of the medium changes in proportion to the net force applied to it. Because the medium is made up of people, who are capable of moving under their own power but not perfectly free to choose how they do so, the net force includes a propulsive term $\mathbf{p}$.

Both $\mathbf{v}$ and $\mathbf{p}$ are vector fields that depend on position and time. Assuming that those two quantities completely characterize the state of the crowd, the researchers wrote down a general model of all the ways that each quantity could deterministically evolve in time. And they found a term that could be key to describing the undulations: $\mathbf{p}$ changes in time in proportion to $-(\mathbf{p} \times \mathbf{v}) \times \mathbf{p}$.

The first part of the vector product, $\mathbf{p} \times \mathbf{v}$, always points perpendicular to the plane of the crowd, but whether it points up or down depends on whether $\mathbf{p}$ is angled to the right or the left of $\mathbf{v}$. The net effect of $-(\mathbf{p} \times \mathbf{v}) \times \mathbf{p}$ is always to rotate $\mathbf{p}$ away from $\mathbf{v}$. But because $\mathbf{p}$ is a force that acts on $\mathbf{v}$, $\mathbf{v}$ also rotates in the same direction. The result, shown in figure 2b, is that $\mathbf{p}$ and $\mathbf{v}$ circle around and around in tandem; whether they circle clockwise or counterclockwise depends on how they were oriented to begin with. That is, the mirror symmetry is spontaneously broken by the initial conditions.

The model agrees well with the undulations observed in the San Fermín crowds, and it's even consistent with the crowd dynamics at the 2010 Love Parade, as inferred from sparser and grainier video footage of that event. The physical origins of the $-(\mathbf{p} \times \mathbf{v}) \times \mathbf{p}$ term are still murky, but the larger insight is that crowd dynamics can be described by a continuum model at all.

"A group of five people is not suitable to be described with these tools," says Bartolo, just like it wouldn't make sense to use the equations of fluid dynamics to model a cluster of five molecules. "What is the magic number of people where the individual interactions get averaged out? We weren't sure that the San Fermín crowds would be over that threshold. But they were, and we're very happy about that."

## Safety in numbers?

The continuum model offers a framework for thinking about what makes dense crowds dangerous in some circumstances but not others. The undulations are an emergent property of the crowd itself, and they're not under the individuals' control. They involve the correlated, collective motion of tens to hundreds of people, with a combined mass of up to tens of thousands of kilograms. If such a group slams into a wall—with no one having the power to stop it—it can exert a crushing force with deadly consequences.

The Pamplona plaza is surrounded by walls. But it also has plenty of side streets that can act as pressure relief valves if too much of the crowd gets compressed near the edge of the square. The Love Parade disaster, in contrast, occurred in the narrow entranceway to the festival grounds, where escape routes were limited. The difference could be key to understanding why the Love Parade ended in tragedy, whereas the San Fermín opening ceremony never has.

Bartolo also points out that crowd undulations start small and grow larger as the crowd density increases. So looking for the undulations while they're still small could be a way for event organizers to tell when a crowd is on the verge of becoming dangerously out of control. "Engineers will have to work hard to develop concrete applications that will prevent real catastrophes," says Bartolo. "But we do think this is potentially useful."

So far, the researchers have limited their study to the average dynamics of the crowd as a whole. They'd like to extend their model to describe how crowds respond to localized stimuli, and they suspect that such an investigation could help clarify the connection between individual behaviors and continuum crowd dynamics. The San Fermín opening ceremony could again provide the setting for a natural experiment, as musicians and security personnel emerge from the city hall and move through the crowd. And every 6 July presents a new chance to gather more data.

**Johanna L. Miller**

## References


1. J. J. Hopfield, "Collective properties of neuronal networks," AIP History Programs, https://www.youtube.com/watch?v=ib6futs-fss&t=777s.
2. F. Gu et al., *Nature* **638**, 112 (2025).
3. A. Sieben, A. Seyfried, *Saf. Sci.* **166**, 106229 (2023).
4. See, for example, I. Echeverría-Huarte et al., *Sci. Rep.* **12**, 2647 (2022).


