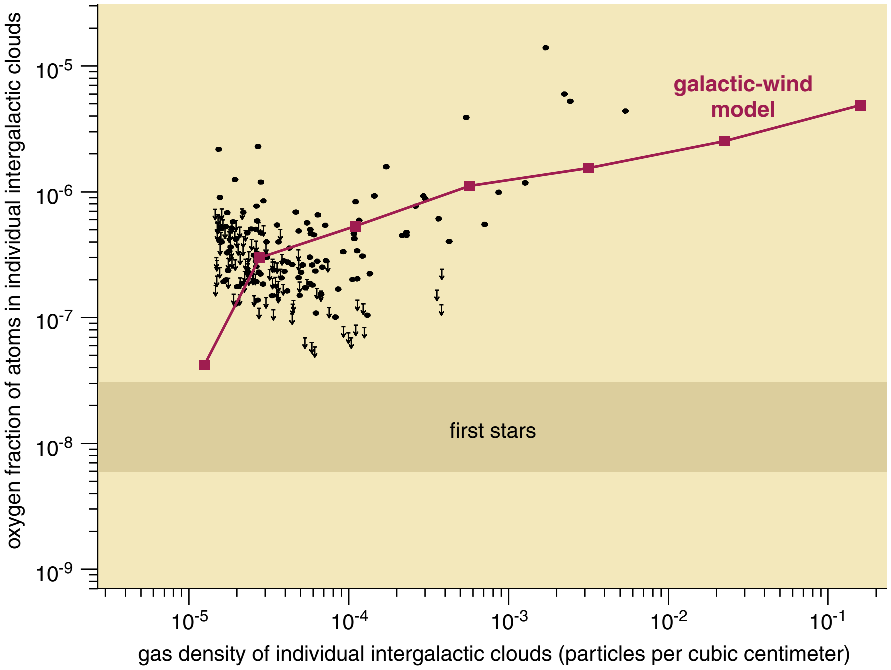

Figure 8. Concentration of oxygen in intergalactic clouds suggests that the heavy elements were scattered into the cosmic web by galactic winds, rather than by the first stars. The oxygen fraction serves as a marker for the distribution of the heavy elements, whereas the relative gas density of each cloud roughly corresponds (inversely) to the cloud's distance from a galaxy. If the first stars were the primary agents of dispersion, then oxygen would have had enough time to mix evenly into the cosmic web, resulting in a low constant value *(tan bar)*. The oxygen concentrations *(dots)* favor the galactic-wind model, but in some instances *(small arrows)* only the upper limits could be assessed. The redshift distance of these measurements reveals the universe 11.2 billion years ago, when the average gas density was about 40 times higher than it is today.

On the theoretical front, the most advanced cosmological simulations are just beginning to incorporate realistic models of galactic winds and the chemical enrichment of the universe. The physics of star formation and galaxy outflows is so complex that even the most sophisticated numerical models must make broad, simplifying assumptions to make the problem computationally tractable. The subject continues to progress rapidly, as both the observations and the theory evolve.

There are, of course, many details to be refined. Exactly how and when did the first stars form? How do galaxies and the intergalactic medium interact? And, perhaps most importantly, what is the nature of dark matter? Yet, when sufficient time has passed to offer a historical perspective, the past decade may well be remembered for the emergence of a standard model of the cosmos that ties all we know about galaxies and the intergalactic medium into a single package.

Bibliography

Cowie, L. L., and A. Songaila. 1998. Heavy-element enrichment in low-density regions of the intergalactic medium. *Nature* 394:44–46.

Padmanabhan, T. 1998. *After the First Three Minutes: The Story of Our Universe*. New York: Cambridge University Press.

Pettini, M., *et al.* 2001. The rest-frame optical spectra of Lyman break galaxies: star formation, extinction, abundances, and kinematics. *The Astrophysical Journal* 554:981–1000.

Rauch, M. 1998. The Lyman alpha forest in the spectra of QSOs. *Annual Review of Astronomy and Astrophysics* 36:267–316.

Simcoe, R. A., W. L. W. Sargent and M. Rauch. 2002. Characterizing the warm-hot intergalactic medium at high redshift: a high-resolution survey for O VI at redshift 2.5. *The Astrophysical Journal* 578:737–762.

Simcoe, R. A., W. L. W. Sargent and M. Rauch. In press. The distribution of metallicity in the intergalactic medium at redshift 2.5: OVI and CIV the spectra of 7 QSOs. *The Astrophysical Journal*.

For relevant Web links, consult this issue of *American Scientist Online:*

http://www.americanscientist.org/