

500+ Xilinx FPGAs Search for Elusive Higgs Boson at 1.5 Terabytes per Second

With an array of more than 500 Virtex and Spartan FPGAs processing 1.5 terabytes of real-time data per second, scientists at the Fermi National Accelerator Laboratory hope to track down the last subatomic particle — the Higgs boson.

WE FIND OURSELVES IN A BEWILDERING WORLD. WE WANT TO MAKE SENSE OF WHAT WE SEE AROUND US AND ASK: WHAT IS THE NATURE OF THE UNIVERSE?

— STEPHEN W. HAWKING, LUCASIAN PROFESSOR OF MATHEMATICS AT CAMBRIDGE UNIVERSITY

by Mark Havener
Science Writer, Software Consultant
havener@inreach.com

Nick Hartl
Gold FAE, Avnet Design Services
nick.hartl@avnet.com

Inside the four-mile long Tevatron, the world's most powerful particle accelerator, protons and antiprotons collide at nearly the speed of light, creating bursts of energy and showers of millions of subatomic particles. If theoretical predictions are correct, over the next five years a million billion collisions (10^{15}) will produce only 120 events with the characteristic pattern most easily recognizable as evidence of the existence of Higgs boson.

Discovery of the Higgs boson will verify the "Standard Model" theory that is the foundation of modern particle physics. Finding a Higgs boson needle in this haystack of particles, however, requires a digital signal processing (DSP) system capable of gathering and processing 1.5 terabytes of data per second.

The scientists at Fermilab (www.fnal.gov) in Batavia, Illinois, investigated multiple ways of capturing this enormous data flow before finally settling on an array of more than 500 Xilinx Virtex™ and Spartan™ FPGAs. The Xilinx components are assembled into a “trigger” – a homemade, massively parallel supercomputer programmed as a multi-level pattern recognition filter. Tracks left behind charged particles are examined, and complex algorithms recognize and discard known patterns. Data about unknown particles are passed on and stored for later analyses. These data could prove that the Higgs boson exists. If it does, the proof will not only extend our understanding of the universe, but it may also earn a Nobel Prize for the physicists at Fermilab.

The Standard Model and the Higgs Boson

Modern theoretical physics describes the world as composed of twelve fundamental matter particles in three generations. First-generation particles are stable and can easily be found in nature, while second- and third-generation particles are extremely unstable and exist for only a tiny fraction of a second before decaying into other particles. Force-carrying particles interact with the matter particles. These particles and their interactions make up the Standard Model of Fundamental Particles and Interactions (Figure 1). (For more information on the Standard Model, see “The Building Blocks of Matter,” www.fnal.gov/pub/inquiring/matter/madeof/, and “The Particle Adventure: Fundamentals of Matter and Force,” particleadventure.org/particleadventure/.)

The four known fundamental force-carrying particles are photons, W and Z bosons, and gluons. First-generation matter particles include up quarks, down quarks, and electrons. Two down quarks and one up quark form a neutron; two up quarks and one down quark form a proton. Protons, neutrons, and electrons combine to form atoms, atoms combine to form molecules, and mol-

ecules combine to form The World As We Know It. All of this is supported by physical evidence gathered from experiments.

Experimental measurements show that most fundamental particles have a very small mass, typically 1 giga electron volt (GeV/c^2) or less. An electron has a mass of 0.511 mega electron volt (MeV/c^2), or 9.11×10^{-31} kilogram. The photon mass is theoretically zero. Indeed, experimental

