New Challenges for Signal Processing for HEP experiments

Geoff Hall

Imperial College London

Introduction & objectives

- First, to identify the challenges, what is the state-of-the-art?
  - then
- How and when did we (HEP) get to where we are now?
  - try to identify some factors making developments successful
  - try to identify trends which are important for the future
- Qualification
  - HEP is a big field
    - I've focused on LHC experiments – which are responsible for much innovation
    - There is an interesting discussion to be had about how things evolve, whether big projects are the right model, how innovation occurs, the role of investment, risk, imagination,…
  - I've chosen projects I am most familiar with for illustration
    - I think they are representative
    - but don't assume they are the best, or only, ones

What are the challenges we face?

- To attempt to answer this (very difficult) question
  - look back, as well as forward
- then focus on a few main areas:
  - silicon sensors
  - ASIC electronics
  - optical links
  - FPGAs, for digital processing
- Why these choices?
  - areas which I know reasonably well
  - which have developed considerably – from different starting points – in the last ~20 years
  - areas which are probably among the most important for particle physics, now and in the future
  - but don't be offended if your favourite is missing – let's discuss later

Silicon sensors

- Originated – for HEP – in the late 1970s
  - in response to the discovery of charm quarks
    - need to acquire higher statistics and precision => electronic detectors capable of measuring short distance decays into multi-particle states
    - alternatives: emulsions, rapid cycling bubble chambers
  - work done for other reasons (nuclear physics, ν beam) offered a solution
    - but not overnight!
  - led to highly segmented, precise, (relatively) low mass, ~2D measurements

- Something which started small, then grew substantially
  - many modest size experiments using these detectors
  - serendipitously profited from growth of microelectronics industry
  - electronics challenges stimulated complementary activity
    - also with large consequences

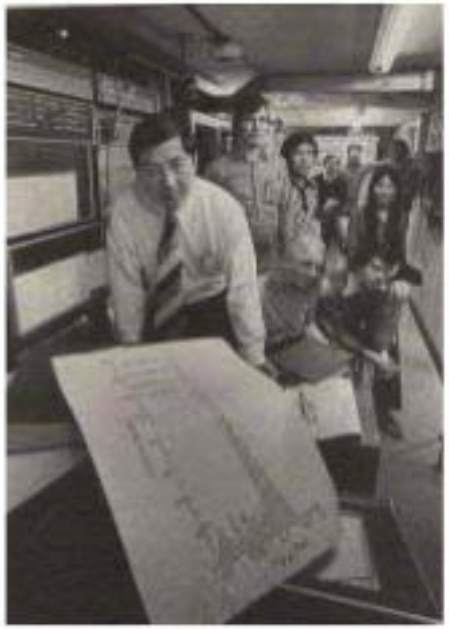

Physics Motivation
November Revolution

J/ψ

11 November 1974

- *from Chris Damerell talk (Snowmass 2001)*

- Gaillard, Lee and Rosner RMP **47** (1975) 277 *Search For Charm*
  'The tracks of charmed particles will be too short to see in bubble chambers, but should definitely be of the order of tens or hundreds of microns: easily detectable in emulsion'.
- Charpak, EPS Conference in Palermo June 1975
  'Drift chambers are the easiest to build, most accurate, cheapest and most convenient detector for localising particles. Whoever is familiar with their operation would be strongly reluctant to use other devices in the planning of a new experiment'.
- ACCMOR collaboration in CERN struggled to see charm hadroproduction (single *e* trigger)
- Succeeded over next 10 years to develop silicon microstrip and pixel detectors (CCDs) as powerful tools for charm physics.

How was this possible?

- Relatively large development budgets were allocated
  - justified by scientific importance and opportunities for experiments
- A number of small scale manufacturers identified
  - Micron, Canberra, SI, Hamamatsu,… & some in-house manufacturing
    - TUM/J. Kemmer et al particularly influential
- Related applications identified
  - photosensors, CCDs, scintillator readout, non-HEP applications
- Large number of experimental studies, as much physics to cover
  - different production, and decays, lifetimes, cross-sections, energy
- New challenges soon identified
  - particularly high-density readout, signal processing, noise,…
  - electronics was an expert niche (still is?)

A fertile and dynamic period
Historical rigour?

What happened next?

