RIDING THE PLASMA WAVE OF THE FUTURE

by Matthew Early Wright

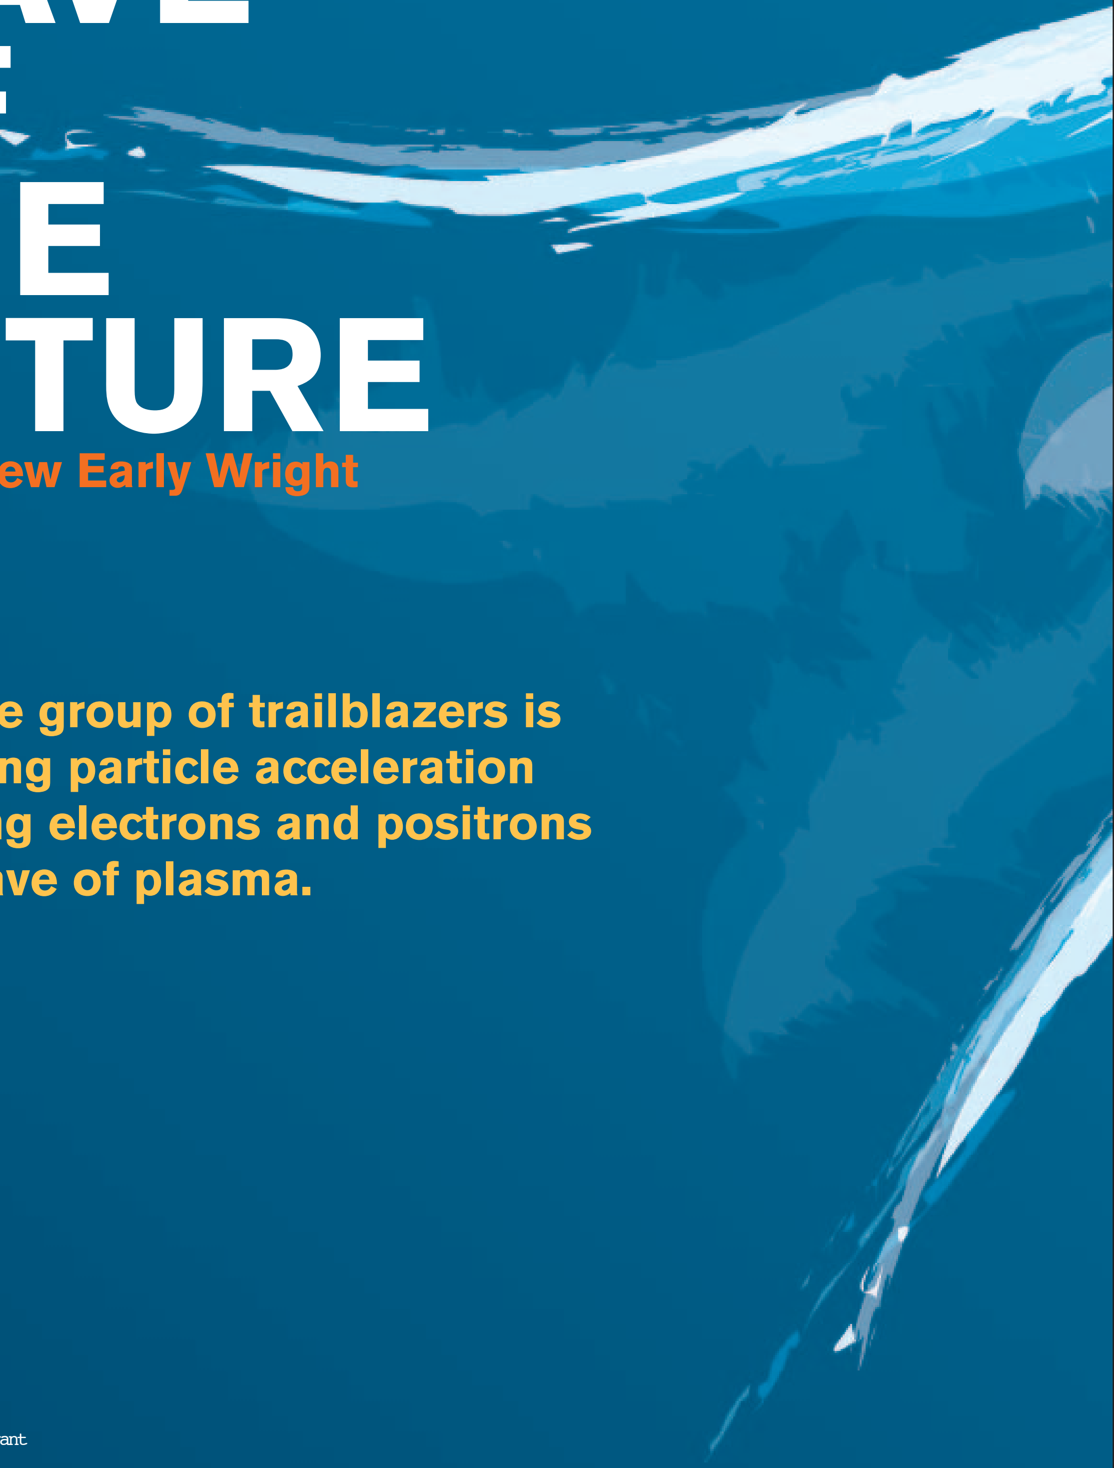

A creative group of trailblazers is reinventing particle acceleration by making electrons and positrons surf a wave of plasma.

Illustrations: Aaron Grant

Particle accelerators have a history of getting larger. Using current technology, it seems the easiest way to produce more powerful machines is to increase their size. But can physicists afford to continue building these enormous machines, or does accelerator technology need a fundamental revision?

Caolionn O'Connell of Stanford Linear Accelerator Center thinks so. "We're working on something that can revolutionize the way accelerators work," she says.

O'Connell is among a handful of physicists pioneering the field of "plasma wakefield" acceleration: coaxing electrons and positrons to surf a wave of ionized plasma created in the wake of a laser burst or electron beam. The technology might one day make compact, tabletop-sized accelerators a reality, and dramatically increase the rate of acceleration possible with traditional radiofrequency machines.

But tabletop accelerators won't appear soon. "There's still some fundamental physics we don't know about," warns Chris Barnes, O'Connell's colleague at SLAC. "We still have some basic engineering questions to answer."

> **“We have demonstrated this to be an amazing technology,” O’Connell says. “But we have a lot of little steps to take before it’s ready for use.”**

The next, next generation

The radiofrequency (rf) technology driving existing particle accelerators is showing its age. The rf-generated electric fields that traditionally accelerate particles are constrained by an absolute upper limit, beyond which the fields become unstable. Worse yet, pushing past this limit may even melt the accelerator’s cavity. As O’Connell puts it, “that would be very, very bad.”