# Plasma channel guides electrons to 10 GeV

Femtosecond-scale laser pulses make a plasma waveguide that helps wakefield-surfing electron bunches keep their energy focused.

High-energy, ultrafast-moving particles have a wide array of both applied and fundamental scientific uses, such as imaging biological tissue in detail and probing questions about the building blocks of matter. But accelerating particles to high energies is not a cheap or easy task. The advanced accelerators being used for that purpose rely on multikilometer-scale facilities. That's because RF cavities—metal chambers that generate the electric fields used to accelerate particles—have fundamental limits that max out their electric field gradients at tens of megaelectron volts per meter. At those gradients, it takes kilometers to get particles up to the tens of gigaelectron-volt energies achieved at facilities such as SLAC and the German Electron Synchrotron (DESY). But there's another approach that can produce electric field gradients that are thousands of times as strong as those produced in RF cavities.

When a laser pulse is sent through an ionized plasma—a sea of electrons and ions—it pushes aside the lighter electrons. The charge rearrangement induces a strong electric field gradient that moves like a wave through the plasma and trails behind the front of the laser pulse, as shown in figure 1. Much like a surfer can ride the wake behind a speedboat, electrons can get trapped in the plasma wave

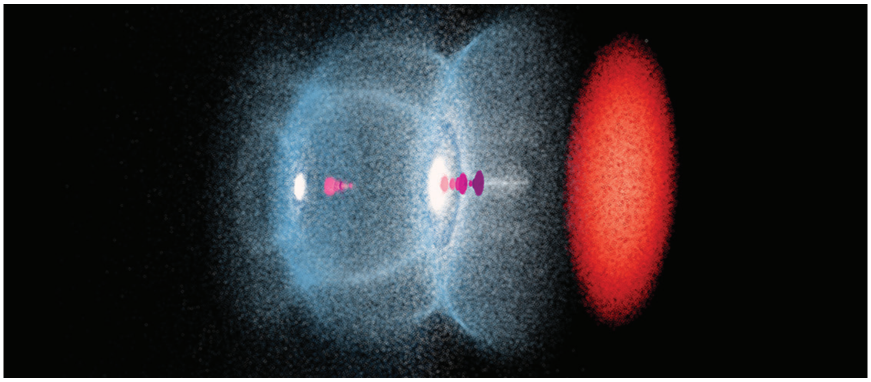

**FIGURE 1. AN ELECTRON WAKE** (blue in this still from a simulation) trails behind a femtosecond-scale laser pulse (red), with a diameter of about 50 μm, traveling from left to right through a plasma. Electron bunches (magenta) that ride the wake are accelerated to high energies over short distances. (Image courtesy of Carlo Benedetti, Lawrence Berkeley National Laboratory.)

and accelerate to high energies. That's the idea behind the aptly named laser wakefield acceleration method (see Physics Today, January 2015, page 11). Researchers hope that the methodology, conceived of in the 1970s[1] and invigorated by promising electron beams produced by three research groups in 2004,[2] could one day produce powerful electron accelerators that are orders of magnitude smaller than today's.

Now Alex Picksley and colleagues at the Berkeley Lab Laser Accelerator (BELLA) Center at Lawrence Berkeley National Laboratory, in collaboration with Howard Milchberg's group at the University of Maryland, have used wakefield acceleration to produce a well-controlled 9.2 GeV electron beam with charge that extends beyond 10 GeV.[3] "It's a really impressive feat," says Rob Shalloo, an experimental plasma accelerator physicist at DESY. "And they've done it with incredible diagnostics of what's happening throughout the acceleration process. That allows us to learn how we can improve our plasma accelerators and gives us a road map as to where to go from here."