- µstrips predominated
  - telescopes vs active targets, CCDs established
- technologies (and companies) became established
  - product portfolios, adaptable and re-usable (profit essential!)
- innovations from HEP community
  - double-sided sensors, Si drift detectors, optimised photodiodes, DEPFETs, biasing schemes (FOXFET, p-stops,..), pixels, CCD development, 3D, ..
  - ASICs!
    - essential to overcome space and power issues, esp. at colliders
- We also learned about
  - quality and reliability, manufacturing, technical challenges (esp. noise and cooling), costs, why 2D is harder than 1D,…
  - radiation damage
  - signal processing fundamentals

A decade later… SSC & LHC

- Radiation damage was the **biggest** issue
  - speed and data volumes also significant – as well as system sizes
- From ~1989, gradual ramp-up of dedicated studies
  - intensive R&D programmes set up (US & Europe)
    - significant funding, many eager participants, great optimism despite challenges
- The outcome – eventual success with silicon
  - large body of data and improved understanding
    - although radiation damage is a very complex subject, still today
    - careful studies of details: bulk & surfac edamage, $C_{interstrip}$, ageing, CCE, …
  - "conservative" µstrip sensor designs based on incremental improvements
    - V_{dep} : ~50V -> ~500V, single-sided, thickness, finish: manufacturer's input vital
    - gradual increase in wafer size (2-3-4-6 inch)
  - but, in parallel, success with innovative sensors, like APDs, drift devices, and later SiPMs

What have we learned from this?

- Importance of collaboration – in HEP, and with industry
  - are more “standard” devices needed, or possible?
- Large volume vendorS needed for large-scale projects
  - specialised HEP needs are subject to commercial interest, or lack of it
- Long term commitments, and specialised infrastructure
  - how to maintain it?
- Innovations are possible, but time is needed
  - On the horizon: greater use of electronic sensors

ASICs

- Original drive for collider opportunities and constraints
  - LEP & SLC mainly, later Tevatron
- and willingness from experimenters to learn new technologies
  - long tradition of exploiting electronics for HEP

- Evolution over ~20 years
  - hand-crafted designs by amateurs or novices =>
  - professional, expert design teams using (close to) standard commercial state-of-the-art processes
- Increasingly complex, large scale chips
  - longer (?) development times, larger budgets needed
    - but lower risk, and methods for cost-sharing
    - just don't take this for granted

History to ~1990

- First HEP ASIC ~1984
  - Microplex @ Mark II Walker, Parker, Hyams et al
  - Notable features: NMOS, 5µm, 14mW/channel
- CMOS successors within 4-5 (!!) years for LEP
  - ~3µm feature size, 1-2mW/channel, τ ~ µsec

	Mark II	**Delphi**	**Aleph**	**Opal**	**CDF**
Chip name	Microplex	MX3	CAMEX64A	MX5	SVXD
Ref	Walker 1984	Stanton 1989 Bingefors 1993b	Buttler 1988	Allport 1993	Kleinfelder 1988
Technology	5µm NMOS	3µm CMOS	3.5µm CMOS	1.5µm CMOS	3µm CMOS
No. channels	128	128	64	128	128
Area [mm^2]	6.3 x 5.4	6.4 x 6.9	6.4 x 5.0	6.3 x 6.7	6.3 x 6.8
Bond pitch [µm]	47.5	44	100	44	48
ENC [e] (C_{in} in pF)	$280+97C_{in}$	$670+55C_{in}$	$335+35C_{in}$	$325+23C_{in}$	$350+58C_{in}$
Power/channel [mW]	~14	~0.5	~1	~2	~1.3

LHC electronics… from 1990

- We did not know how to build detectors for such a hostile environment
  - expected radiation levels were unprecedented ~10Mrad in tracker
  - customised electronics was essential
  - approach used for space systems completely impossible
    - shielding plus qualification of small numbers of COTS components
- Several space-qualified processes, mostly of military origin
  - SoS, SoI, bipolar, GaAs, hardened CMOS,…
  - at least as many companies…
- Technology and other issues
  - highly regulated by governments
  - expensive
  - hard to design in, especially analogue
  - slow turn-round for small, inexperienced customers
  - relatively antiquated: ~1.2 µm feature size

By 1997…

- …the number of available processes had diminished greatly
  - evolution of world political situation
  - modest economic value of LHC business
  - technical performance challenges
    - - noise and power
  - sufficient radiation hardness not guaranteed