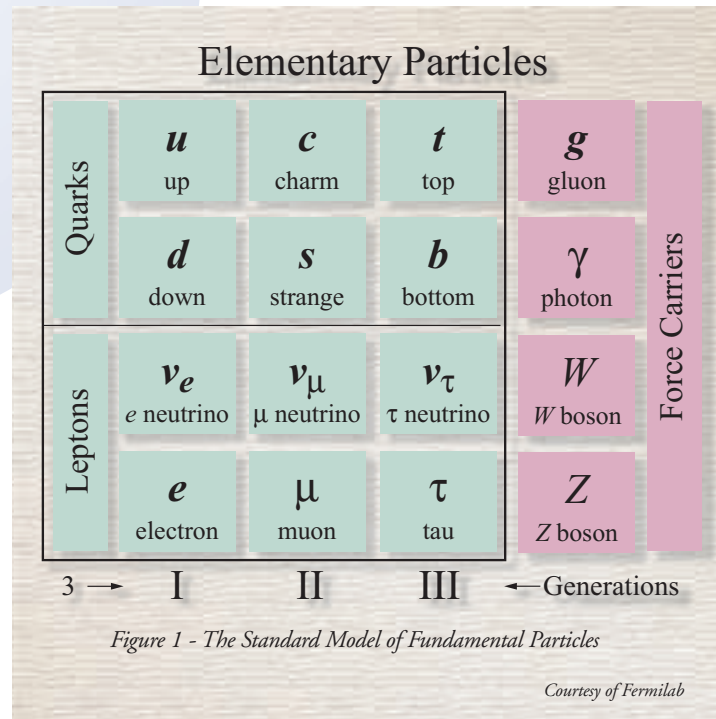
1964, British physicist Peter Higgs postulated the existence of an invisible field that permeates the universe and is responsible for endowing all matter with mass. (To find out more on the Higgs theory, see “The Higgs Boson,” www.jlab.org/~cecire/higgs.html.) According to theory, when a subatomic particle, such as an electron or quark, moves through the Higgs field, the particle acquires mass. The existence of a fundamental force-carrying particle – the Higgs boson – supports the simplest theory that would explain the large masses of the W and Z bosons. The Higgs boson exists as both a field and particle, because matter and force exist as both fields and particles, according to Quantum Theory.

In 1971, Glashow, Salam, and Weinberg included an *ad hoc* Higgs mechanism in calculations that predicted the massive W and Z bosons. These predictions were beautifully confirmed a decade later with their discovery by Carlo Rubbia’s group of experimenters at CERN (European Organization for Nuclear Research, public.web.cern.ch/Public/). The prediction and discovery led to the award of three Nobel Prizes.

If the Standard Model is correct, high-energy collisions in the Tevatron will produce Higgs bosons. Each Higgs boson will exist for only a fraction of a second before decaying, but measurement of the angle and velocity of the resulting decay particles will provide proof of its existence. Discovery of the Higgs boson will provide additional confirmation of the Standard Model and expand physicists’ understanding of mass.

Subatomic Collisions

When protons and antiprotons collide, force particles and unstable second- and third-generation matter particles are created. Because these particles are so short-lived, the only evidence of their existence is the tracks they leave as they decay, as well as the tracks left by the other particles created in the decay process. The Fermilab



evidence indicates that the photon mass can’t be greater than $10^{-33} \text{ GeV}/c^2$, in excellent agreement with theoretical prediction. However, the W boson and Z boson masses have been measured as $80.4 \text{ GeV}/c^2$ and $91.187 \text{ GeV}/c^2$, respectively.

The large masses of these bosons, and their observed interactions with known elementary particles, create a curious inconsistency in the mathematical equations that describe the behavior of matter and force. The equations predict the probability of two very high-energy particles colliding is greater than one. It’s like knowing you’ll always win the lottery, and you won’t even have to buy a ticket. This would be nice, but it’s impossible.