Strictly speaking, 9.2 GeV is not a wakefield record; in 2024, researchers at the Texas Petawatt Laser Facility reported that they had generated 10 GeV electrons in much larger bunches.[4] The feat was achieved by injecting lots of electrons into the wake with the use of nanoparticles, but it also required about six times as much laser pulse energy as the BELLA group's result did and produced a wider energy spread. The accelerator community's enthusiasm about the BELLA result is not only because of the magnitude of the electron energy achieved but also because of the control exerted over several aspects of the process.

## Narrowing the channel

A few key challenges plague the wakefield acceleration method. One is to keep the energy spread of the electrons narrow, an important requirement of precision electron-beam applications. Electrons can enter the wake by being injected intentionally and by being pulled in from the surrounding plasma. Electrons that enter the wake a little bit ahead of or behind the ideal position will accelerate to different energies and increase the spread. Another challenge is to keep the drive beam—the laser pulse that creates the wake—focused so that it can impart maximum energy to the plasma wake. Laser pulses can be focused with lenses, but they don't stay focused for long; their radial expansion detracts from the efficiency of the process.

One way to keep the drive beam narrow is to use a plasma waveguide, an approach that has been in use for decades. (See the article by Wim Leemans and Eric Esarey, Physics Today, March 2009, page 44.) In 2018, researchers at the BELLA Center generated 8 GeV electrons,[5] a near-doubling of the record energy for the method at that time. That energy was achieved with a plasma waveguide created by sending nanosecond laser pulses into a plasma-filled tube about the length of a pencil. The pulses heated the plasma, which expanded to produce a shaft of lower-density plasma down the center of the tube that was

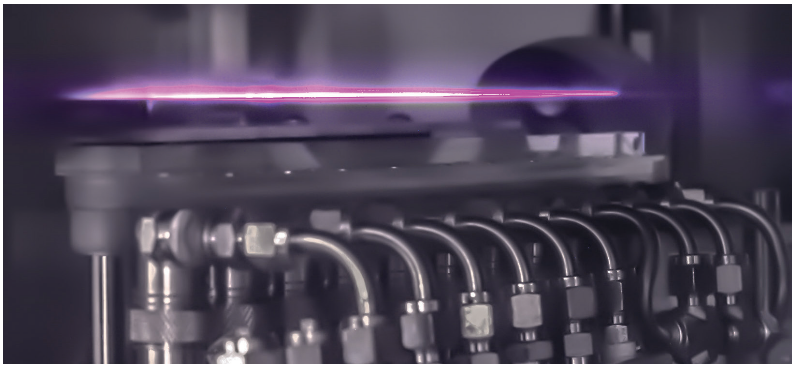

**FIGURE 2. A PLASMA CHANNEL**, 30 cm long, acts as a conduit for a laser pulse used to create a plasma wakefield that accelerates electrons to high energies at the Berkeley Lab Laser Accelerator Center. The plasma is formed when a 40 fs laser pulse is shot through a sheet of supersonic gas that streams upward out of an aluminum slit. The pink color is a result of the hydrogen gas used to generate the channel. (Photo courtesy of Alex Picksley and Anthony Gonsalves, Lawrence Berkeley National Laboratory.)

surrounded by cooler, denser plasma. The low-density channel helped to keep the laser focused.

But laser heating, which relies on electron–ion collisions, becomes ineffective at the low plasma densities that are ideal for acceleration. So the BELLA team turned to another type of plasma guide, one made with a shorter-duration laser pulse and no tube: a hydrodynamic optical-field-ionized (HOFI) plasma channel,[6] shown in figure 2. (See the Quick Study by Bo Miao, Jaron Shrock, and Howard Milchberg, Physics Today, August 2023, page 54.)