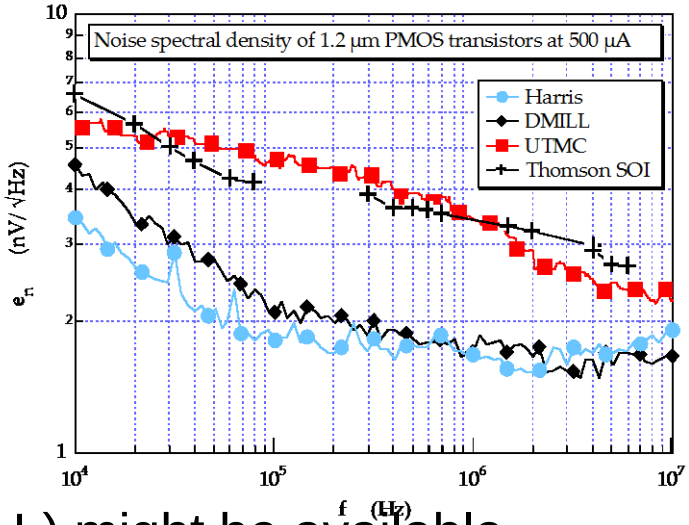

- It seemed likely that only one process (DMILL) might be available
  - custom process developed for HEP
  - not fully qualified for yield, cost and radiation behaviour
- Mostly unexpected news in September 1997 [A. Marchioro et al CERN]

0.25µm CMOS 1997

- Standard commercial processes shown to be unexpectedly promising
  - previously many inconsistent results
  - >>10Mrad achievable with careful transistor design
  - not fully proven that low noise analogue design was possible
    - …but looked worth exploring…

