Chapter 2

Solid State Tracking Detectors

Maurice Garcia-Sciveres

Lawrence Berkeley National Lab
MS 50B-5239, 1 Cyclotron Rd., Berkeley, CA 94720, (510) 486-7354
mgs@lbl.gov
http://physics.lbl.gov/garcia-sciveres

2.1 What Is a Tracking Detector?

In a high energy particle experiment, a beam of particles (electrons, protons, muons, etc.) is focused onto a target and the scattered particles are measured using a detector. The properties of the target can be reconstructed from measurements of the scattered particles. It can be insightful to compare the problem to conventional microscopy where a small object is imaged using scattered light rays. There are important differences from optical microscopy:

1. The target being studied can be the quantum vacuum rather than a solid object, in which case the incident beam of particles is focused onto another incident beam of particles: this is the particle collider configuration.
2. The scattered particles need not be the same as the incident beam particles. Larger numbers and types of particles typically emerge. Energy is conserved, but if the incident beam energy is higher than the rest mass of some particles, those particles may be produced in the collision.
3. There is no lens that can focus scattered particles as done with light. The characteristic ray tracing of optics reconstructs straight line rays from two points: the lens (common to all rays) and the position on the image sensor. The analogous ray tracing for particles is done by measuring two

 https://doi.org/10.1142/9789819801107_0002

or more points along the trajectory of each particle and inferring the trajectory from those points. This is called track reconstruction or tracking.

When considering a single particle emerging from a target and traveling in vacuum, measuring two points would suffice to infer the straight line trajectory. However, there are hundreds or thousands of particles emerging at the same time, which introduces ambiguities when connecting the two measured points. In addition, the particles have electric charge, which means they will not travel in a straight line; rather they will follow a helix due to the applied magnetic field. Finally, the particles do not travel in a vacuum, as both accelerator and detector contain significant material. To address these conditions, significantly more than two points per trajectory must be measured. The detector that makes such measurements is call a tracker.

Practical trackers measure points along the trajectories of electrically charged particles only. The interaction cross-section of charged particles with matter peaks at low values of energy loss [1] so that an energetic charged particle traversing material will gradually lose energy along its path. Thus, small amounts of material can be used to extract enough energy for a detection signal from a traversing charged particle with high efficiency, yet without stopping or significantly changing the trajectory of the particle. Neutral particles (gamma rays, neutrons, and neutral kaons), on the other hand, lose energy in larger, discrete interactions, either being completely stopped or deflected, making it unfeasible to sample their trajectories without disturbing them.

2.2 Why Solid State?

Gas, liquid, and solid materials have all been used to build trackers. Earlier trackers all used internal gain, meaning that the energy lost by a traversing particle is amplified though some physical process within the detection medium before being measured by electronic or optical means. Bubble chambers used a superheated liquid as the energy loss medium, in which traversing particles cause nucleation, leading to local boiling along the trajectory, leaving a trail of bubbles that could be photographed. This is a relatively slow process and not compatible with continuous operation. In gaseous detectors, such as drift chambers, a high electric field near a sense electrode leads to avalanche breakdown in the presence of ionization from a traversing particle. This process is much faster than bubble formation but still too slow (the recovery in particular) for contemporary colliders [2]. Furthermore, the position resolution of gaseous trackers is limited by

ionization statistics. The density of a gas is low enough that the distance between ionizing interactions of a passing charged particle[1] is Poisson distributed with mean of order 100 μm, and this limits how well the track position can be known from the ionization [3].

Solid state trackers, which overwhelmingly use silicon as the detection medium, are both faster and more precise than gaseous trackers. The high density of solids results in sufficient ionization that no intrinsic gain is required in order to measure the signal from a traversing particle. This is also thanks to low noise electronics enabled by integrated circuits and high density interconnection methods (wire bonding and bump bonding). The capacitance of interconnects enters into the electronic noise as will be seen later. Detection without intrinsic gain is faster because one does not need to wait for the intrinsic gain process and its recovery to take place, and this enables the very high rate capability of silicon trackers. Ionization statistics no longer limits resolution (for the time being), as the mean free path between energy deposits of minimum ionizing particles in silicon is of order 1 μm.

While resulting in high ionization which is good for fast, precise detection, the high density of solids is also a liability because it leads to scattering of the particles whose trajectory is being measured. Fundamentally, a tracker must extract energy from traversing particles in order to measure them while at the same time not extracting energy in order to leave their trajectory undisturbed. This is clearly an optimization problem for the thickness and separation of the solid state sensors in a tracker. Early silicon trackers all used 300 μm thick silicon sensors because this happens to be the standard thickness for commercial silicon wafers and results in a signal that can be readily measured, but contemporary trackers use ever thinner sensors, as other factors enter into the optimization, in particular radiation damage to the sensor.

2.3 Strips, Pixels, and Monolithic Pixels

A tracker design involves the optimization of many parameters and tracker designs have evolved in response to changing parameters. Some parameters are driven by the experiment and science requirements, while others are due to available technology and practical matters, such as cost. In broad strokes, the historical trends go from small to large systems (both size

[1] We are mainly concerned with relativistic particles in the so-called minimum ionizing regime [1].

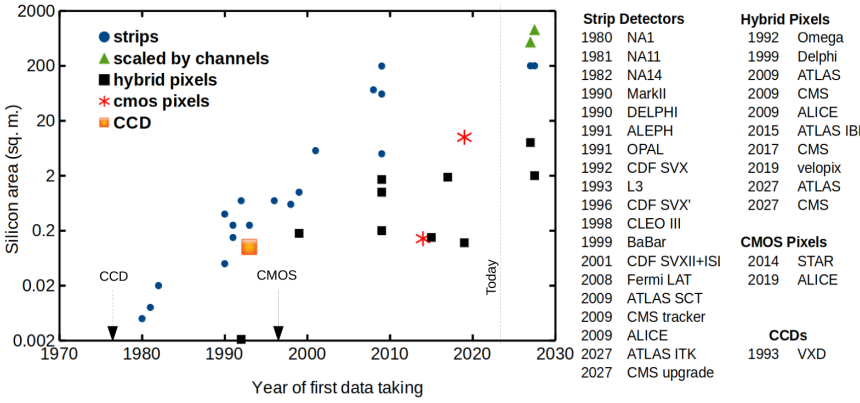

Figure 2.1. Evolution of silicon tracking detectors for particle and nuclear physics, separated into the three main types. All of them are on ground-based experiments except for Fermi-LAT, which is a space mission. The green triangles for strip detectors are duplicate points scaled by channel count rather than area because area is constrained by the existing solenoid magnets the detectors must fit into. The solid square on the horizontal axis (Omega) is just $2\,\mathrm{cm}^2$ (so off scale). The dashed arrow indicates when the CCDs first appeared in consumer digital cameras, while the solid arrow marks when the monolithic CMOS sensors first appeared in webcams. The vertical dotted line shows the time of this writing.

and channel count), from slow to fast, and from less to more radiation tolerant. Figure 2.1 shows the historical area trends for three main types of silicon trackers: strips, hybrid pixels, and monolithic. In all trackers, one can identify a basic unit called a *module* which is replicated many times. A module is like a tile in a tiled floor or wall. A tracker consists of surfaces (which are often cylindrical) tiled with modules. The main difference between the three types of trackers is what the modules are made of.

