Lawrence Berkeley National Laboratory

Recent Work

Title

A review of advances in pixel detectors for experiments with high rate and radiation.

Permalink

<https://escholarship.org/uc/item/1zn1j5x9>

Journal

Reports on progress in physics. Physical Society (Great Britain), 81(6)

ISSN

0034-4885

Authors

Garcia-Sciveres, Maurice
Wermes, Norbert

Publication Date

2018-06-01

DOI

10.1088/1361-6633/aab064

Peer reviewed

REVIEW

A review of advances in pixel detectors for experiments with high rate and radiation

To cite this article: Maurice Garcia-Sciveres and Norbert Wermes 2018 *Rep. Prog. Phys.* **81** 066101

View the [article online](#) for updates and enhancements.

Related content

- [Prototype ATLAS IBL modules using the FE-I4A front-end readout chip](#)
The ATLAS IBL collaboration
- [4D tracking with ultra-fast silicon detectors](#)
Hartmut F-W Sadrozinski, Abraham Seiden and Nicolò Cartiglia
- [The phase II ATLAS Pixel upgrade: the Inner Tracker \(ITk\)](#)
T. Flick

IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

Review

A review of advances in pixel detectors for experiments with high rate and radiation

Maurice Garcia-Sciveres¹ and Norbert Wermes² ¹ Lawrence Berkeley National Laboratory, Berkeley, United States of America² University of Bonn, Bonn, GermanyE-mail: mgs@lbl.gov and wermes@uni-bonn.de

Received 25 May 2017, revised 30 October 2017

Accepted for publication 19 February 2018

Published 29 May 2018

CrossMark

Corresponding Editor Professor Steve Ritz

Abstract

The large Hadron collider (LHC) experiments ATLAS and CMS have established hybrid pixel detectors as the instrument of choice for particle tracking and vertexing in high rate and radiation environments, as they operate close to the LHC interaction points. With the high luminosity-LHC upgrade now in sight, for which the tracking detectors will be completely replaced, new generations of pixel detectors are being devised. They have to address enormous challenges in terms of data throughput and radiation levels, ionizing and non-ionizing, that harm the sensing and readout parts of pixel detectors alike. Advances in microelectronics and microprocessing technologies now enable large scale detector designs with unprecedented performance in measurement precision (space and time), radiation hard sensors and readout chips, hybridization techniques, lightweight supports, and fully monolithic approaches to meet these challenges. This paper reviews the world-wide effort on these developments.

Keywords: pixel detectors, semiconductor detectors, tracking detectors, radiation damage, radiation hard CMOS sensors, hybrid pixels, high energy physics

(Some figures may appear in colour only in the online journal)

Contents

1. Pixel detectors at the heart of particle physics experiments: demands, challenges and concepts	2	3. Pixel sensors and hybridization.....	9
2. Space-time point resolution	3	3.1. Planar pixel sensors.....	9
2.1. Demands and current directions.....	3	3.1.1. Use of 8'' wafers and CMOS foundries.....	11
2.2. Space point and direction measurements.....	3	3.2. 3D-Si sensors	12
2.2.1. Point measurements.....	4	3.3. Bonding techniques.....	13
2.2.2. Space-point reconstruction methods.....	4	3.3.1. Solder bumps and bonding.	13
2.2.3. Limitations to the space point resolution of a pixel detector.....	5	3.3.2. In-In bonding.....	14
2.2.4. In-pixel decoding.	5	3.3.3. Adhesive bonding.	14
2.2.5. Directional information from hit clusters.....	6	3.3.4. Cu-Cu direct bonding.....	15
2.3. Time measurement.....	6	3.3.5. Oxide-oxide direct bonding.....	15
2.3.1. LGAD structures.....	7	3.3.6. Solid-liquid interdiffusion bonding (SLID).....	15
		3.4. Wafer thinning.....	16
		3.5. 3D-integration and through-silicon vias	16
		3.5.1. TSV fabrication.	16
		3.5.2. First HEP experiences with TSVs.	17

4. Readout integrated circuits, data transmission, and power distribution.....	18
4.1. CMOS ROICs	18
4.2. Increasing circuit density through 3D integration	20
4.3. Radiation tolerance of readout integrated circuits.....	20
4.4. Total dose radiation damage in CMOS transistors	21
4.4.1. Oxide charge and its effects.....	21
4.4.2. RINCE.....	22
4.4.3. RISCE.....	22
4.4.4. Dependence on voltage and temperature.....	23
4.5. Single event upsets and mitigation.....	23
4.6. Analog front end and ADC	24
4.7. Readout architecture	26
4.8. Input hit rates and output data transmission (electrical).....	28
4.9. Optical and wireless transmission.....	30
4.10. Power distribution	30
4.11. Future development.....	31
5. Advanced materials: new possibilities for supports and cooling.....	32
5.1. Cooling.....	32
5.2. Supports	33
6. CMOS active sensors: towards monolithic pixels.....	34
6.1. From MAPS to DMAPS	34
6.1.1. Depleted (D)MAPS.....	35
6.1.2. DMAPS capacitances.....	36
6.1.3. Fill-factor choices.....	36
6.1.4. Substrate resistivity.....	36
6.2. Designs and technology variants.....	36
6.2.1. Hybrid active CMOS pixels.....	37
6.2.2. DMAPS in HV/HR technologies with large fill-factor.....	37
6.2.3. CMOS on SOI.....	38
6.2.4. DMAPS with small fill-factor.....	38
6.3. Outlook on DMAPS pixels	38
7. Summary and conclusions	39
Acknowledgments.....	39
References.....	39

1. Pixel detectors at the heart of particle physics experiments: demands, challenges and concepts

This article reviews recent advances and future directions in pixel detectors to measure high energy charged particle trajectories in high rate and radiation experiments. A general treatment of pixel detectors can be found in [1]. This review focuses on the on-detector elements, and does not include data acquisition or downstream data processing such as track reconstruction. As there are recent reviews on radiation damage of silicon sensors [2] and on mechanics and cooling of such detectors [3] these topics receive an abbreviated treatment here. The current state of the art is represented by the detectors in operation or under construction at the large Hadron collider at CERN, in the ATLAS, CMS, LHC-b, and NA62 experiments.

The ATLAS experiment installed a 3-layer hybrid pixel detector in 2007 [4] and an additional layer at lower radius, inserted within the original detector envelope and therefore called insertable B-layer (IBL), in 2014 [5]. The IBL pioneered the use of 3D silicon sensors (in a limited acceptance range) and introduced a new readout integrated circuit (ROIC or readout chip) called FE-I4 [6] with several of the features needed for future high rate detectors.

The CMS experiment installed a 3-layer hybrid pixel detector in 2008 [7], and has replaced it with a new, 4-layer detector in 2017 [8]. The new detector has lower inner radius, significantly lower mass which improves precision, and higher data rate capability needed to cope with increased accelerator luminosity. The upgrade baseline ROIC was redesigned for the outer 3 layers, replacing analog signal readout with on-chip ADCs and digital readout at higher rate [9]. A different ROIC has been designed specifically for the inner layer to handle higher rate and radiation, while keeping to the original footprint. To reduce the mass of the services, radiation hard voltage regulators [10] have been introduced just outside the acceptance, but still inside the CMS inner barrel volume.

The LHC-b experiment is replacing their strip vertex detector with a 26 plane pixel detector (fixed target geometry) [11]. The detector will use triggerless readout at the full 40 MHz collision rate of the LHC. The high data volume will be handled by up to four 5 Gbps serial outputs per ROIC [12], transmitted over high bandwidth copper cables outside the physics acceptance.

The NA62 experiment has installed a fixed target geometry hybrid pixel detector called Gigatracker to measure timing with high resolution (200 ps) as well as position in each pixel [13]. This is a small area detector, but is pioneering the use of per-pixel timing in particle tracking.

The challenges that drive ongoing development can be broadly categorized into scale, intensity, and performance. Particle physics programs demand pixel detectors with larger outer radius to cope with higher rates, and greater length along the beam direction to increase acceptance. Significant development is therefore focused on how to produce pixel detectors with less effort and lower cost. The two main directions are lower cost production of hybrid pixels and diode sensors (section 3), and development of monolithic technology in CMOS foundries (section 6).

Increasing collider intensity places two main demands on pixel detectors: the ability to store and process greater hit rate per unit area (section 4), and higher radiation tolerance (sections 3, 4 and 6). All hits ‘raining’ on a pixel detector unit area must be time-stamped and stored for a trigger latency interval. This means more memory per unit area in the ROIC, which directly translates to a need for a technology with higher logic density (i.e. smaller feature size, following Moore’s law). Note that the need for higher logic density is not a function of pixel size, but of hit rate per unit area. Smaller pixels are needed to maintain efficiency (to avoid pileup of hits in a single pixel), for resolution (section 2), and for radiation tolerance (to keep the leakage current per pixel small). Higher intensity also means that collision events all ‘look alike’, because every

event has a large number of superimposed low energy scatters (pileup), and may or may not have a hard scatter of interest as well, typically with less energy than the sum of the underlying pileup. As distinguishing events of interest becomes more difficult, triggers must increase rate (for a given signal acceptance more background will pass the trigger) and/or use more information, including tracking. Pixel detectors must therefore output much more data (section 4.8).

Achieving larger scale and coping with higher intensity are necessary to increase the physics reach of particle physics experiments. The basic performance of a detector can also improve the physics reach, and therefore, the challenges of scaling and intensity are compounded by the desire to increase performance. Given signal size, sensor capacitance, and device specific transconductance in the ROIC, the pixel front end achievable performance is determined, and actual ROIC's come close to this limit. From a single pixel perspective, the analog performance limit given a power budget and ROIC and sensor technologies is well understood and achieved by the implemented circuits. New developments to increase performance are therefore looking beyond the basic model of a pixel detector as a collection of individual pixels with an output of 3D space points. Correlations between pixels and measurements with internal degrees of freedom (in addition to the usual spatial coordinates) can be used to improve pattern recognition, suppress background, and even provide input to a trigger system. The main degrees of freedom under study are timing, direction, and cluster shape and charge distribution (see section 2).

2. Space-time point resolution

2.1. Demands and current directions

The basic detection mechanism of silicon detectors is the generation and drift of mobile charges (e/h) in a depleted silicon junction. This charge cloud has a rich spatial and temporal structure with some dependence on incident particle type and trajectory as well as existing electric and magnetic fields in the silicon. Silicon tracking detectors have typically had granularities in space and time greater than or equal to the charge cloud deposit, resulting in one 3D space point per particle crossing a sensor with a fairly coarse arrival time stamp, sufficient to associate the point with a given accelerator collision event. Measurement of the magnitude of the collected charge has also been generally available for pixel detectors, and has been used to improve the 3D space point precision through interpolation as well as for particle identification through specific ionization measurement.

In current detector development we are starting to see increased space and time granularity, in order to measure and make use of the structure of charge deposits. There are many interesting applications of this extra information beyond pure space points. We will give below a few present or anticipated examples: multi-track to cluster association using machine learning, an angular information from cluster length, 3D cluster shapes using charge arrival time, and disentangling of multiple interactions from sub-nanosecond hit timing.

To be fully exploited, the trend towards fine granularity in space and time, to resolve the charge deposit structure and even the sequence of particle impacts within a detector, must be accompanied by an evolution of the pattern recognition and track fitting algorithms used. Interestingly, algorithm evolution seems to lag detector development. Currently used techniques are still largely based on space points, with extra information added for specific tasks after space point reconstruction. In this chapter we review the basics of space point reconstruction and extend towards new space-time measurement directions as far as they are known and published today. But we expect to soon see modifications of the conventional, well known formulas to include angle, direction, and time measurements.

In general, the tasks of pixel detectors in HEP experiments can be listed as follows:

1. Pattern recognition and identification of particle tracks at large background and pile-up levels
2. Measurement of primary and secondary vertices;
3. Multi-track separation and vertex identification in the core of (boosted) jets;
4. Momentum measurement of particles (together with other detectors, like strip detectors);
5. Measurement of specific ionization.

Small pixel size keeps the pixel occupancy down at high particle rates (important for item 1) and also leads to good hit resolution. Space point resolutions in the order of $10\ \mu\text{m}$ or less have been routinely achieved at least in one dimension. Achieving the best possible resolution is of utmost importance to cope with the above challenges.

Recent developments (e.g. [14]) also address processing of multi-hit pixel clusters as complex objects to exploit directional information and improve the two track resolution, in particular inside the core of jets.

Time resolutions at the 10–100ps level have not been realized so far with semiconductor tracking detectors. If achieved this will add another measurement dimension which can be exploited for example to distinguish calorimeter jets coming from a hard interaction from those with large pile-up contributions by using the arrival time distribution (being broad for pile-up jets). Precise timing might also allow coincidence measurements e.g. between tracks in the forward and backward regions of a detector.

2.2. Space point and direction measurements

Figure 1 sketches two typical situations for pixel hit clusters. Case (1) represents the case of a particle impinging close to perpendicularly to the pixel module leading to hit clusters of typically one or two pixels. Case (2) represents tracks impinging at steep angles, thus producing larger hit clusters with some directional information when properly treated by reconstruction algorithms, in particular when such clusters appear in several detector layers along a track. Exploiting the hit information this way could become important for hit assignment and track recognition in high track density at the LHC (see for example [14]).

Figure 1. Pixel hit clusters for tracks under different incident angles.

2.2.1. Point measurements. The precision of a space-point measurement enters the momentum resolution in a track measurement with N detector layers as given by the Gluckstern formula [15]:

$$\frac{\sigma_{p_T}}{p_T} = \left(\frac{p_T}{0.3|z|} \frac{\sigma_{\text{point}}}{L^2 B} \sqrt{\frac{720}{N+4}} \right) \oplus \left(\frac{\sigma_{p_T}}{p_T} \right)_{\text{MS}} \quad (1)$$

where p_T is transverse momentum in GeV c^{-1} , L is the radial length in m, B magnetic field in T, z is the particle electric charge in elementary units, σ_{point} is the point resolution of the detectors in m, and N is assumed to be large in this approximation. Important for a precise momentum measurement is the point resolution, but also (quadratically) the total length L of the tracker and the bending field B . The multiple scattering (MS) contribution for a number of detector layers N can be written as

$$\left(\frac{\sigma_{p_T}}{p_T} \right)_{\text{MS}} = \frac{0.0136}{0.3 \beta B L} \sqrt{\frac{(N-1)x/\sin\theta}{X_0}} \sqrt{C_N}, \quad [L] = \text{m}, [B] = \text{T} \quad (2)$$

where L is the tracker length projected onto the plane perpendicular to the magnetic field, 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. 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; it approaches $C_N = 1.33$ for $N \rightarrow \infty$ (continuous scattering).

Even though the hard collisions at the LHC produce high energy jets, low momentum tracks around and below 1 GeV c^{-1} transverse momentum play an important role, especially in the forward direction. For example, efficient jet tagging—jets in general and b-jets in particular—and suppression of pile-up contributions by primary vertex identification suffer from imprecise detection of low momentum particles. For the HL-LHC, low p_T tracks will become even more important as pile-up increases. This renders low mass (low x/X_0) extremely important.

Similarly, the precision of secondary vertex measurement with an N layer tracker can be expressed by the impact parameter resolution (here in linear extrapolation):

$$\sigma_{d_0} \approx \frac{\sigma_{\text{point}}}{\sqrt{N}} \sqrt{1 + \frac{12(N-1)}{(N+1)} \left(\frac{r}{L} \right)^2} \oplus \theta_0 r_{pv} \sqrt{\frac{N(2N-1)}{6(N-1)^2}} \quad (3)$$

where the first term results from the extrapolation from the tracker to the primary vertex with r/L being the ratio of the

extrapolation distance to the tracker length. The point resolution enters linearly. For a pixel detector with four layers as in ATLAS, at radii between 3.3 cm and 12.3 cm and with a point resolution of about $10 \mu\text{m}$, this yields $\sigma_{d_0} \approx 12.5 \mu\text{m}$ without multiple scattering.

The second term is due to multiple scattering approximated by assuming extrapolation from the first layer to the primary vertex, the slope of which is smeared by multiple scattering, with θ_0 being the multiple scattering angle [16]:

$$\theta_0 \approx \frac{0.0136 \text{ GeV c}^{-1}}{\beta p} \sqrt{x/X_0}, \quad (4)$$

and r_{pv} the distance of the first pixel layer to the primary interaction vertex. For a 4-layer geometry like in ATLAS and a material thickness of typically around $3\% X_0$ the multiple scattering contribution to the d_0 resolution yields

$$\sigma_{d_0}^{\text{scat}} \approx \frac{90 \mu\text{m GeV c}^{-1}}{p}. \quad (5)$$

Spatial resolution also impacts pattern recognition performance in complex interplay with pixel occupancy, material thickness, and layout. However, pattern recognition algorithms are outside the scope of this review.

2.2.2. Space-point reconstruction methods. In designing a pixel detector one selects the (initial) amount of charge sharing. The tuning parameter for a given planar sensor thickness is the ratio of the Lorentz angle (the deviation from perpendicular to the sensor surface of the drift path of charges in a magnetic field [17]), to the tilt angle of the module with respect to perpendicular particle passage. While charge sharing between pixels allows for better spatial resolution by charge interpolation, a signal decrease caused by irradiation during the detector's lifetime demands minimal charge sharing. Whether collected charge causes a hit also depends on the pixel threshold. Noise deteriorates the precision of reconstruction and causes spurious hits.

For the present LHC pixel detectors with relatively large pixels, covering mainly central rapidity and not having too steep incidence angles, the number of hits in a cluster from a single track is rarely greater than 2. At large incidence angles in the forward regions of a barrel pixel detector or in the future with smaller pixel sizes, the cluster distribution can become much larger and this information can also be exploited (see below).

Resolutions achievable with classical space point reconstruction methods can be classified as follows:

- For single hit clusters—independent of having binary (yes/no) or analog origin—resolution is given by the well-known pitch/ $\sqrt{12}$ RMS resolution assuming a flat prior distribution of track position within the pixel (most conservative assumption). For an arbitrary given cluster size and track direction the resolution still is constant and is determined by the RMS of the prior track distribution that can produce a cluster of this size. For large clusters, high spatial resolution is possible even with binary readout, due to the averaging of many pixels.

- Analog hit information can be obtained at the expense of a larger total data volume (see e.g. [4, 18] for ATLAS) or by processing the analog pulse off the module as is done in CMS [7]. In this case the reconstructed hit position x_{rec} can be obtained e.g. by the centre of gravity method

$$x_{\text{rec}} = \frac{\sum (S_i + n_i) x_i}{\sum (S_i + n_i)} \quad (6)$$

where the S_i and n_i are the signal and noise fractional weights, respectively, and x_i is the center of each individual pixel in the cluster. The achievable resolution is:

$$\sigma_x^2 = \sigma_n^2 \left[\left(\sum_{i=1}^N x_i^2 \right) + N \langle x^2 \rangle \right] + \mathcal{O}(\sigma_n^3), \quad (7)$$

where σ_n denotes uncorrelated noise and N is the number of hit pixels [19, 20].

- If two hit clusters is the most common case for pixel detectors the η -reconstruction method [21, 22] is optimal for space reconstruction of Gaussian charge clouds, since detector effects are automatically included. For two adjacent left and right electrodes the function

$$\eta = \frac{S_L}{S_L + S_R}, \quad (8)$$

constructed from measured signals $S_{L,R}$ is position dependent, $\eta = \eta(x)$, and the hit position is obtained from the inverse η^{-1} :

$$x_{\text{rec}} = \eta^{-1} \left(\frac{S_L}{S_L + S_R} \right) = \frac{a}{N} \int_0^\eta \frac{dN}{d\eta'} d\eta' \quad (9)$$

where $dN/d\eta'$ is the probability distribution (with normalization N) of hits versus η' resulting from uniform illumination (random x distribution), and a is the channel pitch. Including noise $n_{L,R}$ (fractional, i.e. in units of the signal) (9) becomes [19, 20]:

$$x_{\text{rec}} = \eta^{-1} \left(\frac{S_L + n_L}{S_L + S_R + n_L + n_R} \right) = \eta^{-1} \left(\frac{\eta(x) + n_L}{1 + n_L + n_R} \right) \\ \approx x + \frac{d\eta^{-1}(s)}{ds} \Big|_{\eta(x)} \cdot \left(n_L(1 - \eta(x)) - n_R \eta(x) \right) \quad (10)$$

with $s = \eta(x) + n_L(1 - \eta(x)) - n_R \eta(x)$. Obviously, a steep response function η yields the best resolution, given by:

$$\sigma_x^2 = \sigma_n^2 \left\langle \frac{1 - 2\eta + 2\eta^2}{\eta'^2} \right\rangle + 2 \langle n_L n_R \rangle \cdot \left\langle \frac{\eta^2 - \eta}{\eta'^2} \right\rangle, \quad (11)$$

where the second term vanishes for uncorrelated noise, $\langle n_L n_R \rangle = 0$.

By using analog information either directly (CMS) or via a Time over Threshold (ToT) digitization (ATLAS), resolutions as shown in table 1 have been obtained (references in table). The ToT method is explained in section 4.6.

Table 1. Spatial resolutions obtained with the ATLAS and CMS pixel detectors under LHC conditions. R/O stands for readout and ToT for time over threshold.

Experiment	R/O method	Pitch (μm)	inc. angle ($^\circ$)	σ_x (μm)
ATLAS [23]	Binary/ToT	250	0	66.5
ATLAS [23]	ToT	50	0	10
ATLAS [4]	ToT	50	10	7
CMS [24]	Analog	150	0	44
CMS [24]	Analog	150	30	21
CMS [24]	Analog	100	0	9.4

2.2.3. Limitations to the space point resolution of a pixel detector. When aiming for ultimate spatial resolutions with hybrid pixel detectors the following limitations are encountered:

- Since for hybrid pixels the chip's pixel area must match the sensor's pixel area the smallest pixel size is determined by the amount of CMOS electronics needed to amplify, discriminate, and process the hit information in the area occupied by the pixel cell. Third generation chips for high hit rate achieve about $50 \times 50 \mu\text{m}^2$ or $25 \times 100 \mu\text{m}^2$, limited by logic density as explained in section 4.
- The signal of a traversing MIP (minimum ionizing particle) spreads due to diffusion according to $\sigma_x = \sqrt{2Dx/v_D}$, where D is the diffusion constant, x the drift distance to the electrode and v_D the drift velocity. For a typical detector thickness of $150\text{--}200 \mu\text{m}$, σ_x becomes $4\text{--}8 \mu\text{m}$. Hence pixel pitches well below this (at the same sensor thickness) would lead to excessive charge sharing.
- Pixel hybridization technology (see section 3) can currently cope with bump pitches in the order of $25\text{--}50 \mu\text{m}$. The technological limit for galvanic or evaporation methods for the coming decade seems to lie in the order of $5\text{--}10 \mu\text{m}$ [25] (see also section 3.3).

2.2.4. In-pixel decoding. It has been proposed [26] to effectively achieve smaller pixel size by encoding and decoding the pixel sensor cell by means of 'smart sensors' when using depletable substrates with a CMOS electronics layer (see also section 6).

Figure 2 shows the principle. Two pixels of the readout chip (e.g. FE-I4 with $50 \mu\text{m} \times 250 \mu\text{m}$ each) match six (sub-) pixels of the smart sensor with sizes $33 \mu\text{m} \times 125 \mu\text{m}$. Three sub-pixels at a time are connected to one of the two FE-I4 inputs via AC coupling, either by a capacitor included in the sensor pixel plus DC connection (bump bonding) to the corresponding FE-I4 chip pixel, or by gluing the sensor chip/wafer to the FE-I4 wafer, in which case the coupling is also capacitive. The digital (voltage) output pulse of the smart sensor is capacitively coupled into the CSA input of the readout chip. The three digital output stages of the sensor pixels can be programmed with three different gains (high, medium, low) such that digital pulses of different height are capacitively coupled to the input of FE-I4 depending on which of the sub-pixels has been hit. Further processing of the pulse by the FE-I4

Figure 2. In-pixel en- and decoding using a ‘smart sensor’ bonded to a pixel readout chip (here showing 6 smaller sensor pixels bonded to two readout pixels). The sensor pixels contain amplifications with adjustable gains (here $\times 1$, $\times 2$, $\times 3$). The gains are set such that the origin of a hit can be decoded from the output pulse height encoded as ToT. (a) Coupling of smart sensor and readout chip. (b) Function principle.