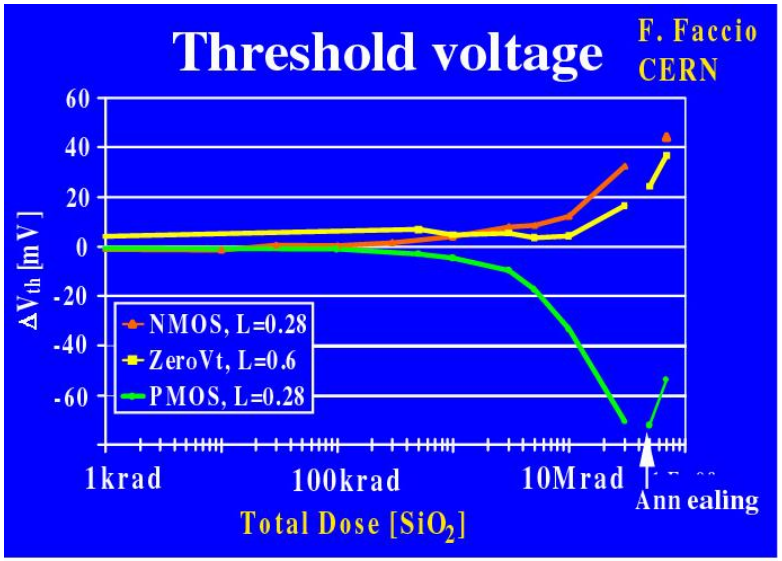

F. Faccio CERN

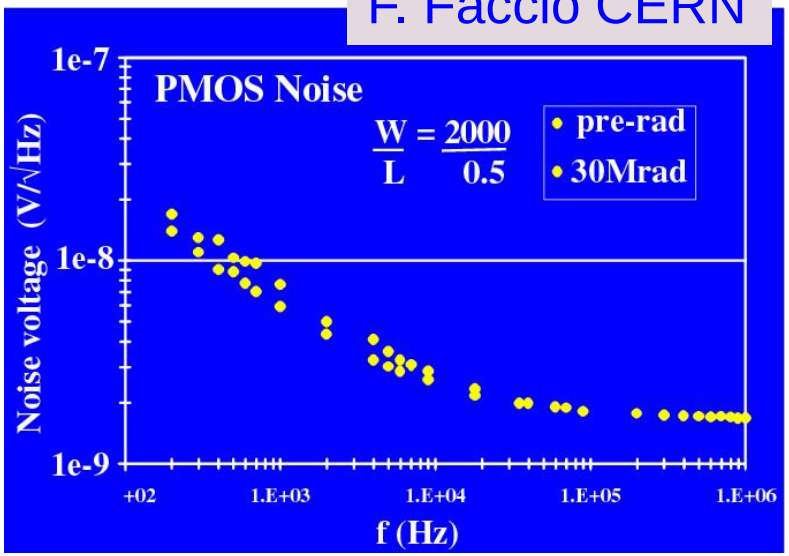

Principles of rad-tol design

(1) Gate oxide scaling

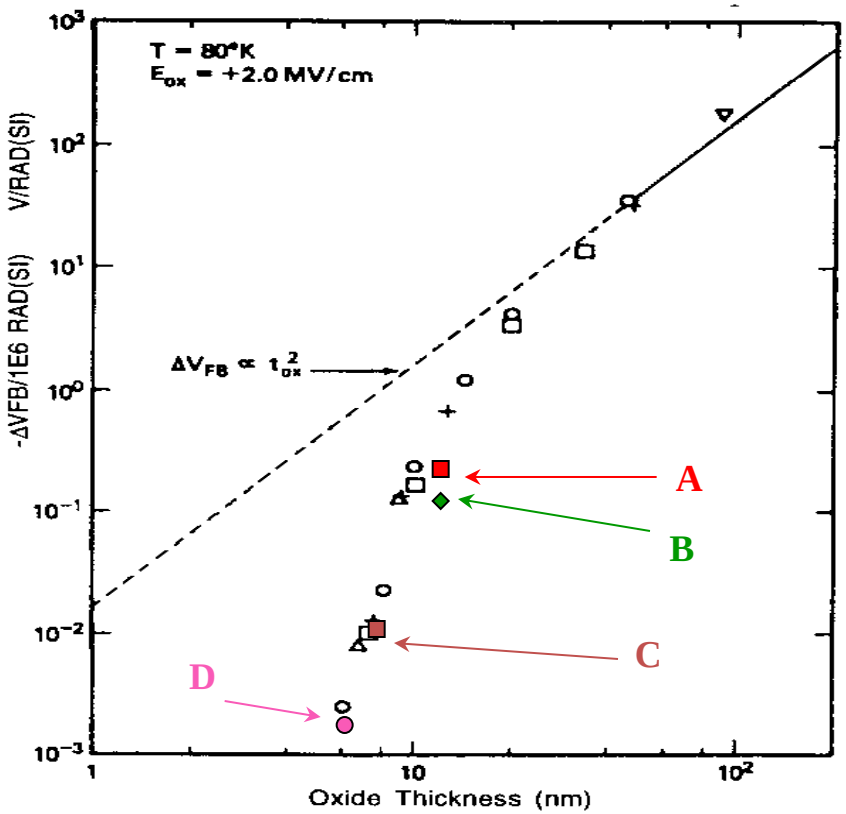

After N.S. Sacks, M.G. Ancona, and J.A. Modolo, IEEE Trans.Nucl.Sci., Vol.NS-31 (1984) 1249

- Electron tunneling neutralizes trapped holes in thin oxides.
- Total dose effects, such as V_t shift, are naturally reduced in deep submicron processes.

M. Letheren CERN

Principles of rad-tol design

(2) Thin gate-oxides + Gate all-around layout

- Min-size NMOS layout
  - Edge-less structure eliminates leakage via parasitic edge transistor.
  - Guard ring eliminates leakage between devices and provides latch-up protection.
  - Higher capacitance of gate all-around structure improves SEU tolerance.
    Further SEU tolerance by circuit design (SEU-tolerant flip-flops) or system design (triple-redundant logic, error detection and correction coding etc.)