Silicon Strips: The modules consist of a silicon sensor and readout electronics (Fig. 2.2(a)). The sensor is an array of parallel line channels or *strips* connected to readout at one end. This is the oldest type of tracker and the simplest to build using more or less generic components. The first such detectors were built with discrete electronics fed by one wire per strip, but area scaling and collider configurations were enabled by integrated circuit electronics and automated wire bonding technology. The first custom integrated circuit for particle physics, the Microplex chip, was developed in 1983 precisely to be able to scale the use of silicon strip sensors to larger areas and collider detector geometry. It is worth noting that the first silicon strip sensors were commercially sold in 1969, so the sensor technology was

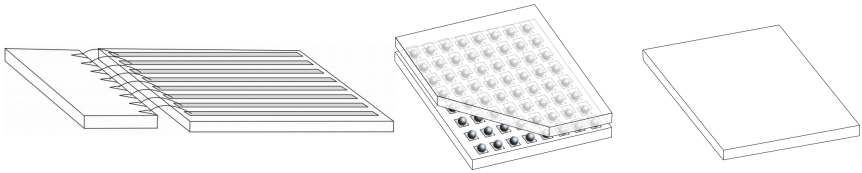

Figure 2.2. Schematic depiction of module types. (a) Strip module with sensor on right and sensor strips connected by wire bonds to the readout electronics on the left. (b) Hybrid pixel module with sensor and readout integrated circuit having matching geometry and connected to each other face-to-face by metal bump bonds. A corner has been cut away to help visualize the assembly. (c) Monolithic module.

mature; it was the development of integrated circuits that made them ideal for particle physics applications.

Hybrid Pixels: The modules consist of a silicon sensor and readout electronics (Fig. 2.2(b)), but the sensor is a matrix of pixels connected to readout electronics via bump bonds (one bump per pixel). Since the sensor and readout integrated circuit must have matched geometries, and because fine pitch bump bonding requires specialized equipment, hybrid pixels are complex to build and require custom designed components. However, a major advantage is that the sensor and readout integrated circuit can be separately optimized using very different fabrication processes. Mature strip sensor technology could be used to produce science grade pixel sensors almost "out of the box." High efficiency at the same time as high readout rate and radiation tolerance could be achieved thanks to the separate optimization.

Monolithic Active Pixel Sensors (MAPSs): A MAPS module can consist of just one MAPS sensor and some interconnects. MAPS are the closest relative of consumer image sensors ubiquitous in every mobile device. While MAPSs are simple to assemble, they are complex to produce, requiring a customized integrated circuit fabrication process. Since integrated circuits are mass produced in large foundries, particle physics-specific customization for producing small volumes is challenging to negotiate. MAPS sensors provide the lowest mass option for particle tracking but have rate and radiation tolerance limitations relative to hybrid pixels. Thus, the highest rate and radiation experiments have not managed to make use of them. Such monolithic technology was investigated very early on [4] (before CMOS imagers replaced CCD imagers in consumer electronics) but was initially disfavored relative to hybrid pixels, which had a mature sensor technology base.

2.4 How Does the Detector Work?

2.4.1 *Hybrid sensor*

A silicon sensor channel (strip or pixel) can be very simply modeled as a parallel plate capacitor in vacuum, in which electric charges magically appear along the trajectory of a particle crossing the gap. A resistor is in parallel to the capacitor and a DC voltage called bias of order 100 V is required across this RC for operation. For silicon sensors, the amount of charge from a passing particle is on average 80 electrons (and 80 positive charges or *holes*) per micron of path length, which for a 300 μm thick sensor traversed perpendicularly is 4 fC. This is the signal. To read out this signal, the charge must be extracted from the sensor channel and turned into a voltage. The capacitance of a silicon strip (pixel) channel is of order 1 pF/cm (100 fF) and the resistor is of order 100 GΩ. So the voltage change on the channel from the added charge is of order 1 mV (40 mV) for a 4 cm strip (pixel). One cannot easily read this voltage that sits on top of the 100 V bias. Instead, the charge is extracted by a readout circuit, which generates a voltage from it.

The sensor is not in reality parallel plates in vacuum but a solid silicon crystal. High resistivity silicon is used, which means it has a low level of doping and hence a low charge carrier concentration. Highly doped implants are used to define the pixels or strips. These implants form PN diode junctions with the high resistivity bulk. For an introduction to and reference on silicon devices, see [5]. Applying a reverse bias voltage to the diodes will grow the depletion region until the entire bulk is depleted of free carriers. The depleted sensor bulk is the vacuum of the parallel plate capacitor model in which charge can magically appear. A traversing high energy particle will lose energy in the silicon crystal and this energy will promote charge carriers from the valence band to the depleted conduction band. The average energy per electron–hole pair promoted to the conduction band in silicon is 3.7 eV. Thus, the 80 e-h pairs per micron correspond to an energy loss of 300 eV per micron by the passing charge particle. This is of course the energy loss by a minimum ionizing charged particle in silicon [1]. As a fun digression, consider what if one really had vacuum instead of silicon. It is still possible to turn energy into charge carriers, but one needs twice the electron mass per electron positron pair created, which would be 80 GeV per micron instead of 300 eV. So other than the 8 orders of magnitude energy difference and a

solid lattice to conserve momentum for the lost energy, a depleted sensor is a good analogy for Dirac's negative energy electron sea filling all of space [6].

A depleted semiconductor is almost like an insulator (or vacuum) but not quite. Imperfections in the crystal have different energy levels than pure silicon and can result in carriers within the forbidden region separating the valence band from the conduction band. Combined with carrier diffusion, this leads to promotion of carriers into the conduction band manifesting as a steady "leakage current." The parallel resistor in the simple model was there to simulate this current. Leakage current, also known as dark current in photodiodes, is a major topic in semiconductor detectors. The first thing to note is that it looks exactly like signal: charge carriers appearing in the depleted conduction band. Consider the 100 GΩ resistor under 100 V vias of the simple model. This results in 1 nA of leakage current, or 1000 fC/s, to be compared with the 4 fC signal. This may sound hopeless, but it just means that the readout electronics need to be fast because the signal appears all at once as an AC pulse, while the leakage current is DC. If the electronics can integrate the signal in 1 μs, for example, then the above leakage contribution becomes negligible. However, the story does not end there because 1 nA is the leakage current of a very high quality, brand new sensor channel, but it can increase by orders of magnitude with radiation damage. A common technique in strip detectors is to capacitively couple each strip to the readout (called AC coupling), which preserves the signal while stopping DC current, but this still leaves the fluctuations in the leakage current to contend with because leakage current is a stochastic process. Furthermore, AC coupling adds processing cost, limits the bias voltage that can be applied, and is not practical for pixels. Two main technologies are used to address leakage current for detectors that must withstand high radiation (which is most of them): (1) sensor material engineering and (2) cooling. An international collaboration called RD50 has been perfecting silicon material for tracker applications for decades [7]. Cooling reduces leakage current because, since leakage current depends on diffusion of intrinsic carriers, it scales like the intrinsic carrier concentration, which is proportional to $T^{1.5}e^{-E_g/2kT}$, where T is absolute temperature, E_g is the bandgap (1.12 eV in silicon), and k is Boltzmann's constant. The use of advanced materials and methods in the low mass mechanical design (see Section 2.8) enables cold operation of silicon trackers, ranging from -20°C to $+10$°C sensor temperature.

