Anatomy of a Supernova

BY STEPHEN M. MAURER AND D. ANDREW HOWELL

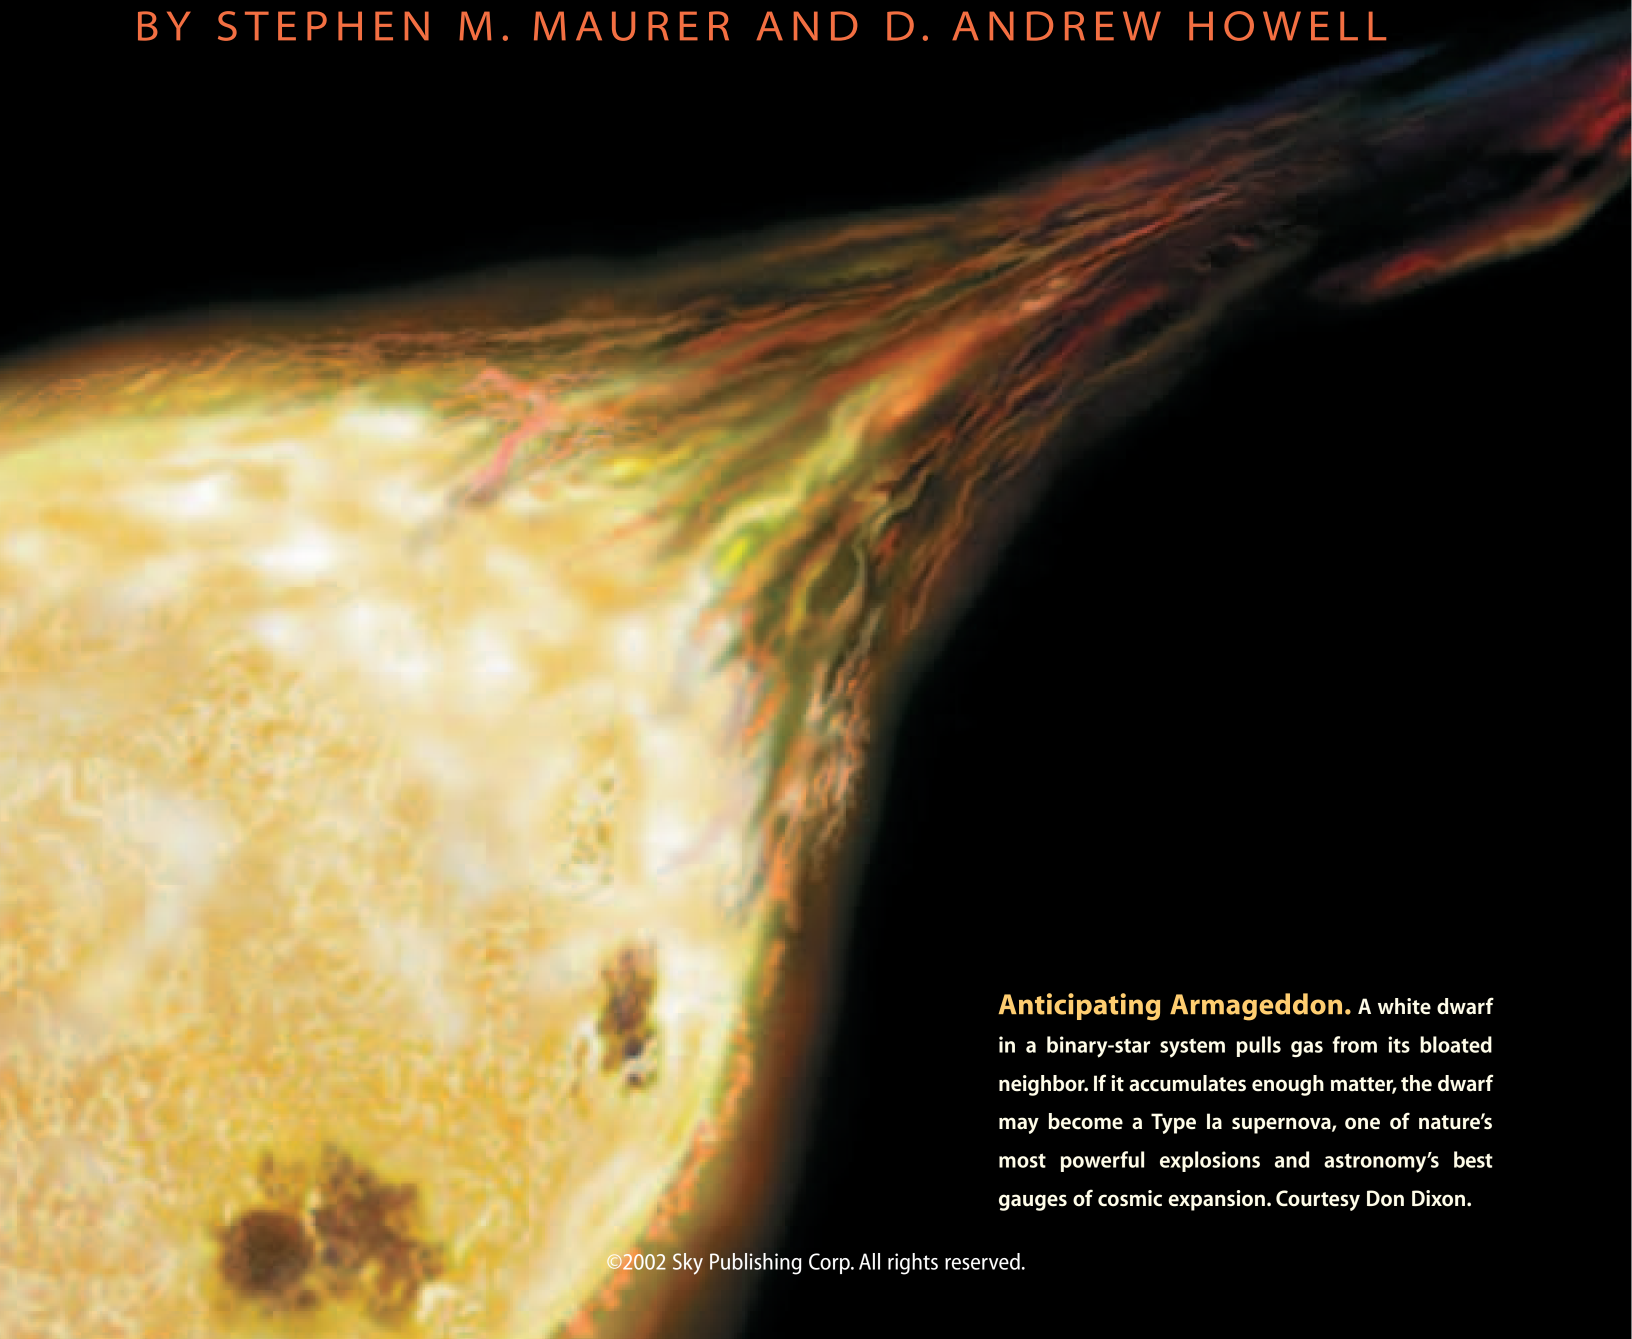

Anticipating Armageddon. A white dwarf in a binary-star system pulls gas from its bloated neighbor. If it accumulates enough matter, the dwarf may become a Type Ia