One way to resolve this theoretical dilemma is to introduce additional particles. In

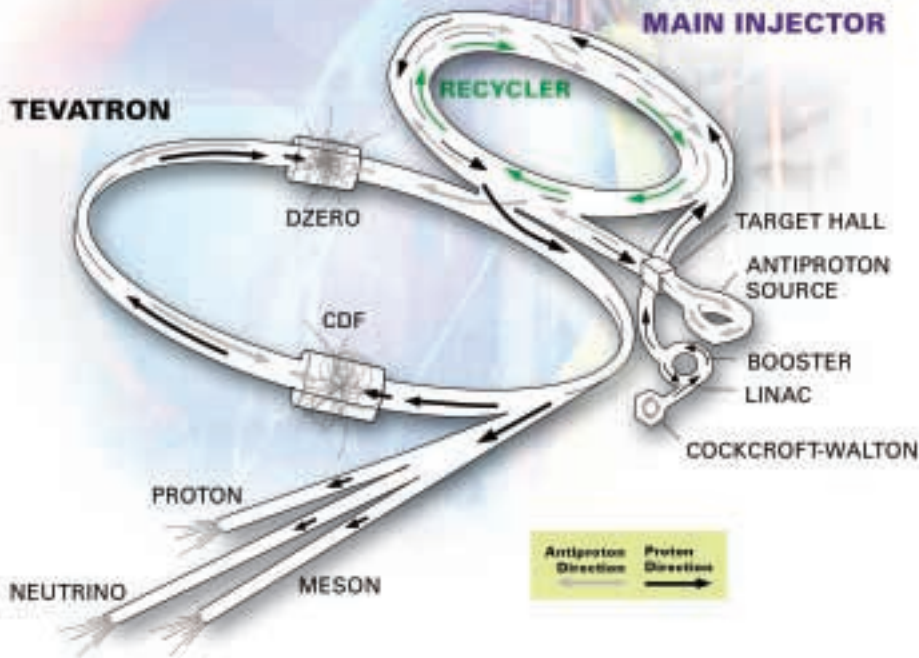


Figure 2 - The Fermilab Accelerator Chain

Courtesy of Fermilab.

physicists use the detectors inside the Tevatron to observe these tracks and the FPGA array to analyze them for proof of the existence of these ephemeral particles.

To get a sense of how difficult this is, imagine a child's Hula Hoop® toy suspended at the 50-yard line of a 100-yard American football field with a machine gun in each end zone. The machine guns aim for the center of the Hula Hoop and fire as fast as they can. Sometimes the bullets collide, sometimes they don't. Some collisions are head-on, and some are indirect glancing blows. Your job is to measure the direction and velocity of the bullet fragments after each collision and then use that data to analyze and re-create the collision.

Just as the imaginary machine-gun bullet collisions are not all identical, the proton and antiproton collisions inside the Tevatron are not all identical. Some collisions are direct, some glancing. The protons and antiprotons travel at slightly different velocities and orientations. The types of particles created – and the directions and velocities in which they are scattered – depend on many factors, with each collision producing a different and distinct “signatures” of particles.

Out of the 10^{15} proton-antiproton collisions expected in Run II of the Tevatron, a very small fraction will result in a top quark and top antiquark meeting head-on. In theory, quark-antiquark collisions can produce a Higgs boson in three distinct ways, each with a distinct particle signature. The current estimate is that all these collisions will produce less than 20,000 Higgs particles – one Higgs boson for every 50 billion collisions. And only 120 of the collisions will yield the characteristic pattern most easily recognizable as Higgs production.

The Tevatron

The Tevatron is the world's most powerful particle accelerator (Figure 2). It uses oscillating magnetic fields to push protons and antiprotons in opposite directions, reaching nearly the speed of light on the four-mile circular path before colliding in one of the two detectors (CDF and DZero).

The process begins with the ionization of hydrogen atoms, creating two electrons and one proton. These particles are accelerated to an energy of 400 MeV and passed through a carbon foil filter that removes the electrons. The protons are further accelerated to 8 GeV and sent into the

Main Injector, where they are yet further accelerated to 120 GeV before some of them are siphoned off and crashed into a fixed nickel target. These collisions produce secondary particles, most of which are ignored and discarded – with the exception of the antiprotons, which are collected and sent back to the Main Injector. About half the size of the Main Accelerator, the Main Injector increases the energy of both the protons and antiprotons to 150 GeV before injecting them into the Main Accelerator.