2.4.2 *Integrated circuit and readout*

The readout integrated circuit (either hybrid or monolithic) is primarily responsible for extracting the signal from the sensor and turning it into communicable information that can be relayed to a data acquisition system, which can be a significant distance away, of order 100 m. The distinct feature of hybrid pixel and strip detectors is that each channel (each individual pixel or strip) is an independent unit with its own electronics, and all channels run in parallel. This is in sharp contrast to consumer image sensors, either CMOS or their CCD ancestors, where all or many pixels on one device are staged onto a common readout channel. The all-in-parallel operation permits trackers to reach the high frame rates required by particle physics experiments, for example, 40 MHz at the LHC. A 40 MHz frame rate is unheard of (and unnecessary) for commercial electronics. If one takes as a figure of merit detector area times frame rate, to cover $1\,\mathrm{m}^2$ of detector at 40 Mfps would require 20 billion mobile phone image sensors since each sensor is about $0.5\,\mathrm{cm}^2$ and capable of 40 fps. Instead, this is done with 2,500 custom integrated circuits in the case of pixels and 10 times fewer in the case of strips. The problem is data volume, as should be clear from the mobile phone comparison. (If everyone on the planet had 10 mobile phones and tried to upload video to the cloud simultaneously for a few years non-stop, this would be a problem.) The central question for tracker readout is: what is the information that must be extracted and how can that be done?

The information content (or information entropy) in a tracker can be calculated [8, 9]. The information entropy, which can be expressed in bits, depends linearly on the number of particles traversing the tracker, which should not be too surprising. But it depends very weakly (logarithmically) on the number channels, and it depends more strongly on the noise performance and timing resolution of the electronics. For example, for High Luminosity LHC pixel detectors (rightmost solid squares in Fig. 2.1), the information content is about 25 bits per particle in each detector layer crossed. This is a measure of the information available. Without lower noise or faster electronics, there is no more useful information to be had. But this does not mean that one will transmit 25 data bits per particle from each detector module to the data acquisition system. This is an information entropy bound, which means that unless information is discarded, it is impossible to move the information off-detector with fewer than 25 bits per particle. But it is actually very challenging to achieve this entropy bound. The High Luminosity LHC pixel detectors will get to within 25%

or so, which is better than ever done before, and with very good reason: the High Luminosity LHC particle rate is extremely large, and moving information out of the pixel detectors (namely, readout cables) is what limits their tracking performance, due to the mass of the cables causing multiple scattering. This illustrates the complexity of tracker design, where apparently disconnected parameters depend on each other, such as the noise target for electronics depending on the amount of data to be moved off detector in two separate ways: first, because it would be useless to generate more information (lower noise = more information) than can be transmitted, and second, because higher precision (lower noise = higher precision) than the multiple scattering smearing due to data cables will not improve tracking performance (Section 2.6).

Viewed as images, the frames from a tracker module will be dark fields with a few bright dots. Transmitting such frames as images would use many bits because of all those dark channels, which are of no interest for tracking yet are preserved in an image. Instead, dark channels are suppressed by applying a threshold (called zero-suppression), so only the bright dots must be read out. The threshold is applied as early as possible, before any digitization. In fact, this threshold discrimination is itself used as the digitization technique. Some detectors simply record which channels are above threshold (also called hits) with no further information, while others store a low precision amplitude value for every hit by counting "Time over Threshold" (ToT).

The readout process is depicted conceptually in Fig. 2.3. Even if the entropy bound were achieved, transmitting all information for every bunch crossing usually results in too much data to handle off-detector. The High Luminosity LHC ATLAS and CMS experiments will see about 10,000 charged particles per beam bunch crossing. If each particle crosses 10 detector layers, that would be 2.5 Mbits per bunch crossing, which leads to 100 Tb/s. Therefore, the trackers need a way to decide which bunch crossings to send off-detector rather than sending all of them. This selection is shown as a trigger signal in Fig. 2.3. Since deciding to keep or discard a bunch crossing takes some time, this means the readout integrated circuit must be a large digital memory to store *all* the data until a trigger arrives. For this reason, tracker integrated circuits are produced in the technology nodes with the smallest transistors that detector builders have access to. Smaller transistors means more memory per unit area to be able to store more hits per unit area. Therefore, Moore's law has been an important enabler of advances in semiconductor trackers.

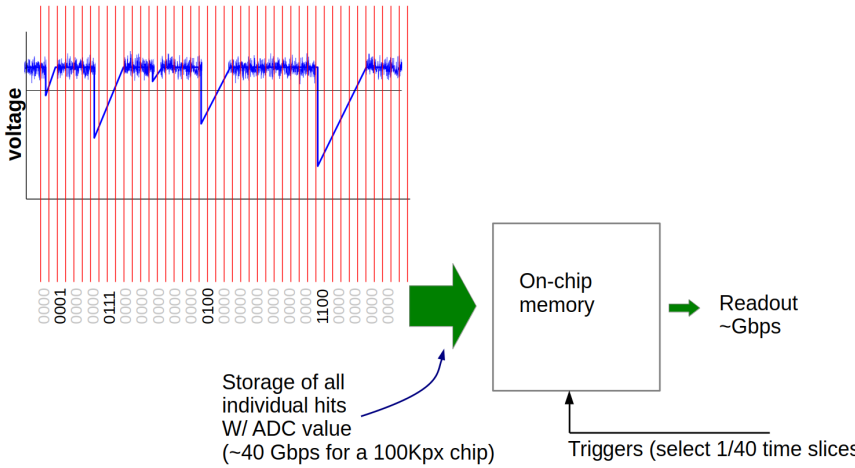

Figure 2.3. Representation of readout integrated circuit data flow with High Luminosity LHC pixel detector values. The waveform shows the analog signal for a single channel, with a noisy baseline and negative-going signal pulses. The horizontal line represents the applied threshold. Only pulses exceeding threshold are further processed, digitized, and stored in on-chip memory. The vertical lines mark the 40 MHz beam bunch crossings, with only one possible digitization per bunch crossing. The trigger signal selects only specific bunch crossings for off-detector readout.

Trigger signals are often supplied from outside the tracker, but a tracker can also be self-triggered, locally deciding which hits to transmit off-detector and which to discard. Hits produced by charged tracks will be correlated from one layer to the next, and the spatial correlation will depend on the track origin and momentum (if there is a magnetic field). Measuring the charge deposition profile in a sensor of finite thickness is sensitive to the incident particle direction but with precision limited by the sensor thickness d, as well as the position precision of the sensor Δx. The angular measurement error will be $\Delta x/d$. Interconnecting two layers some distance apart can greatly increase precision by increasing d far beyond sensor thickness. The CMS experiment is implementing such a direction-sensitive detector with electrically interconnected pairs of sensing layers consisting of one strip and one pixel, which measure correlated hit pairs called stubs. These are 3-D vectors rather than 3-D points in space. The stubs are filtered on-chip, and only those compatible with particle tracks of interest are read out, and combined off-detector to serve as a trigger signal for other subsystems.