supernova, one of nature's most powerful explosions and astronomy's best gauges of cosmic expansion. Courtesy Don Dixon.

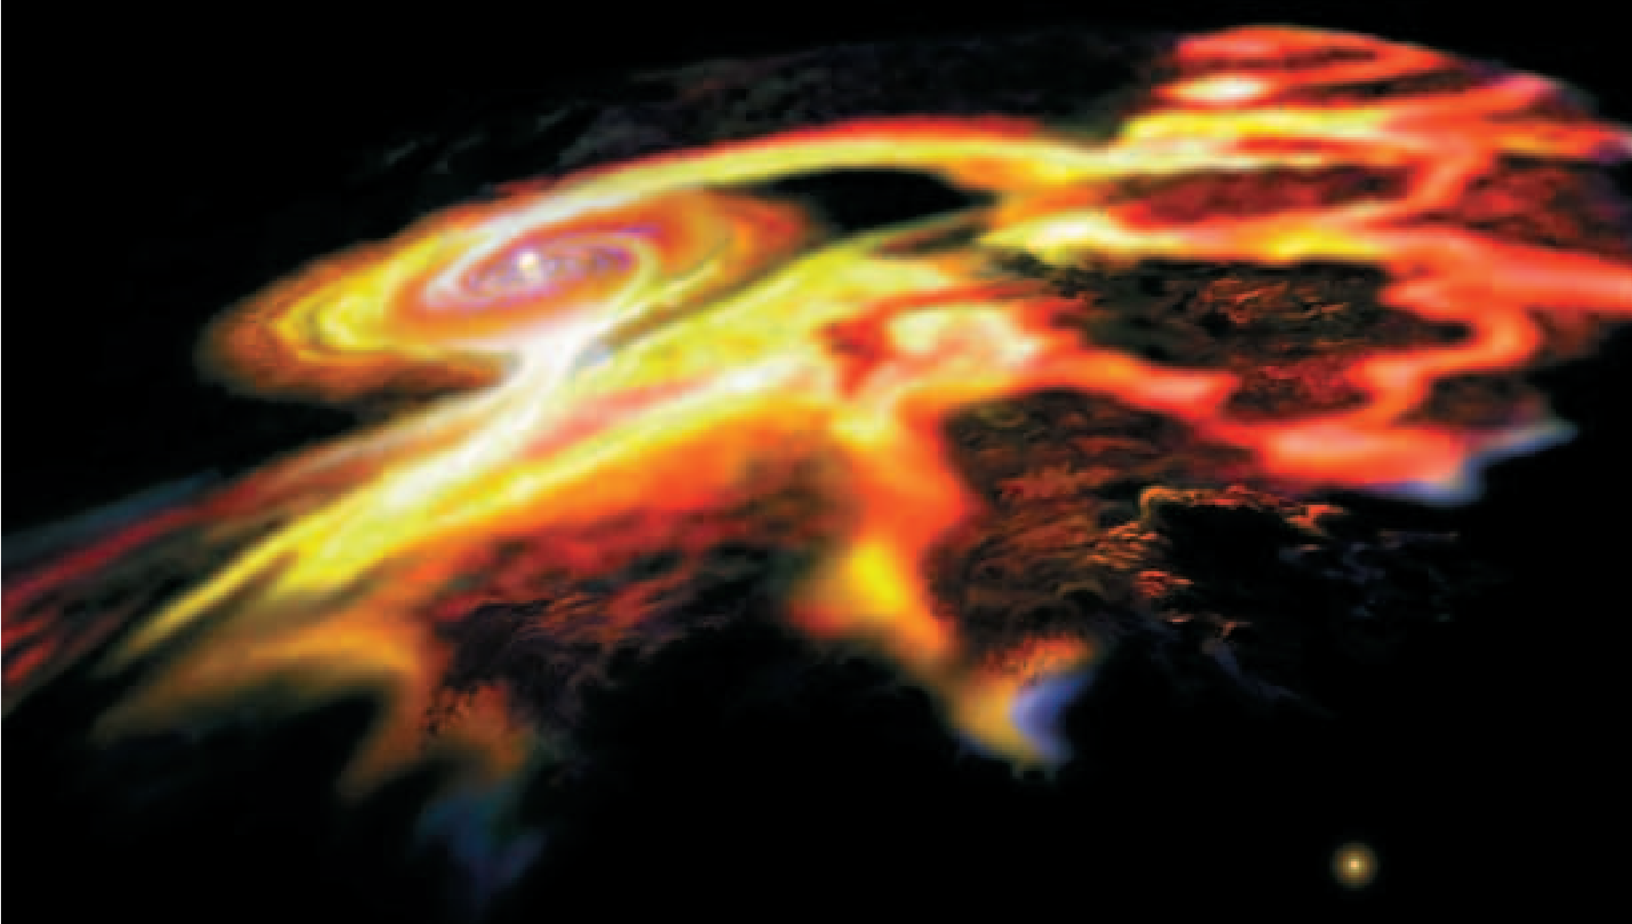

Astronomers use Type Ia supernovae as cosmic yardsticks. But what causes these stellar cataclysms in the first place?