Plasma wakefield accelerators can subvert this problem by filling the accelerator cavity with plasma—a “soup” of ionized gas in which electrons are stripped from atoms. Plasma can handle heavier energy flows, enabling these accelerators to dodge the issue of structural failure.

The acceleration rate of a modern particle accelerator is often measured in MeV (millions of electron volts) per meter. That is, electrons typically feel millions of volts of electric fields accelerating them through every meter of the machine. Most rf machines are pushing up against the hard ceiling on acceleration rate, which tops out at around a few tens of MeV per meter. So the only way to generate more powerful collisions using rf is to make longer, more expensive accelerators.

But by some estimates, plasma acceleration technology could potentially achieve rates in the GeV (billions of electron volts) per meter range. If plasma accelerators deliver on this promise, huge rf accelerators like the Large Hadron Collider (LHC) and the proposed International Linear Collider (ILC) might be the last of their kind.

The improved performance of plasma technology translates to more acceleration in less space. This can mean one of two things. On one hand, plasma machines may one day make the vaunted “tabletop” accelerator a reality. Lengths measured in meters rather than kilometers could bring accelerator labs within the reach of any university or industrial lab.

On the other hand, O’Connell envisions plasma “afterburners” that could be used to upgrade existing rf accelerators. If this application materializes, particle physicists might be able to use facilities like LHC and ILC to perform experiments beyond their imaginings.

If large, superconducting accelerators like LHC and ILC are the “next generation,” then plasma technology offers promise for the “next, next generation.” O’Connell estimates that practical, commercially feasible plasma machines are at least 10 to 20 years away.