Earlier we said the data rate was too high to read out everything without any trigger. However, in the case of the LHCb experiment, the number of particles per beam collision is an order of magnitude less, and additionally the fixed-target detector geometry permits routing cables outside the volume where the particle propagates. Thanks to these features, the LHCb experiment has implemented full triggerless readout of their detectors. This actually simplifies the architecture of the readout chips, as a storage of data while waiting for a trigger is no longer needed. This simplification is offset by the need to move data faster both on-chip and to the DAQ, requiring higher bandwidth circuits. The data must still be filtered off-detector to select events of interest. More sophisticated filtering is possible, as correlations between all detector elements can be exploited but only if enough computing resources can be deployed, which can be a significant cost. Even where triggerless readout is possible, detector design must balance the potential physics advantage from having all data off-detector against the lower cost and processing complexity of filtering data at the source, before it is read out.

Analog Front End and Signal Discrimination: The analog Front End (FE) is the readout circuit element responsible for extracting the signal from the sensor and converting it to a substantial voltage for further processing. It must move an electric charge Q from a significant sensor capacitance C_d, where it only equates to a small voltage $V = Q/C_d$, onto a smaller capacitor C_f where it will result in a larger voltage $V = Q/C_f$. The design and function of integrated circuit amplifiers for silicon detectors are extensive topics and only a superficial overview is given here in order to motivate the different types of devices used; for a more in-depth treatment, see, for example [10]. A plumbing analogy for the FE starts a small amount of water at the bottom of a bathtub, which needs to be moved to a glass in order to measure how much water there is. For the case of the bathtub, this can be accomplished with a pump, with the time taken and power used to transfer the water to the glass depending on the pump characteristics. Furthermore, for unattended operation, one needs a mechanism to empty the glass periodically or continually (through a small hole at the bottom). With a small hole, the water level in the glass will rise as the pump empties the tub and then fall as the glass drains, resulting in a water level vs. time pulse shape depending on how fast the pump pumps and the size of the hole. Silicon detector FEs are most commonly of this kind, where the electrical equivalent of the pump is a high open-loop gain preamplifier, the

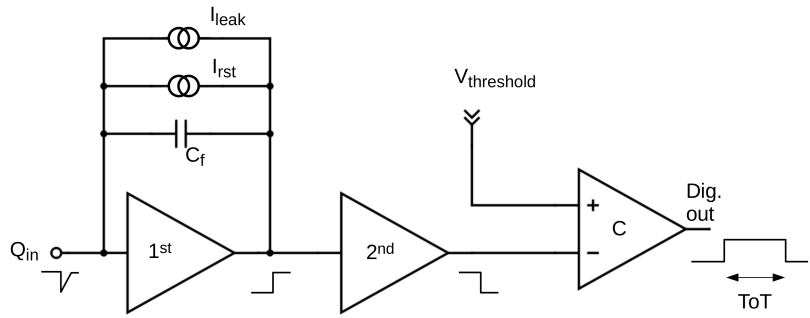

Figure 2.4. Schematic diagram of a typical analog front end (FE). Signal polarities and Time over Threshold (ToT) are indicated along the bottom. The 1st stage or preamp shows a feedback capacitor C_f, a reset current $I_{\rm rst}$, and a leakage compensating current $I_{\rm leak}$. The comparator stage, C, shows a threshold voltage adjustment $V_{\rm threshold}$.

glass is a capacitor placed in the preamp feedback loop (C_f), and the hole is a continuous reset of the capacitor ($I_{\rm rst}$) provided by a resistor, a current source, or a more complex circuit to control linearity and baseline. These elements can be seen schematically on the left of Fig. 2.4.

Figure 2.4 also shows an additional feedback element, $I_{\rm leak}$, a 2nd stage preamp, and voltage comparator (C) with an applied threshold. The 2nd stage is conceptually not a separate functional element but typically required to provide additional gain and/or drive the comparator input. The $I_{\rm leak}$ feedback element is functionally important if the sensor is DC connected to the FE input, such that sensor leakage current will flow through the FE. Sensors can also be capacitively coupled to the FE, in which case the no leakage current circuitry is needed, but this adds cost and complexity to the sensors and is typically only possible for strip detectors.

The comparator carries out the pulse height discrimination, distinguishing pulses above a user-chosen threshold, which ideally are only produced by particles hitting the sensor, from electronic noise (or thermal noise) and from leakage current fluctuations (called shot noise). Optimization of the threshold to be more than 99% efficient for real particle hits while having low firing rate in the absence of real hits is a main challenge of analog design.

Thermal noise can be thought of as fluctuations of voltage in an electronic system. At no point in an electronic system is the voltage a perfectly constant value with time, but it wiggles with some RMS noise value. Therefore, if one samples a voltage (or compares two voltages as

is done in a comparator), the result depends on exactly at what point in time the sampling happens. If a circuit is very slow (low bandwidth), the wiggles are slow and the result does not change much for small changes in sampling time, but if the bandwidth is high, the change is larger. So a slow amplifier will have less noise, but it can't be made too slow because it has to distinguish bunch crossings. A very important point is that thermal noise is a voltage noise, whereas the signal in particle detectors is a charge. Noise in silicon detectors is typically expressed in units of Equivalent Noise Charge (ENC), also known as input-referred noise. We want to know the value of signal-like charge fluctuation that results in a voltage fluctuation equal to the electronic voltage noise. Crucially, this depends on the sensor capacitance for no more complicated reason than the $Q = CV$ relationship. Thus, thermal noise (or more correctly thermal signal to noise, S/N) scales with sensor capacitance! Another important point is that the ENC depends on the amplifier signal gain and charge collection efficiency. Recall that in a continuous reset FE the glass being filled by the pump has a hole it, so the slower the pump fills it, the lower the maximum level the water will rise to. For a resistive reset, the discharge current is V_{out}/R, while the charging current from the amplifier operation is $(Q_0/\tau)e^{-(t/\tau)}$. The current starts at some value Q_0/τ and asymptotes to zero when the entire signal charge Q_0 has been extracted from the sensor. The integral of this extracted current is Q_0 and the time constant, τ, is a characteristic of the amplifier (how fast the pump pumps in the plumbing analogy). The output voltage is $V_{\text{out}} = QC_f$, where Q is instantaneous charge on the feedback capacitor C_f (analogous to the volume of water in glass next to the bathtub):

$$\frac{dQ}{dt} = \frac{Q_0}{\tau}e^{-(t/\tau)} - \frac{Q}{RC_f} \tag{2.1}$$

This has a characteristic pulse shape with a rise time given by τ and a fall time given by RC_f. The peak amplitude Q_{peak} occurs when $dQ/dt = 0$, and it can be readily appreciated from Eq. (2.1) that the faster the fall time (the smaller R), the lower Q_{peak} for a given Q_0. All this is to say that the ENC is not only the voltage noise times the detector capacitance, but it is further increased by a factor Q_0/Q_{peak}.