Inside the Main Accelerator, protons and antiprotons are accelerated with powerful electromagnetic fields, using harmonic oscillation at gigahertz frequencies. As the particles reach higher speeds, additional magnetic force is used to bend the beams into a circular path. Protons travel clockwise and antiprotons counterclockwise, faster and faster, to within 200 miles an hour of the speed of light. At this speed, the energy of the particles approaches a thousand billion electron volts, or one tera electron volt (1 TeV) – and this is where the Tevatron gets its name. The beams are slightly offset from each other, crossing at two points. High energy collisions occur at these intersections, where the detectors are located.

The DZero Detector (www-d0.fnal.gov) uses Silicon Microstrip Tracker (SMT) and Central Fiber Tracker (CFT) subdetectors to record tracks of charged particles produced in the collisions. The CFT is made of scintillating fibers mounted on eight concentric cylinders. A charged particle passing through the fiber produces a tiny amount of light that is converted into an electrical pulse by visible light photon counters. These are small silicon devices with an array of eight photo sensitive areas, each 1 mm in diameter. Recorded electric signals make it possible to reconstruct an accurate three-dimensional image of the particle's path.

The DZero Upgrade

The Tevatron collider began operating in 1983, with continuing improvements and additions over the next 14 years. The original DZero (a.k.a. DØ) Detector was commissioned on Valentine's Day 1992. By any

measure, Run I was a successful experiment, monitoring a few trillion collisions and culminating with the discovery of the top quark in 1995.

In trillions of collisions, the physicists observed only 90 top quark events – events with a signature similar to what the Standard Model would predict if a top and an antitop quark were produced in the collision. Higgs candidates are even more elusive than top quarks, with an expected production of one Higgs boson in every 50 billion collisions. To find a Higgs, many more collisions will be needed. In 1997, the Tevatron was shut down for final installation of the Main Injector and Antiproton Recycler to increase the luminosity of the beams, and for improvements to the particle detectors to monitor the additional collisions.

The improved DZero Detector (Figure 3) was initially designed using commercially available DSP components. Computer models were built, with simulations



Figure 3 - DZero, side view
Courtesy of Fermilab

designed and executed to verify the design's functionality. Running the simulator revealed that the DSP design was wholly inadequate to handle the number of collisions that the increased luminosity would produce. A new approach was needed ... and the final design of the new DZero Trigger relies heavily on Xilinx FPGAs.

The DZero Trigger

The DZero Trigger is a multi-level pattern recognition filter that processes nearly a trillion signals every second. Its job is to identify the collisions that are most likely to produce Higgs particles, and save that data for later detailed analyses. When a charged particle passes through the CFT, the light from the fibers is first converted to an electric signal. Next, the digital signal is sent to a DZero Trigger subsystem called

the Central Track Trigger (CTT), which progressively filters signals (Figure 4 - d0server1.fnal.gov/projects/VHDL/General/ctt-diagram.pdf). Each collision creates multiple particles and each particle creates multiple signals. From the resulting signals, it is possible to reconstruct the paths of the particles involved in the collision. Complex algorithms in the DZero Trigger identify and separate “interesting” signals from “uninteresting” ones.