Figure 3. Improvement in hit cluster assignment and tracking efficiency using artificial neural networks to resolve merged clusters for tracking in dense environments (TIDE). Reproduced from [14]. CC BY 3.0. (a) Cluster assignment efficiency as a function of separation of two pion tracks from simulated $\rho^0 \rightarrow \pi^+\pi^-$ decays at the innermost pixel detector layers. (b) Track reconstruction efficiency inside dense jets for high p_T jets. (a) Cluster to track assignment. (b) Track reconstruction efficiency inside dense jets

amplification and discrimination stages turns different pulse heights at its input into different ToT values at the output. Hence the sub-pixel hit by the MIP is finally decoded from its ToT value. Test beam measurements demonstrated the functionality of the method [27, 28].

2.2.5. Directional information from hit clusters. Multivariate algorithms like artificial neural nets (ANN) have now also entered the area of hit/cluster to track association for improved track reconstruction in high multiplicity environments. The single hit precision is not the only figure of merit to tune. One must also optimize the capability of cluster shape analysis (especially for inclined tracks) and the analysis of pattern information of clusters from dense tracks over several layers. For example in [14] it is shown how successive employment of ANNs can improve (a) the association of merged clusters to tracks, (b) track reconstruction inside the core of (boosted) jets, and (c) heavy flavour and τ -jet identification. The amount of charge per pixel, for instance provided by the ToT technique [4], and a precise knowledge of the pixel coordinates is already sufficient as input for the ANN to identify merged clusters efficiently. The emission of δ -rays and secondary interactions of tracks prevent a perfect ANN performance. The approach benefits from knowing the incident angle of tracks, estimated from the coordinates of modules with respect to the

beam spot, and from correlating information from consecutive layers of the pixel detector [14]. Figure 3(a) demonstrates using simulation the gain in cluster assignment efficiency using this technique on a sample of simulated $\rho^0 \rightarrow \pi^+\pi^-$ decays. Figure 3(b) shows the improvement in tracking efficiency inside dense jet-cores obtained using the ANN-based improved merger identification. Improvement in high- p_T heavy flavour identification of 7–13% (b-tagging and b-jet identification) and of 5% in 3-prong τ -jets above $p_T = 600$ GeV is the result [14].

2.3. Time measurement

The charge collection time, i.e. the time until the arrival of the last electron at the pixel electrode, typically is in the order of 3–10 ns, depending on sensor thickness and E-field. The induced current pulse at the electrode instantly starts with the particle passage and consists of the incoherent sum of induced pulses from each carrier’s path towards (e^-) and away (h^+) from the pixel electrodes. The individual contributions are weighted by the pixel weighting field (see e.g. [1, 19, 29]) leading to the fact that the carriers’ paths close to the pixel implant have the strongest contributions to the signal. In a typical pixel detector the current pulses are integrated by the input stage of the subsequent readout chip, a charge sensitive

Figure 4. Cross section through an LGAD structure using a high ohmic p-type bulk and featuring a metallurgical junction underneath the n^{++} electrode acting as the amplification layer (adapted from [38]). The shape of the electric field strength is indicated on the left. Reprinted from [38], Copyright (2015), with permission from Elsevier.

amplifier (CSA). While such a readout method has low noise performance³ it does not allow timing measurements at or below 100 ps, values at which timing becomes interesting for high energy particle detection.

For a time coordinate measurement to reach equivalent spatial resolutions in the cm to mm regime, time resolutions of the order of a few tens of picoseconds are needed. To achieve this with pixel detectors, where charge collection times are in the ns range, is a real challenge. Recently, amplification structures in silicon have been brought forward to cope with these demands [31–33]. Potential benefits are the suppression of pile-up jets by recognizing the time-wise association of tracks to a primary vertex.

2.3.1. LGAD structures. In order to reach into the picosecond timing regime with silicon detectors, in-silicon amplification can be employed. Avalanche generating silicon based devices have been developed in the context of photodetectors (see for example [34–36]). They can be distinguished by their operation mode, either as linear avalanche photodiodes (APD) operating in a linear amplification mode or as so-called Geiger APDs operating near the breakdown point thus sensing a light pulse by a large breakdown pulse with follow-up recovery time. Timing in Geiger-mode operation will be governed by the time jitter introduced by the breakdown process (multiplication jitter), similar to the case for gaseous detectors where fast timing is addressed for example with RPCs (resistive plate chambers) [37]. Another approach has been chosen by [31], namely by operating avalanche silicon diodes at a low gain operation point, so-called *low gain avalanche diodes*, LGADs.

In order to create precise time stamps, fast and large signals and low noise are needed. The *slew rate* can serve as a figure of merit, defined as

$$\frac{dV}{dt} \approx \frac{\text{signal height}}{t_{\text{rise}}}. \quad (12)$$

Figure 4 shows a principal cross section of an LGAD structure [38]. In a p^- bulk, faced on both surfaces by the usual very highly doped⁴ n^{++} and p^{++} electrodes rendering the

Figure 5. LGAD signal pulse detailing the contributions from electrons and holes before and after amplification following weighting field simulations [39]. The contributions from both electron and hole signals are shown before and after amplification as well as the total signal (adapted from [38]). Reprinted from [38], Copyright (2015), with permission from Elsevier.

structure of the detector, an additional, highly doped p^+ layer ($N_A \approx \mathcal{O}(10^{16} \text{ cm}^{-3})$) is implanted immediately underneath the n^{++} electrode, thus creating a very local high electric field of up to 300 kV cm^{-1} . This is the multiplication layer amplifying all arriving electrons by a factor of the order of 10–20, thereby creating secondary e/h-pairs.

The signal induced on the electrodes therefore has several components (see figure 5):

- Electrons and holes drifting from the point of creation along the track towards their respective collection electrodes. The individual e/h parts of this contribution are small before amplification takes place and end when the last carriers have arrived at their respective electrodes.
- Amplification electrons created in the multiplication layer reach the top electrode (figure 4) almost instantaneously and their contribution to the induced signal current $i_S(t)$ is close to negligible, because according to the Shockley–Ramo theorem [40, 41]

$$i_S(t) = e\vec{v}_D(t) \cdot \vec{E}_w$$

with $\vec{E}_w$ = weighting field and v_D = drift velocity, there is almost no current contribution since the drift duration for electrons is quasi non-existent and ending before the next primary electrons reach the amplification layer.

- By contrast, the multiplied holes drifting from the amplification layer towards the backside contribute to $i_S(t)$ as

³ Typical noise values are on the order of $150 e^-$ for the present LHC detectors [4, 7] and will be around $80 e^-$ for future detectors with smaller pixels [30].

⁴ We here adopt the notation p^{++} for $N_A \gg 10^{16} \text{ cm}^{-3}$, likewise for n^{++} .

Figure 6. LGAD weighting field simulations and measurements: (a) slew rate as a function of detector thickness for different amplification gains (simulation); (b) comparison of time resolutions from weighting field simulations (WF2) and from test beam measurements using constant fraction discrimination. Open triangles are simulation points, crossed circles are measured resolutions without errors given (bars do not indicate the measurement errors). Note the different pad sizes entering the figure. The WF2 simulations for $1 \times 1 \mu\text{m}^2$ are underlined by the dashed line only to guide the eye. The data are taken from [33, 38, 43, 44], and [42]. (a) dV/dt versus thickness. (b) Time resolution.

long as they are drifting. Their contribution adds up, provided that the weighting field is formed such that it does not suppress contributions from distant holes too much. This implies that the pixel/pad implant widths must be large compared to the implant distances and should be in the same order of magnitude as the detector thickness.

As can be seen from figure 5, a large slew rate and hence fast timing requires fast (primary) electron arrival at the amplification zone together with fast (amplified) hole movement away. Short bunching and fast movement in high fields ($\sim 20 \text{ kV cm}^{-1}$) benefit from thin sensors.

The time resolution has several contributions [42]:

$$\sigma_t^2 = \underbrace{\left(\frac{V_{th}}{dV/dt} \right)_{rms}^2}_{\sigma_{\text{time walk}}^2} + \underbrace{\left(\frac{\text{Noise}}{dV/dt} \right)^2}_{\sigma_{\text{noise}}^2} + \sigma_{\text{arrival}}^2 + \sigma_{\text{dist}}^2 + \sigma_{\text{TDC}}^2 \quad (13)$$

The first term represents time walk introduced by different signal pulse heights due to different energy deposits coming from Landau fluctuations in the energy loss process and its resulting pulse height fluctuations at the discriminator input. Minimization of this term is possible e.g. by employing a constant fraction discrimination or a ToT correction architecture in the readout. The second term is noise jitter, i.e. time fluctuations due to noisy signals. Both contributions are made small when large slew rates are achieved. An irreducible contribution comes from fluctuations resulting from non-uniform depositions of charge along the particle track moving towards the amplification junction, hence causing an intrinsic jitter in the arrival time σ_{arrival} (third term). Fluctuations from secondary ionizations and from the amplification process enter here as well. The thinner the detector the less disturbing is this effect. An additional 4th contribution is signal distortion due to non-uniform weighting field regions and variations in (non-saturated) drift velocities. The final term is time fluctuations due to uncertainties in the time digitization. It is

assumed that the latter can be made negligible (below 10 ps) with GHz TDCs.

The key ingredients for fast timing and small time jitter are thin detector substrates ($\mathcal{O}(50 \mu\text{m})$) providing three essential benefits:

- larger slew rate ($\times 1.6$ slew rate increase going from $d = 100 \mu\text{m}$ to $50 \mu\text{m}$ at a gain of 20),
- better bulk radiation hardness,
- smaller arrival time fluctuations due to a shorter charge deposition path.

In figure 6(a) the slew rate is plotted against the diode thickness for different gains assuming an input capacitance of 2 pF [38]. It is evident that large slew rates can be obtained with thin detectors and sufficient gain ($\gtrsim 10$). Much higher gains compromise the S/N ratio, because bulk shot noise is also amplified by the amplification structure (called excess noise) such that there is an optimal gain for maximum signal-to-noise (see for instance [36] or [42]).

Prototype LGAD structures have been fabricated [38, 43] with relatively large pads (from $8 \times 8 \text{ mm}^2$ to $1 \times 1 \text{ mm}^2$), rather than pixels, mainly to validate the underlying models and simulations of the achievable time resolution. Figure 6(b) shows the predicted time resolution by simulations [39] as a function of the detector thickness. At the right hand side measured data from $300 \mu\text{m}$ thick LGAD pad detectors are compared with simulations. Extrapolation to thin detectors yields time resolutions in the order of 30 ps. Very thin detectors have only recently become available. In a test beam measurement with 180 GeV pions, a $50 \mu\text{m}$ LGAD sensor with $1.2 \times 1.2 \text{ mm}^2$ pads has been characterized, achieving 34 ps time resolution at a bias voltage of 200 V [44], in excellent agreement with predictions from weighting field simulations (figure 6(b)).

A major concern for the use of LGADs at the HL-LHC is their performance in high radiation areas. Apart from the usual effects caused by radiation in silicon detectors like leakage current increase and deteriorated charge collection efficiency,

effective doping changes play a sensitive role. Since the metallurgical junction providing the amplification gain requires p-doping concentrations of order of 10^{16} cm^{-3} , radiation induced acceptor removal will have detrimental effects upon the gain. Detailed radiation studies are currently ongoing [42], including very thin sensors ($\lesssim 50 \mu\text{m}$) and devices with gallium substituted as p-dopant in the amplification implant instead of boron, to reduce interstitial capture that scales with atomic mass [45]. Another proposed method is to use carbon-enriched wafers with (abundant) carbon being trapped by silicon interstitials rather than boron.

Currently LGADs are not yet discussed for pixel tracking detectors but rather as precision timing layers, aiming to distinguish primary vertex positions (in z) from each other (see e.g. [46]). CMS-TOTEM is planning to use LGADs in roman pot stations [47].

3. Pixel sensors and hybridization

‘Hybrid pixels’ including a sensor (electronically passive) and a readout chip as separate entities (figure 7) currently are the state-of-the-art technology for large scale pixel detectors in most particle physics experiments [1]. For both parts, the radiation requirements at LHC-type experiments constitute a challenge which is addressed by dedicated R&D and process engineering techniques. In this chapter we describe passive sensor types used in hybrid pixel assemblies, to be contrasted to CMOS active monolithic sensors that integrate electronics circuitry and sensor, described in section 6. The mechanical and electrical mating of sensor and readout chip, called hybridization, requires state-of-the-art processing technologies, also described in this chapter (section 3.3). The hybridization mating partner, the readout chip is described in detail in section 4.

After the success of planar sensors (section 3.1) in LHC run-1, having planar pixel implants at one side of the sensor and being fabricated using 4" sensor wafer technologies, the desire for large detector modules at affordable cost has fueled the move to 6" and possibly even 8" wafer sensor technologies. In line with this a better understanding of radiation damage effects and radiation hard sensor device engineering became a research branch on its own, coordinated within the CERN RD50 collaboration [45].

Radiation hard device engineering includes using silicon with increased oxygen content supplied in the growth process [48–50] and operation at low-temperature, both employed to reduce the damage’s impact on the detector performance. Furthermore, so-called 3D-silicon sensors [51] (see section 3.2) have been developed having vertically structured electrodes fabricated within the volume of the silicon substrate (figure 11). Their main feature is a better ratio of ionization thickness to charge collection distance than for planar sensors. After an operation performance demonstration in the first pixel upgrade of ATLAS (Insertable B-Layer, IBL [5]), an innermost pixel barrel layer at a distance of 3.5 cm from the interaction point, equipped with 25% 3D-Si sensors (at high η) and 75% planar pixel sensors, both sensor types now

compete for the HL-LHC upgrade as pixel sensors for the radiation hot areas near the interaction point.

A challenge for both sensor types imposed by the radiation fluences of $>10^{16} n_{\text{eq}} \text{ cm}^{-2}$ is the small foreseen sensor thickness (100–150 μm) and the correspondingly increased handling and hybridization difficulty.

Another advance in cost reduction of sensor wafers has been made possible by the availability high quality p-substrate material ($\gtrsim 2 \text{ k}\Omega\text{cm}$). With n-in-p type sensors cost-efficient single sided processing suffices for fabrication while maintaining electron collection (n-type pixels). During the initial LHC pixel development period (late 1990s) high ohmic p-type sensor wafers were not favoured, largely for reasons of historical development⁵, and n^+ -in- n planar pixel sensors have been used instead, requiring processing steps on both sides of the wafer.

With CMOS technology vendors becoming more open to smaller market customers, the use of high ohmic 8" wafers fabricated in (much lower cost) CMOS processing lines become options for (planar) sensor options (see section 3.1.1).

Currently, therefore, the trend in hybrid pixel module development goes to large readout chips and thin sensors, preferentially using p-type bulk material, all being beneficial for cost, radiation tolerance, and detector mass, while the handling and hybridization demands increase. The latter together with the increased interest for large area coverage using pixel detectors also revived collaboration with industry to develop lower cost hybridization techniques with thinner wafers/chips.

3.1. Planar pixel sensors

The proven standard of planar pixel sensors at the LHC experiments has n^+ pixel implants on n -type substrate material as shown in figure 8(a). The fabrication demands double-sided processing. The negative bias potential on the backside is brought down to zero potential at the sensor edge by a many guard ring structure implanted on the backside. The pn-diode is initially on the unstructured backside of the sensor, changing to the electrode side after radiation has turned the bulk’s n^- doping into effectively p-type (type inversion). For typical substrate resistivities of 2–5 $\text{k}\Omega\text{cm}$, this happens after fluences of some $10^{12} n_{\text{eq}} \text{ cm}^{-2}$. After type inversion the depletion zone grows from the electrode side into the bulk guaranteeing that signals induced by moving charges can be sensed on the pixels even if the substrate is no longer fully depleted, i.e. not depleted at the side opposite to the pixels. To prevent an electron accumulation layer from short circuiting the n^+ -in- n pixels before type inversion, shallow p -doping is used in between the pixel electrodes, with a smoothly changing doping profile preventing high field corners. Details are given in [1].

With increasing radiation fluence the bias voltage needed for sufficient depletion must be increased, eventually reaching

⁵ Single sided Si sensor development had begun using n-substrate material with hole collection on p-electrodes. For radiation tolerance reasons electron collection became favored leading to n^+ -in- n structures with double sided processing. For n-in-p sensors it then has taken time to develop n-side multi-guard-ring structures and reliable HV bias isolation from the electronics before high quality p-type sensors were fabricated.

Figure 7. Hybrid pixel detector: (a) Layout of an individual pixel cell having a sensor and an electronics cell in 1–1 correspondence; (b) a hybrid pixel matrix; sensor and electronics chip have pixels of the same size, bonded to each other by means of bump contacts. (a) Hybrid pixel. (b) Pixel matrix.

Figure 8. Planar sensors: (a) LHC conventional n^+ -in- n design requiring 2-sided wafer processing; (b) n^+ -in- p design with single-sided processing (bump) bonded to a readout chip. Note the small distance between readout chip and guard ring implants; (c) n^+ -in- p sensor hit efficiency as a function of bias voltage for different fluences and various thicknesses [52]. Reprinted from [52]. Copyright (2016), with permission from Elsevier. (a) n^+ in n design: 2-sided. (b) n^+ in p design: 1-sided. (c) Hit efficiency.

values of more than 500 V at fluences around and beyond 10^{15} neq cm^{-2} . For the innermost layers close to the interaction point this corresponds to detector lifetimes of only a few years at current LHC luminosity.

The development of planar sensors for the LHC high luminosity upgrade has therefore concentrated on tailored designs guaranteeing high fields and sufficiently large depletion depths after fluences of up to and above 10^{16} neq cm^{-2} . This has been achieved using n^+ pixel implants in p -substrate material and thin (100–150 μm) sensors operated at bias

voltages of 500–700 V [52]. While providing smaller signals per MIP, the benefits of thin sensors are higher electric fields as well as shorter and faster electron collection for a given bias voltage and hence better radiation tolerance. Thin 6'' or even 8'' sensor wafer production is enabled by techniques employing SOI or Si-Si handling wafers, or by thinning (e.g. by cavity etching) and forming a back side ohmic contact at low temperature after the front side processing is complete. Currently the limit in thickness is considered to be around 50 μm [53].

Figure 9. Hybrid pixel module using passive CMOS pixel sensors (schematic). Depleted sensor employing CMOS technology with 1–2 metal planes that can be used for (i) AC coupling and (ii) rerouting. The insert details these features. The yellow area in every pixel denotes the charge collection node.

The guard ring designs, bringing down the potential from the bias implants (HV) to the pixel implants (0V), play an important role in optimizing the breakdown behaviour of sensors after irradiation. The number of rings, implant distance and size, as well as metal overhangs for vertical field suppression are the parameters to optimize. Optimal performance regarding radiation tolerance has been obtained [53] with 10–11 rings with metal overhangs on both sides of the implant over a total length of $350\ \mu\text{m}$. The distance from the edge to the active part of the sensor is $\gtrsim 400\ \mu\text{m}$. Other designs trying to minimize this distance make use of etched and doped edges, called active edges, reaching values down to about $50\ \mu\text{m}$ [54].

Figure 8(a) shows a planar pixel design as used in present LHC hybrid pixel detectors ($200\ \mu\text{m}$ thick, n^+ -in- n design) compared to a thin n^+ -in- p design (figure 8(b)). Note that, if cost saving single sided processing is to be used, the guard ring structure needs to be placed on the electrode side and is thus in very close proximity ($\lesssim 20\ \mu\text{m}$) to the readout chip, a challenge for the design. Dedicated passivation with BCB (benzocyclobutene) or parylene, however, appears to be able to solve this issue [55]. Figure 8(c) shows that high hit efficiencies can be achieved after irradiation with thin sensors for bias voltages in excess of 500 V .

3.1.1. Use of 8" wafers and CMOS foundries. It is also possible to fabricate planar sensors on 8" wafers using high ohmic substrates either in dedicated sensor fabrication facilities or by employing CMOS foundries, as has been successfully demonstrated for strip and pixel sensors [56, 57]. This has the following advantages compared to the standard small volume sensor fabrication:

- Large volume production lines can be used with price advantages, suited in particular for large area detectors, e.g. in outer tracking layers at collider experiments.
- Wafer sizes of 8" or 12" are a commercial standard allowing large sensors when stitching over reticle boundaries is applied. Stitching is a method to cross the boundaries between reticles of a wafer with metal lines and implant areas, widely applied in the digital imaging industry [58]. Current reticle sizes are about $25 \times 25\ \text{mm}^2$. By putting individual building blocks of a design

Figure 10. Hit efficiency measured in 3.2 GeV electron test beams [60] of passive CMOS pixel sensors as a function of bias voltage, unirradiated and for two fluence levels. For low threshold, unirradiated, there is only one point for each thickness at minimal bias voltage. Reproduced from [60]. CC BY 3.0.

separately (rather than together) on the reticle/mask one can join them together over the reticle's boundary by special lithographic programming, thus allowing IC fabrication of sensor sizes much larger than the reticle area of the CMOS process.

- The wafers can be purchased with solder bumps already provided by the CMOS vendor (e.g. so-called C4 bumps [59], pitch $\gtrsim 150\ \mu\text{m}$).
- One or two metal layers can be exploited for AC coupling of the sensor cell to the amplifier and for connection redistribution, e.g. to connect amplifier inputs to pixels in areas not covered 1:1 by readout chip cells (regions in between readout chips).

A sketch illustrating some features employing passive CMOS sensors is shown in figure 9. Measurements on passive CMOS pixel sensors, $100\ \mu\text{m}$ and $300\ \mu\text{m}$ thick, irradiated to fluences of $1.1 \times 10^{15}\ \text{neq cm}^{-2}$ have shown lab and test beam performance equal to those of planar sensors fabricated in dedicated sensor production lines [60]. Figure 10 shows the mean hit efficiency as a function of bias voltage measured with 3.2 GeV electrons for DC and AC coupled passive CMOS sensors bonded to the ATLAS readout chip FE-I4. In particular the AC-coupled devices show excellent performance without any

Figure 11. 3D-Si sensors: (a) Design (single sided) with columns going completely through the sensor bulk [62]. Reprinted from [62], Copyright (2005), with permission from Elsevier; (b) double sided design with columns entering from both sides, but not reaching through (adapted from [66]). Reprinted from [66], Copyright (2008), with permission from Elsevier. (c) Thin design optimized for HL-LHC (adapted from [67]) with two top view sketches for $50 \times 50 \mu\text{m}^2$ and $25 \times 100 \mu\text{m}^2$ pixel sizes, respectively. Reprinted from [67], Copyright (2016), with permission from Elsevier. (a) 3D-Si: penetrating columns. (b) Partial 3D-Si (double sided). (c) HL-LHC design.

efficiency losses. The DC devices are a bit less efficient due to the area taken by the punch-through dot for biasing of the pixel implants.

3.2. 3D-Si sensors

So-called 3D-silicon sensors have been developed since the late 1990s [51, 61] featuring columnar electrode implants driven into the Si substrate perpendicular to the sensor surface (figure 11). The electrode distance is made smaller ($50 \mu\text{m}$) than the typical sensor thickness ($200\text{--}250 \mu\text{m}$), thus rendering a shorter average drift distance for particles impinging on the sensor face than in the case of planar sensors (compare figure 7(a) to figure 11(a)). In addition, high drift fields are obtained with still moderate bias voltages. Both these facts result in an increased radiation tolerance due to a reduced trapping probability.

The 3D-Si technique has been developed over many years. The first structures were fabricated at Stanford (later also at Oslo) [62] using single sided processing with columns reaching completely through the bulk (called ‘full-3D’ and shown in figure 11(a)). Further development by CNM [63] and FBK [64] of sensors used in the ATLAS IBL detector, resulted in double-sided 3D designs with columns entering the bulk from both sides, either in full-3D or in partial-3D (shown in figure 11(b)). The process fabricates about $10 \mu\text{m}$ diameter columns by etching, followed by a $1 \mu\text{m}$ polysilicon layer covering the inside of the etched holes, then passivated by a wet oxide [63]. In addition the sensor edges can be fabricated with active edge implants thus rendering sensors with

an unrivaled active area fraction [65]. More details on etching holes into silicon can be found in section 3.5.