The shot noise arises because charge is quantized. Integrating leakage current for a period of time equates to counting a number of charges and like any counting is subject to a Poisson fluctuation of $\sqrt{N}$ for large N. The larger the N (or the integrated leakage current Q_{leak}), the smaller

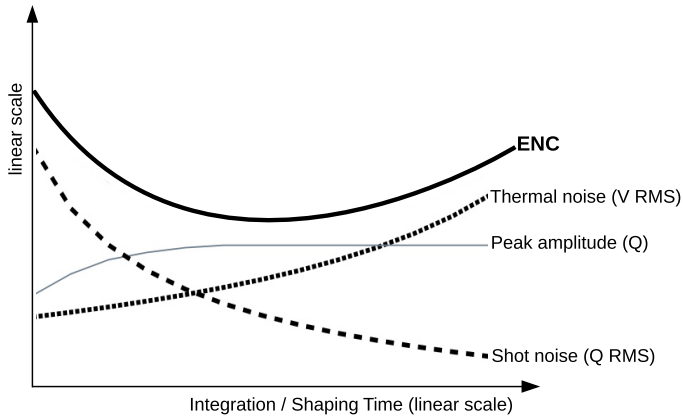

Figure 2.5. Depiction of contributions to Equivalent Noise Charge (ENC) from thermal (voltage) noise and leakage current shot noise. The ENC is the quadrature sum of the thermal noise converted to equivalent charge, including peak amplitude (gain) correction as explained in the text, and the shot noise, which already has units of charge. The peak amplitude decrease on the left of the figure signals that the shaping time is becoming comparable to the amplifier rise time (the hole in the glass is becoming too large in the plumbing analog of the text).

the relative error $\sqrt{N}/N$. Thus, shot noise decreases for longer integration time (larger value of R). Note that it makes no difference if the sensor is AC or DC coupled for shot noise, since the fluctuations are an AC signal. However, since shot noise is already a charge fluctuation, it does not depend on detector capacitance. Unlike shot noise, thermal noise increases with increasing R because the bigger time window allows more voltage wiggles to be sampled. The amplifier shaping (RC_f in this simple case) can thus be optimized to minimize ENC (maximize S/N), as shown in Fig. 2.5.

There are many important conclusions that can be drawn from the simple noise analysis presented. Strip detectors have significant sensor capacitance and leakage current, simply because the area of each strip is relatively large. They need more sophisticated amplifiers with high gain due to the large sensor capacitance and cannot operate at very high bandwidth because of leakage current. Hybrid pixel detectors, on the other hand, have exactly the same signal size as strip detectors but lower capacitance and leakage currents; they are just like extremely small strips. So they can have a simpler, lower gain FE and can operate at higher bandwidths. MAPSs have even smaller capacitance and therefore can use an even simpler FE or no FE

at all: the sensor capacitance is analogous to the size of the bathtub in the plumbing example, so when the bathtub becomes the size of a glass, one no longer needs a pump to collect the water into a measuring glass. However, MAPSs also have smaller signal size, and the development of that signal can be slower than in a standalone sensor, so they cannot necessarily exploit the low capacitance to achieve high speed; they are generally significantly slower devices than hybrid pixels. This already hints at the importance of shot noise for MAPSs, since the signal is small, the integration time is relatively long, and shot noise does not care that the detector capacitance is tiny.

For detectors with charge readout, the fact that a comparator follows the FE of silicon detectors is universally exploited to digitize the signal amplitude with the technique to Time over Threshold (ToT), as mentioned in Section 2.4.2. This kind of digitization requires minimal circuitry (just a counter) by making use of the already available comparator output and the beam crossing clock. The comparator output is a digital pulse whose duration is equal to the time that the FE output exceeded the comparator threshold. This time is measured by using that pulse to gate a counter that counts beam crossing clock cycles. Clearly, this only works if the FE output pulses are typically longer than one beam crossing, which tends to be the case because the amplifier rise time is chosen to be just fast enough to distinguish beam crossings, and the fall time must be longer to minimize ENC, as per Fig. 2.5.

We can now return to to the hit discriminating function of the comparator in more detail. The ENC must be small enough (equivalently S/N large enough) that one can choose a threshold which achieves the required 99% hit efficiency with a low rate of noise hits. This is depicted in Fig. 2.6. In the ideal case of the figure, the threshold can be chosen where there is zero response probability, resulting in zero hit rate in the absence of signal and 100% efficiency for signal. In practice, there is typically some overlap between tails of the two distributions. The distributions can't simply be added to determine the minimum probability point suitable for the threshold because they have different normalizations. The particle response function is normalized by the real hit rate impinging on the detector, while the ENC distribution is normalized by the maximum hit rate that the FE can respond to, which depends on FE speed and shaping. The ENC hit rate will rise exponentially as the threshold is decreased, up to the maximum at zero threshold. The operating threshold can be chosen to either reduce noise hit rate or increase signal efficiency. An interesting

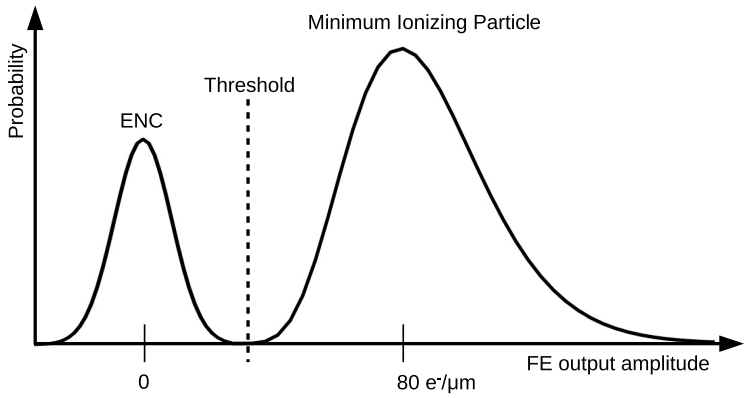

Figure 2.6. Depiction of FE response probability distribution functions and the suggested threshold setting. In the absence of signal, the FE response is given by the ENC, a Gaussian centered at zero (fluctuations can be negative or positive). When a minimum ionizing particle ionizes the detector, a random charge value obeying a Landau distribution with most probably value of $80\,e^-/\mu$m appears at the FE input, leading to a response at the output given by Landau, convoluted with the ENC (since noise is always present whether or not there is a signal). In this ideal example, there is zero response probability between the two distributions, where the threshold is shown.

observation is that for high rate experiments such as those at the High Luminosity LHC, the real particle hit rate will reach values of $3\,\mathrm{GHz/cm^2}$ and therefore very high noise hit rates (in the $\mathrm{MHz/cm^2}$ range) could in principle be tolerated with minimal impact to the data readout. However, experiments typically opt for running with much lower noise hit rates for a variety of reasons, including the ability to perform noise-free calibrations.

A final consideration about the threshold is that it must be the same (when referred to the input) for all channels in a chip and also stable with time so that the hit efficiency is uniform and, in the case of ToT, the charge measurement is uniform. Uniform response of analog circuits is not automatic within integrated circuits. Such uniformity in the static case is typically achieved with a programmable channel-by-channel threshold adjustment. But threshold fluctuations in time can be more challenging to control. For example, changes in chip power consumption can lead to threshold shifts, and processing of hits can affect power consumption. In pixel detectors, this tends to lead to a maximum occupancy beyond which operation may become unstable. For these reasons, the threshold is better thought of as a band of some thickness rather than the thin line in Fig. 2.6.