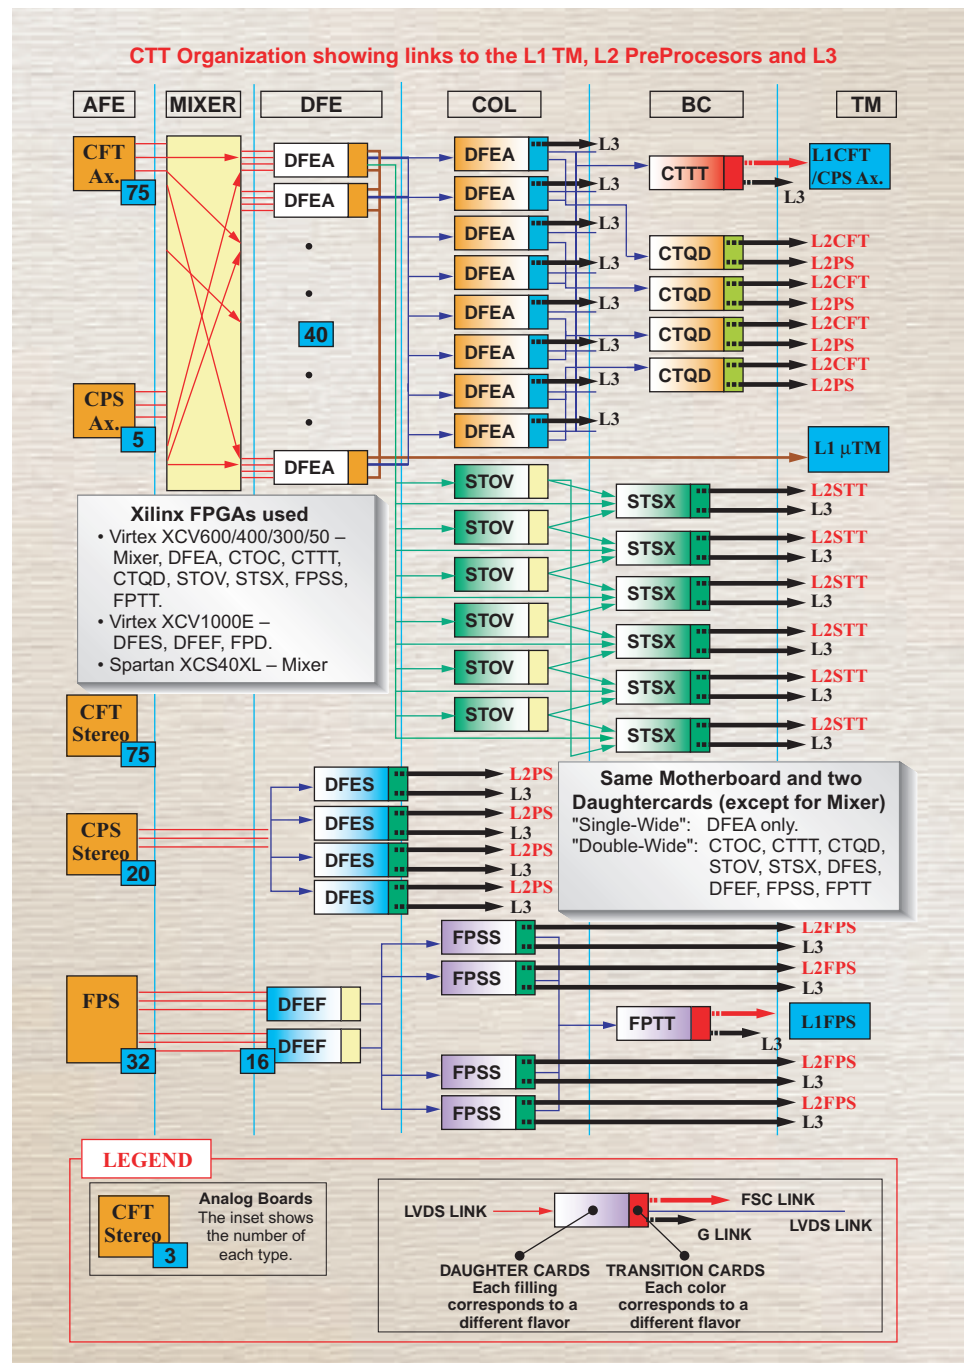


Figure 4 - The DZero CTT FPGA Array

Courtesy of Fermilab

Each second the DZero Trigger must look at 7 million collisions and decide in real time which ones to save. Only a fraction of the collisions can be saved, so recognizing and saving the ones that are most likely to show important events (like the production of a Higgs boson) is the key to success.