Within the ATLAS IBL detector 3D-Si pixel sensors have been proven to operate well in a running experiment [5]. After two years of operation the performance of 3D-Si pixel modules in terms of operation characteristics (signal, noise, threshold settings, in-time efficiency) are on par with planar pixel modules operated with significantly higher voltages [68].

Current developments motivated by HL-LHC demands [63, 69–70] target the following goals to optimize radiation hardness, granularity, material budget, and processing costs [70]: (a) thin sensors ($\sim 100 \mu\text{m}$) on 6μ wafers, (b) narrower electrodes ($\sim 5 \mu\text{m}$), (c) shorter electrode spacing ($\sim 30 \mu\text{m}$), and (d) very slim ($\sim 50 \mu\text{m}$) or active edges. Single sided processing is preferred providing cost benefits. An advanced design [70] is shown in figure 11(c). A thin, highly resistive (p-type) sensor wafer is supported by a low ohmic (p^{++}) handle wafer that can be backside thinned after processing. While the p^{+} columns are deep etched through to the handle wafer where they receive their electric potential, the n^{+} columns stop about $15 \mu\text{m}$ short from the handle wafer. In addition to cost and yield advantaged of single-sided processing, studies have shown that the trade-off between signal efficiency and breakdown performance favors partial depth n-columns (not extending all the way through the thickness) [71]. At the top surface isolation of the n-columns is achieved by a p-spray layer preventing the electron accumulation layer underneath the oxide from creating shorts. Sensors are designed to meet the currently planned pixel area sizes of $50 \times 50 \mu\text{m}^2$ or $25 \times 100 \mu\text{m}^2$, as shown on

Figure 12. 3D-Si sensors: hit efficiency as a function of bias voltage for different fluences and design variants. Reproduced from [69]. © IOP Publishing Ltd. All rights reserved.

the right of figure 11(c). The performance of such designs has been shown to yield high breakdown voltages before and after irradiation [72].

The hit efficiency obtained with 3D-Si structures designed by CNM [69] is demonstrated in Figure 12 for fluences of up to 0.9×10^{16} neq cm⁻² [69]. A >97% efficiency plateau is reached with comparatively low bias voltages of 150V even at the highest fluences, with the missing 3% being largely due to tracks with straight incidence into the column structures, a case not possible for tracks from actual LHC collisions. Smaller implant pitches with lower voltages for the same field strength also reduce power dissipation due to leakage currents after irradiation. This creates some safety margin against thermal runaway, not completely negligible even at the moderate bias voltages of 3D-Si sensors [73].

3D-Si sensors therefore are a strong contender for hybrid pixel modules for the innermost pixel detector layers at the HL-LHC. It should be noted, however, that the fabrication process does is currently low volume, making it unlikely to cover large areas (>1 m²) with this technology choice.

3.3. Bonding techniques

Wafer- and chip-bonding as well as 3D-integration technologies are in widespread demand in the semiconductor industry sector of IC module assembly and stacking. Many techniques exist that are suitable for a large variety of applications and given fabrication constraints. A large amount of literature is available, reviewed and summarized for example in [74–76].

For hybrid pixel detectors, the mating of readout chips to substrate (sensor) plates or substrate wafers is the main application of bumping/flip-chipping and 3D-integration techniques. The overriding demands in these applications are small capacitance additions to the preamplifier by the bond connection, good yield with $<10^{-4}$ defect rate (open or short), good contacts (<100 mΩ), and robustness against temperature cycling (−40 °C to +60 °C). The need for small pitches between bond connections imposed a very high demand on industrial standards when pixel R&D for the LHC started. At the time the required pitch of 50 μm was about 15 years ahead

of industrial demands. This pitch has become a standard today and requirements for current hybrid pixel R&D often target pitches less than 50 μm.

Bump or under bump metalization (UBM, see below) application usually is a wafer-scale process. Flip-chipping, on the other hand, is conventionally employed to mate bumped readout chips to (bumped) sensor plates⁶. For the existing detectors at the LHC experiments and likely also for the upcoming upgrades, *eutectic soldering*⁷ [77] and *In–In* thermocompression bonding are the methods that have been preferred over other techniques (see e.g. [1]) including anodic bonding, fusion bonding, and adhesive bonding (see below). In current techniques ICs are usually bonded to sensor plates or wafers (chip-to-wafer bonding, C2W). Wafer-to-wafer bonding (W2W) is a cost attractive future possibility for pixel detectors that might become interesting for applications in connection with further advances in vertical electrical interconnection in so-called 3D integration techniques. In particular *through silicon vias* (TSV) open the possibility to provide elegant and space efficient electrical contacts for hybrid pixel module fabrication when employing W2W or C2W bonding techniques. TSVs also allow reaching through to the backside of a chip (or sensor). This can be exploited to use the chip’s backside metal for redistribution of readout or service lines (redistribution layer, RDL).

3.3.1. Solder bumps and bonding. One can perhaps subdivide the bonding technologies relevant in this review into (a) technologies requiring intermediate media to perform the bonding like for example eutectic bonding by means of solder or adhesive bonding using glue layers and (b) direct bonding, either metal-to-metal or silicon oxide-to-oxide. When intermediate materials are used no requirements on special surface treatments exist except for flatness. After UBM (and bump) fabrication, oxide removal is needed before flip-chipping can be done. The bonds are strong and post-processing—like thinning—is often possible. The lack of electrical connections (in the case of adhesive bonding) or price and pitch constraints (for solder bonding) are drawbacks.

The bonding technique currently used most often is the electrochemical application of solder micro bumps to either of the mating parts: chip or sensor (figure 13). A description of the method and the process steps is given in [1]. Solder bumping was introduced in 1969 by IBM in the C4-process (controlled collapse chip connection) [59]. The process is still in use by CMOS vendors and is often offered with CMOS wafer production as an add-on. The pitch of the applied bumps is limited to 170 μm or larger, often too large for fine-pitch, high granularity pixel detectors in particle physics. However, price and reliability arguments render C4-bumping still very attractive for low cost, large area applications at LHC upgrades, as is for example addressed in section 3.1.

⁶ Depending on the techniques one of the mating partners can have bumps as well, e.g. for *In–In* flip-chipping, or can just have UBM for the mating process, e.g. for solder bump flip-chipping.

⁷ Eutectic systems are solid mixtures that form a superlattice by striking a unique atomic percentage ratio and thus have the same melting point.

Figure 13. Eutectic solder bumps and the flip-chipping process: (a) Solder bump attached to a readout chip is bonded to the under bump metalization of the pixel sensor; (b) SEM micrograph of a SnAg bump. Reproduced with permission from Fraunhofer IZM Berlin. (a) Solder bumping and flip/chip. (b) Photograph.

Figure 14. In–In bonding: indium bumps produced before (a) and after (b) reflow to a sphere [1] (2006) Springer-Verlag Berlin Heidelberg. With permission of Springer. (a) Indium bumps (flat). (b) Photograph.

Eutectic solder bumping has meanwhile been developed with high connection density down to pitches of $25\ \mu\text{m}$ [78] and bump dimensions of $(15\ \mu\text{m})^3$. High yield is obtained with Sn/Pb as well as with lead free Sn/Ag or Au/Sn alloys [79]. The ultimate pitch limit is believed to be around $5\text{--}10\ \mu\text{m}$ [25]. Readout ICs typically have Al or Cu pads for bump connection bounded by a passivation layer for chip protection. Typically a thin (100 nm) Ti:W layer is first sputtered onto the wafers to promote adhesion and provide a barrier to prevent Cu atoms from diffusing into the pad metal (Al). Then a Cu or Au plating base ($>150\text{ nm}$) is applied for the electrochemical contact followed by a well wettable Cu layer ($1\text{--}5\ \mu\text{m}$). This layer stack is the UBM and is applied to both mating parts. In a wafer-based process cylinders of an electrodeposit are now grown onto the UBM of one part and turned into a spherical shape by reflow. Bonding of both parts is so far mostly done after dicing of bumped readout ICs and sensor plates or wafers supplied with UBM in a flip-chipping process and subsequent re-melting of the solder bumps. Note that this reflow process provides self-alignment of the mated parts.

Eutectic solder bumping is regarded as present day's work-horse bonding technique for hybrid pixel detectors. For the ATLAS pixel detector very high bond yields of more than 99% have been obtained (i.e. number of modules not rejected due to bonding issues).

3.3.2. In–In bonding. In–In bonding has also been used in LHC pixel detectors for a large quantity of pixel modules. The technique usually employs vapour deposition⁸ of indium through openings in the lift-off masks deposited on the wafer (see e.g. [1]). The mask is pulled off from the wafer by a wet lift-off process, requiring the bumps to be fairly flat and thin ($\approx 10\ \mu\text{m}$, figure 14(a)). Also a UBM is necessary for this process, usually Ti–Pt–(Au) is used. The applied indium bumps are then bonded by thermocompression at about $100\ ^\circ\text{C}$. The advantages of the technique are its ease of application and the low temperature requirements. The bonds are, however, comparatively fragile, and the fabricated modules tend to have a lower mechanical damage threshold than solder bumped assemblies. In the Run-1 ATLAS pixel detector about 50% of the modules were produced using In–In bonding, showing high reliability with acceptance yields around 90%. For the CMS pixel detector a reflow step has been introduced after bump deposition [80], turning the In bumps into spherical shapes (figure 14(b)). The two spheres are merged in a second reflow at $180\ ^\circ\text{C}$. This provides a larger distance between chip and sensor, but it introduces self-alignment and stronger bonds.

3.3.3. Adhesive bonding. Recently adhesive wafer-to-wafer bonding has been discussed in the context of active CMOS sensors, capacitively coupled to readout chips where the capacitive coupling is provided between the metal pads of sensor and chip, respectively, separated by the adhesive bonding

⁸ Electroplating of indium is also possible.

Figure 15. Process steps of solid-liquid interdiffusion bonding (SLID): (a) structure with metal stacks before bonding, (b) alloy forming step under pressure at 240–320 °C, (c) resulting bond connection.

layer. Adhesive bonding has been claimed to reduce production cost by avoiding the UBM, bump deposition, and reflow steps. This has, however, not yet been demonstrated and is not at all obvious. Planarity requirements are stricter and requirements for alignment for chip-to-chip/sensor placement are the same as for standard bump bonding, the latter being the cost driver, not to mention the need for (cost intensive) electrical connections of the mated parts if provided by TSVs.

In the context of capacitive coupling of ‘smart sensors’ to readout ICs (CCPD) adhesive bonding has been investigated regarding thickness control and uniformity of the adhesive as well as radiation tolerance [81]. Good glue uniformity at the micrometer level has been achieved, however further evaluation of its competitiveness is needed, in particular whether electrical connections can be provided by cheap techniques like e.g. wire bonding, thus, however, abandoning the goal of compact module shapes.

3.3.4. Cu–Cu direct bonding. Metal-metal adhesion (Cu–Cu or Au–Au) of flat and polished surfaces has been known for a long time. The bonding forces principally involve capillary, Van der Waals, and electrostatic forces, as well as solid bridges caused by impurities and hydrogen bonds between OH groups.

Two mirror-polished wafers are put into contact and held together by adhesive forces without any intermediate material [74]. The process completes by Cu atom diffusion between the two Cu layers. Thermocompression (diffusion) bonding is simplest, but the required temperatures are usually too high for typical sensors and CMOS chips used in pixel detector applications.

Especially interesting is surface-activated Cu–Cu bonding as it works at ambient temperature. Surface-activation in this context means increasing the bonding force by surface treatment, e.g. by augmenting the number of hydrogen bonds (hydrophilicity) or by generating new types of chemical bonds. The treatment methods comprise wet chemical processes and (oxygen) plasma etching. The bonds are electrically conductive, an advantage for any wafer-to-wafer but also chip-to-wafer bonding project.

Requirements for good bonding are native oxide and other remnant removal, excellent chemical-mechanical polishing (CMP), and highly planar parts (wafers) with sufficient total Cu fraction. Very fine pitch (<4 μm) and low capacitance contacts are possible. The demands on surface cleanliness and

flatness are the biggest drawbacks. Wafers must be processed soon after fabrication. Wafers must be very planar and sufficiently stiff which might compromise very thin solutions. At present the technique is also still fairly expensive.

3.3.5. Oxide-oxide direct bonding. Under room temperature conditions silicon wafers are covered with an oxide layer that can be used for direct wafer bonding, a method commonly used for SOI wafer production. Most common is ‘hydrophilic’ wafer bonding in which the wafers are, after cleaning, rinsed or stored in deionized water. Water is bonded via Si–OH groups on the silica surface. The bonding can be done at room temperature but the bonding strength increases with temperature (and time). With the addition of an electric field one speaks of anodic bonding, which can tolerate rougher surfaces.

Required is high quality CMP polishing and extremely good surface cleaning to avoid large bonding voids. The bonds are non-electrical. Hence TSVs or other electrical connections are needed for pixel assemblies. For this process wafer-to-wafer bonding is much easier than chip-to-wafer or chip-to-chip.

3.3.6. Solid-liquid interdiffusion bonding (SLID). SLID bonding uses an intermetallic alloy formation and represents an alternative fine-pitch bonding technique based on thin (often eutectic) Cu–Sn connects [74, 82]. Between two Cu layers of several micron thickness a thin (3 μm) Sn layer is applied on one of the Cu layers (figure 15). They are brought in contact and heated to a temperature of around 240–320 °C. At this temperature the tin diffuses into the copper forming the Cu–Sn alloy. As the melting point of this alloy is around 600 °C multiple layers can be stacked and connected without melting the previously formed SLID connections. Furthermore, the process is also a flux-free bonding alternative with pitches much below 20 μm, whereas the number of processing steps is about the same as for solder bonding (the reflow step is not needed). Technically, for SLID, pressure and temperature are required during the bonding step. This renders sequential bonding of multi-chip pixel modules necessary (chip after chip), whereas bonding by reflow does not require pressure (see page 13) and can be done in parallel for a number of chips bonded to a sensor plate.

Despite their similarity SLID and soldering are fundamentally different processes with distinct and unique properties. A

Figure 16. 3D-integration principle: (a) TSVs connecting different electronics device layers [75] John Wiley & Sons. Copyright © 2000-2018 by John Wiley & Sons, Inc., All rights reserved. (b) bonding of two wafers with electrical SuperContacts to the outside [83] John Wiley & Sons. Copyright © 2000-2018 by John Wiley & Sons, Inc., All rights reserved. (a) Device layers connected by TSVs. (b) Bonding of two wafers with electrical connection via a SuperContact TSV.

major difference is that soldering is reversible, whereas SLID is irreversible and only melts at temperatures much higher than the SLID process temperature. This implies also that no reworking is possible. Some process similarities to metal-metal thermocompression bonding also exist. A detailed comparison is given in [74]. SLID is particularly suited for wafer-to-wafer and chip-to-wafer bonding.

3.4. Wafer thinning

Thin wafers are a goal not only for low-mass pixel detectors in particle detection, but also for a number of integration processes, in particular for through-silicon vias. While thick wafers are better for handling, other process steps like for example etching benefit from small wafer thickness. However, dedicated ‘handling wafers’ are needed to deal with flatness and bowing issues. Often handling wafers are of the SOI (silicon-on-insulator) type, because the Si–SiO₂ interface transition offers a sharp etch stop with the oxide also acting as a sacrificial layer after ion etching.

Wafers are thinned by (backside) grinding. The active (front) side is first protected by tape. Grinding in two steps (coarse/fine) is performed using grinding wheels with diamond grain sizes from 1–8 μm (fine) to 20–80 μm (coarse). After the thinning process, breaking strength, warpage, and bow are characterization parameters that depend on the grinding tool characteristics and the grinding speed. More details are given in [75].

3.5. 3D-integration and through-silicon vias

Extending electronics integration into the third dimension is—apart from feature size shrinking—regarded as the second route to ever increasing circuit density. 3D-stacking of several electronic device layers (tiers) is thus an eminent field of industrial research. Further advantages are reduced power consumption due to smaller connections and smaller involved capacitances as well as larger I/O bandwidth and more functionality at lower cost. Vertical vias running through layers of silicon (TSVs) are a key ingredient for 3D-stacking [75]. Tiers can be different

CMOS layers but also electronics and sensor layers or layers interfacing to an optical signal transport. Figure 16(a) is a sketch illustrating the principle of 3D-interconnection via TSVs.

Besides TSV fabrication, wafer thinning and aligned wafer bonding are the key follow-up processes for 3D-integration. Direct bonding, especially Cu–Cu fusion or SLID (see above) are the preferred bonding techniques for 3D-stacking, providing strong fine-pitch bonds that are also conductive.

For future pixel detectors in particle physics, 3D-integration and TSVs are of interest if (a) wafer scale production becomes practical and/or (b) if TSVs allow better connectivity of detector modules by avoiding space consuming wire bonds and providing better buttable modules. In addition TSVs can reach through to the backside of a chip, allowing redistribution of electrical connections on the chip’s backside metal is addressed. This in turn saves the material normally added by flex interconnect circuits.

3.5.1. TSV fabrication. In IC fabrication all processes before the first wiring metal are called *front end of line* (FEOL), whereas the *back end of line* (BEOL) begins with this first metal processing and ends with the last IC processing step. Different types of TSVs are distinguished depending on the point in the process flow at which they are fabricated [86], i.e.

- *via-first* or FEOL via (polysilicon) applied before or during FEOL fabrication.
- *via-middle* or BEOL via (metal) applied after FEOL during BEOL fabrication
- *via-last* or post-BEOL via (metal) applied after IC fabrication post-BEOL. Here one distinguishes also *front side via-last* as in figure 17(right) and *backside via-last*, depending on whether the via is etched from the front through the BEOL stack (difficult) or from the back through (thinned) silicon substrate.

Note that the conducting material for FEOL TSVs must be doped polysilicon for reasons of thermal and material compatibility, having the disadvantage of higher resistivity compared to metals. For BEOL and post-BEOL vias tungsten and

Figure 17. Different types of vias. (Left) *via-first*: fabrication before or during FEOL, (middle) *via-middle*: fabrication after FEOL during BEOL, (right) *via last* fabrication after IC completion (post-BEOL). Reproduced with permission from [84]. Semiconductor Industry Association (SIA).

Table 2. Typical TSV parameters (2017) [75, 86, 87].

Via type	Smallest diameter (μm)	Smallest pitch (μm)	Typical aspect ratio
Via-first	3–5	6–10	10 : 1
Via-middle	3–5	6–10	5–10 : 1
Front side via-last	10–20	20–40	5 : 1
Backside via-last	5–20	10–40	5–12 : 1

copper are most common. Figure 17 shows these distinct TSV types schematically.

Important TSV specifications are diameter, pitch, and aspect ratio (length-diameter ratio). Present typical ranges are given in table 2.

The main process for TSV production is *deep reactive ion etching* (DRIE, invented by Bosch [85]). The process employs a directional repetitive sequence of ion etching and wall passivation resulting in anisotropic etching of the silicon bulk. The principle is shown in figure 18. The alternating cycles (typically lasting about 5 s) are: etching under a bias voltage with SF_6 or NF_3 in argon atmosphere to form gaseous SiF_x^+ products—passivation of the surfaces with C_3F_6 , C_4F_8 , or CHF_3 in argon forming a protecting fluorocarbon polymer surface layer on sidewalls and bottom. The bias voltage provides a directional orientation of the bombardment. The resulting via has a typical scalloped surface with undercuts that are, however, substantially reduced by advanced DRIE parameters.

3.5.2. First HEP experiences with TSVs. A characteristic of TSV application for HEP pixel detectors is the fact that read-out chips or sensors are still comparatively thick (100–200 μm) for via fabrication. On the other hand the required via pitch and density often is relaxed, for example when only service lines need to be via-connected or when I/O pads can have large pitches.

Therefore, the first successful HEP pixel module operation employing TSVs [88] was obtained with small aspect ratio ($\sim 1:1$) tapered vias as shown in figure 19, fabricated at Fraunhofer IZM, Berlin. Chip wafers (FE-I3) of the ATLAS pixel detector production were thinned to about 80 μm

Figure 18. Deep reactive ion etching (DRIE) in the Bosch process. Reproduced with permission from [85]. US Grant US5498312A. From top to bottom a sequence of ion etching and passivation is executed (see text).

thickness. The I/O pads were accessed through TSVs and redistributed on the backside of the chip. The via geometry can be seen from figure 19(a): 110/45 μm diameters (top/bottom), $\sim 70^\circ$ tapering angle, 150 μm via pitch. After wafer dicing a pixel module was assembled by bonding a readout chip containing vias to a pixel sensor using solder bonding. The pixel module was operated and characterized with radioactive sources without performance loss compared to modules connected via wire bonds instead of vias.

Figure 19. Tapered via demonstrated on an ATLAS FE-I3 pixel module: (a) cross section detail; (b) SEM photograph. Reproduced from [88]. © IOP Publishing Ltd. All rights reserved.

Figure 20. Straight via applications used for HEP pixel detector prototypes: (a) IZM straight via (here in a MEDIPIX3 chip). Reproduced with permission from [87]; (b) LETI/CEA. Reproduced with permission from [92].

Larger via depths and aspect ratios are required to connect large area chips like the ATLAS FE-I4 chip [6] having about $2 \times 2 \text{ cm}^2$ area. Without handling supports, such chips show large bows and the thickness must be optimized for handling, bonding, and TSV yields. A successful approach was achieved using straight copper vias (Fraunhofer IZM, Berlin) with an aspect ratio of 2.7:1 applied to $160 \mu\text{m}$ thick FE-I4 wafers [89, 90] (figure 20(a)). Similar efforts are carried out at LETI/CEA [91] (figure 20(b)). A $2 \times 2 \text{ cm}^2$ pixel module was produced along similar lines containing about 25 500 pixels, and characterized using radioactive sources. To increase the via yield to close to 100% a new approach using chip wafers thinner than $100 \mu\text{m}$ with handling wafer support using straight Cu-coated vias is under way.

4. Readout integrated circuits, data transmission, and power distribution

The ROIC provides the functionality of a hybrid pixel assembly. From the start, the purpose of the hybrid approach has been to allow independent optimization of ROIC and sensor [93]. The ROIC technology development has been enabled by Moore's law (the doubling of logic density every 2 years in commercial electronics) towards deep submicron commercial CMOS, in order to meet rate and radiation tolerance demands. The alternative of stacking circuit layers in the 3rd dimension without reducing feature size has also been explored, but has so far been less successful. Data transmission and power distribution are system challenges concerning the ROIC as well as other elements such as optical components, encoding/decoding protocols, cables, and cooling.

Table 3. Hit storage memory and output bandwidth for each generation of pixel readout chip.

	1st gen.	2nd gen.	3rd gen.
Hit data storage density	$< 1 \text{ Gb/s/cm}^2$	5 Gbps	40 Gb/s/cm^2
Chip output bandwidth	$40\text{--}160 \text{ Mb s}^{-1}$	$0.3\text{--}1.2 \text{ Gbps}$	$2\text{--}20 \text{ Gb s}^{-1}$

4.1. CMOS ROICs

It is useful to classify CMOS ROICs in terms of generations, which have a rough correlation with the feature size used. First generation pixel chips are typified by those in the original ATLAS and CMS detectors [94, 95] and Medipix [96], using $0.25 \mu\text{m}$ feature size CMOS with custom layout techniques for radiation tolerance. Second generation chips have been developed and fabricated using 130 nm CMOS and are running in current experiments or devices [6, 97, 98]. Third generation chips are under development, focusing on 65 nm CMOS [30, 99, 100]. Table 3 shows the evolution of memory and readout bandwidth between generations.