“We have demonstrated this to be an amazing technology,” O’Connell says. “But we have a lot of little steps to take before it’s ready for use.”

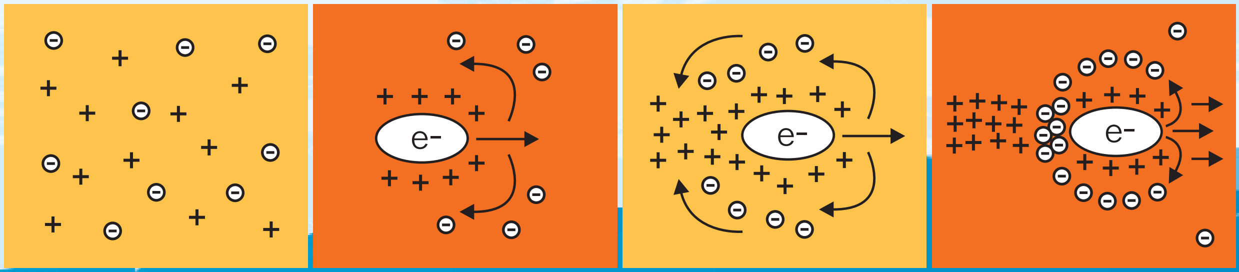

Picturing the plasma acceleration

1. A plasma, made of positive ions and free electrons before an electron bunch enters.
2. The electron bunch enters the plasma, repelling all the free electrons from its path, and attracting the positive ions. The moving electron bunch leaves a wake of positive ions behind it as it passes.
3. The displaced free electrons are now attracted to the mass of positive ions behind the electron bunch.
4. The free electrons in their new position give the electon bunch an acceleration.

Surf's up

For O'Connell and Barnes, thinking about surfing is part of their work. And for good reason: it's a fitting analogy to use in describing plasma accelerators. More accurately, plasma acceleration is a lot like "wake surfing." Unlike catching natural waves at the beach, wake surfing depends on a specially designed, rear-ballasted boat to make waves large enough to surf. As the boat speeds across the surface of a lake with its bow tilted skyward, it leaves a powerful, rippling wake behind it. The surfer, initially towed by a rope, can eventually let go and free-surf in the wake.

To extend the metaphor to plasma accelerators, imagine a mass of plasma as the lake. A driving force, such as an electron pulse or a laser pulse, acts as the boat. The particle bunch is the surfer, gaining energy as it is pushed along by the plasma wake.

The two types of "boats" used to generate the plasma wake—either electrons or lasers—involve different physics. Accelerator physicists tend to specialize in one form or another, partly because they require access to very different facilities. O'Connell and Barnes perform electron-beam driven experiments at SLAC.

The system they have developed uses a single electron bunch as both the boat and the surfer. "The leading electrons drive a wave in the plasma that's analogous to the wake behind any boat," Barnes explains. "The trailing electrons then surf on the wake, gaining substantial amounts of energy at the expense of the first particles." But this method requires a pre-existing beam source, much like the one at SLAC, to provide the energy needed to drive the process. The presence of the plasma increases the efficiency of the acceleration and prevents overloading of the beam cavity.

"Basically, the system acts as a transformer where we take energy from the head of the beam, and give it to the tail," says O'Connell. This one-bunch system is effective, but O'Connell believes it can be improved upon. Eventually, their research group plans to transition to a two-bunch system, with a leading bunch used exclusively to drive the wake.

Laser driven acceleration is similar in principle to beam driven acceleration. But instead of depending on the leading electrons in the bunch to generate the plasma wake, a short, 10 femtosecond pulse (10 quadrillionths of a second) from a high-energy laser is used instead. A recent set of publications describes the results of three separate groups working on laser acceleration.

Unlike in beam accelerators, for which the electron bunch is also the accelerating force, injecting the electron bunch into the plasma wake at the right time and place has been a challenge for laser accelerators. Zulfikar Najmudin, who works with laser accelerators at Imperial College, London, explains that "like the surfer, the particles must have some velocity in the direction of the wave, otherwise the wave will just wash over them." But each of the research groups independently demonstrated that if the wave is large enough, electrons can be caught up in the wave "just like the froth that is sometimes carried forward on the crest of breaking ocean waves."