GIVE AN ASTRONOMER A "STANDARD CANDLE" — that is, tell her how bright a star or galaxy is — and she can figure out its distance. Naturally, bright candles work best. Starting in the 1930s, astronomers dreamed of using the ultraluminous explosions called supernovae to survey the cosmos. Unfortunately, each new supernova seemed to have a different brightness. During the 1960s, astronomer Fritz Zwicky resolved part of the confusion by showing that the word "supernova" included five different kinds of explosions. Over the next 20 years, astronomers used better spectra to reorganize Zwicky's categories (see page 33). The final piece of the puzzle fell into place in 1993 after observatories traded in their photographic plates for electronic detectors. Better brightness data from the Chile-based Calán/Tololo Supernova Survey showed that one of the subcategories — Type Ia — was predictable enough to serve as a rough, if imperfect, candle. Just as important, bright Type Ia's last longer than dim ones. Astronomers used this fact to develop a "stretch correction" that turns each Type Ia's duration into a customized brightness estimate.

In 1989 scientists at the Lawrence Berkeley National Laboratory began to develop semi-automated techniques for finding supernovae. This allowed them to monitor huge amounts of sky; their program eventually became the Supernova Cosmology Project. Later, a second group called the High-Z Supernova Search Team joined the game. By 1998 each group had found a handful of Type Ia's that let them measure the distance to remote galaxies (*S&T:* September 1998, page 38). Astronomers were dumbfounded. Instead of slowing down, as generations of cosmologists had been taught, the universe's expansion rate actually was accelerating. Of course, Einstein's equations of general relativity were never in danger — Willem de Sitter had found accelerating solutions to them as far back as 1917. For these solutions to work, however, the universe had to be filled with a previously unsuspected "dark energy."

Physicists have plenty of theoretical models for dark energy. These include the vacuum energy from Einstein's original cosmological constant, a new kind of energy called quintessence, and parallel universes called branes. In principle, the history of our universe should tell us which model is right. For now, though, astronomers' best measurements of that history still aren't good enough because of uncertainties in the "wattage" of their Type Ia standard candle. So what *do* astronomers really know about these supernovae?

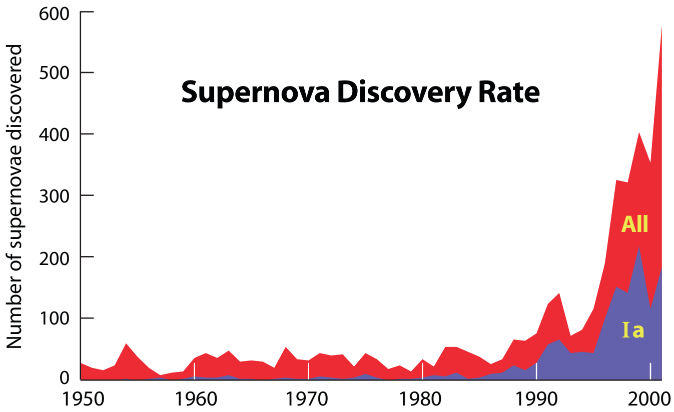

The use of automated search techniques produced a flood of supernova discoveries in the 1990s, including many Type Ia's. Courtesy Stephen M. Maurer and D. Andrew Howell.

Making a White Dwarf

The story of a Type Ia begins with an ordinary star and its lifelong struggle against collapse. In order to counteract gravity, a middle-aged star has to generate enormous internal pressure, which it gets from nuclear fusion. This process of hydrostatic equilibrium works like a diesel engine. As gravity tries to crush the star, it acts like a descending piston. The downstroke forces energy into the gas, which raises its temperature. Higher temperatures, in turn, make the star's hydrogen fuel "burn" faster (that is, fuse into helium). In a well-behaved engine (or star), this rapid burning generates enough extra pressure to make the piston change direction and start its upstroke. As the gas expands anew, it cools. Combustion slows, and the cycle repeats. Most main-sequence stars, including our Sun, balance gravity and pressure so exquisitely that these oscillations are tiny and damp themselves out.

In 1994 the Hubble Space Telescope recorded a Type Ia supernova on the outskirts of NGC 4526, a 10th-magnitude spiral in Virgo. By the time this image was acquired, the supernova had already faded from its peak. At maximum, a Type Ia supernova can outshine an entire galaxy. Courtesy NASA and the High-Z Supernova Search Team.

Eventually the star's core starts to run out of hydrogen. Fusion slows, and the star's center contracts. This drives up temperatures until the core can burn a new fuel, helium, turning it into carbon and oxygen. Meanwhile, the superheated core blasts the star's outer layers with radiation and makes them puff up. The result is a red giant. For stars like the Sun, the process ends here because temperatures never get high enough to ignite other elements, including carbon and oxygen. Instead, the new elements simply accumulate in the star's core.

A star's final years are violent. It sheds its outer layers through pulsations or powerful stellar winds, creating a planetary nebula. Eventually the bare core is exposed to space and becomes a white dwarf. Now the formerly Sun-size star is smaller than the Earth. In fact, white-dwarf matter is so compressed that a teaspoonful weighs more than a car.

Theorists have spent the last 70 years trying to understand how such unfamiliar stuff behaves. Fortunately, they have a clue. Atomic physics says that it is hard to crowd electrons together. Pack the electrons tightly enough — that is, force them into what physicists call a degenerate state — and they push back. This electron-degeneracy pressure supports the white dwarf against collapse. Since degeneracy pressure doesn't need fuel it can, if left alone, support the star forever. Eventually, the star's carbon and oxygen cool enough to form one huge crystal — in effect, a degenerate diamond. This is the fate of our Sun.