The basic elements and organization of a hybrid pixel readout chip are common to all generations and applications. The defining characteristic is the use of a dedicated amplification and discrimination channel per pixel (*front end*), with the ability for all these front ends to operate in parallel, followed by parallel processing of the output signals. Section 4.6 reviews the front end characteristics. The output of the front end is most commonly digitized in some way on the pixel, but can also be sampled on an analog memory for later readout as an analog level [95]. On-chip digitization allows for higher

Figure 21. Schematic representation of a conventional CMOS inverter (left) and the triple well process implementation (right).

readout bandwidth with pure digital readout and is expected for all third generation chips. The preferred digitization method has so far been ToT (see section 4.6). We normally refer to a pixel *firing* when the amplified signal exceeds the discriminator threshold. Managing power transients and noise coupling when all pixels are simultaneously *live* (having the ability to fire at any time) is a main system challenge of pixel ROIC design. Many chips are designed for a given maximum occupancy (number of pixels firing at once) and therefore only guarantee system stability in applications that do not exceed this occupancy.

While the front end architecture remains similar from one generation to the next, processing of the information produced by the front end, and the relationship of the front end to the rest of the chip, have evolved significantly. First generation chips were mainly analog circuits, with some logic and memory to manage hit buffering and readout. The most salient new feature of second generation chips was the use of synthesized logic side-by-side with analog front ends in the entire pixel matrix, but with an organization in columns and small digital blocks stepped and repeated. The third generation will go further, implementing an almost entirely digital chip, with a large number of pixels within a synthesized logic basic unit. The front ends will be embedded as data sources within a digital fabric entirely surrounding each analog unit.

Logic density is critical to pixel ROIC design, unlike in most other physics detector applications, where analog performance tends to be of the highest importance. Achievable logic density is intimately tied to rate capability as will be seen later. The 1st generation chips had a mainly analog pixel matrix, with digital operations implemented using full custom circuits (produced with analog design methods). Any digital processing and data buffering were done in the periphery. This resulted in a relatively large periphery area and the need for high bandwidth data transfer from pixels to periphery (see section 4.7). The most prominent new development in 2nd generation chips was the use of synthesized logic within the pixel matrix, enabled by a leap in logic density (see section 4.3). Thus the matrix became a mix of analog columns designed with conventional analog methods and digital columns synthesized with digital design tools. Digital processing and storage in the columns meant the periphery area could shrink and the hit rate capability increased (section 4.7). The 3rd generation chips are now expected to be essentially digital, with embedded analog amplifiers kept to a minimum and logic complexity rivaling commercial microprocessors. This opens the door to unprecedented functionality within the front end chip.

A key element in the transition to essentially digital chips has been isolation of circuits within the same chip. Digital switching produces current spikes. Even if analog and digital circuits use different power supplies and grounds, return currents induced in the silicon substrate can inject parasitic signals into sensitive analog circuits. This can be mitigated by isolating circuits from the bulk substrate, using options provided in many deep submicron processes. Figure 21 compares schematically a logic inverter in conventional CMOS with the triple well version. It can be seen that the NMOS device and the local ground are isolated from the substrate in the triple well process. Some processes offer more sophisticated implant structures that allow further isolation of the N wells, involving separately biased additional wells and/or insulating trenches. Silicon on insulator (SOI) processes have not been popular for high radiation tolerance because they contain a thick oxide layer under the circuitry that accumulates trapped charge with radiation dose, giving rise to undesirable back-gating effects in fully depleted devices (FDSOI). The format of isolation structures exploited varies with vendor and technology. Designs also vary between isolating only analog circuits, only digital circuits, or both from the true bulk substrate. While such isolation strongly attenuates noise coupling, it does not completely eliminate it. Therefore, in addition to isolation, controlling power consumption transients or changes during chip operation is increasingly important. Isolation implants have also been exploited for monolithic active pixel designs, covered in section 6.

Logic density is most important in a triggered system (see section 4.7). This is the case for the highest rate detectors in a collider geometry, where the raw hit volume is too high to allow for read-all operation with an acceptable mass of data cables. In a triggered system the ROIC function is to remember the position, arrival time, and charge of all hits until a trigger arrives. The rate of hits depends on the rate of particles impinging on the detector and, since each hit requires a certain number of bits for digital storage, it can be characterized as a specific bandwidth in bits per second per unit area, as was done in table 3. Note that this does not depend on pixel size to first order, but on particle flux (one incident particle makes one pixel cluster regardless of pixel size, and with the use of region architectures (section 4.7) the effect of cluster size is small). Given a particle flux, the total memory needed depends on the trigger latency, which is typically a few microseconds. The achievable memory per unit area is limited by the CMOS process logic density. Thus, for a triggered system, higher hit rate (and/or longer latency) requires higher logic density. This requirement, not pixel size or radiation,

Figure 22. Conceptual comparison of (left) typical linear transistor layout and (right) an enclosed layout transistor (ELT).

has driven pixel ROICs for high rate and radiation to ever deeper submicron CMOS processes (e.g. 65 nm instead of 130 nm). This is not immediately evident, because pixel size also scales down with increasing hit rate. However, simply making smaller pixels would not require a new process with smaller transistors (smaller feature size), since the analog functionality uses a small number of relatively large transistors. The requirement that does demand smaller transistors is increased logic density. The reason pixel size must also be reduced with increasing hit rate is to avoid in-pixel pileup. After registering a hit, a pixel needs a certain recovery time before being able to record another hit. To avoid losing efficiency to in-pixel pileup, the average time between hits in any given pixel (which scales inversely with pixel area) has to be much longer than the recovery time. An alternative solution is synchronous operation (where every pixel resets before each beam collision). This has been investigated, but so far it has not been popular because, as such a front end must be faster, it consumes more power.

4.2. Increasing circuit density through 3D integration

The use of TSVs in hybrid assemblies has been discussed in section 3.5 in the context of assembly integration. Very small TSVs can also provide an alternate path to increase logic density (in a given footprint) by stacking and vertically interconnecting multiple tiers of circuitry. Stacking of circuit layers this way results in a linear increase in logic density with number of layers, whereas feature size reduction results to first order in a quadratic increase with feature reduction factor. In the commercial sector, industry has so far preferred Moore's law scaling over 3D integration. Only for future system-on-chip architectures, below the 10 nm feature size or equivalent, do industry leaders favor 3D integration [101]. One big advantage of 3D integration is reduction of power. This is because power scales with distance between switching elements, and in 3 dimensions more switching elements can be placed in close proximity than in 2 dimensions. In fact, the main actual and growing use of 3D integration in commodity complex devices is to put memory as close as possible to processors, by integrating a 3D memory stack directly on top of a CPU chip [102]. This increases transfer speed and reduces power. But since high volume introduction of TSV in 2006, the fastest growth of 3D integration has been in DRAM memory and image sensors, where either simple and regular structures (DRAM) or heterogeneous structures (image sensors) are being stacked. The solid state RAM use case is obvious: it

Figure 23. Relative density of radiation hard (RH) and commercial logic in different CMOS technology nodes. In each node, we compared a memory latch from a commercial 7-track standard cell library to the smallest memory latch in the radiation hard library used for LHC pixel chips. The points have no error bars since each corresponds to a single latch.

is not an alternative to Moore's law scaling, but an augmentation of it. The DRAM levels being stacked already have the highest available 2D logic density and 3D integration is the only way to fit even more into a given footprint. In high energy physics, the use of 3D integration for higher logic density ROIC fabrication was explored (with 130 nm feature size CMOS), in parallel to the exploration of plain (not 3D) 65 nm feature size. Working devices were produced and good results were obtained [103, 104], but so far, as in industry, going to a smaller feature size process has been the more effective path to high logic density.

4.3. Radiation tolerance of readout integrated circuits

First generation radiation hard ROICs were only possible in commercial 0.25 μm CMOS processes with custom enclosed layout transistors (ELT). They achieved radiation tolerance above 50 Mrad [105]. In an ELT the gate completely surrounds the source (drain) and in turn the drain (source) surrounds the gate (figure 22). This avoids some channel edge effects (discussed later) since the channel is edgeless, but uses more real estate than a traditional linear transistor. ELT in 0.25 μm CMOS represented a 2.5-fold increase in logic density over 0.8 μm feature size military grade radiation hard technology used in the 1990's, and this enabled sufficient logic density to produce the triggered ROICs for ATLAS and CMS. Early development of pixel ROICs using 0.8 μm radiation hard processes did not succeed in incorporating all the needed functionality in a uniform pixel matrix; the logic density was simply not high enough. Because of the ELT area penalty, radiation hard 0.25 μm lagged behind the Moore's law scaling of their commercial counterparts (figure 23). Starting with the 130 nm node, the minimum size linear transistor standard cell logic used commercially exhibited high total dose radiation tolerance. The ability to use linear transistors translated into a faster than Moore's law increase in radiation hard logic density, which went hand-in-hand with the transition to 2nd generation chips with synthesized logic in the pixel matrix. In

addition to high logic density, the good performance of linear transistors enabled the out-of-the-box use of commercial logic libraries. These have been perfected and extensively validated by large scale commercial manufacturers of consumer electronics and come well integrated with powerful design and simulation tools. However, the use of linear transistors comes with some risks and side effects, which, once again, make radiation hard circuits in 65 nm lag their commercial counterparts in logic density.

The radiation hard design approach of 1st and 2nd generation ROICs was to make or find devices (transistors) that tolerate the required dose without significantly changing their response. The design process typically included increased margin on the specifications, but otherwise relied on the device models available from the manufacturer. After fabrication, the ROIC was tested to measure the actual radiation tolerance. This approach is still followed for analog circuits. Optimization of analog performance would suffer if one had to consider a wide range of properties for each transistor in the circuit, so working with fixed properties is desirable. Fortunately, since for analog circuits the transistor dimensions and geometry are fully under the designer's control, it is possible to choose transistors that suffer little change with radiation, which means larger than minimum dimensions in the case of deep submicron CMOS (section 4.4).

On the other hand, for high density synthesized logic, designers do not have the freedom to select devices that show little change with radiation. Instead, digital design must increasingly rely on accurate modeling of radiation damaged transistors. Conventional digital design already offers several different models for the same transistor: a *typical* model plus so called *corner* models, which simulate operation with different temperatures, voltages, and/or variation of fabrication process parameters within an allowed range. The same transistor will be faster or slower depending on temperature, and so on. The synthesis tools produce circuits to function in all the selected corners. The new design approach is to use additional corner models that represent the radiation damaged transistors. Such models can be custom made by parameterizing measurement data. In this way, logic can be synthesized to work both before and after radiation damage, even though logic gate propagation delays, setup and hold times, etc. can change very significantly: a factor of 2 or more. This is to be compared to changes of just tens of percent between the operating temperature extremes. Such large changes will limit the achievable clock speed, but most pixel ROIC applications require clocks of order 100 MHz, rather than the typical GHz speeds of microprocessors in the same technology. Radiation damaged transistor models are also used to simulate analog and mixed signal circuits in order to confirm the design prior to fabrication and radiation testing, which still provides the ultimate validation.

4.4. Total dose radiation damage in CMOS transistors

Radiation damage in CMOS circuits is entirely due to charge carriers generated by ionization in the dielectric layers of the process, and not to bulk damage of the silicon lattice. Ionizing dose is delivered at hadron colliders by a combination of

Figure 24. Conceptual diagram of a MOS transistor of channel length L and width W , not to scale. A PMOS (NMOS) transistor would have N-type (P-type) substrate and P^+ (N^+) implants, where the $+$ sign indicates high doping. The gate conductor (not shown) would fill the region between the spacers. Spacers are used to control the doping at the ends of the channel. The spacers and the STI (shallow trench isolation) oxides are responsible for radiation damage effects. The STI closing off the front of the figure is not shown.

minimum ionizing particles (mainly pions) and background x-rays. The doping concentrations in CMOS transistors are high (10^{15} cm^{-3} and higher), compared to which the defect density introduced by bulk radiation damage is negligible [106] (below 10^{14} cm^{-3} for HL-LHC inner layers after 3000 fb^{-1}). However, there are many dielectric structures in a modern CMOS process and each one leads to its own radiation effect due to ionizing dose. 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 becomes necessary. The importance of each effect also depends on transistor geometry and size. In this respect, the 130 nm CMOS technology node represented a ‘sweet spot’ for which commercial logic libraries could be used out-of-the-box up to doses well in excess of the 250 Mrad requirement (double or perhaps triple). In contrast, in the 65 nm node it is necessary to select or customize logic cell designs depending on the desired radiation tolerance, effectively trading off radiation tolerance for logic density. If the expected radiation dose is low, the out-of-the-box commercial logic can be used, while for higher expected doses, lower density logic cells must be substituted (The logic density point shown for 65 nm in figure 23 was appropriate for an expected dose of 500 Mrad). These general statements should be seen as rough trends, and must be tempered by process specificity and environmental effects discussed later. Beyond 65 nm, which is at present well characterized, new research in high energy physics instrumentation is focused on understanding radiation damage in 40 nm and 28 nm feature processes [107], with detailed investigations expected over the coming years.

4.4.1. Oxide charge and its effects. Three oxide structures have been identified as producing radiation damage effects in the bulk CMOS processes used so far: the gate oxide, the

shallow trench isolation oxide (STI), and the gate spacers, all shown in fig 24. Ionizing radiation generates charge carriers within these oxides, which can lead to an accumulation of positive static charge, because hole mobility in SiO₂ is 6 orders of magnitude lower than electron mobility. While charge in the gate dielectric would be most problematic, because it would directly cause a threshold shift, gate oxide in 0.25 μm feature size and below cannot charge up. The gate oxides are thin enough (6 nm and below) that quantum mechanical tunneling is an important effect, resulting in an effective electrical conductivity. Tunneling also means that there is a gate leakage DC current (unrelated to radiation), which is what eventually caused industry to move away from silicon dioxide gates in nodes below 45 nm feature size. The gate leakage in the nodes used so far for pixel chips is not yet large enough to result in significant standing power consumption, but it is visible in bias networks, which are high impedance nodes feeding large gate areas.

STI and spacers are both thick oxide structures and therefore will accumulate positive static charge. However, not all the positive charge is truly static. Charge held in so-called deep traps is static at normal temperatures, but charge in shallow traps will gradually drift in an applied electric field (as present during chip operation) and will thus reach the oxide-silicon interface. This is a complex region which undergoes manipulation during the fabrication process [108]. This manipulation involves the passivation of so-called dangling bonds, which would otherwise manifest themselves as states within the energy band gap and spoil the semiconductor properties. Positive charge reaching the interface reverses some of the passivation, bringing back some of the traps that were painstakingly neutralized during fabrication. However, these will then trap negative charge, and that will actually compensate some of the oxide static charge. Remarkably, we would not be able to use commercial CMOS processes for ROICs, were it not for this fortunate cancellation of two detrimental phenomena: oxide charge and interface traps.

At relatively low dose (up to a few Mrad) the net results can be complex depending on the time constants for positive charge drift, the activation of interface traps, and their subsequent trapping of negative charge. These time constants can be vendor-specific, depending on oxide structure formation chemistry and steps, and are certainly affected by temperature. At high total dose the movement of charge and activation of traps have equilibrated, leading to an approximate, but not perfect cancellation of interface charge. A positive net charge always remains, and this degrades transistor performance. Since the sign of this net charge at the oxide-silicon interface is always positive, it affects NMOS and PMOS transistors differently. The net positive interface charge results in parasitic lateral gating in the case of STI and in a modification of source and drain in the case of spacers. The former is known as radiation induced narrow channel effect (RINCE) [109], while the latter is known as radiation induced short channel effect (RISCE) [110].

4.4.2. RINCE. RINCE or lateral gating affects the sides of the channel (see figure 24), causing them to be conductive in the case of an N-channel and non-conductive in the case of a P-channel. This does not have a significant impact on a

wide transistor, because a small fraction of the total current flows near the sides. But in a very narrow transistor all the current is near the sides and the effect can be large, hence it is called narrow channel effect. ELT's were used in 1st generation ROICs to avoid this effect, as the channel in an ELT has no sides (figure 22). This was necessary in 0.25 μm technology, because the radiation-induced parasitic standing current in linear NMOS transistors was measured to be too high for a useful design. But note that, while an ELT has no sides, it also cannot have a very narrow channel. The width of the channel must be at least the perimeter of the source or drain, and so has to be many times the minimum feature size. Therefore, as long as standing parasitic current is tolerable, making wide conventional transistors is just as good as making ELT's, with the advantage that conventional transistors are standard and extremely well modeled. Thus, in 130 nm and below, ELT's have not been used. This led to some surprises at low total dose in the case of the ATLAS IBL detector [111]. The parasitic standing current in NMOS transistors, while completely tolerable at high radiation dose, can go through a transient period of being higher than tolerable depending on vendor and on temperature. Accurate modeling of this varying leakage current behavior has been now developed [112].

RINCE affects PMOS and NMOS transistors in opposite ways. NMOS transistors develop parasitic standing current, but this does not interfere with the transistor action, it simply adds to it. PMOS transistors are parasitically turned off near the sides, which hinders the transistor action. A critical question for PMOS is how far away from the sides is the channel affected by interface charge. The effect can be visualized as a radiation dependent width change, such that the effective width is $W_{\text{eff}} = W_{\text{layout}} - 2\Delta W_{\text{RINCE}}(\text{TID})$, where W_{layout} is the width as drawn, $\Delta W_{\text{RINCE}}(\text{TID})$ is how far away from the STI-channel interface is the channel affected, and TID is total ionizing dose. From the observation that minimum width PMOS devices in 65 nm feature size are highly degraded at 500 Mrad, while those in 130 nm feature size are mildly degraded, we can say that $\Delta W_{\text{RINCE}}(500 \text{ Mrad}) \approx 30 \text{ nm}$.

4.4.3. RISCE. RISCE affects PMOS and NMOS in a similar way, by impeding charge flow between source or drain and channel. It can be roughly modeled as 'adding' a certain length to the channel (a longer transistor conducts less current than a shorter one). One may therefore think that making very short transistors would be a good strategy against RISCE, because this would be a way to compensate. But the opposite is true: longer transistors are less affected by RISCE, because the relative change in effective length is small if the original device is long to begin with. Once again we can write $L_{\text{eff}} = L_{\text{layout}} + 2\Delta L_{\text{RISCE}}(\text{TID})$, where we are now adding effective length rather than subtracting effective width. The magnitude of the effect is about twice as large in PMOS than in NMOS. A 60 nm long PMOS (NMOS) will experience a 70% (30%) reduction in full-on current after 500 Mrad. Since transconductance scales as $1/L$, a 70% (30%) reduction in current is equivalent to a factor of 2.5 (1.4) increase in length. For the original channel length of 60 nm this implies $\Delta L_{\text{RISCE}}(500 \text{ Mrad}) \approx 45 \text{ nm}$ ($\approx 12 \text{ nm}$) for PMOS (NMOS).

Figure 25. Current drive variation in 65 nm technology PMOS devices of minimum length but varying width, as function of (a) total ionizing dose at -15°C and (b) annealing duration and temperature under worst case bias. Reproduced from [113]. CC BY 3.0.

4.4.4. Dependence on voltage and temperature. The above $\Delta W_{\text{RINCE}}(500 \text{ Mrad})$ and $\Delta L_{\text{RISCE}}(500 \text{ Mrad})$ should be seen as rough sketches of a more complex underlying behavior, and were given to provide intuition about the magnitude and sense in which the transistors are affected, and why the effects become more important with decreasing feature size: effective width and length should not be regarded as a real model.

The RINCE and RISCE effects are modulated by transistor bias and by temperature. In general, both effects occur only when transistors are powered, which means there are electric fields in the STI and spacer oxides. The larger the field (which depends on transistor bias conditions) the greater the effect, though there are quantitative differences between NMOS and PMOS. The effective width and length given above are for worst case bias. Unpowered devices suffer little or no damage regardless of temperature—an important point to keep in mind when estimating how long a detector will last.

For powered devices, high temperature increases damage during irradiation, and can decrease (anneal) or increase damage after irradiation. Figure 25 shows results of a study [113] in which different width PMOS devices were first irradiated at low temperature and then annealed at different temperatures, always under power. All transistors are minimum length. When powered, devices can suffer very large damage in a matter of hours at 100°C , but the same damage at 0°C would take over 10 years. This suggests that activation of deep traps is responsible. As an oxide can contain multiple traps with different activation energies, the effect can be complex and process-dependent, which is confirmed by observations. Since the combination of power and high temperature is explicitly excluded by interlocks in all pixel detectors, the main impact of high temperature annealing damage is to complicate qualification. The high integrated dose on a detector is delivered gradually over years. Past qualification protocols emulated this with high dose rate irradiation followed by high temperature annealing. However, application of such protocols is no longer straightforward in 65 nm. Activation energies must be precisely determined before thermal acceleration tests can be confidently applied.

4.5. Single event upsets and mitigation

In addition to long term degradation due to accumulated dose, energy loss by ionizing particles leads to instantaneous soft errors called single event upsets (SEU). 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 ROICs and 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 ROICs 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 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 versus LET is typically fit with a Weibull function, resulting in a threshold and saturation cross section as shown in figure 26. In submicron technologies, the typical LET

Figure 26. Conceptual plot of single event upset (SEU) cross section versus linear energy transfer (LET) for a typical memory cell.

threshold pretty much regardless of memory cell type is of order $1 \text{ MeV} \cdot \text{cm}^2 \text{ mg}^{-1}$. Saturation cross section varies with cell design, but is of order 10^{-7} – 10^{-8} cm^2 for common SRAM, latches and flip-flops. However, an energetic proton (or pion) has a LET of order $0.01 \text{ MeV} \cdot \text{cm}^2 \text{ mg}^{-1}$, 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 [114]. There is thus a kinematic threshold depending on the nuclei in the material, rather than a LET threshold. Typical SEU cross sections for protons are of order 10^{-13} cm^2 [114]. 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.

SEU hardening of memory cells exploits redundant connections or storage separated by some distance. This is more effective for heavy ions, which deposit energy very locally, but still useful for relativistic protons (or pions) which can produce extended deposits [115]. A type of memory called dual interlocked cell (DICE) [116] is extensively used as its state can only be changed by switching two physically separate voltages in coincidence, which gives it a lower SEU cross section than common latches. Special layout techniques, such as *interleaving* can be used to carefully separate sensitive elements and prevent charge sharing [117]. Table 4 compares the SEU cross sections of common structures [118] in a high energy proton beam, for 65 nm technology.

One is not always free to choose the memory type with the lowest SEU cross section. The almost 4 orders of magnitude gain from a standard latch to maximal use of triple redundancy in table 4 comes with an associated footprint increase of a factor of 10. Such cells can therefore not be used in the high density logic of the pixel matrix. Suppose every pixel has 8 bits of configuration and the acceptable fraction of corrupted pixels during an 8 h data run is 1%, with 1 GHz cm^{-2} particle flux. The run-integrated fluence is $2.5 \times 10^{13} \text{ cm}^{-2}$ and therefore the pixel SEU cross section must be $1\%/(2.5 \times 10^{13} \text{ cm}^{-2})$ or less. As there are 8 bits per pixel, the single bit cross section must be $5 \times 10^{-17} \text{ cm}^2$ or less. Even without considering

Table 4. SEU cross sections of different memory structures in 65 nm technology with 24 GeV protons [118]. TR stands for triple redundant and L for latch.

Cell type	Cross section in cm^2
Standard latch	2.8×10^{-14}
DICE latch with interleaved layout	3.1×10^{-15}
TR standard latch	1.2×10^{-16}
TRL with error correction and control triplication	2.3×10^{-17}
TRL with error correction, control triplication and separation of sensitive nodes	6.8×10^{-18}

added design margin, table 4 shows that this requires triple redundancy, which may not fit due to its large footprint. A new solution to this problem in 3rd generation ROICs is to implement a control protocol that allows continuous writing of configuration data during data taking, instead of configuring in a distinct operation mode. If configuration values are constantly being refreshed from outside, one does not need to rely on ‘long term memory’ within the chip. For example if the time between consecutive writes is 1 s instead of the duration of the data run, this relaxes the SEU cross section requirement by 4 orders of magnitude. It should be noted that hit data have never required SEU protection, because the time that data bits spend in chip memory before readout is very short, given by the trigger latency. Thus the relevant time scale for hit data is of order $10 \mu\text{s}$ for triggered, 3rd generation ROICs. Even with a significant error amplification from data encoding⁹, the fraction of data lost due to SEU will be negligible with storage in standard latches¹⁰.