2.5 Tracker Performance

In the ideal case of a helical trajectory in a magnetic field, the momentum resolution of a tracker with N layers is given by the following Gluckstern formula [11]:

$$\frac{\sigma_{p_T}}{p_T} = \left(\frac{p_T}{0.3|z|}\frac{\sigma_{\text{point}}}{L^2 B}\sqrt{\frac{720}{N+4}}\right) \oplus \left(\frac{\sigma_{p_T}}{p_T}\right)_{\text{MS}} \tag{2.2}$$

where p_T is transverse momentum in GeV/c, L is the radial length in m, B is the magnetic field in T filling the tracker volume, z is the particle electric charge in elementary units, σ_{point} is the resolution of the detector measurements in m, and N is assumed to be large in this approximation. L has the largest effect in the equation, but note that the whole of L has to be filled with magnetic field, and the stored energy in a solenoidal field scales like B^2L^2, so making L very large is not trivial. It may therefore seem that improving point resolution may be the most economical way to enhance tracker performance. However, this has to compete against the second term in quadrature. Point resolution also has to compete against alignment precision, which is not considered in Eq. (2.2).

The second term in Eq. (2.2) is the multiple scattering (MS) contribution for a number of detector layers N. It can be written as

$$\left(\frac{\sigma_{p_T}}{p_T}\right)_{\text{MS}} = \frac{0.0136}{0.3\,\beta\,BL}\sqrt{\frac{(N-1)x/\sin\theta}{X_0}}\sqrt{C_N}, \quad [L] = \text{m},\ [B] = \text{T} \tag{2.3}$$

where L is the radial track length and $(x/\sin\theta)/X_0$ is the total material thickness traversed by a particle incident with polar angle θ with respect to the beam, in units of the radiation length. One radiation length is the thickness of material that reduces the energy of impinging relativistic electrons by $1/e$, or equivalently that pair-converts $e^{-7/9}$ (45%) of impinging energetic gamma rays. C_N is a factor depending on the number of layers: $C_N = 2.5$ for the minimum of three layers to measure a circle and approaches $C_N = 1.33$ for $N \to \infty$ (continuous scattering). Equation (2.3) shows that the low momentum performance of a tracker is limited by its mass. The radiation length of 1 mm thick silicon, copper, and carbon composite, is 1%, 7%, and 0.3%, respectively. State-of-the-art hybrid (monolithic) trackers achieve under 2% (1%) of a radiation length per layer.

2.6 How Does Tracking Work?

Track reconstruction, or tracking, is the process of converting the hit coordinates that were read out from the detector, known as space points, into momentum vectors of the charge particles that produced them. The results of Section 2.5 only apply after tracking has done a good job. Multiple operations are involved, which we group into four main categories: pattern recognition, track fitting, alignment, and simulation. One might here ask why not train an artificial neural network to perform this complex transformation from space points to momentum vectors? This is a very active development area but does not affect the division into these four categories. Machine learning is being applied within each of the four individual categories, but a solution to infer final vectors directly form hits in a single step (called end-to-end), that is, competitive in terms of precision and performance, is not considered feasible.

Pattern Recognition: Pattern recognition is the process of dividing all the space points from an event into track candidates, where each track candidate is just a group of space points. This process is very sensitive to the number and spacing of the layers in a detector and it is the greatest consumer of computing time in modern experiment track reconstruction. Space points in different layers are linked together by extrapolating "seeds," where a seed is any pair or triplet (different techniques use pairs or triplets) of space points that may or may not belong to the same track. Seeds are typically created combinatorially and can therefore reach astronomical numbers, but most of them will be spurious combinations that will be gradually rejected for having no additional points in their extrapolations. It is instructive to analyze the linking of hits of a hypothetical three-layer arrangement. In order to distinguish true tracks from random combinations, the probability of finding random combinations that align into tracks should be small. This "fake" probability is straightforward to calculate. Even when a particle track is known up to a certain layer, the extrapolation of that track to the next layer has an angular uncertainty α, with a lower bound given by multiple scattering (the multiple scattering angular error can be found in [1]). This uncertainty projects a circle of radius $x\alpha$ onto the next layer, where x is the distance to that next layer. The number of random coincidence hits within this circle is the area of the circle times the hit density in that layer, ρ. The hit density is given by the track density at that layer, so it is known regardless of detector details. Actually, to the extrapolation area, one must add a position resolution area, because the

measured hit positions have their own uncertainty, given a resolution d (for simplicity, we assume equal resolution in two dimensions, but in practice, detectors can have different resolution in each dimension (notably strip detectors have excellent resolution in one dimension only). The probability of finding a random hit in the extrapolation of a track is therefore

$$P_{12} = \pi(\alpha^2 x^2 + d^2)\rho \tag{2.4}$$

where P_{12} denotes the probability of the extrapolation from layer 1 to layer 2 linking a fake hit, not related to the track in question. In the case of a three-layer tracker with total length L, two successive extrapolations (from layer 1 to layer 2 and from layer 2 to layer 3) must each find a random hit coincidence in order to result in a fake track. The fake probability from layer 1 to layer 2 was given in Eq. (2.4). The fake probability from layer 2 to 3 (P_{23}) is also given by the same equation but replacing the distance x (which was the distance from layer 1 to 2) by $L - x$. The probability of the two successive fakes is simply the product $P_{12}P_{23}$:

$$P_{\text{fake}} = P_{12}P_{23} = \pi^2\rho_2\rho_3(\alpha_1^2 x + d_2^2)(\alpha_2^2(L-x)^2 + d_3^2) \tag{2.5}$$

where the subscripts now refer to the layers since different layers can have different masses and resolutions. Note the multiple scattering in the extrapolation source layer and the position resolution in the destination layer determine the uncertainty circles. Figure 2.7 shows the relative probability for fake tracks vs. the position of layer 2, ranging from very close to layer 1 to very close to layer 3 (the separation between layers 1 and 3 is always the constant L), where we assumed for simplicity α and d are the same for all layers. Two cases are plotted: the extrapolation uncertainty at distance $L/2$ is dominated by the hit resolution or dominated by multiple scattering (dominated means twice as large). It can be appreciated that when multiple scattering is small (which is the case for high momentum particles) there is a slight pattern recognition advantage to evenly spaced layers, whereas if multiple scattering is important (the case for low momentum particles), there is a significant advantage to placing some layers close together (often called doublets, as in the self-triggered CMS tracker discussed earlier). This makes it clear that there is no one-size-fits-all tracker geometry, but the geometry must be optimized depending on the detailed requirements.

Track Fitting: Section 2.5 assumed particle trajectories are perfectly helical, the magnetic field is perfectly uniform, multiple scattering smears

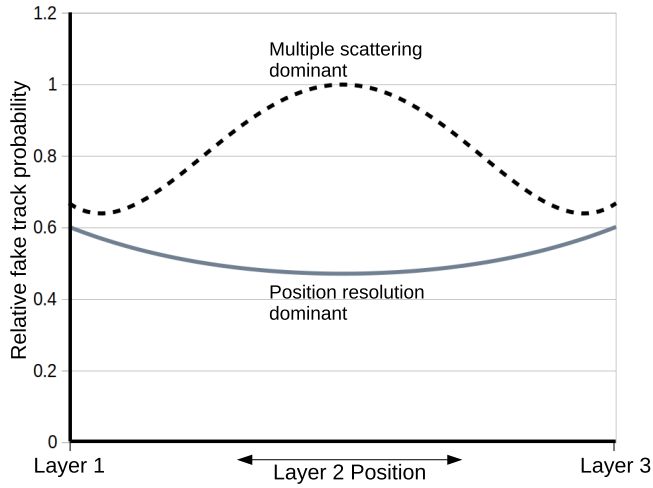

Figure 2.7. Relative fake track probability in three-layer tracker discussed in the text and given in Eq. (2.4). The positions of layers 1 and 3 are at the left and right of the plot, respectively, with constant separation L. The effect of varying the position of layer 2 between layers 1 and and 3 is plotted for the cases where the layer position resolution is the dominant uncertainty (solid) and multiple scattering is the dominant uncertainty (dashed).