The trigger has three separate levels called L1, L2, and L3. Each level has a progressively finer filter. The data output rate at each level is lower than the data input rate. The difference between the rates determines how much data is rejected.

The DZero FPGA Array

The complete trigger consists of 582 Xilinx FPGAs ranging from Spartan FPGAs to Virtex 300s to Virtex-E 1000s. The 582 FPGAs are assembled into 21 unique designs that are repeated to make multiple data channels. The common footprint of the Virtex family allowed a design utilizing a single printed circuit board that could be populated with different chips. This was a great advantage, because the single common board could be customized by placing different numbers and sizes of FPGAs on it to create the various subsystems used in the trigger.

Once the hardware design was completed, the next major challenge was programming the chips. The function of the DZero Detector and the data it produces had to be fully understood and incorporated into an algorithm that would save the correct data. The most difficult task was creating an algorithm that operated in the minimum amount of time. Because data cannot be discarded until the system reaches a save/don't save decision, and because there is a finite data buffer, it is important that the calculations be completed before the buffer is overwritten. Completing this task in the limited time available proved to be very challenging.

Choosing the Right FPGAs

After careful consideration, the DZero team chose Xilinx FPGAs. As Jamison Olsen, principal EE on the project, explained: "The common footprints used for the Virtex family allowed us to lay out one board that could be populated by a variety of different size chips with no change to the board. This was a great



Figure 5 - Single-wide daughtercard

Courtesy of Fermilab

advantage, because we could design a common printed circuit card that we could customize by placing different numbers and sizes of FPGA to create the various subsystems used in the trigger."

Olsen continued, "Other considerations that led us to Xilinx were very fast fitting of the devices, a good price-to-performance ratio, and several Virtex features, including the flexible RAM architectures."

Common Footprint

The DZero Trigger system architecture uses a base carrier card to take care of backplane I/O and to carry one or two daughtercards with the Xilinx FPGAs that do the actual work (Figures 5 and 6). This architecture allowed the team to build a variety of subsystems on common hardware. The three different tiers of processing within CTT have different processing requirements, so the CTT was built with the appropriate components at each level. Because the Xilinx components all share a common footprint, the base printed circuit boards can all be identical, providing both an initial cost savings and a much more efficient store of replacement parts.

The common footprint also provides the ability to boost performance by reconfiguring with more powerful devices as they become available during Run II.

Memory

The Xilinx components have more on-board RAM than competing devices. The Virtex-E FPGAs have as much as 1 Mb of internal configurable distributed RAM and up to 832 Kb of synchronous internal block RAM. Data cannot be discarded until the trigger system reaches a decision whether or not to save it, so the generous RAM provides a buffer to store the data while the system completes its calculations. Even with this much RAM, completing the calculations in as few clock cycles as possible before the buffer was overridden was a challenging task. The Virtex flexible RAM architecture also came into play. The hierarchical memory system LUTs are configurable as 16-bit RAM, 32-bit RAM, 16-bit dual ported RAM, or 16-bit shift register, with fast interfaces to external high-performance RAMs.



Figure 6 - Double-wide daughtercard

Courtesy of Fermilab

Performance and Bandwidth

In the development phase, the scientists decided the initial DSP design was unacceptably slow. They selected Xilinx components because the FPGAs were faster. The Virtex FPGAs operated at system speeds as fast as 200 MHz, and the Virtex-E parts achieved more than 311 MHz. The DZero processing is extremely I/O bound. With more than 1.5 trillion events per second, the amount of

data flowing into the array is staggering. In the Virtex-E family, I/O performance in each component is 622 Mb/s using source-synchronous data transmission architectures.

Configurability

Xilinx FPGAs provided design flexibility that allowed the scientists to connect all the data paths early in the design – and figure out what to do with the data later. The physicists were able to implement their original algorithms and begin the experiments – and if necessary, they will be able to reconfigure and upgrade the FPGAs during the course of Run II.