4.6. Analog front end and ADC

The analog front end elements are shown schematically in figure 27. The front end design has universally consisted of a charge sensitive amplifier (CSA), followed by a 2nd stage to provide additional voltage gain, and a comparator (C) to carry out the pulse height discrimination. The CSA has a 1st stage amplifier with capacitive feedback (C_f), a continuous reset (I_{rst}) and a low bandwidth feedback to compensate for sensor DC leakage current (I_{leak}). This is needed because hybrid pixels have so far been DC coupled to the readout and sensor leakage current can be significant after irradiation (for development of AC-coupled pixels see section 3.1).

On-chip signal amplitude digitization has used ToT of the comparator output. ToT is a simple digitization method that counts clock cycles while the comparator is high (meaning above threshold). The most important function of the pixel front end has been to discriminate true hits from noise with the correct timing to within one bunch crossing of the accelerator, which for the LHC means 25 ns. This requires a fast leading edge response (high slew rate), which can be achieved with high input gain. The input charge to output voltage gain

⁹ For example, a single bit flip corrupts an entire 64 bit packet.

¹⁰ Since no detector is 100% efficient, data losses well below the 1% level are insignificant.

Figure 27. Schematic diagram of a typical pixel analog front end. Signal polarities and time over threshold (ToT) are indicated along the bottom.

is given by the inverse of the feedback capacitor C_f in figure 27 and the characteristics of the first stage amplifier. On the other hand, the charge transfer efficiency from the sensor to the integrator is given by $(C_{\text{det}} + G \times C_f) / C_{\text{det}}$, where C_{det} is the detector load capacitance and G is the 1st stage open loop gain. Clearly, a high open loop gain is needed in order to have both high gain (small C_f) and good charge transfer efficiency. However, the specific transconductance (g_m / I_D , where I_D is the drain current) of CMOS transistors is only of order 10 S A^{-1} , which means that a simple (i.e. single transistor) inverting amplifier can only achieve $G \sim 10$ for $1 \mu\text{A}$ current and 1 V supply, since $G \approx V_D g_m / I_D$, where V_D is the supply voltage. Two simple amplifiers in series would achieve an open loop gain proportional to g_m^2 (so $G \sim 100$), but with low bandwidth. The solution in pixel ROICs has always been to use a cascode topology for the 1st stage, which achieves an open loop gain proportional to $g_m(M1) \times g_m(M3)$ without sacrificing bandwidth. The labels M1 and M3 refer to figure 28, where two cascode variants are shown. In the straight cascode configuration (a), the current source M4 is larger than the current source M2, and their combined current passes through the input transistor M1. A small change in the input will cause a change in the M1 drain voltage according to the transistor's gain. This in turn shifts the source of M3, leading to a significant V_{gs} change for M3, which amplifies the change in the output voltage. The M3 gate voltage is held at a fixed bias, V_{casc} . In contrast, in the regulated cascode (b), the current source M4 is smaller than the current source M5, and the additional transistor M2 shifts the M3 gate voltage in response to a change of M1 drain voltage, leading to additional output gain (bigger V_{gs} change for M3).

While the front end configuration and 1st stage cascode architectures have remained, there have been important quantitative changes between 1st and 3rd generation ROICs. The specific transconductance of CMOS transistors has nearly doubled going from 16 S A^{-1} in the $0.25 \mu\text{m}$ process to 22 S A^{-1} in 130 nm to 28 S A^{-1} in 65 nm . The higher hit rate requirements have led to smaller pixels (with lower capacitance per pixel) and faster shaping to avoid in-pixel pileup. While some 1st generation ROICs used PMOS input transistors for better substrate isolation and better $1/f$ noise performance,

with the faster pulse shaping and common substrate isolation techniques, NMOS has become the input device of choice for amplifiers in 130 nm CMOS and below [119]. The decrease in pixel capacitance (by about a factor of 4) and increase in transconductance have translated into higher gain and speed. The effect of time-walk is therefore significantly reduced in 3rd generation ROICs, even before making use of digital processing to correct the time of hits near threshold¹¹. While so far all detectors have used ROICs with continuous reset as shown in figure 27, there has also been R&D on front end designs that reset before every bunch crossing (these are referred to as synchronous front ends) [120]. In this case there is by construction no time-walk, but the power consumption is higher due to the fast shaping and reset settling needed. So far this approach has not been adopted for pixel detectors.

The need for very high speed front ends with sub-ns time stamp resolution capability has been spearheaded by NA62 and its Gigatracker detector. The Gigatracker ROIC [121], implemented in 130 nm CMOS, uses a straight cascode front end with RC feedback. The RC time constant is 5 ns , which gives the device its 5 ns peaking time. Two differential gain stages are then used in front of a constant fraction discriminator to achieve a 200 ps time stamp resolution. Future fast timing detectors discussed in section 2.3 will continue to push this type of design to finer time resolution.

A pixel ROIC with free-running front ends is an inherently metastable circuit. Every pixel can fire at any time, and the act of firing switches logic that would not otherwise switch. Any coupling from digital switching to analog front end can be a positive feedback mechanism. For low enough threshold and/or enough pixels firing at the same time, this positive feedback can set off a chain reaction causing all pixels to fire. For a given instantaneous hit occupancy (fraction of pixels firing at the same time due to an external stimulus), there will be a minimum stable threshold. Each ROIC is different in this respect, but in general the minimum stable threshold was around 2500 e^- in 1st generation ROICs, whereas it will be around 500 e^- for the 3rd generation. This reduction has been deliberate: required by decreasing input signal values. Large pixels ($2 \times 10^4 \mu\text{m}^2$), thick sensors ($>200 \mu\text{m}$), and moderate sensor radiation damage for 1st generation detectors translated into expected signals of order 10 ke^- , while small pixels ($0.25 \times 10^4 \mu\text{m}^2$), thinner sensors ($100 \mu\text{m}$), and heavier sensor radiation damage will lead to signals as low as 2 ke^- at

¹¹ Time-walk is the variation in relative delay between front input pulse and comparator firing, as a function of pulse amplitude above threshold.

Figure 28. Cascode circuit variants with NMOS input transistor, as typically used for 1st stage pixel amplifiers. The bias V_{cas} and the gates of M2 and M4 in (a) and of M4 and M5 in (b) are held at DC voltages, typically generated by bias current mirrors (not shown). (a) Straight cascode. (b) Regulated cascode.

the HL-LHC. The minimum stable threshold is reduced by exactly the same factor, since the important figure of merit for pixel detectors is not signal to noise ratio, but rather signal to threshold ratio. The front end noise does set a lower bound to the threshold, but does not determine how far above this lower bound is the minimum stable threshold. Assuming a noise occupancy of 10^{-4} is acceptable, this corresponds to the Gaussian 1-sided tail fraction for 3.7σ . Inverting the question, for a 500 e^- minimum threshold, the noise must be less than $500/3.7 = 135\text{ e}^-$ equivalent input charge. But this is just an upper bound. How far below this bound the noise needs to be depends on how the threshold varies with time and how from pixel to pixel.

One can express the front end noise requirement as

$$\text{ENC} < \sqrt{(T/3.7)^2 - T_{\text{RMS}}^2(\vec{x}) - T_{\text{RMS}}^2(t)} \quad (14)$$

where ENC is the equivalent input charge noise, T is the threshold, and T_{RMS} is the threshold variation as a function of position ($\vec{x}$) or time (t). $T(\vec{x})$, or the pixel to pixel threshold variation, is known as dispersion. The main cause of dispersion is fabrication process mismatch between transistors and passives of identical design. While the mismatch amount is process dependent, and the translation from mismatch to dispersion is design dependent, a typical dispersion value is in the range of $300\text{--}500\text{ e}^-$. As this is clearly unacceptable, all ROICs so far have compensated for mismatch by programming a different threshold voltage in each pixel, such that in units of input charge all pixels have an equalized (or *tuned*) threshold value. With such tuning, a value of 40 e^- dispersion is typically achieved. The cost of this technique is circuit area in the pixel in order to implement the needed DAC and memory. However, tuning only controls $T_{\text{RMS}}(\vec{x})$, and not $T_{\text{RMS}}(t)$. The main source of $T_{\text{RMS}}(t)$ has been coupling of transients from other circuits to the front end, which has in turn been the main limitation on minimum stable threshold. This has been mitigated with circuit isolation, and control of power transients, as already explained. Still other sources of $T_{\text{RMS}}(t)$ may gain importance in the future. In particular, short term

radiation damage can spoil tuning and cause threshold shifts [122]. If the radiation dose in between tuning cycles is high enough, $T_{\text{RMS}}(t)$ could become dominant for reasons unrelated to circuit isolation. This has not been a problem until now, but will need attention as accelerator intensity increases and windows of opportunity for tuning decrease. Techniques for near real-time equalization, which would compensate both $T_{\text{RMS}}(\vec{x})$ and $T_{\text{RMS}}(t)$ simultaneously, include periodic capture of the baseline level in each pixel, or auto-zeroing [123], and self-adjustment in the pixel to equalize to the same noise occupancy [124].

4.7. Readout architecture

Section 4.8 gives an overview of the hit rate impinging onto the ROIC as well as the bit rate coming out of the data links. Readout architecture refers to the transfer of data from the pixels, where it is generated, to the ROIC output(s), including storage and buffering along the way. There has been much evolution in this area from 1st to 3rd generation. The most significant requirement for the readout architecture is whether the system is full readout or triggered readout. In a full readout ROIC every hit from every pixel must eventually make its way to the data output. In the language of queuing theory this is a single server queue with random customer arrival time (each pixel is a *customer*, while the ROIC data output stage is the single *server*), which is characterized by equation (15), discussed below. A full readout system may not require high logic density, because, provided the data output stage has high enough bandwidth, there is no need to store hit data on chip for any length of time. The problem with a full readout system is that it may not be possible to implement sufficiently high output bandwidth to handle the pixel hit rate. In a triggered system, on the other hand, all pixel hits must wait a pre-defined *latency* until a trigger decision determines which ones to read out. In this case the data output bandwidth only has to match the trigger rate, and is decoupled from the primary pixel hit rate. But the full incoming hit rate must be stored during the trigger latency, and storage is limited by the process logic density.

Figure 29 compares schematically three readout architectures, where the buffered storage of hit data is represented by W for wait, and the motion of hit data by T for transfer. In the classic *column drain* architecture of 1st generation chips (figure 29(a)), all pixels in a column share a data bus that only one pixel at a time can use. As soon as a pixel is hit, it tries to grab the bus and transmit its data. The bus is arbitrated so each hit pixel must wait its turn. This column bus is itself a single server queue, just like the case of full chip readout. The bandwidth of the bus must exceed the incoming hit rate for this to be viable. Equation (15) shows the probability for waiting a time greater than t in such a simple single server queue¹²,

$$P(W > t) = \frac{\lambda}{\mu} e^{-(\mu - \lambda)t} \quad (15)$$

¹² In Kendall notation [125] this is an M/M/1 queue.

Figure 29. Diagram of (a) classic column drain readout with triggered readout, (b) shift register implementation of column drain with full readout, and (c) local buffer triggered readout. W stands for wait and T stands for transfer.

where λ is the input hit rate and μ is the output bandwidth. For a pixel hit rate r , $\lambda = Nr$, where N is the number of pixels served by the bus. Hits must be transferred within a short time, t_s , as there is no place to store them within the pixel array. Therefore, to avoid hit loss, $P(W > t_s) \approx 0.001$, which means that $(\mu - \lambda)t_s \approx 6$. Since t_s is small, one must have $\mu \gg \lambda$: the well known condition that column drain needs a high bandwidth. Implementation of a trigger buffer in 1st generation chips was at the chip bottom, requiring all hits to be transferred out of the pixel matrix. Only after storage in the periphery was the trigger selection applied.

The LHC-b Velopix readout chip implements a read-all architecture using a shift register (SR) for transferring hit data instead of a bus. This is represented by figure 29(b). The SR not only transfers data, but is also temporary storage. The output bandwidth μ of the SR is simply the clocking speed (assuming one pixel hit is read out per clock cycle). But unlike an arbitrated bus, each pixel sees a different output bandwidth. The very first pixel feeding the SR sees the full bandwidth, but the next pixel sees less bandwidth because occasionally it will be blocked by data from the previous pixel already in the SR. Thus, pixel i sees an output bandwidth $\mu_i = \mu - \sum_j r_j$, where r_j is the hit rate of the j th pixel. Because the SR stores hits, the short transfer time requirement of column drain is removed, and one just needs $\mu_i > r_i$ (as opposed to $\gg r_i$). If all pixels have the same hit rate r , one needs $\mu > \lambda$, where $\lambda = Nr$, as before. Since the last pixel sees the smallest output bandwidth, its waiting time is the longest, so one can choose μ such that the last pixel waiting time is less than the time between hits. This architecture is particularly well suited to the highest rate chips in the Velopix detector, because the hit rate is *not* the same in all pixels. The pixel nearest the beam line has the highest rate, and also happens to be the pixel farthest from the readout—that is, the first pixel in the SR. The last pixel in the SR has the lowest hit rate. Thus, the pixel dependent hit rate r_i nicely follows the pixel dependent output bandwidth μ_i .

Starting with the 2nd generation triggered ROICs, high logic density permitted the implementation of local hit storage within the pixel matrix (figure 29(c)). The input rate to the column bus is reduced as $\lambda' = \lambda \times h \times t_{BX}$, where h is the trigger rate and t_{BX} is the bunch crossing period. If the readout

Figure 30. Block diagram of 4-pixel region logic on FE-I4 ROIC. N stands for neighbor, D for front end discriminator output, T for readout token. Reproduced from [18]. © Copyright 2009 IEEE.

time can be long, then one can have $\mu \gtrsim \lambda'$, which makes for a very relaxed column bus transfer bandwidth requirement. In addition to higher hit rates, this relaxed requirement has been exploited to make larger ROICs (more pixels per column bus). Instead of being limited by the transfer rate, the hit rate that can be handled is limited by the amount of memory that can fit in the pixel matrix. All hits must be stored for the entire trigger latency, so the memory needed scales like $\lambda \times t_L$. The longer the trigger latency the smaller the hit rate that can be handled.

The local storage of hit data does not make efficient use of memory when done at the individual pixel level. This is because the mean time between single pixel hits is longer (but not much longer) than the trigger latency. For example, for 50kHz single pixel hit rate and 5 μ s trigger latency, the average number of hits per latency period is 0.25. Using Poisson statistics, one would need to store up to 3 hits per pixel in order to keep 99.9% of hits. We refer to this number (3 in this example) as the *buffer depth*. Instead, grouping pixels into regions with shared storage significantly reduces memory needs. Consider grouping 4 pixels together. If the 4-pixel region average hit rate were 1 hit per latency period, then, using Poisson statistics, one would need a buffer depth of 5 region-hits per latency period in order to keep 99.9% of hits.

Figure 31. Layout detail from RD53 Collaboration illustrating the concept of islands of analog circuitry (blue) embedded in a ‘digital sea’ of synthesized logic (green).

That is only $5/4 = 1.25$ buffer locations per pixel instead of 3 as before. But the advantage is even greater, because in a pixel detector the hits among neighbor pixels are correlated (charged particles produce clusters). The correlation between pixels means that the region hit rate is less than the single pixel hit rate times the number of pixels per region. The amount of correlation depends on sensor and detector geometry. Charge information, on the other hand, must be stored for every hit pixel regardless of how the pixels are grouped, and this reduces the advantage of having many pixels per region. The simplest approach is to store a fixed number of charge bits per pixel times the buffer depth. Suppose in the above example one stores 10 bits of region hit arrival time and some flags, plus 4 bits per pixel charge information. If storage is independent per pixel, each pixel must store 14 bits per hit times buffer depth of 3, or 42 bits per pixel. Conversely, in a 4 pixel region, one must store $10 + 4 \times 4 = 26$ bits per region hit times buffer depth of 5, or 32.5 bits per pixel. The FE-I4 ROIC architecture is based on a 4-pixel region along these lines [18], shown schematically in figure 30.

To take full advantage of larger regions (or to efficiently store higher precision charge information) one needs a more sophisticated storage mechanism that allocates memory only to hit pixels, and does not store zeros for pixels below threshold in the same region. Effectively, hit information must be transferred as soon as it is produced from the pixels to a central storage. This should evoke the column drain architecture of figure 29(a). A large region with arbitrated storage can be labeled as ‘region drain’ and can be analyzed the same way as column drain. The main problem of column drain was the high bandwidth requirement for the column bus, $\mu \gg \lambda$. But given a bus bandwidth μ , one can always make λ small enough by reducing the number of pixels on the bus, N , since $\lambda = Nr$ (r being the single pixel hit rate). In other words, column drain can accommodate high hit rate for a small enough ROIC. Therefore, a large, high hit rate ROIC could be built as a matrix of ‘little chips’, each with column drain architecture. It is interesting that the architecture evolution moved away from column drain because it could not scale to larger size (or to higher hit rate for fixed size), in favor of region architecture, for which size and hit rate are decoupled. But, evolving to large regions eventually leads to the same point as replicating

a column drain architecture many times, as a way to achieve both large total size and high rate.

Third generation ROICs will continue to expand region architectures. While the basic 4 pixel region described above probably represents a sweet spot for efficient memory usage, the use of larger regions is being explored [126]. Already in the 2nd generation an important change took place, in which a ROIC core went from being a collection of pixel circuits, stepped and repeated, to a collection of regions, stepped and repeated. A 3rd generation ROIC will be a collection of digital cores, stepped and repeated. A core is not the same as a region and can in fact contain many regions. A core is simply the smallest stepped and repeated instance of digital circuitry. A relatively large core allows one to take full advantage of digital synthesis tools to implement complex functionality in the pixel matrix, sharing resources among many pixels as needed. Large cores can have 2 dimensional digital connectivity, removing the constraint on all previous ROICs that communication could only take place up and down pixel columns, but not along rows. Figure 31 shows a layout detail from the RD53 Collaboration in which identical 4-pixel analog front end islands can be seen completely surrounded by synthesized logic. The core logic has been dubbed ‘digital sea’ to stress that it results in a different and variable environment surrounding each analog island (depending on where synthesis tools place gates and connections). This is a radical departure from the single pixel step and repeat, perfectly symmetric environment of the 1st generation, raising potential concerns about systematic variations within the pixel matrix. The FE65-P2 prototype [122], which implemented 4 by 64 pixel cores, has shown that with modern isolation techniques (see section 4.6) excellent uniformity can be achieved within a large synthesized core.

4.8. Input hit rates and output data transmission (electrical)

Hit rates and output bandwidth increase by an order of magnitude or more between 2nd and 3rd generation ROICs. For both past and planned detectors, output data are transmitted electrically for at least the first meter away from the ROIC (optical links are covered in section 4.9). The LHC-b experiment is implementing a full triggerless readout system in the upgraded VELO detector using the Velopix IC [12].

Figure 32. Geometry of LHC-b VELO detector. Beams circulate bottom left to top right and vice-versa, colliding near the center. The detector planes (dark grey and red) are perpendicular to the beam, and shown in the open configuration: during collisions they will close onto the beam leaving a small square opening for the beam to pass through. Only interactions boosted forward (towards top right) will be in the detector acceptance. Reproduced from [11]. CC BY 3.0.

They are able to do this because the hit rate, while very high, is not extreme, and the experiment has a fixed target geometry, allowing data cables to be placed outside of the physics acceptance. The ATLAS and CMS experiments, on the other hand, must contend with an order of magnitude higher hit rate per ROIC and have nearly 4π physics acceptance, so most data cables must pass through the active volume. They will therefore use triggered readout. However, the trigger rates will be an order of magnitude higher than today, which combined with busier events will make for a 20- to 40-fold readout bandwidth increase for 3rd generation ROICs. The LHC-b hit rate is projected to be 1.6 GHz cm^{-2} for the most occupied pixel, but as the rate decreases almost quadratically with radius and the modules are perpendicular to the beam, the average hit rate in the most occupied ROIC is 0.3 GHz cm^{-2} . In contrast, for the ATLAS and CMS upgrades all pixels (and therefore all ROICs) in the innermost barrel layer have similar hit rates, projected to be near 3 GHz cm^{-2} for a 3.2 cm layer radius.

Figure 32 shows the geometry of the LHC-b VELO upgrade. The fixed target geometry allows for high bandwidth data transmission with almost no cable mass in the active volume. The Velopix IC supports 4 parallel differential outputs at 5.12 Gbps each, fed by the shift readout architecture described in section 4.7. The data use a fixed frame custom format, where each frame contains a time stamp and a bit map for a group of 8 pixels (no charge information is recorded, just hit/no hit). Reading individual pixels instead would use 30% more bandwidth, because it would not take advantage of the clustered nature of pixel detector hits.

A theoretical discussion of lossless data compression in pixel detector readout can be found in [127], which explains that the information content (entropy) is roughly proportional to the number of clusters and the achievable compression

Figure 33. Conceptual sketch of attenuation versus frequency in an electrical transmission line (solid line). The attenuation is given by the signal amplitude at the output divided by the input signal amplitude (solid line, no specific functional form assumed). The horizontal axis represents the frequency of a pure sine wave signal. The shaded areas represent the frequency content of a DC-balanced transmission if done at low (left) or high (right) bit rate. The dashed lines represent the total attenuation of cable plus equalization, which must be flat across the frequency band of interest to avoid signal distortion.

depends on cluster size. The 30% compression afforded by multi-pixel frames in Velopix is not the maximum possible compression, but is adequate for LHC-b needs. The projected data packet rate for the most occupied chip is 520 MHz , while the Velopix output consisting of 4 links at 5.12 Gbps each will saturate at a 640 MHz packet rate.

The ATLAS and CMS upgrades plan an output bandwidth of 5.12 Gbps per ROIC, which is actually less than 25% of the LHC-b maximum because the planned ROIC size is slightly larger. A critical difference is that ATLAS and CMS must have longer cable runs and at the same time minimize data cable mass, which means tolerating a high transmission line loss (meaning signal attenuation, not data loss). Figure 33 shows conceptually the attenuation of a signal passing through a lossy transmission line. Two vertical gray bands represent low and high bit rate transmission. These bands have a lower cutoff, rather than extending to zero frequency, because DC-balance has been assumed. DC balance avoids low frequencies by guaranteeing at least one transition every n bits (e.g. for the 8b/10b encoding $n = 5$). If the attenuation curve within a gray band is flat, that means the signal shape will not be distorted: a good condition for error-free transmission. This makes it obvious why DC balance improves transmission: the narrower the band, the smaller the variation of the attenuation within the band. However, for the band at high bit rate, even with DC balance there is a large variation of the attenuation within the band. This situation can be improved with equalization, indicated by the horizontal dashed lines. Equalization attenuates lower frequencies more than higher frequencies (the opposite of what cables do) to achieve a flat response within the band. All this is well known textbook material [128], but is only now playing a critical role in pixel detector design. The telecommunications industry and consumer electronics have long ago optimized performance in lossy transmission situations. As a consequence, particle physics designs are moving towards more sophisticated commercial protocols and solutions. State of the art equalization

included in commercial FPGA's can achieve reliable transmission with line losses as high as 28 dB. While prior ROICs (for example FE-I4) have used 8b/10b encoding, the RD53A ROIC will use an open source commercial protocol implementation of 64b/66b encoding, including a multilane version for balancing data over four 1.28 Gbps outputs. Equalization will be an integral part of future systems: RD53A drivers will have pre-emphasis capabilities, which boost high frequencies at the transmitter (sending a purposely distorted signal to counteract the cable distortion) and receivers with different types of equalization are also being developed [129, 130]. Receiver equalization has the advantage that it can be more easily made adaptive, so that it does not have to be manually tuned for every link or every time environmental conditions change. Even with these techniques, because of various constraints including radiation damage, the ATLAS experiment aims to keep line losses at 20 dB or less.