all measurements equally, etc. But none of these conditions hold exactly. To obtain the best results, track fitting is not a simple least squares fit to an analytic formula. Starting from the initial momentum we wish to know, a particle's position in the first tracker layer is only smeared by the material entering the tracker (like the collider beam pipe), while at the last layer, it is smeared by a random walk through the entire tracker. The fitting algorithm that correctly unfolds these progressive errors, universally used in track reconstruction, is the Kalman filter [12]. An integral part of this method is comparing each hit position to the expectation of where it should be based on all prior hits. The track parameters are updated hit by hit. This requires a model of the detector and machinery to predict, from the track parameters, where the next hit should be. This is called track propagation. Track propagation performs the transport of track parameters and the associated covariance matrices through the detector geometry, taking into account interactions with the material (using the detector model) and the magnetic field (using a magnetic field map). Accurate simulation is an absolutely integral part of tracking.

2.7 Radiation Considerations

Sensors are fabricated from high purity single crystal silicon so that a significant depth can be depleted of charge carriers. Impurities and lattice defects degrade the sensor performance. Nuclear interactions from impinging hadronic particles can damage the crystal lattice, a mechanism called Displacement Damage (DD) or bulk damage. DD is quantified by Non-Ionizing Energy Loss (NIEL) in terms of the flux of 1 MeV neutrons that cause the same damage (abbreviated 1 MeV n eq.). DD leads to increased leakage current (by orders of magnitude), increase of reverse bias voltage needed to deplete the sensor, and reduction of free carrier lifetime (which reduces charge collection efficiency). Through fabrication techniques known as defect engineering, silicon sensors are currently being made to withstand NIEL doses of 2×10^{16} 1 MeV n eq.

Radiation damage in CMOS circuits is entirely due to charge carriers generated by ionization in the dielectric layers of the process and not displacement damage. Ionizing dose is delivered at hadron colliders by a combination of minimum ionizing particles (mainly pions) and background X-rays and termed Total Ionizing Dose (TID). The doping concentrations in CMOS transistors are so high ($10^{15}\,\mathrm{cm}^{-3}$ and higher) that compared to them, the defect density introduced by bulk radiation damage is negligible [13] (below $10^{14}\,\mathrm{cm}^{-3}$ for HL-LHC inner layers after $3000\,\mathrm{fb}^{-1}$). This means NIEL damage has no effect on CMOS electronics. However, there are many dielectric structures in a modern CMOS process and each one leads to its own radiation effect due to TID. It is not by accident that radiation tolerance requirements have kept pace with the logic density evolution in the ROIC generations. The reason is that both hit rate and radiation dose scale with particle flux. Required radiation tolerance went from 50 Mrad for the 1st generation, to 250 Mrad for the 2nd, to 1 Grad in the 3rd. 1 Grad corresponds to about 50 minimum ionizing particles crossing every Si lattice cell. Not all effects from charge generation in the dielectrics are equally important. As radiation dose increases, understanding and managing previously negligible effects become necessary.

In addition to long-term degradation due to accumulated dose, energy loss by ionizing particles leads to instantaneous soft errors called Single Event Upsets (SEUs). The most common SEU is the flipping of a stored bit in a memory. SEU also can produce voltage transients on signal or control lines that can result in accidental operations (for example, a single-level asynchronous line to reset logic or memory would be vulnerable to

SEU). Protection against SEU involves hardening of memory cells, avoiding designs with vulnerable control signals or hardening control signals where their use cannot be avoided, and circuit triplication. These techniques have been in use since 1st generation readout integrated circuits and they have not seen significant changes in the 2nd generation. However, as collider rate continues to increase and higher logic density translates into lower deposited charge needed for upset, these techniques will no longer be sufficient. An approach being introduced in 3rd generation readout integrated circuits is to design for reliable operation while a significant level of upsets is taking place. Fundamentally, this is abandoning the idea of circuit hardening as a solution to the SEU problem and instead designing all functions such that SEU is not a problem to begin with. In practice, a combination of hardening and SEU-friendly functionality will be used.

Extensive literature and experience exist on SEU of memory cells in the context of electronics used in space. This is not directly applicable to particle physics pixel detectors but is nevertheless a good starting point. SEU of a given circuit, like an SRAM cell, latch, or flip-flop, depends on the amount of energy deposited by an impinging ion, which is characterized by a Linear Energy Transfer (LET). It is important that this is meant to describe non-relativistic ions, which lose energy approximately uniformly along their path through electromagnetic interactions. The upset rate is characterized by the cross-section for causing a bit flip (SEU cross-section). Cross-section vs. LET is typically fit with a Weibull function, resulting in a threshold and saturation cross-section, as shown in Fig. 2.8. In submicron technologies, the typical LET threshold is of order $1\,\mathrm{MeV{\cdot}cm^2/mg}$ pretty much regardless of

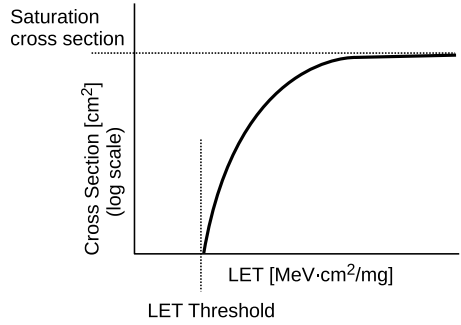

Figure 2.8. Conceptual plot of Single Event Upset (SEU) cross-section vs. Linear Energy Transfer (LET) for a typical memory cell.

memory cell type. Saturation cross-section varies with cell design but is of order 10^{-7}–$10^{-8}\,\mathrm{cm}^2$ for common SRAM, latches, and flip-flops. However, an energetic proton (or pion) has an LET of order $0.01\,\mathrm{MeV{\cdot}cm}^2/\mathrm{mg}$, which immediately signals that it cannot upset memory cells by the same energy loss mechanism as ions (it is far below the LET threshold). Upsets in this case are due to nuclear interactions. This can be seen from the fact that SEU cross-sections are about the same for energetic neutrons and protons [14]. There is thus a kinematic threshold depending on the nuclei in the material rather than an LET threshold. Typical SEU cross-sections for protons are of order $10^{-13}\,\mathrm{cm}^2$ [14]. At relatively low energies, an adequate model has the proton imparting momentum to a nucleus which then becomes a traditional high LET heavy ion. But at the GeV energies of the LHC, inelastic collisions can produce showers of high LET particles, affecting a large area of silicon. This is important for hardening techniques.