But degeneracy pressure will not support stars much bigger than the Sun. As the late physicist Subrahmanyan Chandrasekhar deduced purely from atomic theory, it works only for stars having less than 1.4 solar masses. This stellar boundary is known as the Chandrasekhar limit.

Death Spiral

Most stars belong to binary systems. Suppose that one star becomes a white dwarf. When the second star evolves into a red giant, its gaseous envelope expands to fill a region roughly an astronomical unit — the distance between the Earth and Sun — in radius. If the dwarf star is close enough to its swollen companion, its gravity grabs some of this gas, which rains onto the dwarf. The resulting drag gradually shrinks the dwarf's orbit so that gas capture continues. After millions of years, stolen material may push the white dwarf's mass up toward the Chandrasekhar limit. If the mass of the precursor star was a little larger than the mass of the Sun, the dwarf can accrete enough matter to go past the 1.4-solar-mass limit and collapse into a neutron star — a high-density object only 15 to 20 miles (25 to 32 kilometers) across, composed almost entirely of neutrons. This collapse does not eject matter or produce a bright flash, so it has never been seen directly. However, these "silent supernovae" must happen, because the galaxy is littered with neutron-star corpses with masses near the Chandrasekhar limit.

If the white dwarf contains carbon, it never gets the chance to reach the Chandrasekhar limit. That's because the carbon starts to burn, leading to a runaway reaction that destroys the star. Why is carbon burning so explosive? Remember the diesel. If you want an engine to explode, jam the piston so that it can't move. This is what happens in a Type Ia. Degenerate matter is so incredibly stiff that it stays practically the same size over a huge range of pressures. So instead of expanding gradually as the carbon burns, the white dwarf explodes.

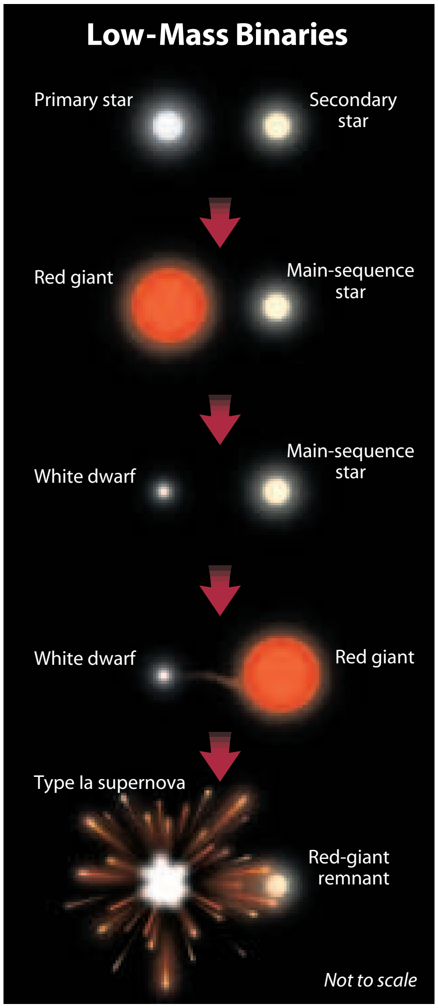

S&T: GREGG DINDERMAN

If the stellar components of a low-mass binary are close, their evolution will ultimately lead to the destruction of one of the stars. After one star evolves into a white dwarf, it begins pulling material from its companion, even as the second star proceeds along its own evolutionary path. Eventually the gas-gobbling white dwarf accretes too much matter and explodes as a Type Ia supernova.

Burning Down the Star

Astronomers have struggled for a half century to explain the explosion of a white dwarf with a runaway fusion reaction at its core. In principle, the answer can be computed from nuclear-physics data gained in atom-smasher experiments. In practice, that's like saying that anyone who knows chemistry can predict what a forest fire will do. The challenge is to develop computer models that start with nuclear physics and predict observed light curves and spectra.

Fortunately, astronomers have a starting point. Mathematical combustion theory and terrestrial experiments show that there are two basic types of explosion. The first is called *deflagration.* Deflagration waves are subsonic and act like fast-moving fires. The second — *detonation* — is a supersonic wave and happens when combustion produces an ultrafast shock that compresses and ignites the material ahead of it. Whether a particular explosion produces detonation or deflagration depends on factors like density in the preshock environment. On the other hand, waves that start in one density regime can travel to another; laboratory experiments routinely show detonation waves turning into deflagration waves and vice versa.

During the 1960s astronomers' first computer simulations used detonation waves to destroy a carbon-rich white-dwarf star undergoing a runaway nuclear reaction. The problem with these early models was that the supersonic waves moved so rapidly that the stellar material couldn't expand before it burned. Since combustion at very high densities and temperatures is extremely thorough, everything that can burn, does. All that's left are iron, cobalt, and nickel — the so-called iron-peak elements. Burning stops at these elements because making still-heavier ones would absorb more energy than is released.

Astronomers realized early on that detonation waves could not be the whole story. In the spectra of real-world supernovae, strong iron-peak lines take weeks to appear. This shows that the heaviest elements are concentrated near the center. Outer layers, on the other hand, are dominated by lighter elements like silicon, sulfur, and calcium. Obviously, something limits the maximum temperature and stops combustion in these regions. That something is expansion. In the 1970s and early 1980s astrophysicists began developing one-dimensional simulations that featured subsonic deflagration waves. Unlike a detonation wave, a deflagration wave gives the star time to expand. But this model also had problems. Some Type Ia's eject elements like calcium at nearly one-tenth the speed of light. It was hard to see how a subsonic flame could do this.