It is clear from figure 33 that there is a limit beyond which the signal is completely lost and no amount of equalization can recover it (recall the 28 dB mentioned above). The only remaining option to further increase the data rate on such a cable is to use multiple logic levels at a given frequency. For example with 4 logic levels instead of two, one can send two bits per clock cycle instead of one. The problem is that the available output voltage must then be divided among the logic levels. The voltage difference between two levels is $V_{\text{out}}/(n-1)$, where n is the number of levels. Noise can cause confusion between levels, so one can use multiple levels as long as $V_{\text{noise}} \ll V_{\text{out}}/(n-1)$. This also applies to standard equalized signals with two logic levels ($n=2$), which means that one cannot go too far to the right in figure 33. But one can go further to the right for $n=2$ than for $n=4$, so whether multi-level transmission is advantageous or not depends on how fast the cable transmission cuts off. What should be clear from this discussion is that using higher voltage at the driver (for a given noise level) will increase the achievable data bandwidth. Commercial IC interface logic operates at a higher voltage than the core transistors, thanks to thicker gate oxide. But such thick oxide devices are much less radiation tolerant and therefore not adequate for pixel detector applications. For 65 nm the maximum operating voltage of core transistors is 1.2 V. Future development could include multi-level transmission using stacked core transistors, where the voltage difference between all levels remains around 1 V instead of $1.2\text{ V}/(n-1)$. A final problem facing high speed drivers in pixel ROICs is that high speed logic is more vulnerable to radiation damage than low speed logic, because switching speed depends on transistor gain. This could be a pixel-specific reason favoring the use of multiple logic levels at lower clock speed over the standard two levels at higher clock speed.

While not strictly an ROIC issue, the quality of available cables is a critical point for data transmission. Commercial data cables are not developed with low mass requirements and contain dielectric materials that are not radiation hard. Therefore, development of custom cables for low mass, radiation hard, high speed transmission is an ongoing activity. This includes flexible printed circuits [131], short run

unshielded twisted pairs, long run shielded twisted pairs, and twin-axial [132].

4.9. Optical and wireless transmission

The technical design report for the original ATLAS pixel detector [133] specified optical links located on the module, less than 1 cm away from the closest ROIC, with optical fibers forming part of the module interconnect. Such an arrangement was never realized and designs have instead moved further away from this model. There are three reasons why optical conversion is located some distance away from pixel modules: radiation, geometry, and reliability. Fibers and active optical elements both suffer radiation damage [134, 135] and are not currently suitable for environments exceeding 100 Mrad. For the laser drivers, but not the fibers, there have been promising developments in silicon photonic systems using Mach-Zehnder interferometers rather than switching lasers, following the trend in commercial high speed optical links [136]. But customization of such devices for particle physics has so far not been widely successful. Geometrically, pixel detectors are compact, with highly constrained cable routing to achieve hermetic coverage. Routing of fibers to all modules in a pixel detector would require complex manipulation of unjacketed fibers, which are fragile and have a minimum bend radius that must not be exceeded. Electrical cables are much more forgiving and can tolerate any bend radius.

4.10. Power distribution

The 1st generation ATLAS and CMS pixel detectors achieved a power distribution efficiency of approximately 25%. This means that 75% of the power delivered by rack power supplies was dissipated in cable IR loss and voltage regulators. Since the only current return path is thorough the detector, the power delivery efficiency is given by V_L/V_{sup} , where V_L (load voltage) is the voltage across the detector elements connected to each single power supply channel, and V_{sup} is output voltage at the power supply. Inefficient power delivery was recognized as a problem, but nevertheless the 1st generation detectors could be built using simple, direct power delivery to each module within a material budget considered acceptable. The planned high luminosity upgrades, on the other hand, face a much greater power delivery challenge. Not only will the detectors be larger, but the deep submicron electronics used will operate at half the voltage yet with slightly more power per unit area. This leads to a total detector current requirement an order of magnitude higher than present detectors.

However, the current that can be delivered from rack power supplies to the ATLAS or CMS inner detectors is limited by external factors. The resistance of the cable plant supplying the inner detector has a lower practical limit, R_{min} , imposed by mass and space considerations. At the same time, the IR-loss heat load on the subdetectors traversed by these cables is limited to P_{lim} so as not to degrade their performance. Therefore, the rack power supplies are only allowed to deliver at most $I_{\text{sup}} = \sqrt{P_{\text{lim}}/R_{\text{min}}}$. Approximate values for ATLAS

Figure 34. Schematic comparison of serial (a) versus parallel (b) powering of a set of 4-chip modules. Each module has 4 ROICs that are always in parallel within the module, while the modules themselves can be placed in series (a) or in parallel (b). The load voltage for the parallel configuration is equal to the ROIC operating voltage V , whereas in the configuration with n modules in series the load voltage seen by the remote power supplies is $n \cdot V$. The voltage drop in the cables is shown above discrete resistors of value R (same R in both cases).

are $P_{\text{lim}} \approx 10$ kW and $R_{\text{min}} \approx 0.5$ m Ω (round trip), yielding $I_{\text{sup}} \approx 5$ kA. The only free parameter left to control the power supplied to the detector is the load voltage V_L . For an expected 3rd generation ROIC power of order 0.6 W cm $^{-2}$, a 10 m 2 pixel detector would need a load voltage of 12 V to deliver the 60 kW total power using 5 kA of rack supply current. Since the ROIC operating voltage $V_{\text{ROIC}} = 1.2$ V, directly powering the ROICs is out of the question: power conversion at the load is required.

From the above discussion, the voltage conversion ratio required for HL-LHC pixel detectors is of order 10 . Two conversion methods have been extensively investigated: DC to DC converters and serial power. The DC–DC conversion technologies investigated include magnetic switching using air core inductors [137, 138]¹³, piezo-electric converters [139], and capacitor charge pumps [140]. While discrete DC–DC converters have now been developed to be compatible with LHC experiments [138], their use on pixel detector modules would result in too much added mass due to the high current per unit area required. Fully on-chip DC–DC conversion has been investigated [141] as this would not add any mass, but has not yet matured enough for consideration in the construction of HL-LHC detectors.

Serial power achieves power conversion without added mass by connecting devices in series and operating with constant current instead of constant voltage power supplies. Figure 34 compares 4-chip modules in serial versus parallel power configurations. The serial configuration only needs two conductors and one power supply channel to serve n modules. Because $V_L = n \cdot V_{\text{ROIC}}$, the power delivery efficiency (V_L/V_{sup}) can be increased as needed by increasing n . Furthermore, constant current supplies are better suited to delivering power over resistive cables than the more common constant voltage supplies. Constant current powering is used commercially with very long, resistive cables, for example in undersea communications lines [142]. However, the requirements of pixel detector operation are very different and a custom serial power

technology solution had to be developed. Serial power has been found to be reliable to implement and compatible with pixel [143–145] (and strip [146]) detector operation, and is the baseline for both ATLAS and CMS HL-LHC pixel detector designs. Implementation requires specialized constant current regulators, which have been developed [147] and included in the FE-I4 ROIC and the RD53A ROIC design. Prior efforts used simpler shunt voltage regulators [94].

4.11. Future development

Even as monolithic solutions mature, hybrid technology with special purpose ROICs will continue to be necessary for the highest rate capability and the implementation of new functionality, such as fast timing. Looking to industry we find that the monolithic CMOS sensors in modern smartphones are actually 3D integrated devices. To the extent that such technology becomes open to low volume third party customization, monolithic devices for particle physics could also consist of multiple layers with processing and readout in small feature size CMOS exactly as conventional hybrid ROICs. In the coming years ROIC R&D will explore smaller feature sizes, with 28 nm being the likely candidate to follow 65 nm for 4th generation ROICs within a ten year horizon [148].

The 4th generation ROICs will add even more sophisticated digital functionality with minimal, but high speed analog front ends to enable fast timing. The operation of the individual pixels will likely be continually monitored and adjusted in real time with no need for storage of configuration data or lengthy calibrations. Improvements in communication and data compression will continue. Single lane output switching frequency will likely not increase beyond 5 Gbps due to issues with radiation tolerance of high speed logic, but data rate could be increased by making use of multi-level logic, essentially following industrial trends to maximize data rates in existing infrastructure originally designed for lower bandwidth. The point in particle physics is not to fit into existing infrastructure but to use the lowest possible mass cables. It is also possible that improvements in data cables themselves will be realized (higher bandwidth for

¹³ Ferromagnetic cores cannot be used in the strong tracker solenoid fields.

Figure 35. ATLAS IBL stave support with 1.7 mm diameter Ti pipe for CO₂ evaporative cooling. Full stave view (a) and cross section (b), showing thermally conductive foam interior and 100 μm wall thickness pipe.

lower mass), for example by replacing copper conductors and/or aluminum shields with carbon nanotube ‘rope’ or graphene foil, respectively. Some use of wireless technology is possible, if nothing else to facilitate testing by being able to communicate during detector construction/integration without physical connections. Finally, on-chip power management will surely increase, dynamically managing core voltage levels and possibly adding on-chip DC–DC conversion with high ratio (4 or more) providing an alternative to serial power.

5. Advanced materials: new possibilities for supports and cooling

The original ATLAS and CMS pixel detectors achieved an amount of material around 3.5% of a radiation length per layer at normal incidence [149, 150]. Of this, 0.5% was due to silicon (sum of ROIC and sensor). Since then, significant advances in mechanical support (section 5.2) and cooling (section 5.1) have taken place, and future ATLAS and CMS detectors are projected to be more than 50% lighter, between 1% and 1.5% per layer, out of which 0.3% will be silicon. Still lower mass is achieved with monolithic instead of hybrid technology, of order 0.3% per layer [151, 152]. The reason is not only less silicon, which can at most save 0.3%, but less stringent cooling requirements for monolithic sensors with small pixels and small depleted volume (and therefore small leakage current and no thermal runaway issues). However, monolithic active sensors compatible with high rate and radiation are still under development (see section 6).

Some critical developments have been CO₂ evaporative cooling, new composite materials, new methods for structure design and fabrication, and serial power distribution. Serial power distribution, which is an electronic system development, has been covered in section 4.10.

5.1. Cooling

The high power density and large extent of pixel detectors at the LHC demands a coolant that can transport a large amount of heat with low mass flow. Since evaporation can remove more heat for the same mass flow as single phase liquid cooling, it is preferred. Gas phase (air) cooling has been used for the STAR heavy flavor tracker [153], but that detector is compact (20 cm active length and 0.2 m² active area) and based on CMOS sensors that do not need to be kept cold. Nevertheless, 0.2 W cm⁻² was removed with 10 m s⁻¹ simple air flow (no fins or

Figure 36. Photograph of a prototype mechanical support built with braided carbon pipes, embedded in carbon foam.

heat transfer enhancement features). An important advantage of evaporation is the absence of a temperature gradient along cooled structures. With an evaporative cooling system, the coolant temperature is fixed, regardless of heat load (for load less than maximum capacity). The heat is absorbed by the liquid to gas phase transition, so the difference between inlet and outlet of a cooled structure is the liquid fraction (high at the inlet and lower at the outlet). The maximum load capacity is reached when all the liquid has evaporated—a condition called dry-out, which should never be reached in a properly working system.

CO₂ is an ideal evaporative coolant for tracking detectors because it offers heat transfer coefficients an order of magnitude higher than traditional refrigerants, and its high evaporation pressure (around 50 bar) means small produced vapor volume, resulting in small diameter tubing. However, pressure safety is one of the challenges of using CO₂, requiring a rating of 200 bar throughout. The use of evaporative CO₂ for tracker cooling was pioneered by the AMS experiment [154], and in colliders by the LHC-b experiment. It was later adopted for the ATLAS IBL upgrade (figure 35) and the CMS pixel upgrade. For a review see [155]. The original ATLAS and CMS detectors used other coolants that were better established at the time and did not need very high pressure plumbing. ATLAS used evaporative C₃F₈ [156], while the original CMS pixel detector used monophasic cooling with C₆F₁₄.

While aluminum has been used for cooling pipes in more conventional (fluorocarbon) evaporative systems [156], it is very difficult to prevent corrosion in aluminum pipes during detector construction [157], and weakening of pipe walls due to corrosion can be fatal for a high pressure system. Therefore, for many present CO₂ cooling systems the pipe material of choice is titanium, which is corrosion resistant and high strength (so

Figure 37. Photograph of a micro channel cooled substrate (right), with diagram of cross section (top left), and photo of sectioned substrate showing micro channels (bottom left). Reprinted from [158], Copyright (2015), with permission from Elsevier.

Figure 38. Simple model of thermal behavior of hybrid pixel devices. The specific thermal resistance is known as thermal figure of merit (TFM).

can withstand high pressure with thin walls), yet relatively low mass. The main reduction in mass from CO₂ cooling comes not from the pipe material, but from the very small diameter pipes that can be used. Other materials have also been explored, most notably carbon fiber, for which braided fiber tubes are a common industrial product and so it is possible to produce braided pipes, but so far this has not been adopted. Figure 36 shows a prototype support with 2 mm diameter braided pipes embedded in carbon foam. For a review of pixel detector cooling and thermal management materials see [3].

A more recent development has been the use of silicon microchannels with CO₂ cooling, in order to achieve very high cooling capacity with even lower mass than possible with separate mechanical supports cooled with metal pipes. This is suitable for fixed target geometries where mass can be placed immediately outside active elements, and has been pursued in the NA62 Gigatracker [158] (figure 37) and LHC-b [159] upgrades.

For hybrid pixels, detector lifetime and cooling performance are closely coupled. The irradiated sensor leakage current scales with temperature (rule of thumb is that current doubles every 7 °C), while the temperature depends on the power dissipated in the sensor, which in turn depends on leakage current. This leads to the well known condition of thermal runaway above a certain temperature [160]. The behavior can be captured by a simple model with the cooling performance of mechanical supports represented by a specific thermal resistance ($\text{cm}^2 \times ^\circ\text{C W}^{-1}$), also referred to as thermal figure of merit (TFM) [161], as shown in figure 38. With an evaporative cooling system, the coolant temperature is fixed, so the sensor temperature rises linearly with power dissipation, and the proportionality constant is the TFM. Thus the TFM defines positive feedback that leads to thermal runaway. High TFM means large positive feedback and therefore early thermal runaway (short lifetime), while zero TFM would be ideal. The mechanical supports of 1st generation ATLAS and

Figure 39. Photographs of carbon composite pixel support prototypes: (a) ‘I-beam’ geometry coupled layer and (b) ‘Alpine’ design with thermally conductive foam pedestals to hold modules out of plane. Reproduced from [168]. CC BY 4.0. (a) I-beam. (b) Alpine.

CMS pixel detectors as built achieved a TFM of $30 \text{ }^\circ\text{C cm}^2 \text{ W}^{-1}$, while support structures prototyped for the HL-LHC upgrade detectors have achieved a TFM as low as $10 \text{ }^\circ\text{C cm}^2 \text{ W}^{-1}$ (can expect $15 \text{ }^\circ\text{C cm}^2 \text{ W}^{-1}$ for as built detectors). This impressive advance has been obtained by utilization of new, more performant carbon composite materials, such as thermally conductive carbon foam [162]. High thermal conductivity, low mass foam allows one to take advantage of the CO₂ cooling by providing a way to efficiently collect heat from a large area module and couple it to a small diameter tube. Commonly used values of foam density and thermal conductivity are 0.2 g cm^{-2} and 40 W/m/K , respectively¹⁴.

5.2. Supports

Support structures must combine the excellent cooling capacity discussed in section 5.1 with low mass and mechanical stability high enough to maintain module positions to within $10 \text{ }\mu\text{m}$. Carbon composite designs and methods have enabled substantial mass reduction. A successful design approach for inner pixel layers has been the use of coupled layer structures. Single layer supports spanning long distances (tens of cm) suffer gravitational and vibrational deformations that elements are too large without auxiliary stiffening, such as shells or frames, and the mass of these elements also contributes to the detector radiation length. By coupling layers in pairs, structures can be made stiff enough over long spans without

¹⁴ Foam thermal conductivity scales approximately with density, which can be adjusted in production.

Figure 40. (a) Conventional monolithic active pixel sensor (MAPS) with charge collection in an epi-layer mainly by diffusion. An *n-well* acts as the charge collection node. Other *n-wells* in the circuitry must be shielded (here by a *p-well*). (b) DMAPS (Depleted MAPS) structure. Multiple wells on high resistive substrate allow complete embedding of the CMOS electronics layer in a charge collecting *deep n-well*. (a) MAPS. (b) DMAPS.

auxiliary support structures, thanks to the large moment of inertia or the coupled layers. The ALICE and STAR experiments used box beam structures for this purpose [152, 163]. In the case of STAR the approximately 30cm long structures were held cantilevered from one end only. A further development in this direction, providing similar mass and stiffness with additional usable space, higher nesting freedom, and monolithic construction uses an I-beam rather than a box beam shape [164] (figure 39(a)). Other large moment of inertia structures use truss-like assemblies [151].

New materials such as foam have also enabled development of supports with more complex geometry that position modules with some tilt angle relative to the colliding beam direction [165]. An example of a support called Alpine [166] using foam pedestals is shown in figure 39(b). Endcap pixel detectors have seen similar evolution, with CO₂ cooling and carbon composite developments allowing for lower mass and greater design freedom [8]. Planned high luminosity upgrades of ATLAS and CMS will have greater acceptance coverage and therefore larger endcap systems [165, 167].

6. CMOS active sensors: towards monolithic pixels

The many advantages of hybrid pixel detector technology have been highlighted throughout this report. However, the hybrid choice also has some notable disadvantages. Hybrid pixels constitute a relatively large material budget, typically more than 1.5% X_0 per detector layer (ATLAS IBL), distributed among the module components and the cooling and support structures. The module production including bump-bonding and flip-chipping is complex and laborious, leading to a large number of production steps. Consequently, hybrid pixel detectors are comparatively expensive.

Modern CMOS imaging sensors instead make use of 3D integration to combine high resistivity and fully depleted charge collection layers with high density CMOS circuitry, in order to achieve high speed and high collection efficiency (for low light operation). Such a combination of fully depleted high resistivity silicon with CMOS readout sounds like a requirement from particle physics, not from consumer electronics, but smartphone image sensors have independently evolved in this direction for different reasons and with

different optimization. For example, the pixels are very small ($<20 \mu\text{m}^2$) and the depleted layer is very thin (few μm). Nevertheless, their out-of-the-box use for radiation and particle detection is being explored [169]. Such high end fabrication processes are currently not available to be customized for particle physics.

The following sections address current developments towards fully depletable, radiation tolerant, and high rate capable monolithic pixel devices, inheriting in part from successful adaptation of CMOS camera type sensors with thick epitaxial layers. The development of fully depletable devices becomes feasible by exploiting advances in the CMOS industry, in particular the availability of multi-well technologies for both low and high resistivity substrate wafers.

6.1. From MAPS to DMAPS

Employing commercial CMOS technologies to produce a monolithic (rather than hybrid) pixel detector in which pixel sensor and electronics circuitry form one entity, has first been proposed [170] and realized [171] in the early 1990s. Some years later, monolithic active pixel sensor (MAPS) detectors were introduced [172, 173], exploiting as the sensing volume an epitaxial layer often grown on top of the lower quality substrate wafer and hosting the CMOS circuitry. The thickness of this epi-layer typically is in the range of 1–20 μm , where thicker layers are often used in processes addressing CMOS camera applications (figure 40(a)).

For particle detection the charge deposited in the epi-layer can be as large as 4000 e^- for a typical thickness of 15 μm . Since the epi-material usually has low resistivity and the allowed biasing voltages are low in CMOS technologies, the epi-layer usually is depleted only very locally around the charge collection node. The deposited charge of a traversing particle therefore is mostly collected by diffusion rather than by drift. This renders the signal generation slow and incomplete (not all charges arrive at the collection node) and is, besides the lower radiation hardness, the main reason why this original MAPS technology is not suited for high rate applications as needed in LHC pp-experiments. Furthermore, other *n-wells*, e.g. those hosting PMOS transistors, act as competing nodes for charge collection. The latter can and must be

Figure 41. (a) Capacitance contributions of a typical CMOS pixel cell to the amplification node. Apart from capacitances to the backside and between pixels a non-negligible contribution comes from the close-by wells, deep n-well and p-well. (b) Through the capacitance between p- and n-well digital signals can couple into the sensor bulk if no particular care is taken to prevent this. (a) DMAPS capacitances. (b) Coupling.

Figure 42. Two principal variants of CMOS cell geometries: (a) Large fill-factor: the charge collecting deep n-well encloses the complete CMOS electronics. (b) Small fill-factor: the charge collection node is placed outside the CMOS electronics area.

cured by additional deep p-well protections as is also shown in figure 40(a).

MAPS pixel detectors have been successfully used in lower than LHC rate and low radiation experiments, such as the STAR experiment at RHIC [174, 175]. Also the ALICE ITS Upgrade [151] has chosen MAPS pixels based on the 180 nm CMOS node offered by TowerJazz with 6 metal layers [176]. Note, however, that the expected radiation level for HL-LHC ion collisions is 700 krad and $10^{13} \text{ n}_{\text{eq}} \text{ cm}^{-2}$, respectively, i.e. 1500 times lower than for LHC pp-experiments.

6.1.1. Depleted (D)MAPS. In order to further develop CMOS pixel detectors for LHC type rate and radiation applications, improved development lines have been pursued leading to *depleted monolithic active pixel sensors* (DMAPS) [26, 177]. The goal of this new development is to employ commercial CMOS technology with some modifications to obtain sufficient signal and fast timing in monolithic CMOS designs, while maintaining charge collection via charges drifting in an electric field inside the chip's substrate. The technology and the sensing properties must survive the radiation environment at the HL-LHC, at least in the outer layers, far enough from the interaction point, such that the radiation levels are similar to those presently encountered at the inner layers, i.e. 100 Mrad ionisation dose and $2 \times 10^{15} \text{ n}_{\text{eq}} \text{ cm}^{-2}$ particle fluence, respectively.

The development of such detectors much relies on recent advancements and freedom in CMOS technologies, offered in particular by vendors interested in market corners away from mass IC production and offering process add-ons or modifications. The goal is to achieve some (50–100 μm) depletion depth

$$d \sim \sqrt{\rho V} \quad (16)$$

where ρ is the substrate resistivity and V is the bias voltage, yielding a reasonably large signal ($\sim 4000 \text{ e}^-$) with fast and in-time efficient charge collection, while avoiding long collection paths on which charges can be trapped after irradiation. At the same time full CMOS functionality shall be maintained, i.e. equal and unconstrained usage of PMOS and NMOS transistors and no or little interference of electronics signals and detector signal pulses.

DMAPS detectors in particular exploit the following CMOS technology features:

- High voltage technology add-ons (from automotive and power management applications) that increase the voltage handling capability and create a depletion layer in a well's pn-junction of depth in the order of 10–15 μm .
- Medium to high ($>100 \text{ } \Omega\text{cm}$) resistivity 8'' silicon substrate wafers, accepted and qualified by the foundry. A depletion layer develops due to the high resistivity with only moderate bias voltages applied from the electronics side or a (specially processed) backside contact.
- Multiple nested wells (see also section 4) that can be used to isolate transistors and shield deep well potentials in order to optimize charge collection. The foundry must accept some process or design rule changes in order to optimise the design for HEP applications.
- Backside processing add-ons allowing for example a backside biasing contact acting as an additional field shaping potential of the device.

Interest in CMOS pixels for HL-LHC has been aroused by its potential for low cost and the feasibility of large area monolithic devices in outer tracker layers when stitching is employed, but also out of intellectual curiosity about whether a one-piece pixel detector can be developed for HL-LHC environments demands.

6.1.2. DMAPS capacitances. Figure 41 shows the main capacitances in a DMAPS design that contribute to the total amplifier input capacitance of a pixel cell. In addition to the pixel-to-pixel (C_{pp}) and pixel-to-backside (C_b) capacitances, also present in any other sensor design, inter-well capacitances to the well sides (C_{SW}) and between deep n-well and p-well (C_{WW}) play a significant role. In particular, if the deep n-well is large (large fill-factor, see below) C_{WW} can achieve significant values (as large as 100 fF for 10000 μm^2 pixel area) increasing the total capacitance. This is to be compared to typical hybrid pixel capacitances of $C_D \approx 120$ fF (planar) and $C_D \approx 180$ fF (3D-Si) [178]. A large amplifier input capacitance C_D directly enters the thermal noise (and also 1/f noise) figures of a pixel detector with CSA readout and shaping as well as the detector response time τ_{CSA} :

$$\text{ENC}_{\text{thermal}}^2 \propto \frac{4 kT C_D^2}{3 g_m \tau_{\text{sh}}} \quad \text{and} \quad \tau_{\text{CSA}} \propto \frac{1}{g_m} \frac{C_D}{C_f} \quad (17)$$

where C_f is the feedback capacitance, g_m the transconductance, τ_{sh} the shaping time, and kT is the Boltzmann constant times temperature.