2.8 System Aspects Overview

The mechanical and electrical system designs of a tracker are what ultimately limits its practical performance, not only due to cooling but also due to material that particles must pass through and to the alignment precision achieved. Achieving low mass, or more precisely low multiple scattering, is critical for tracker and heavily influences the system design. The multiple scattering detector mass is quantified in terms of electromagnetic radiation interaction length: the mean distance over which a high-energy electron loses all but $1/e$ of its energy by bremsstrahlung. Hybrid trackers achieve masses as low as 1–2% of a radiation length per layer, while MAPS trackers can be as low as 0.3–0.5%. One percent of a radiation length corresponds to a thickness of about 1 mm of either silicon or aluminum, or 0.16 mm of copper, or 3 mm of carbon fiber composites. This should convey the difficulty of producing detectors with 1% of a radiation length per layer or less, including micron-level mechanical precision mechanical structures, power cables, cooling, and readout. Simply in terms of energetics, the most power intensive hybrid systems can dissipate over $0.5\,\mathrm{W/cm}^2$, which for a $10\,\mathrm{m}^2$ tracker means 50 kW, yet this same tracker must operate at $-10^{\circ}\mathrm{C}$ to control sensor leakage current and readout integrated circuits radiation damage. Given almost zero thermal mass, an exquisite balance must be maintained between 50 kW in through electrical wires and 50 kW out though circulating coolant. It is therefore not surprising that trackers make use of the highest performance to mass ratio materials and systems available.

All modern tracker structures are built of carbon composites. Composites allow choosing structural properties different from the intrinsic properties of individual bulk materials, such as thermal expansion coefficient or thermal conductivity. Furthermore, properties can be chosen to be different in different directions. Tracker structures have a near-zero thermal expansion coefficient so that trackers can be assembled at room temperature and maintain micron-level precision at their operating temperature 50°C colder.

Cooling, electrical power delivery, and data transmission are all system aspects that require custom solutions to meet tracker low mass requirements. Additionally, these elements must be radiation hard to different degrees, but this configuration is determined by the mass requirement. The technology of evaporative CO_2 cooling was developed for HEP trackers, as the first application that absolutely needed the extreme performance it offers. This is based on a typical two-phase refrigeration cycle, where the refrigerant is condensed and cooled outside of the active volume and evaporated though a pressure change inside the active volume. Heat is removed mainly through the latent heat of evaporation of the refrigerant. Liquid CO_2 has almost twice the density of water, which means a large mass flow can circulate in a small pipe, and more than twice the latent heat of evaporation as typical refrigerants. It has one significant challenge, however, which is that it only has a liquid phase at pressures above around 50 bar (depending on temperature). The implementation of systems that pumps cold, high-pressure liquid to the detector and extracts high pressure gas to be condensed outside, has allowed building of large low mass trackers with significant power consumption, not possible with any other cooling technology. The same cooling performance achieved with few mm diameter CO_2 pipes would require cm diameter conventional refrigerant pipes and several cm diameter pipes for simple water (single phase) cooling.

Just as important as removing the heat is supplying all that electrical power with low mass. Power supply to detector distance scales tend to be of order 50–100 m if nothing else because off-the-shelf electronics cannot be sited too close to a particle accelerator. State-of-the-art integrated circuit electronics operate at voltages below 2 V. Supplying 50 kW from 100 m away with cables that do not themselves produce more than 50 kW of heat due to their resistance (for a total wall power of 100 kW instead of 50 kW) would require 180 kg of copper per meter of cable length. Clearly, not good for making a low mass tracker. The well-known solution for delivering electrical power without massive cables, present everywhere around us in our power

grid, is to do so at high voltage. The problem for trackers then becomes one of transforming "high" to low (under 2 V) voltage without significant mass, in a strong magnetic field, and with radiation tolerance. Two technologies have been developed: radiation hard, low mass DC-DC converters that operate in a high magnetic field and serial power. The former works with 10–12 V input and 1.5–2.5 V output, multiplying the input current by the voltage ratio. DC-DC conversion is adequate for relatively low power density detectors, such as hybrid strips, as it adds an amount of mass proportional to the power density (twice as much power in the same area needs twice as many DC-DC converters and therefore twice as much mass). Serial power can be implemented with readout integrated circuits designed to operate from a constant current rather than constant voltage power supply. Such devices can be chained in series and with a special type of power regulation [15] also placed in parallel. The total voltage drop of a serial chain is $N_d V_d$, where V_d is the voltage across a single device (of order 2 V). The mass of N_d devices is the same whether they are chained in series or powered in parallel, so serial power adds no mass to the tracker, making it suitable for high power density trackers, such as hybrid pixels. But serial power does add complexity in terms of operation, communication, and sensor bias. Values of N_d as high as 14 are being used in current trackers. Both DC-DC and serial technologies achieve an efficiency of 70–80%, which means an increase of the heat to be removed by the cooling system by 25%. This is acceptable thanks to the high capacity of CO_2 cooling.

Once power delivery and cooling mass have been minimized with the above techniques, the most significant mass contributions can be due to readout cables. While commercial high speed links can move vast data volumes with compact formats, a tracker presents a different readout problem, not well matched to the commercial solutions. Aside from radiation tolerance, a tracker contains a large number of data sources spread throughout a m^3 volume or larger. This is different from the commercial problem of a concentrated data source such as a processor. A 100 Gbps point-to-point link is not very useful for reading a tracker because that volume of data is produced by order 100 different devices over a 1 m length scale. There isn't a common solution to this problem: each detector finds a different optimization, depending of the data to be moved and geometrical constraints. Common components have been developed, notably a radiation hard, low mass optical data link capable of 10 Gbps. Custom electrical cables, including flexible circuit strip lines, aluminum conductors, and twin-axial transmission lines, are commonly produced for trackers.

References

[1] R. L. Workman *et al.* Review of particle physics. *PTEP*, 2022:083C01, 2022. Section 34: Passage of particles through matter. pdg.lbl.gov.
[2] M. Aleksa *et al.* Rate effects in high-resolution drift chambers. 446(3):435–443, 2000.
[3] W. Blum and L. Rolandi. *Particle Detection with Drift Chambers*. Springer-Verlag, 1993.
[4] W. Snoeys, J. Plummer, S. Parker, and C. Kenney. PIN detector arrays and integrated readout circuitry on high resistivity float zone silicon. 41:903–912, 1994.
[5] S. M. Sze, Y. Li, and K. K. Ng. *Physics of Semiconductor Devices*. John Wiley, 1969, 1981, 2007, 2021.
[6] P. A. M. Dirac. A theory of electrons and protons. *Proc. R. Soc. Lond. A*, 126(801):360–365, 1930.
[7] RD50 Collaboration. http://rd50.web.cern.ch/.
[8] M. Garcia-Sciveres and X. Wang. Data encoding efficiency in binary strip detector readout. 9(04):P04021, 2014.
[9] M. Garcia-Sciveres and X. Wang. Data encoding efficiency in pixel detector readout with charge information. 815:18–22, 2016.
[10] H. Spieler. *Semiconductor Detector Systems*, Vol. 12. Oxford University Press, Oxford, 2005.
[11] R. L. Gluckstern. Uncertainties in track momentum and direction, due to multiple scattering and measurement errors. 24:381, 1963.
[12] R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(Ser. D):35–45, 1960.
[13] R. Radu *et al.* Investigation of point and extended defects in electron irradiated silicon–dependence on the particle energy. *Journal of Applied Physics*, 117:164503, 2015.
[14] P. Roche *et al.* A commercial 65 nm CMOS technology for space Applications: Heavy ion, proton and gamma test results and modeling. 57(4):2079–2088, 2010.
[15] M. Karagounis *et al.* An integrated shunt-LDO regulator for serial powered systems. In *Proceedings of ESSCIRC '09*, 2009.