In the late 1980s a new picture combined the strengths of both models. Today scientists believe that a Type Ia supernova starts with a deflagration wave at the heart of the star. Here, confining pressures are so large that deflagration waves can burn at the densities needed to make iron-peak elements. As the wave travels outward, pressures fall and burning is less complete. Now the wave makes only relatively light elements like magnesium, silicon, and sulfur. The final act comes when the explosion reaches the star's upper layers. Somewhere around 0.7 stellar radius, the environment is no longer dense enough to support deflagration. The explosion goes into detonation mode, accelerating combustion products like calcium to supersonic velocities. This elegant theory explains most real-world observations, but astronomers need more computing power to

Astronomers have struggled for a half century to explain the explosion of a white dwarf with a runaway fusion reaction at its core.

refine it. Most current simulations rely on simplified two-dimensional calculations or study just part of the explosion. Full end-to-end simulations are still a decade away.

Deconstructing Starlight

Amazingly, the explosive combustion phase described above ends within just a few seconds. Now the hapless dwarf is filled with heat, momentum, and radioactive isotopes. Because the debris is so energetic, gravity can no longer contain it, and the star flies apart. You might think that the debris is a maelstrom. In fact, computer simulations and nuclear-weapons experience suggest that the explosion is remarkably orderly. Even though all the debris parcels expand, they never overlap. Instead, the supernova grows like a sliced onion seen at increasingly close range. Geometrically, this can happen only if each parcel travels faster than the one below it.

A few days after the explosion, the supernova is still so hot and dense that astronomers can't see past the shattered star's outer layers. Spectra show violently Doppler-shifted absorption lines from backlit gases. Later, as the fireball thins and cools, astronomers can peer deeper into the wreckage. New emission and absorption lines start to appear; each set of lines reveals the chemical composition of another layer.

Strangely, most of the supernova's light doesn't come from the heat of combustion. Instead, the explosion makes huge quantities of a radioactive isotope called nickel-56. By the time the burning wave flickers out, roughly 40 percent of the Type Ia supernova — six-tenths of a solar mass — has been converted into nickel-56! Nickel-56 decays to cobalt-56 with a half-life of 6.1 days. Cobalt-56 is also radioactive and decays to iron-56 with a half-life of 77.1 days. Initially energized, the cobalt and iron nuclei emit copious gamma rays. After imparting much of their energy to neighboring electrons, those gamma rays are absorbed by atoms in the expanding supernova shrapnel. This heats the exploded star's outer envelope to a temperature of about 10,000° Kelvin, making it glow. The visible light emitted by this energized envelope slowly scatters outward, making the supernova's light curve peak about 20 days after the explosion.

Intergalactic Lighthouses

Astronomers don't study Type Ia's just to learn why they go "boom" — they also want to use them to better measure cosmic distances. This second goal took center stage after the "accelerating universe" was announced in 1998. That year astronomers found that distant supernovae were unexpectedly dim — 20 or 30 percent fainter than expected. This implied they were 10 or 15 percent more distant than they would have been had the universe's expansion slowed as much as conventional cosmology predicted. Instead, it appeared that the universe's expansion was actually accelerating. Some astronomers tried to explain the finding by invoking dust that would make distant supernovae look dimmer, or an evolutionary scenario in which the average luminosity of Type Ia's has changed over time. Fortunately, acceleration has an important signature. Einstein's equations show that acceleration — and the extra dimming it causes — must switch off somewhere around a redshift of $z = 0.5$. Recent observations show that neither dust nor evolutionary theory can account for the brightness of high-redshift supernovae (*S&T:* July 2001, page 20).

Today, most astronomers agree that the expansion of the universe is accelerating. But they need better distance estimates to clinch the case and to test competing theories. Unfortunately, typical Type Ia luminosities vary by about 30 percent. Present-day calibration techniques can correct for supernovae that are intrinsically overly bright or overly dim, since more luminous supernovae take longer to fade, while feebler ones disappear quickly. This cuts the uncertainty in half. But this "stretch correction" is only as good as the data from which it comes. Finding supernovae in nearby galaxies — galaxies whose distances can be measured by other means — will dramatically improve matters.