While it may be possible to cope with additional input capacitance as a noise factor, e.g. by increasing g_m (which increases power), the inter-well capacitance also couples the CMOS electronics with the sensor volume as illustrated in figure 41(b). This requires careful circuit design and special measures to prevent digital activity from coupling to the sensor part, thus faking particle signals. An alternative is to place the digital logic away from the active sensor part (n-well), at the price of a smaller fill-factor (see below).

6.1.3. Fill-factor choices. Figure 42 shows two principal variants of a pixel cell arrangement. While in figure 42(a) the entire CMOS electronics is enclosed in the deep, charge collecting n-well, in figure 42(b) the deposited charge is collected at a small n-well located outside the electronics area. A large fill-factor provides good charge collection properties over the entire pixel area with on average shorter travel distances and hence smaller trapping probabilities after irradiation, but it suffers from a comparatively large inter-well capacitance contribution (~ 100 fF) as discussed above leading to larger noise figures and slower timing. Increased power (compared to hybrid pixels) for the same pixel area is needed to cope with the larger capacitance. Most developments so far have chosen this variant to minimize radiation hardness issues, the main challenge at the HL-LHC. Note that with the advancement of CO₂ cooling a change in the power bill does not linearly translate into material. Thermal conductivity of the components is the key parameter (see also section 5.1). For example, for a cooling tube embedded in a carbon foam,

Figure 43. Effective space charge concentration measured after neutron irradiation for different starting materials varying from 10 k Ωcm to 2 k Ωcm as a function of the neutron fluence. Reproduced from [180]. CC BY 3.0. The insert shows the same curves at low fluences.

increasing (decreasing) the power by a factor of 2 leads to a material increase (decrease) of 10 or 20% [179].

The small fill-factor variant on the other hand promises node capacitances of only 5–20 fF and hence excellent noise and timing performance. However, the radiation tolerance is an issue, given that on average the drift distances of signal charges are longer for same cell size. Small pixel sizes are therefore beneficial for a small fill-factor design at the expense of power density.

6.1.4. Substrate resistivity. According to equation (16) high resistivity is a means to obtain adequate signal sizes by sufficient depletion depth at moderate bias. Figure 43 [180], however, shows that independent of the starting substrate resistivity, over a large range of resistivities (10 Ωcm –2 k Ωcm) one observes that after a fluence of about $10^{15} \text{ n}_{\text{eq}} \text{ cm}^{-2}$ the effective space charge concentrations approximate each other at a level of about $N_{\text{eff}} \approx 10^{14} \text{ cm}^{-3}$ corresponding to about 100 Ωcm . This effect is commonly attributed to radiation induced acceptor removal, setting in earlier at higher p-doping concentrations and eventually leading to similar N_{eff} concentrations. Therefore, for high radiation applications, the starting resistivity is not very critical, and in fact an intermediate (100 Ωcm –1 k Ωcm), not too high initial resistivity may be advantageous as this will lead to a fairly constant charge collection throughout the lifetime.

6.2. Designs and technology variants

The high energy physics community has targeted different prototyping designs with various foundries providing standard CMOS technology add-ons necessary to cope with the given demands. The emphasis at the start was either on dedicated high voltage technologies [26] or on technologies accepting high resistivity substrate wafers for processing [181]. It turned

Figure 44. (a) Pixel layout in 180 nm HV technology (AMS). The triple well technology contains both transistor flavours in a deep n-well acting as the charge collecting node at the same time. (b) DMAPS realisation employing SOI technology. A buried oxide layer (BOX) separates the depleted sensor substrate from the electronics layer. Vertical vias connect electronics to the bulk. Trench isolations shield individual transistors. Partially depleted (thick film) n-well and p-well regions prevent the ‘back-gate’ effect. (a) CMOS w/ HV add-on. (b) CMOS on SOI.

out that both, mid to high resistivity ($\lesssim 1 \text{ k}\Omega\text{cm}$) as well as sufficient bias ($\gtrsim 150 \text{ V}$), are needed for good performance under irradiation. The general approach of different groups has been in three prototyping steps.

1. Simple prototypes to characterize technology features and charge collection performance.
2. Large pixel arrays with stand-alone readout as well as readout via a dedicated pixel readout chip (usually FE-I4 [6]) bonded to it.
3. Large, fully monolithic CMOS pixel matrices including on-chip digital readout architecture for rates and occupancies expected at the HL-LHC.

The development is currently progressing fast. One can say that steps 1 and 2 have been successfully carried out, and different designs for step 3 have been fabricated and are under characterization, but not yet published at the time of writing.

6.2.1. Hybrid active CMOS pixels. A variant under step 2 exploits usage of a ‘smart CMOS sensor’ coupled to a readout chip like FE-I4 [6]. In this case the CMOS sensor provides a first stage of analog pulse processing (usually a CSA preamplifier plus discriminator) leading to an output voltage pulse which is DC or AC coupled to the FE-I4 analog input. This approach is commonly called ‘capacitively coupled pixel detector’, CCPD [182]). While it allows for some extra freedom to explore new functionality such as the implementation of subpixel decoding described in section 2, it is basically an alternative hybridization approach, and so must be compared to traditional hybrid pixels in terms of material budget and power performance.

6.2.2. DMAPS in HV/HR technologies with large fill-factor. CMOS pixels with some depletion depth were first implemented in [183] using a dedicated HV technology (called HVC MOS) with 350 nm and later with 180 nm feature sizes (AMS H35 and H18) allowing for up to three nested wells. Both PMOS and NMOS transistors sit in a large deep n-well

which at the same time acts as the charge collection node (large fill-factor, see figure 44(a)). Because the PMOS transistors’ n-well is identical to the collection node, care must be taken in its usage in the electronics design limiting somewhat the CMOS functionality. The achievable depletion depth is around $15\text{--}30 \mu\text{m}$ for low resistivity substrates around $10 \Omega\text{cm}$. Higher resistivities up to about $1 \text{ k}\Omega\text{cm}$ are under study. HVC MOS prototypes have been tested to TIDs well above 100 Mrad and fluences up to $5 \times 10^{15} \text{ n}_{\text{eq}} \text{ cm}^{-2}$ [184]. The most critical performance metric is the so-called in-time efficiency of hits after irradiation which includes both radiation damage characterizations as well as rate demands at the LHC. The term *in-time* means that a hit must be time-stamped within one LHC bunch-crossing window of 25 ns . Test beam characterization of prototype matrices coupled to the FE-I4 pixel chip via adhesive bonding yielded efficiencies of above 97% (99.7% up to $1 \times 10^{15} \text{ n}_{\text{eq}} \text{ cm}^{-2}$) and in-time efficiencies comparable to those of FE-I4 hybrid pixel modules used in the same test beam ($>95\%$ within three beam bunch crossings of 25 ns) [184].

Prototype matrices were tested up to fluences of $1 \times 10^{15} \text{ n}_{\text{eq}} \text{ cm}^{-2}$ and TIDs of 50 Mrad [27] using an approach that differs from the above mainly in two aspects, (i) use of a quadruple well 150 nm process (LFoundry LF15A) as in figure 40(b), and (ii) use of high resistivity substrate wafers ($2\text{--}3 \text{ k}\Omega\text{cm}$). This technology allows for bias voltages in the $150\text{--}200 \text{ V}$ range. The devices performed well in test beams, showing high time integrated efficiencies, but not yet with fast enough timing (1 bunch crossing in-time efficiency $<91\%$) owing to the large capacitance inherent in the large fill-factor approach. The depletion depth of these devices has been measured for increasing radiation fluence [180] (up to $8 \times 10^{15} \text{ n}_{\text{eq}} \text{ cm}^{-2}$, Figure 45) using edge TCT (*transient current technique*, see for example [185]), which measures the transient current generated by a laser pulse parallel to the CMOS sensor surface, entering the bulk through a diced edge. As a function of bias voltage, the square root shape of equation (16) is observed, clearly showing that depletion depth values of $30\text{--}50 \mu\text{m}$ are maintained even for fluences beyond $10^{15} \text{ n}_{\text{eq}} \text{ cm}^{-2}$.

Figure 45. Depletion depth measured using edge TCT for high resistivity (2–3 kΩcm) CMOS pixel devices. Reproduced from [180]. CC BY 3.0. Plotted is the extracted depletion depth as a function of bias voltage, in seven increasing irradiation steps (up to $8 \times 10^{15} \text{ neq cm}^{-2}$). Open symbols denote devices with and full symbols without backplane contact.

Fully monolithic versions in these technologies (AMS and LFoundry) have been fabricated and are currently being characterized [186, 187].

6.2.3. CMOS on SOI. The SOI technology provides a buried oxide layer (BOX) separating the CMOS electronics from the substrate layer. Both parts are connected by vertical via structures reaching through the BOX and leading to an n-implant which acts as the charge collecting node. Monolithic SOI-based pixel structures have been developed for some time using the fully depleted (FD) SOI technology [188], invented for high speed CMOS electronics with reduced (parasitic) capacitances. The CMOS electronics layer is embedded in depleted silicon. The developments in this FD-SOI technology, however suffer from effects inherent to the BOX oxide layer, most notably the so-called back-gate effect and from radiation effects [188] that can be compensated to some degree [189] for moderate radiation doses, but not for the radiation conditions expected at the HL-LHC.

Thick film SOI, however, featuring trench isolation and doped, only partially depleted regions underneath CMOS transistors, thus shielding the transistors (see figure 44(b)), is free of these difficulties and can also sustain much larger radiation doses. SOI CMOS pixel detectors have been realised [190] and characterised [191] in the lab and in test beams. The devices show impressive TID tolerances tested up to 1 Grad. Also here the substrate resistivity initially increases with hadron fluence due to acceptor removal. The measured in-time efficiency also needs modest improvement in order to cope with LHC demands.

While the CMOS on SOI approach towards monolithic pixel detectors remains very interesting, not the least because of the attractive approach of separating sensing and electronics volume by a buried oxide layer, it is currently not the main focus of the HL-LHC developments.

6.2.4. DMAPS with small fill-factor. A recent development aims to achieve sufficient radiation hardness with a

Figure 46. Small fill-factor design of a pixel cell in a 180nm quadruple well technology (TJ 180nm). Reproduced from [192]. © IOP Publishing Ltd. All rights reserved. Note that the electronics circuitry is fully enclosed in a large deep p-well while charge is collected by a small n-well collection node.

small fill-factor design, promising low noise and fast timing due to the small resulting input capacitance. The development inherits from the design of the ALICE ITS pixel chip (ALPIDE) development [176, 193] introducing some technology improvements for better charge collection after irradiation [194]. Figure 46 shows a principal cell layout. The charge collection is obtained in a high ohmic ($>1 \text{ k}\Omega\text{cm}$) epitaxial layer. The quadruple well features are very similar to those in figure 40(b). The main difference is the layout of the charge collection node placed far outside the electronics area with minimal geometric size. The capacitance is very low: 3–10 fF [195], depending on the geometry and pitch of the collection node. The pixel area must be small ($\lesssim 40 \mu\text{m}$) in order to limit the maximum path length for charge collection.

In a first prototype chip (TJ Investigator) the main features of such a design have been investigated [192] confirming the low capacitance and the good timing precision ($\tau_{\text{rise}} \sim 16 \text{ ns}$). To obtain a fully depleted volume the TJ process was modified replacing the epitaxial p-layer by a planar deep n^+-p junction [192]. A fully monolithic chip including readout architecture has been recently submitted for fabrication with expected noise figures of 16–20 e^- and similar timing precision.

6.3. Outlook on DMAPS pixels

The development of depleted CMOS active pixels currently has a large momentum in the high energy physics community. Many groups are active in developing DMAPS sensors and also DMAPS based modules. The specific developments discussed above are not an exhaustive survey and we are aware of other manufacturers and competing efforts that have not yet published results. Significant new material is therefore expected within 1 year of this review. Whether the development can be brought to a level of maturity sufficient to be adopted for (parts of) an HL-LHC pixel detectors remains to be seen. Regardless of the outcome for HL-LHC, researchers in this field expect

that CMOS active pixels will be the preferred technology for future, large area, high rate and radiation trackers

7. Summary and conclusions

With the planned high luminosity upgrades of the LHC machine and experiments, pixel detectors operating in high rate and radiation environments are facing yet another extreme challenge, surpassing that of the current LHC by roughly an order of magnitude. While integrated (monolithic) pixel detectors have been developed already for some time (and are in fact used in detectors under less hostile experimental conditions), as of today hybrid pixel detectors still constitute the state of the art in tracking detector technology close to the interaction point of such high rate and radiation experiments.

In this paper we have reviewed the fundamentals of track space-time point and direction measurements and described development routes and choices to be made for the next generation HL-LHC experiments. These address all components of the hybrid pixel technology, most prominently the sensor and the readout chip, but also interconnection and 3D-stacking techniques, and lightweight support and cooling structures. Beyond the traditional space point measurement, ways to extract and use additional information from pixel detectors are being explored, such as space vector information or precision timing information with the promise of so-called 4D tracking.

Silicon sensors capable of good performance after fluences of $10^{16} \text{ n}_{\text{eq}} \text{ cm}^{-2}$ and beyond have been developed (sections 3.1 and 3.2). Both 3D-Si sensors and conventional planar sensors have been shown to be able to cope with the demands. 3D sensors after further optimization following their successful performance in the IBL pixel upgrade of ATLAS [196], and planar sensors after reducing thickness and optimizing guard ring structures.

A major challenge is the development of a next generation of pixel readout chips which must cope with very harsh radiation conditions (total doses up to 1 Grad), with particle rates $>3 \text{ GHz cm}^{-2}$, and with MHz trigger rates leading to data output bandwidths of $\mathcal{O}(10 \text{ Gb s}^{-1})$. The development of a close to 10^9 transistor readout chip is jointly addressed by ATLAS and CMS through the RD53 collaboration [30] using deep submicron 65 nm technology. Narrow and short transistor channel effects lead to complex radiation damage behavior that demands sophisticated modelling to be properly addressed (section 4.4). An approach dubbed ‘analog islands in a digital sea’ is employed to process the very large rates on chip. The RD53 chip will serve as the workhorse for the HL-LHC hybrid pixel detectors of the experiments.

Advances in evaporative cooling and in using lightweight yet stiff composite materials render possible a material budget reduction per pixel layer from previously above 3% to below 1%, even for power and cooling intensive high rate and radiation applications (section 5). Coolant choices have narrowed down to CO_2 as today’s state of the art. Very small diameter pipes (micro channels) are an attractive development path for high cooling capacity.

Finally, the monolithic pixel module ansatz is being revived by employing multiple well technologies on high

resistivity substrates that have become more commercially accessible to the HEP community and now offer good performance even under HL-LHC requirements (section 6). This has also launched R&D using CMOS technology lines for (passive) sensors in hybrid pixel detectors (section 3.1.1), offering advantages as low cost, capacitive coupling of electrodes as well as efficient routing between sensor and ROIC pixels by exploiting CMOS metal layers.

Acknowledgments

The authors are grateful to many colleagues for their excellent published work and ideas, and for many discussions and other contributions. This work was supported by the Office of High Energy Physics of the U.S. Department of Energy under contract DE-AC02-05CH11231, by the Deutsche Forschungsgemeinschaft DFG, grant number WE 976/4-1, and by the German Ministry for Research BMBF under grant number 05H15PDCA9.

ORCID iDs

Maurice Garcia-Sciveres <https://orcid.org/0000-0002-5800-4210>

Norbert Wermes <https://orcid.org/0000-0001-9971-0077>

References

- [1] Rossi L, Fischer P, Rohe T and Wermes N 2006 *Pixel Detectors: from Fundamentals to Applications* (Berlin: Springer)
- [2] Moll M 2017 Displacement damage in silicon detectors for High Energy Physics *Proc. of Radecs 2017* (Geneva, Switzerland)
- [3] Viehhauser G 2015 Thermal management and mechanical structures for silicon detector systems *J. Instrum.* **10** P09001
- [4] Aad G *et al* 2008 ATLAS pixel detector electronics and sensors *J. Instrum.* **3** P07007
- [5] Hügging F 2011 The ATLAS pixel insertable B-layer (IBL) *Nucl. Instrum. Methods A* **650** 45–9
- [6] Garcia-Sciveres M *et al* 2011 The FE-I4 pixel readout integrated circuit *Nucl. Instrum. Methods A* **636** 155
- [7] Kastli H *et al* 2007 CMS barrel pixel detector overview *Nucl. Instrum. Methods A* **582** 724
- [8] Dominguez A *et al* 2012 CMS technical design report for the pixel detector upgrade *Technical Report CERN-LHCC-2012-016*, CMS-TDR-11
- [9] Hits D and Starodumov A 2015 The CMS pixel readout chip for the phase 1 upgrade *J. Instrum.* **10** C05029
- [10] Affolder A *et al* 2011 DC–DC converters with reduced mass for trackers at the HL-LHC *J. Instrum.* **6** C11035
- [11] Collins P *et al* 2013 LHCb VELO upgrade technical design report *Technical Report CERN-LHCC-2013-021*, LHCb-TDR-013
- [12] Poikela T *et al* 2015 VeloPix: the pixel ASIC for the LHCb upgrade *J. Instrum.* **10** C01057
- [13] Aglieri Rinella G *et al* 2017 The NA62 gigatracker *Nucl. Instrum. Methods A* **845** 147–9
- [14] ATLAS Collaboration 2015 The optimization of ATLAS track reconstruction in dense environments. ATLAS note: ATL-PHYS-PUB-2015-006, CERN

- [15] Gluckstern R 1963 Uncertainties in track momentum and direction, due to multiple scattering and measurement errors *Nucl. Instrum. Methods* **24** 381
- [16] Olive K *et al* 2014 Review of particle physics (RPP) *Chin. Phys. C* **38** 090001 (An updated RPP is published every 2 years. The online access to the current and previous issues is available from <http://pdg.lbl.gov/>)
- [17] Bartsch V *et al* 2002 Lorentz angle measurements in silicon detectors *Nucl. Instrum. Methods A* **A478** 330–2
- [18] Hemperek T *et al* 2009 Digital architecture of the new ATLAS pixel Chip FE-I4 *Proc., 2009 IEEE Nuclear Science Symp. and Medical Imaging Conf.* pp 791–6
- [19] Kolanoski H and Wermes N 2016 *Teilchendetektoren: Grundlagen und Anwendungen* (Berlin: Springer)
- [20] Fischer P 2015 Comments on the reconstruction of hit positions in segmented detectors Private Communication
- [21] Belau E *et al* 1983 Charge collection in silicon strip detectors *Nucl. Instrum. Methods* **214** 253
- [22] Turchetta R 1993 Spatial resolution of silicon microstrip detectors *Nucl. Instrum. Methods A* **335** 44–58
- [23] ATLAS Collaboration 2016 IBL efficiency and single point resolution in collision events *Technical Report ATL-INDET-PUB-2016-001*, CERN, Geneva
- [24] Charchyan S *et al* 2014 Description and performance of track and primary-vertex reconstruction with the CMS tracker *J. Instrum.* **9** P10009
- [25] Soussan P, Majeed B, Le Boterf P and Bouillon P 2014 Development and evaluation of lead free reflow soldering techniques for the flip chip bonding of large gaas pixel detectors on Si readout chip *Proc. EPTC 2014, the 16th Electronic Packaging and Technology Conf. (Orlando, USA 2014)* pp 453–7
- [26] Peric I *et al* 2014 High-voltage pixel sensors for ATLAS upgrade *Nucl. Instrum. Methods A* **765** 172–6
- [27] Hirono T *et al* 2016 Characterization of fully depleted CMOS active pixel sensors on high resistivity substrates for use in a high radiation environment (arXiv:1612.03154 [physics. ins-det])
- [28] Miucci A *et al* 2014 Radiation-hard Active Pixel Sensors for HL-LHC Detector Upgrades based on HV-CMOS Technology *J. Instrum.* **9** C05064
- [29] Spieler H 2005 *Semiconductor Detector Systems* (Oxford: Oxford University Press)
- [30] RD53 Collaboration 2015 RD53A integrated circuit specifications *Technical Report CERN-RD53-PUB-15-001*
- [31] Sadrozinski H F W *et al* 2013 Ultra-fast silicon detectors *Nucl. Instrum. Methods A* **730** 226–31
- [32] Cartiglia N *et al* 2014 Performance of ultra-fast silicon detectors *J. Instrum.* **9** C02001
- [33] Cartiglia N *et al* 2017 Tracking in 4 dimensions *Nucl. Instrum. Methods A* **845** 47–51
- [34] Buzhan P *et al* 2001 An advanced study of silicon photomultiplier *ICFA Instrum. Bull.* **23** 28–41
- [35] Buzhan P *et al* 2003 Silicon photomultiplier and its possible applications *Nucl. Instrum. Methods A* **504** 48–52
- [36] Stokes T 2005 Avalanche photodiodes: theory and applications, www.photonicsonline.com/doc/avalanche-photodiodes-theory-and-applications-0001
- [37] Riegler W and Lippmann C 2004 The physics of resistive plate chambers *Nucl. Instrum. Methods A* **518** 86–90
- [38] Cartiglia N *et al* 2015 Design optimization of ultra-fast silicon detectors *Nucl. Instrum. Methods A* **796** 141–8
- [39] Cenna F *et al* 2015 Weightfield2: a fast simulator for silicon and diamond solid state detectors *Nucl. Instrum. Methods A* **796** 149–53
- [40] Shockley A 1938 Currents to conductors induced by a moving point charge *J. Appl. Phys.* **9** 635
- [41] Ramo S 1939 Currents induced by electron motion *Proc. I.R.E.* **27** 584–5
- [42] Cartiglia N *et al* 2016 The 4D pixel challenge *J. Instrum.* **11** C12016
- [43] Sadrozinski H-W *et al* 2016 Ultra-fast silicon detectors (UFSD) *Nucl. Instrum. Methods A* **831** 18–23
- [44] Cartiglia N *et al* 2017 Beam test results of a 16 ps timing system based on ultra-fast silicon detectors *Nucl. Instrum. Methods A* **850** 83–8
- [45] RD50 Collaboration CERN 2018 <https://rd50.web.cern.ch/rd50>
- [46] Mannelli M, Pöschl R and Seiden A 2017 Energy and time measurements with high-granular silicon devices AIDA-2020-NOTE-2017-005 (arXiv:1704.01304)
- [47] Arcidiacono R *et al* 2017 Test of UFSD silicon detectors for the TOTEM upgrade project *JINST* **12** P03024
- [48] Lindstrom G *et al* 2001 Radiation hard silicon detectors—developments by the RD48 (ROSE) Collaboration *Nucl. Instrum. Methods A* **466** 308–26
- [49] Harkonen J *et al* 2004 Radiation hardness of Czochralski silicon, float zone silicon and oxygenated float zone silicon studied by low energy protons *Nucl. Instrum. Methods A* **518** 346–8
- [50] Pintilie I, Lindstroem G, Junkes A and Fretwurst E 2009 Radiation-induced point- and cluster-related defects with strong impact on damage properties of silicon detectors *Nucl. Instrum. Methods A* **611** 52–68
- [51] Kenney C, Parker S and Segal J 1997 3d—a proposed new architecture for solid-state radiation detectors *Nucl. Instrum. Methods A* **395** 328–43
- [52] Macchiolo A, Nisius R, Savic N and Terzo S 2016 Development of n-in-p pixel modules for the ATLAS Upgrade at HL-LHC *Nucl. Instrum. Methods A* **831** 111–5
- [53] Macchiolo A 2016 Thin planar silicon sensors private communication
- [54] Terzo S, Andricek L, Macchiolo A, Moser H G, Nisius R, Richter R H and Weigell P 2014 Heavily irradiated N-in-p thin planar pixel sensors with and without active edges *J. Instrum.* **9** C05023
- [55] Terzo S 2015 Development of radiation hard pixel modules employing planar n-in-p silicon sensors with active edges for the ATLAS detector at HL-LHC *PhD Thesis TU Munich* <https://publications.mppmu.mpg.de/2015/MPP-2015-291/FullText.pdf>
- [56] Dragicevic M *et al* 2014 Results from a beam test of silicon strip sensors manufactured by Infineon Technologies AG *Nucl. Instrum. Methods A* **765** 1–6
- [57] Bergauer T, Dragicevic M, König A, Hacker J and Bartl U 2016 First thin AC-coupled silicon strip sensors on 8-inch wafers *Nucl. Instrum. Methods A* **830** 473–9
- [58] Schinella R D and Chao K 1997 Use of reticle stitching to provide design flexibility *US Patent* 5,652,163
- [59] Miller L F 1969 Controlled collapse reflow chip joining *J. Res. Dev.* **13** 239–50
- [60] Pohl D L, Hemperek T, Hügging F, Janssen J, Krüger H, Macchiolo A, Owtscharenko N, Sierra I C, Vigani L and Wermes N 2017 Radiation hard pixel sensors using high-resistive wafers in a 150 nm CMOS processing line *JINST* **12** P06020
- [61] Da Via C *et al* 2013 3D active edge silicon sensors: Device processing, yield and QA for the ATLAS-IBL production *Nucl. Instrum. Methods A* **699** 18–21
- [62] Da Via C *et al* 2005 3D silicon detectors status and applications *Nucl. Instrum. Methods A* **549** 122–5
- [63] Pellegrini G *et al* 2013 3D double sided detector fabrication at IMB-CNM *Nucl. Instrum. Methods A* **A699** 27–30
- [64] Giacomini G *et al* 2013 Development of double-sided full-passing-column 3D sensors at FBK *IEEE Trans. Nucl. Sci.* **60** 2357–66
- [65] Dalla Betta G-F *et al* 2013 3D silicon sensors: irradiation results *PoS Vertex2012* 014