The researchers make the HOFI plasma channel by sending a femtosecond-scale laser pulse through a sheet of supersonic gas. The pulse optically ionizes a string of plasma that radially explodes into the surrounding gas to form a cylindrical shock wave. The resulting plasma channel lasts for a matter of nanoseconds and appears as a flash of light, tinged with different colors depending on the gas used to generate it. "It has the same operating principle as an optical fiber," says Shalloo, "but is immune to laser damage, as a fresh plasma waveguide is formed each time the laser is fired."

## A look inside

With the HOFI channel in place, a second laser pulse acts as the drive beam. The HOFI channel brings multiple benefits: It is narrower than the laser-heated channels, and thus better matches the width of the drive beam; it has a lower-density plasma; and because it is formed in an open space, researchers can easily change its length by changing the length of the gas sheet used to create it. That last property was used to probe details of how the drive beam was behaving inside the plasma accelerator. Although the full channel length was 30 cm, the BELLA team adjusted the channel length for each pulse to see how the laser pulse moved through the full length.

"We were able to peer into the interaction, seeing how the laser pulse evolves as a function of distance," says Anthony Gonsalves, a member of the BELLA team. "It shows us the way, experimentally, that simulations have been showing us for a long time."

With a look inside the process, the researchers could see that the shape of the laser pulse was causing a loss of energy. To maximize energy output, crystal-based petawatt lasers are designed to produce pulses with a top-hat intensity profile—a uniform full-power circle. But the optimal intensity profile for plasma wakefield generation is a Gaussian distribution that is tapered at the edges. Simulations show that with an optimized laser profile and the same energy pulse, the BELLA setup should be able to achieve more than 13 GeV. "Now we have motivation for laser builders," says Gonsalves.

The other HOFI channel advantage—lower-density plasma—was crucial to reaching the 10 GeV range. That's because the energy reached is inversely proportional to the plasma density. As a laser travels through plasma, the refraction of light slows it down. If the laser pulse slows down enough, the wakefield-surfing electrons, which are moving at nearly the speed of light, can catch up to the drive beam and start to decelerate. Lower-density plasma means a faster-moving laser pulse. Additionally, denser plasma leads to more self-trapping of electrons, in which the electrons get pulled from their optimal position in the bubble behind the drive beam and lose energy.

A localized addition of nitrogen gas to the plasma was used to precisely inject electrons at the optimum position in the wake bubble. The controlled injection, which was facilitated by the supersonic gas sheet configuration, resulted in a narrow energy spread and minimized self-trapping. "The guiding technology enabled us to boost the electron energy and control the quality of the accelerated beam a little bit more," says Picksley.

The results offer a clear path for making further advancements in laser wakefield acceleration. Improvements to lasers are part of that path. In addition to a petawatt laser with a Gaussian intensity profile, a laser with a higher pulse repetition rate and improved energy efficiency would make the method more viable for regular use. The team at BELLA sees promise for those qualities in fiber lasers. Another strategy on the table is staging multiple lasers together to combine their power. Although wakefield accelerators aren't yet in common usage, many labs are already exploring their applications. Lawrence Berkeley National Laboratory has used the technology to power a free-electron laser, and DESY has plans to use wakefield acceleration to feed electrons into its PETRA synchrotron light source. The latest result brings tabletop particle accelerators a little bit closer to fruition.

**Laura Fattaruso**

## References


1. T. Tajima, J. M. Dawson, *Phys. Rev. Lett.* **43**, 267 (1979).
2. S. P. D. Mangles et al., *Nature* **431**, 535 (2004); C. G. R. Geddes et al., *Nature* **431**, 538 (2004); J. Faure et al., *Nature* **431**, 541 (2004).
3. A. Picksley et al., *Phys. Rev. Lett.* **133**, 255001 (2024).
4. C. Aniculaesei et al., *Matter Radiat. Extremes* **9**, 014001 (2024).
5. A. J. Gonsalves et al., *Phys. Rev. Lett.* **122**, 084801 (2019).
6. R. J. Shalloo et al., *Phys. Rev. E.* **97**, 053203 (2018).