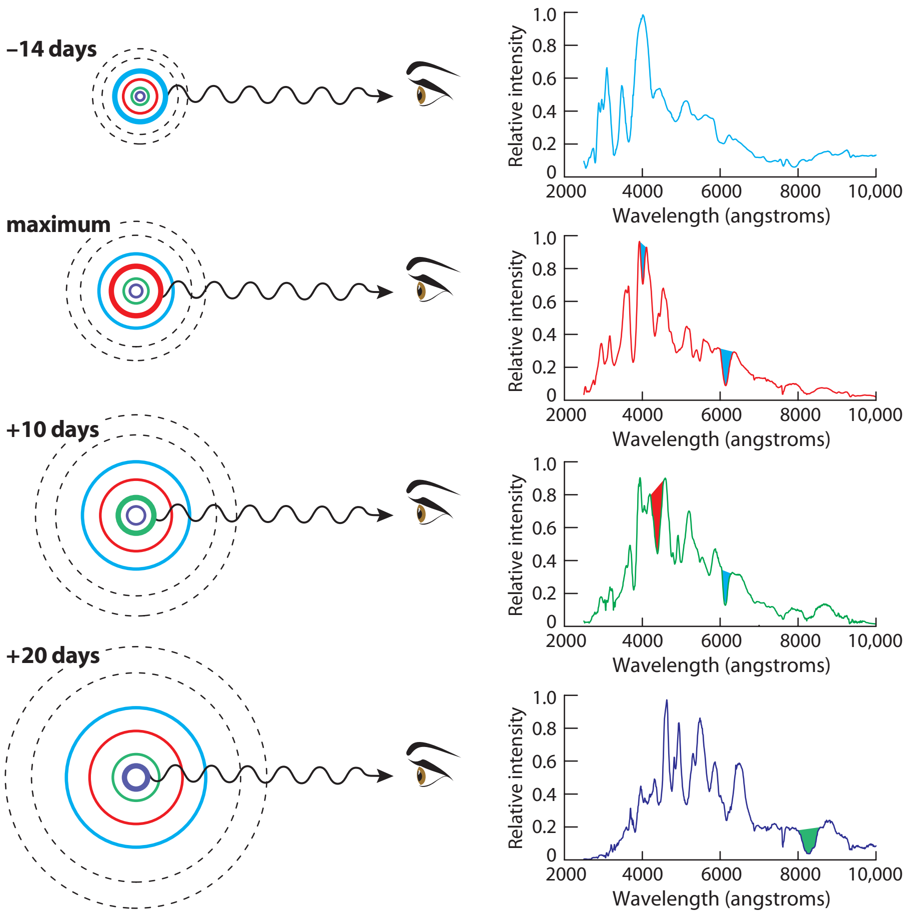

At first, days before maximum brightness, supernova spectra are dominated by light from shallow parts of the explosion (cyan). Later, photons from progressively deeper layers (red, green, and dark blue, respectively) escape and add features to the spectrum. As we observe these deeper layers, dominant elements from the shallower layers are backlit and seen as absorption lines (here filled with colors indicating the layers in which they dominate). The changing pattern of emission and absorption lines reveals what each layer is made of. Courtesy Saul Perlmutter (Lawrence Berkeley National Laboratory).

Astronomers also are looking for other ways to estimate the intrinsic luminosities of these "standard bombs." For example, hotter explosions produce more light. In principle, the supernova's temperature can be inferred from its color, or by comparing the strengths of different spectral lines.

Loose Ends

Despite 70 years of supernova research, plenty of mysteries remain. Thirty percent of all Type Ia's are peculiar. For example, some explosions have been wildly overluminous (SN 1991T) or underluminous (SN 1991bg). Current theory probably can accommodate subluminous supernovae if the deflagration-to-detonation transition comes late enough. However, astronomers are still not sure about overluminous supernovae. The basic problem is figuring out how any explosion can make enough nickel-56 — nearly one solar mass — to power such a bright supernova.

Astronomers also want to know what the progenitor star is made of. Simulations show that fairly pristine samples should survive in the supernova's outermost layers. In principle, ultraviolet spectra can detect this material, but because our atmosphere absorbs ultraviolet light, observers will have to use the Hubble Space Telescope to get their data. They'll also have to catch the supernova very, very early — probably within its first week. After that, light from the star's interior swamps the signal.

The companion star's identity is also a mystery. The most likely candidates include middle-aged, Sun-type stars and red giants. However, some theorists, including Icko Iben (University of Illinois), have argued that the companion might be another white-dwarf star that collides with its primary. While many astronomers doubt this scenario — they think that the merged star would immediately exceed the Chandrasekhar limit and collapse — no one really knows. In principle, better observa-

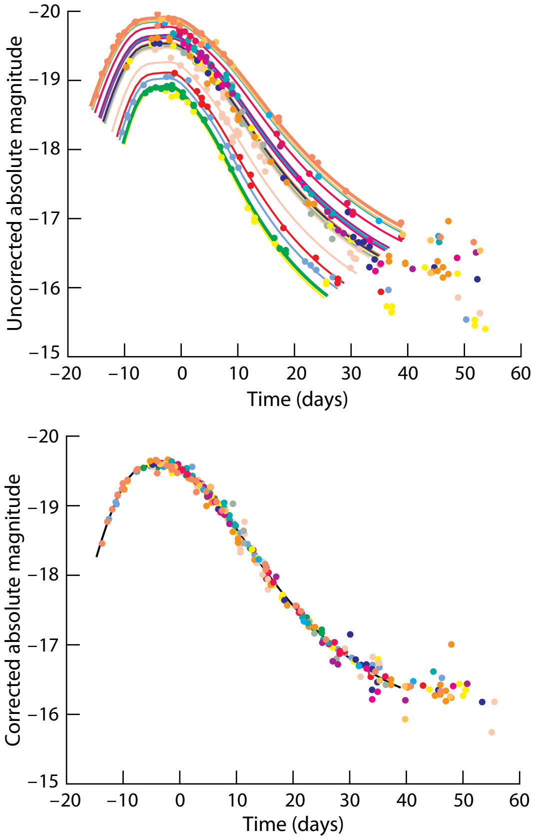

Top: **These light curves show how Type Ia supernovae differ in their peak luminosities, with the most luminous being those that fade most slowly.** ***Bottom:*** **Correcting for the "brighter is broader" relationship makes a nearly universal template that will indicate whether newfound Type Ia's are under- or overluminous (and by how much). Courtesy the Supernova Cosmology Project.**

A Supernova Bestiary

Despite revolutionary advances in knowledge, astronomers still classify supernovae using the same basic system that pioneering Caltech astronomer Fritz Zwicky invented during the 1960s. The central idea is to separate supernovae whose spectra have hydrogen (Type II) from those that don't (Type I). Since the 1980s, astronomers have further subdivided the Type I category by separating spectra with strong silicon (Type Ia) or helium lines (Type Ib) from the rest (Type Ic). Type II's, for their part, now have five different subcategories.

Except for Type Ia's, all supernovae are caused by core collapse. If a star starts out life weighing more than eight solar masses, nuclear fusion continues past oxygen and carbon until a degenerate iron core starts to form. But degenerate matter has an Achilles' heel: destroy the electrons and the star's supporting pressure vanishes. When the iron core gets large enough, central pressures rise so high that electrons merge with protons, making neutrons. The star collapses. According to this picture, Type II supernovae (such as SN 1987A) are produced when a massive star's core collapses. Type Ib's are caused by core-collapse events in stars that have previously shed their hydrogen envelope. Finally, Type Ic's are caused by "naked core collapse" events in stars that also have lost their helium.

The fact that Type Ia's are powered by thermonuclear burning makes them different from all other supernovae. Type Ia's occur only in binary systems, and their fusion reactions stop at iron. Only a core-collapse supernova can make heavier elements such as gold.