- [66] Pellegrini G *et al* 2008 First double-sided 3D detectors fabricated at CNM-IMB *Nucl. Instrum. Methods A* **592** 38–43
- [67] Dalla Betta G-F *et al* 2016 Development of a new generation of 3D pixel sensors for HL-LHC *Nucl. Instrum. Methods A* **824** 386–7
- [68] La Rosa A 2016 The ATLAS insertable B-Layer: from construction to operation *J. Instrum.* **11** C12036
- [69] Lange J *et al* 2016 3D silicon pixel detectors for the high-luminosity LHC *J. Instrum.* **11** C11024
- [70] Dalla Betta G-F *et al* 2016 Development of new 3D pixel sensors for phase 2 upgrades at LHC <http://inspirehep.net/record/1501528/files/arXiv:1612.00608.pdf>
- [71] Dalla Betta G-F 2017 3D-Si sensor design private communication
- [72] Sultan D *et al* 2016 First production of new thin 3D sensors for HL-LHC at FBK *Proc., 18th Int. Workshop on Radiation Imaging Detectors (Barcelona, Spain, July, 2016)*
- [73] Lange J 2017 3D-Si sensor design private communication
- [74] Ramm P *et al* 2012 *Handbook of Wafer Bonding* (Weinheim: Wiley)
- [75] Garrou P, Bower C and Ramm P (ed) 2012 *Handbook of 3D Integration vol 1&2* (Weinheim: Wiley)
- [76] Garrou P, Koyanagi M and Ramm P (ed) 2014 *Handbook of 3D Integration vol 3* (Weinheim: Wiley)
- [77] Oppermann H and Hutter M 2012 Au/Sn solder *Handbook of Wafer Bonding* ed P Ramm *et al* (Weinheim: Wiley) pp 119–38
- [78] Dietrich L, Töpper M, Fritzsche T, Ehrmann O and Reichl H 2007 Prospects and yield of electrochemical wafer plating for bumping and signal routing *Proc. EMPC 2007, the 16th European Microelectronics and Packaging Conf. (Oulu, Finland)* pp 458–66
- [79] Klein M *et al* 2008 Development and evaluation of lead free reflow soldering techniques for the flip chip bonding of large gaas pixel detectors on si readout chip *Proc. ECTC 2008, the 58th Electronic Components and Technology Conf. (Orlando, USA)* pp 1893–9
- [80] Broennimann C *et al* 2006 Development of an indium bump bond process for silicon pixel detectors at PSI *Nucl. Instrum. Methods A* **565** 303–8
- [81] Alimonti G *et al* 2016 Pixel hybridization technologies for the HL-LHC *J. Instrum.* **11** C12077
- [82] Munding A *et al* 2008 Cu/Sn solid liquid interdiffusion bonding *Wafer Level 3D ICs Process Technology* (New York: Springer) pp 131–70
- [83] Patti R 2008 *3D Integration at Tezzaron Semiconductor Corporation* (New York: Wiley) pp 463–86
- [84] Semiconductor Industry Association 2009 The international technology roadmap for semiconductors, SEMATECH, Austin TX (2010 update)
- [85] Laermer F and Schilp A 1996 Method for anisotropic plasma etching of substrates *US Patent* 5,498,312
- [86] Wolf M 2017 Through Silicon Vias private communication
- [87] Fritzsche T 2017 Straight through silicon vias private communication
- [88] Barbero M *et al* 2012 A via last TSV process applied to ATLAS pixel detector modules: proof of principle demonstration *J. Instrum.* **7** P08008
- [89] Sarajlic M *et al* 2016 Development of edgeless TSV x-ray detectors *J. Instrum.* **11** C02043
- [90] Vermes N *et al* 2016 TSVs in FE-I4 *Presentation at the AIDA-2020 Annual Meeting (Hamburg, Germany)*
- [91] Kühn S *et al* 2016 Compact pixel module with through-silicon vias *IEEE NSS/MIC Conf. (Strasbourg, France)*
- [92] CEA/LETI 2017 Private Communication
- [93] Heijne E and Jarron P 1989 Development of silicon pixel detectors: an introduction *Nucl. Instrum. Methods A* **275** 467–71
- [94] Peric I *et al* 2006 The FEI3 readout chip for the ATLAS pixel detector *Nucl. Instrum. Methods A* **565** 178–87
- [95] Barbero M *et al* 2004 Design and test of the CMS pixel readout chip *Nucl. Instrum. Methods A* **517** 349–59
- [96] Llopert X and Campbell M 2003 First test measurements of a 64k pixel readout chip working in single photon counting mode *Nucl. Instrum. Methods A* **509** 157–63
- [97] Ballabriga R *et al* 2013 The Medipix3RX: a high resolution, zero dead-time pixel detector readout chip allowing spectroscopic imaging *J. Instrum.* **8** C02016
- [98] Poikela T *et al* 2014 Timepix3: a 65 K channel hybrid pixel readout chip with simultaneous ToA/ToT and sparse readout *J. Instrum.* **9** C05013
- [99] Valerio P P *et al* 2014 A prototype hybrid pixel detector ASIC for the CLIC experiment *J. Instrum.* **9** C01012
- [100] Campbell M *et al* 2016 Towards a new generation of pixel detector readout chips *J. Instrum.* **11** C01007
- [101] Or-Bach Z 2015 EETimes: Intel calls for 3D IC www.eetimes.com/author.asp?doc_id=1325921
- [102] Merrit R 2015 EETimes: chip stacks at 30000 feet www.eetimes.com/author.asp?doc_id=1326936
- [103] Deptuch G *et al* 2010 VIPIC IC-design and test aspects of the 3D pixel chip *IEEE Nuclear Science Symp. Conf. Record (Orlando, FL, USA)* pp 1540–3
- [104] Deptuch G *et al* 2016 Fully 3D integrated pixel detectors for x-rays *IEEE Trans. Electron Devices* **63** 205–14
- [105] Nowlin R *et al* 2005 A new total-dose-induced parasitic effect in enclosed-geometry transistors *IEEE Trans. Nucl. Sci.* **52** 2495–502
- [106] Radu R *et al* 2015 Investigation of point and extended defects in electron irradiated silicon: dependence on the particle energy *J. Appl. Phys.* **117** 164503
- [107] Mattiazzo S S *et al* 2017 Total ionizing dose effects on a 28 nm Hi-K metal-gate CMOS technology up to 1 Grad *J. Instrum.* **12** C02003
- [108] Grove A *et al* 1965 MOS C-V model *Solid-State Electron.* **8** 145
- [109] Faccio F and Cervelli G 2005 Radiation induced edge effects in deep submicron CMOS transistors *IEEE Trans. Nucl. Sci.* **52** 2413–20
- [110] Faccio F *et al* 2015 Radiation-induced short channel (RISCE) and narrow channel (RINCE) effects in 65 and 130 nm MOSFETs *IEEE Trans. Nucl. Sci.* **62**
- [111] Dette K 2016 Total ionising dose effects in the FE-I4 front-end chip of the ATLAS pixel IBL detector *J. Instrum.* **11** C11028
- [112] Backhaus M 2017 Parametrization of the radiation induced leakage current increase of NMOS transistors *J. Instrum.* **12** P01011
- [113] Menouni M *et al* 2015 1 Grad total dose evaluation of 65 nm CMOS technology for the HL-LHC upgrades *J. Instrum.* **10** C05009
- [114] Roche P *et al* 2010 A Commercial 65 nm CMOS technology for space applications: heavy ion, proton and gamma test results and modeling *IEEE Trans. Nucl. Sci.* **57** 2079–88
- [115] Menouni M *et al* 2008 Design and measurements of SEU tolerant latches *Electronics for Particle Physics. Proc., Topical Workshop (Naxos, Greece, 15–9 September 2008)* pp 402–5
- [116] Calin T, Nicolaidis M and Velazco R 1996 Upset hardened memory design for submicron CMOS technology *IEEE Trans. Nucl. Sci.* **43** 2874–8
- [117] Amusan O *et al* 2007 Single event upsets in a 130 nm hardened latch design due to charge sharing *45th Annual IEEE Int., Reliability Physics Symp., Proc.*
- [118] Menouni M 2017 Private Communication

- [119] Manghisoni M *et al* 2007 130 and 90 nm CMOS technologies for detector front-end applications *Nucl. Instrum. Methods A* **572** 368–70
- [120] Fahim F 2014 Fermi CMS Pixel (FCP130) test ASIC *Technical Report FERMILAB-CONF-14-435-PPD*
- [121] Martin E *et al* 2010 The 5 ns peaking time transimpedance front end amplifier for the silicon pixel detector in the NA62 Gigatracker *IEEE Nuclear Science Symp. Conf. Record*
- [122] Carney R *et al* 2016 Results of FE65-P2 pixel readout test chip for high luminosity LHC upgrades *Proc. 38th Int. Conf. on High Energy Physics (Chicago, USA)* p 272
- [123] Monteil E *et al* 2016 A prototype of a new generation readout ASIC in 65 nm CMOS for pixel detectors at HL-LHC *J. Instrum.* **11** C12044
- [124] Heim T and Garcia-Sciveres M 2017 Self-adjusting threshold mechanism for pixel detectors to appear in *Nucl. Instrum. Methods A* (<https://doi.org/10.1016/j.nima.2017.06.040>)
- [125] Kendall D 1953 Stochastic processes occurring in the theory of queues and their analysis by the method of the imbedded Markov chain *Ann. Math. Stat.* **24** 338
- [126] Gaioni L, De Canio F, Manghisoni M, Ratti L, Re V and Traversi G 2016 65 nm CMOS analog front-end for pixel detectors at the HL-LHC *J. Instrum.* **11** C02049
- [127] Garcia-Sciveres M and Wang X 2016 Data encoding efficiency in pixel detector readout with charge information *Nucl. Instrum. Methods A* **815** 18–22
- [128] Proakis J and Salehi M 2000 *Digital Communications* (New York: McGraw-Hill)
- [129] Wallängen V and Garcia-Sciveres M 2017 Decision feedback equalization for radiation hard data link at 5 Gbps *J. Instrum.* **12** C01067
- [130] Bates R *et al* 2017 High speed electrical transmission line design and characterization *J. Instrum.* **12** C02002
- [131] McFadden N, Hoeferkamp M and Seidel S 2016 Radiation-tolerant, low-mass, high bandwidth, flexible printed circuit cables for particle physics experiments *Nucl. Instrum. Methods A* **A830** 461–5
- [132] Shahinian J *et al* 2016 High speed data transmission on small gauge cables for the ATLAS Phase-II Pixel detector upgrade *J. Instrum.* **11** C03024
- [133] ATLAS Collaboration 1998 ATLAS pixel detector: technical design report *Technical Report ATLAS-TDR-11*, CERN-LHCC-98-013
- [134] Hall D, Huffman T and Weidberg A 2012 The radiation induced attenuation of optical fibres below -20°C exposed to lifetime HL-LHC doses at a dose rate of 700 Gy(Si)/hr *J. Instrum.* **7** C01047
- [135] Mazza G G *et al* 2013 The GBLD: a radiation tolerant laser driver for high energy physics applications *J. Instrum.* **8** C01033
- [136] Seif El Nasr-Storey S *et al* 2015 Effect of radiation on a Mach-Zehnder interferometer silicon modulator for HL-LHC data transmission applications *IEEE Trans. Nucl. Sci.* **62** 329–35
- [137] Michelis S *et al* 2007 Inductor based switching DC-DC converter for low voltage power distribution in SLHC *Proc. Electronics for Particle Physics*
- [138] Faccio F *et al* 2015 FEAST2: a radiation and magnetic field tolerant point-of-load buck DC/DC converter *IEEE Radiation Effects Data Workshop*
- [139] Lui C *et al* 2006 Design and modeling of the step-down piezo transformer *Proc. Euro. Particle Accelerator Conf.* p 26
- [140] Denes P *et al* 2009 A capacitor charge pump DC-DC converter for physics instrumentation *IEEE Trans. Nucl. Sci.* **56** 1507
- [141] Krieger A *et al* 2014 Design of fully-integrated DC-DC converters for high energy physics applications *Presented at IEEE NSS/MIC (Seattle)*
- [142] Chesnoy J (ed) 2015 *Undersea Fiber Communication Systems* (Amsterdam: Elsevier)
- [143] Stockmanns T *et al* 2003 Serial powering of pixel modules *Nucl. Instrum. Methods A* **511** 174–9
- [144] Ta D *et al* 2006 Serial powering: proof of principle demonstration of a scheme for the operation of a large pixel detector at the LHC *Nucl. Instrum. Methods A* **557** 445–59
- [145] Filimonov V *et al* 2017 A serial powering pixel stave prototype for the ATLAS ITk upgrade *J. Instrum.* **12** C03045
- [146] Matheson J 2011 Progress and advances in serial powering of silicon modules for the atlas tracker upgrade *J. Instrum.* **6** C01019
- [147] Karagounis M *et al* 2009 An integrated Shunt-LDO regulator for serial powered systems *Proceedings of ESSCIRC*
- [148] Zhang C *et al* 2016 GigaRad total ionizing dose and post-irradiation effects on 28 nm bulk MOSFETs *IEEE Nuclear Science Symp.*
- [149] ATLAS Collaboration 2015 Studies of the ATLAS inner detector material using $\sqrt{s} = 13\text{ TeV}$ *pp* collision data *Technical Report ATL-PHYS-PUB-2015-050*, CERN, Geneva
- [150] ATLAS Collaboration 2010 ATLAS insertable B-layer: technical design report *Technical Report ATLAS-TDR-19*, CERN-LHCC-2010-013
- [151] Abelev B *et al* 2014 Technical design report for the Upgrade of the ALICE inner tracking system *J. Phys. G: Nucl. Part. Phys.* **41** 087002
- [152] Schambach J *et al* 2014 The STAR heavy flavor tracker (HFT) *Proc., 20th Int. Conf. on Particles and Nuclei (Hamburg, Germany, 24–29 August 2014)* pp 659–64
- [153] Wieman H *et al* 2009 STAR PIXEL detector mechanical design *J. Instrum.* **4** P05015
- [154] Delil A, Woering A and Verlaet B 2002 Development of a mechanically pumped two-phase CO_2 cooling loop for the AMS-2 tracker experiment *SAE Technical Paper* (Warrendale, PA: Society of Automotive Engineers)
- [155] Verlaet B and Colijn A-P 2009 CO_2 cooling developments for HEP detectors *PoS VERTEX2009* 031
- [156] Attree D *et al* 2008 The evaporative cooling system for the ATLAS inner detector *J. Instrum.* **3** P07003
- [157] Garcia-Sciveres M 2009 Post-installation status of the ATLAS pixel detector *J. Instrum.* **4** P03021
- [158] Romagnoli G *et al* 2015 Silicon micro-fluidic cooling for NA62 GTK pixel detectors *Microelectron. Eng.* **145** 133–7
- [159] de Aguiar Francisco O *et al* 2015 Evaporative CO_2 microchannel cooling for the LHCb VELO pixel upgrade *J. Instrum.* **10** C05014
- [160] Kohriki T *et al* 1996 First observation of thermal runaway in the radiation damaged silicon detector *IEEE Trans. Nucl. Sci.* **43** 1200–2
- [161] Sharma A 2017 The phase-II ATLAS pixel tracker upgrade: layout and mechanics *Proc., IEEE Nuclear Science Symp. and Medical Imaging Conf. (Strasbourg, France, 29 October 2016)*
- [162] Miller W *et al* Allcomp, Inc. 2017 www.allcomp.net/
- [163] Pepato A *et al* 2006 The mechanics and cooling system of the ALICE silicon pixel detector *Nucl. Instrum. Methods A* **565** 6–12
- [164] Hartman N *et al* 2013 Novel fabrication techniques for low-mass composite structures in silicon particle detectors *Nucl. Instrum. Methods A* **732** 103–8

- [165] Smart B 2017 ATLAS pixel detector design for the HL-LHC *J. Instrum.* **12** C02011
- [166] Delebecque P *et al* 2013 Alpine pixel detector layout *Technical Report* ATL-UPGRADE-PUB-2013-009, CERN, Geneva
- [167] Migliore E 2016 CMS pixel detector design for HL-LHC *J. Instrum.* **11** C12061
- [168] ATLAS Phase-II Upgrade Scoping Document 2015 *Technical Report* CERN-LHCC-2015-020. LHCC-G-166, CERN, Geneva
- [169] Whiteson D *et al* 2015 Observing ultra-high energy cosmic rays with smartphones (arXiv:[1410.2895](https://arxiv.org/abs/1410.2895))
- [170] Parker S 1989 A proposed vlsi pixel device for particle detection *Nucl. Instrum. Methods A* **275** 494–516
- [171] Kenney C, Parker S, Peterson V, Snoeys W, Plummer J and Aw C H 1994 A prototype monolithic pixel detector *Nucl. Instrum. Methods A* **342** 59–77
- [172] Dierickx B, Meynants G and Scheffer D 1998 Near 100% fill factor CMOS active pixel In *Proc. SPIE* **3410** 68–76
- [173] Turchetta R *et al* 2001 A monolithic active pixel sensor for charged particle tracking and imaging using standard VLSI CMOS technology *Nucl. Instrum. Methods A* **458** 677–89
- [174] Contin G *et al* 2016 The MAPS-based vertex detector for the STAR experiment: lessons learned and performance *Nucl. Instrum. Methods A* **831** 7–11
- [175] Schambach J *et al* 2015 A MAPS based micro-vertex detector for the STAR experiment *Phys. Proc.* **66** 514–9
- [176] Mager M 2016 ALPIDE, the monolithic active pixel sensor for the ALICE ITS upgrade *Nucl. Instrum. Methods A* **824** 434–8
- [177] Vermes N 2016 Depleted CMOS pixels for LHC proton–proton experiments *Nucl. Instrum. Methods A* **A824** 483–6
- [178] Havranek M, Hügging F, Krüger H and Vermes N 2013 Measurement of pixel sensor capacitances with sub-femtofarad precision *Nucl. Instrum. Methods A* **714** 83–9
- [179] Andersen E 2016 Cooling of pixel support structures Private Communication
- [180] Mandic I *et al* 2017 Neutron irradiation test of depleted CMOS pixel detector prototypes *J. Instrum.* **12** P02021
- [181] Havranek M *et al* 2015 DMAPS: a fully depleted monolithic active pixel sensor-analog performance characterization *J. Instrum.* **10** P02013
- [182] Peric I, Kreidl C and Fischer P 2010 Hybrid pixel detector based on capacitive chip to chip signal-transmission *Nucl. Instrum. Methods A* **617** 576–81
- [183] Peric I 2007 A novel monolithic pixelated particle detector implemented in high-voltage CMOS technology *Nucl. Instrum. Methods A* **582** 876–85
- [184] Benoit M *et al* 2016 Testbeam results of irradiated AMS H18 HV-CMOS pixel sensor prototypes *JINST* **13** P02011
- [185] Kramberger G *et al* 2009 Investigation of electric field and charge multiplication in silicon detectors by edge-TCT *Proc., 2009 IEEE Nuclear Science Symp. and Medical Imaging Conf. (Orlando, Florida, 25–31 October 2009)* pp 1740–8
- [186] Wang T *et al* 2017 Development of a depleted monolithic CMOS sensor in a 150nm CMOS technology for the ATLAS inner tracker Upgrade *J. Instrum.* **12** C01039
- [187] Peric I *et al* 2015 Overview of HVCMOS pixel sensors *J. Instrum.* **10** C05021
- [188] Arai Y *et al* 2011 Development of SOI pixel process technology *Nucl. Instrum. Methods A* **636** S31–6
- [189] Miyoshi T *et al* 2014 SOI monolithic pixel detector *J. Instrum.* **9** C05044
- [190] Hemperek T, Kishishita T, Krüger H and Vermes N 2015 A monolithic active pixel sensor for ionizing radiation using a 180 nm HV-SOI process *Nucl. Instrum. Methods A* **796** 8–12
- [191] Fernandez-Perez S *et al* 2015 Radiation hardness of a 180 nm SOI monolithic active pixel sensor *Nucl. Instrum. Methods A* **796** 13–8
- [192] Pernegger H *et al* 2017 First tests of a novel radiation hard CMOS sensor process for depleted monolithic active pixel sensors *J. Instrum. JINST* **12** P06008
- [193] Aglieri G *et al* 2013 Monolithic active pixel sensor development for the upgrade of the ALICE inner tracking system *J. Instrum.* **8** C12041
- [194] Gao C *et al* 2016 A novel source-drain follower for monolithic active pixel sensors *Nucl. Instrum. Methods A* **831** 147–55
- [195] van Hoorne J 2015 Study and development of a novel silicon pixel detector for the Upgrade of the ALICE inner tracking system *PhD Thesis* CERN CERN-THESIS-2015–255 p 61
- [196] Darbo G 2015 Experience on 3D silicon sensors for ATLAS IBL *J. Instrum.* **10** C05001

Maurice Garcia-Sciveres is Senior Scientist in the Physics Division of the Lawrence Berkeley National Laboratory (LBNL), USA, and a fellow of the American Physical Society. He obtained his PhD in 1994 with a measurement of the weak mixing angle, V_{cb} , in the CLEO experiment at Cornell University. He joined LBNL after graduate school, where he focused on the development of vertex detectors at hadron colliders. The devices he worked on enabled studies of the top quark and discoveries including Bs meson oscillations (CDF experiment) and the Higgs boson (ATLAS experiment). He had leading roles on the development of various integrated circuits (IC) for both strip and pixel silicon

detector readout, and is currently co-spokesperson of the CERN RD53 collaboration, charged with producing a pixel readout IC for the ATLAS and CMS high luminosity upgrades.

Norbert Vermes is Full Professor of Physics at the University of Bonn, Germany. Since completing his PhD in 1982 he has been researcher at Stanford University (SLAC) and CERN (Geneva) before becoming a Professor at the University of Heidelberg in 1989. In 1992 he accepted a chair position at the University of Bonn where he carries out research and teaches since then. His and his group's research focuses on the physics of elementary particles, especially Higgs physics and heavy quarks, and on detector development, especially semiconductor pixel detectors and ASIC readout electronics.