Building a better surfboard

While the underlying principle is the same for laser- and beam-driven accelerators, the technical challenges of each are very different. Najmudin thinks part of this is due to a historical gap. “Electron beam facilities have more than 60 years of development, whilst the first laser beam was only demonstrated 40 years ago,” he explains. For now, this means that electron beam accelerators are more reliable and reproducible than laser accelerators, he says.

Csaba Toth of Lawrence Berkeley National Laboratory (LBNL) believes this will improve as lasers become more compact and efficient. “The current systems are still multi-tabletop, complex arrangements,” he says. “We expect major changes in the next five to six years on this front as diode-pumped and high-efficiency lasers become more widespread and economical.”

The nature of the physical interaction between a laser pulse and a mass of plasma also presents a problem. Unlike an electron pulse, which tends to stay focused in plasma, a laser pulse will spread out as it travels farther, spreading the wake with it. The longer the plasma chamber, the more of a problem this becomes. The introduction of more efficient lasers will help, but there is still a limit beyond which the problem must be attacked from other angles.

Toth and principal investigator Wim Leemans, with the rest of their group at LBNL, have devised a possible solution. They used a preformed channel of dense plasma to guide the laser beam, in much the same way an optical fiber can guide light. This strategy extended the length of the tightly focused plasma wake by an order of magnitude.

Even with such modifications, laser acceleration is so far only possible over minuscule distances. Laser-driven setups such as those at LBNL and Imperial College function on the millimeter scale. In contrast, the beam-driven plasma chambers O’Connell and Barnes use in their work are 10-30 centimeters long. To extend the acceleration distance of laser machines, Toth believes the answer lies in linking together, or “staging” several laser units.

Barnes enjoys the advantages in working with beam-driven systems. “The problem is lasers don’t stay focused, so the net energy gain is very small,” he says. “Electron bunches want to stay focused, so we can exploit the plasma, rather than fight against it.”

This is one reason why O’Connell chose to do her doctoral work at Stanford. “SLAC has a unique combination of a high-energy facility and a high-quality beam,” she says. “Both of these are absolutely necessary. This is an experiment that has to happen at SLAC.”

If rf-boosting “afterburners” become a reality, these will almost certainly be of the beam-driven type. Since these units depend on preformed electron bunches produced by standard rf acceleration, relatively short plasma units could be inserted directly into the beamline to give a substantial boost. O’Connell estimates that this strategy could double the energy of a facility like SLAC.

“For the time being, we’re still in the proof-of-principle stage,” Barnes says. “Just proving that we’ve substantially accelerated particles for a macroscopic distance is our major result.”

“This work requires some wildly disparate skills–we need to know a little about a lot of different things.”

Over the horizon

Everyone working in the field agrees that it will be decades until either beam- or laser-driven plasma wakefield accelerators are used to probe the nature of matter. But other applications may be just around the corner, comparatively speaking.

“Beams with energy on the GeV scale are expected to be generated in the next few years,” says Leemans. This could open up a number of uses for plasma acceleration, such as generating light sources for fine-scale imaging, producing radioactive isotopes for medical use, and creating gamma-ray and terahertz radiation for materials testing.

Eventually, the laser and electron beam camps may not be divided along such distinct lines. Laser-driven machines offer compact size, stand-alone capability, and higher energy outputs. On the other hand, beam-driven systems are reliable, tightly focused, and can generate a plasma wake over substantial distances.

Leemans believes that the two approaches are equally important to pursue. “Developing both technologies and understanding the similar physics between both beam- and laser-driven systems is important and complementary,” he says.

Barnes agrees: “It might be that a hybrid is more effective. But that’s a ‘pie-in-the-sky’ scenario, way off in the future.”

For now, the plasma researchers all seem to enjoy being one of a handful of surfers on a particularly choice wave. “This work requires some wildly disparate skills—we need to know a little about a lot of different things,” says Barnes. “But I like being way in front of theory.”

O’Connell knows a transition is coming, from proof-of-principle to practical application and meeting the needs of particle physicists. “But for now the field is so new, we can publish every piece of data we get,” O’Connell says with a satisfied grin. “It’s so fringe, and so cool, with so many possibilities. How could you not love it?”