Tool Set

The high-level software tools available allowed the DZero team to go from knowing nothing about programming FPGAs to building some of the most sophisticated DSP devices in the world. Their learning curve included understanding a new computer language and mastering all the new tools that go with it. In less than a year, the physicists and engineers were able to program in VHDL, adopt the tools, and use them at the level of very experienced digital designers.

“While our people are very talented, I think the fact that we were able to learn and master the new language, tools, and art of high-level digital design so quickly also speaks very well about the ease of use of Xilinx development tools,” said Levan Babukhadia, who led the team in developing the VHDL firmware. The DZero team used various releases of Xilinx ISE, as well as Aldec Active-HDL™, Synopsys FPGA Express™ (Xilinx Edition), and Synplicity software.

Vendor Support

Avnet Design Services was able to provide all necessary training. Nick Hartl, an ADS Gold FAE, taught part of an intense five-day introduction to VHDL and Active HDL. He also arranged for two days of Aldec instruction. For many of the physicists, this was their first exposure to digital design. Additionally, Hartl worked with Fermilab on product selection, and he provided consulting and information on core integration, design optimization, and system-level architecture choices.

Conclusion

The DZero team has built an ultra high-bandwidth real-time supercomputer out of off-the-shelf Xilinx components to search for the Higgs boson. As powerful as this system is, it still is not able to monitor every collision and record every event. Within two years, the Fermilab Tevatron is going to ramp up its luminosity to a higher level. The ramp up will require refinements to the track finding and other algorithms – and much more powerful FPGAs.

Certain parts of the algorithms are easiest to implement in software, yet the team cannot afford to give up the raw power of parallel processing in FPGA hardware. The natural next step appears to be to marry the software and hardware by migrating to Xilinx Platform FPGAs, such as the Xilinx Virtex-II Pro™ series. Virtex-II Pro Platform FPGAs offer as many as four embedded IBM PowerPC™ 405 cores – and as many as 10 million system gates.

Meanwhile, the Fermilab scientists will continue to refine and improve their equipment, looking not only for the Higgs boson, but also searching for supersymmetry, extra dimensions, and other new phenomena. Babukhadia concluded, “We are on the way to exciting physics, with the first results coming soon, and exciting years ahead!” ❧

Glossary

- **Energy:** Since energies in the world of elementary particles are so tiny compared to our everyday, macroscopic experience, they are typically given in units of electron volts (eV). One eV is the amount of energy one electron would acquire having passed through a +1 Volt potential difference. Or perhaps in more familiar energy units of food ratings, it is equal to about 3.8×10^{-23} (food) calories. Because particle accelerators collide beams of particles of very high energies, these energies are usually given in billions of electron volts, or GeV.
 - MeV - million electron volts (mega electron volts)
 - GeV - billion electron volts (giga electron volts)
 - TeV - trillion electron volts (tera electron volts)
- **Mass:** Owing to Einstein’s celebrated relation $E=mc^2$, describing the equivalence of mass and energy, mass of fundamental particles is typically given in units of energy. A convenient unit turns out to be GeV/c^2 , or billions of electron volts divided by the speed of light (in a vacuum) squared. For example, in these units the proton mass is approximately $1 \text{ GeV}/c^2$ or, equivalently, about 1.78×10^{-27} kilograms. With the speed of light further set to unity, mass is often given simply in units of GeV.
- **Luminosity:** This is the “brightness” of the particle beam. Measured in particles per square centimeter per second, luminosity determines how many collisions can occur. The higher the luminosity, the higher the collision rate.

The ultimate goal of the DZero team might best be described with another quote from Stephen Hawking: *Ever since the dawn of civilization, people have not been content to see events as unconnected and inexplicable. They have craved an understanding of the underlying order in the world. Today we still yearn to know why we are here and where we came from. Humanity’s deepest desire for knowledge is justification enough for our continuing quest. And our goal is nothing less than a complete description of the universe we live in.*