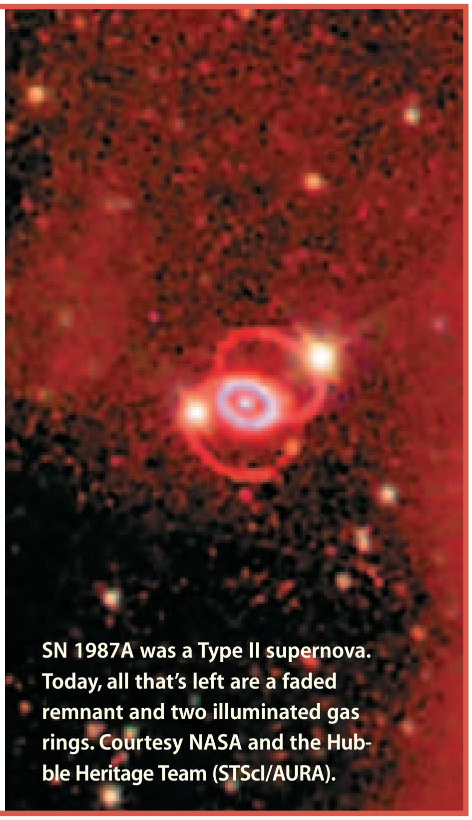

SN 1987A was a Type II supernova. Today, all that's left are a faded remnant and two illuminated gas rings. Courtesy NASA and the Hubble Heritage Team (STScI/AURA).

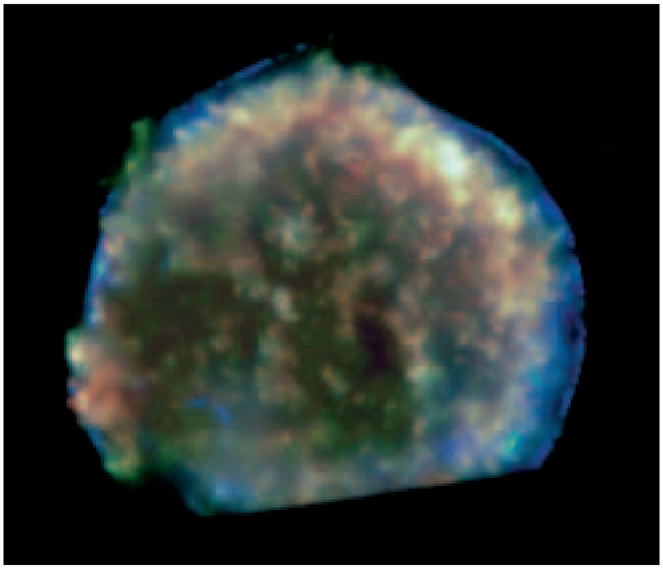

NASA / CXC

In A.D. 1572 Danish astronomer Tycho Brahe observed a supernova in the Milky Way; today astronomers think it may have been a Type Ia. Although the precursor star remains undiscovered, debris from the explosion is visible as an expanding shell of gas now 24 light-years across. This new view of the Tycho supernova remnant by the Chandra X-ray Observatory reveals interactions between the dead star's wreckage and the interstellar medium. The colors red, green, and blue represent low, medium, and high X-ray energies, respectively.

tions can settle the question. If the companion is a conventional star, the explosion will blow off part of its envelope. This should add traces of hydrogen to the supernova's spectrum.

Surprisingly, the companion itself should survive — minus its outer layers and swollen with excess heat. Although astronomers don't understand the details, this scorched remnant should occupy a region of the color-magnitude diagram that ordinary, internally heated stars never reach. Survivors are also expected to lie deep inside the debris field with a velocity vector that points back to the center. Robert Fesen (Dartmouth College) has looked for survivors with HST's Wide Field and Planetary Camera 2. Fesen's search focused on the Milky Way's only suspected Type Ia candidates from the past

The hunt for Type Ia's has also moved into space.

1,000 years — SN 1006 in Lupus and Tycho's supernova of 1572. "So far the answer is no," Fesen admits, "but the search isn't airtight. SN 1006 is so close that it's about 30 arcminutes across, which is a big area to search. Tycho is farther away but has a lot more extinction. That makes our color-data information hard to interpret."

The Quest Continues

High-redshift supernovae are still the name of the game. The High-Z Supernova Search Team and the Supernova Cosmology Project both have launched second-generation searches that will bag hundreds of distant Type Ia's over the next five years. Unlike earlier surveys, most of these programs feature rolling searches that revisit the same patch of sky over and over again. This strategy guarantees repeated brightness measurements for every supernova. Exposure times are set by the desire to find Type Ia's at redshifts near $z = 0.5$ — stars that exploded around the time that the universe's acceleration seems to have kicked in. Cosmologists need such data to choose between Einstein's cosmological constant, quintessence, and other models.

Finding nearby Type Ia's is less glamorous but just as essential. Astronomers need better spectra and luminosity data to understand and calibrate the Type Ia candle. The task is enormous because researchers have to monitor huge amounts of sky. This is partly due to perspective; nearby galaxies are scattered all across the heavens. Also, the total volume of nearby space is, by definition, limited. A Franco-American collaboration called the Supernova Factory is attacking these obstacles with automation. On a typical day, observations from the previous night, made at the 1.2-meter reflector at Palomar and the 1.2-meter Air Force Telescope at Haleakala, Hawaii, arrive in the morning and are processed by lunch. The group is also building an automated spectrograph to follow up its discoveries; it will be placed on the 2.2-meter University of Hawaii reflector on Mauna Kea.

Finally, astronomers are taking advantage of a new generation of telescopes to push Type Ia science into the near in-

This illustration shows the devastating effect on a red-giant star of one solar mass when its white-dwarf companion becomes a Type Ia supernova *(frames 1 and 2)*. Approximately 3 hours after the white dwarf explodes, ejected material reaches the red giant *(frame 3)* and begins stripping its outer atmosphere. Six hours later a shock wave has been driven into the giant star's envelope *(frame 4)* and is slowly making its way around the star. Two days after the explosion the shock has passed through the giant, heating and distorting what little remains of its outer envelope *(frame 5)*. After a month of bombardment, the former red giant has lost 98 percent of its outer atmosphere, but the hot degenerate core survives with about half the star's original mass *(frame 6)*. Courtesy Aaron McEuen, Hansen Planetarium.

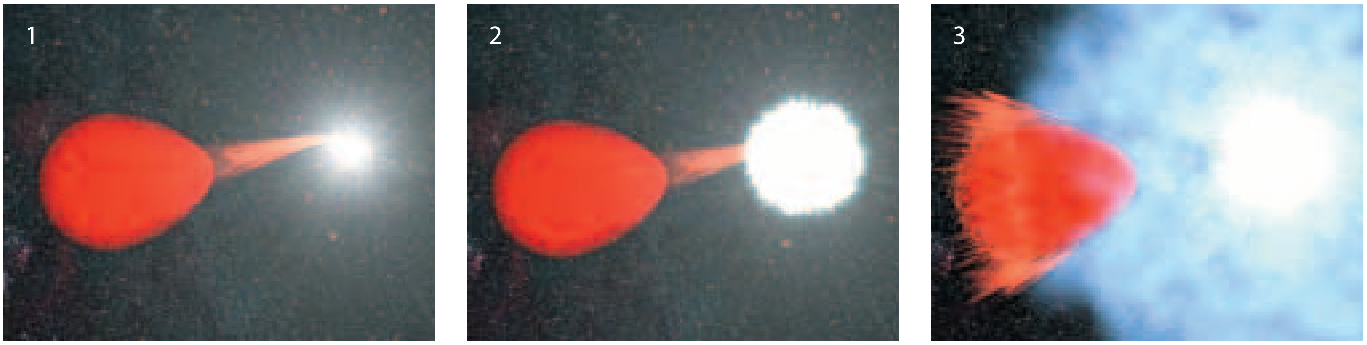

Above: **Located in the nearby elliptical galaxy NGC 5018 in Virgo, Supernova 2002dj was caught almost two weeks before its maximum brightness — one of the earliest recorded observations of a stellar explosion. Courtesy Nicholas Suntzeff, R. Chris Smith, Kevin Krisciunas, and Pablo Candia (NOAO/AURA/NSF).**

Below: **Glimmering at only 18th magnitude and invisible in this image (it's located within the circle at the top of the photograph), the recurrent nova U Scorpii is usually not much to look at; for a few days every decade or so, it erupts to 8th or 9th magnitude. But sometime in the next 700,000 years the star, located at the head of Scorpius, may become a Type Ia supernova and briefly outshine Venus. Courtesy Akira Fujii.**

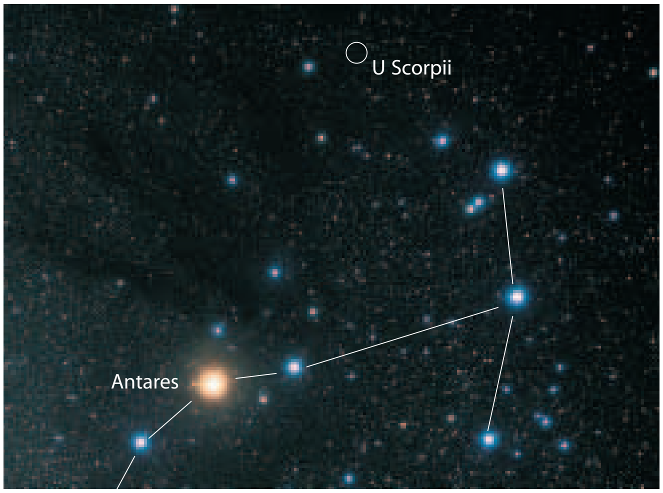

frared. Most chemical spectral lines are well separated and easy to identify in the infrared. Furthermore, infrared radiation lets astronomers see deeper into the post-supernova wreckage. Various groups are planning campaigns on the 6.5-meter Magellan Telescope at Chile's Las Campanas, the Gemini South Telescope at Cerro Tololo (also in Chile), and the Hobby-Eberly Telescope at McDonald Observatory in Texas.

The hunt for Type Ia's has also moved into space. The water vapor in Earth's atmosphere prevents ground-based astronomers from getting accurate spectra and brightness data from Type Ia supernovae whose light has been redshifted into the infrared. This includes supernovae at the cosmologically interesting redshifts beyond $z = 1$. Instead, they routinely turn to HST for follow-up observations. HST is also searching for new supernovae. Adam Riess (Space Telescope Science Institute) and his collaborators used Hubble Deep Field images to study the farthest supernova found thus far, at $z = 1.7$ (*S&T:* July 2001, page 20). Riess now is using HST in hopes of finding six to eight additional high-redshift Type Ia's. With luck, the new data will probe the era when the universe was still decelerating. "The goal," Riess explains, "is to find an epoch when the universe was dominated by dark matter instead of dark energy."

HST may be good, but it is not the last word. Since its field of view is tiny, rolling searches are impractical, so observers are forced to study one supernova at a time. Also, its sensors are not optimized for the infrared wavelengths needed to measure Type Ia luminosities and spectra above $z = 1$. A group at Lawrence Berkeley National Laboratory wants to do better. Their SuperNova Acceleration Probe (SNAP) proposal would send a 2-meter telescope into orbit by the end of the decade. Because SNAP's field of view will be 200 times bigger than HST's, its rolling searches will monitor 20 square degrees at a time. Furthermore, its deep-field exposures should discover supernovae one-tenth as bright as today's best ground-based searches. Finally, SNAP's detectors will be optimized for Type Ia's in the cosmologically interesting redshift range between $z = 0.5$ and $z = 1.7$. All in all, SNAP should discover and obtain detailed follow-up data for almost 2,000 Type Ia's per year.

Today at least a dozen groups are working to find the truth about Type Ia's. Make no mistake: The "golden age" of supernova science is about to arrive. Ten years from now, astronomers will understand Type Ia's much better than they do today. In the process they may also solve the dark-energy problem once and for all.

Stephen M. Maurer *is a San Francisco–based science writer. His articles have appeared in* Nature, Science, *and in the August 2001 issue of* Sky & Telescope. D. Andrew Howell *is a researcher at the Lawrence Berkeley National Laboratory and a member of the Supernova Cosmology Project. He also edits the astronomy Web site starstuff.org.*