M. Letheren CERN

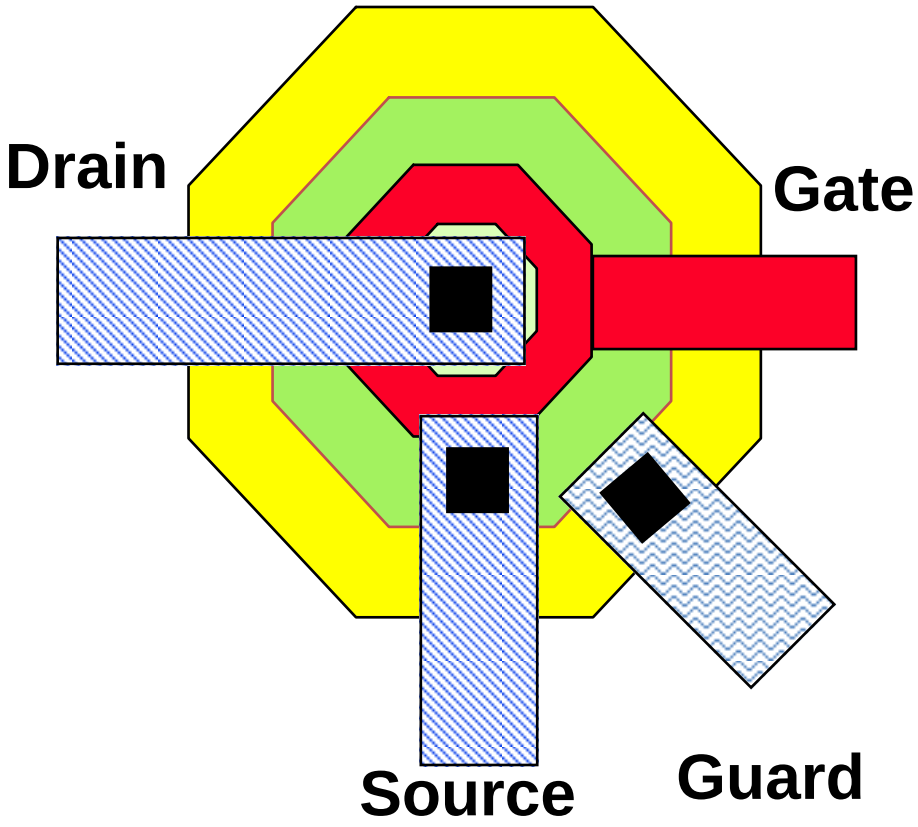

What next for ASICs?

- ASICs are mainly for FE readout – includes some ADCs and digital pipelines, plus data transmission (optical) links
  - elsewhere, avoid customisation to gain flexibility
- More advanced technologies?
  - 250 nm for LHC first generation. 130 nm established, 65 nm in use
    - commercial evolution is constant – but not easy to ride the technology wave
  - Does HEP need the finest feature sizes?
    - Small size components save space and allow more digital processing (intelligence)
    - but radiation tolerance may worsen, design complexity increases
    - NRE cost increases substantially (with cost of foundries & processing/masks)
  - More complex designs are possible with lower risk
    - sophisticated design and simulation tools but longer design times for larger chips
    - evaluation with sensors & real environment is crucial – “perfection” takes time
  - Nevertheless, trend is towards more intelligence on the module

The sensor-ASIC boundary

- Sensors and ASICs use very similar technologies
  - manufacturing cost per unit area very similar
  - why not use ASICs to produce sensors?
- We do! MAPS (monolithic APS) see next slide
  - limited radiation hardness and readout speed – not fundamental?
- but
  - ASICs are limited in area by reticle (masking steps)
  - Yield is high but not 100%, so wafer scale detectors are difficult
  - Sensors usually require 100% coverage
  - Data transfer across modules remains challenging (& increases)
- Hybrid technologies assemble sensors and ASICs
  - came of age during LHC experiment construction
  - bonding is still a costly step

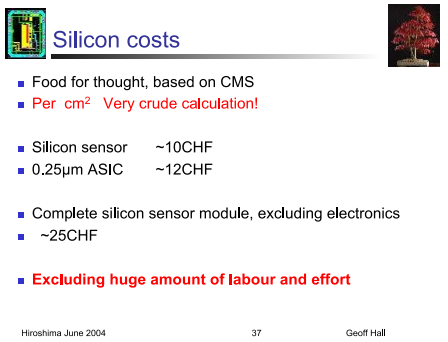

ALICE proposal

A MAPS detector

New ITS Layout

25 G-pixel camera
(10.3 m^2)

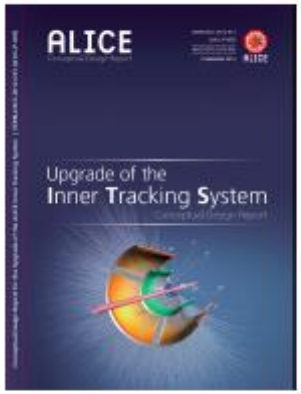

CDR – LHCC-2012-013

7 layers of MAPS

Outer layers

Middle layers

Inner layers

Beam pipe

Radial coverage
22 – 406 mm

700 krad/ $1x10^{13}$ 1 MeV n_{eq}
Includes safety factor 10

R Yarema 2007

Physicist's Dream

- Physicists have long dreamed of integrating sensors and readout electronics.
- Pixel designs have progressed from hybrid designs to MAPS.
- There is now an opportunity to provide further improvements with newer technologies (SOI detectors[1], and 3D).

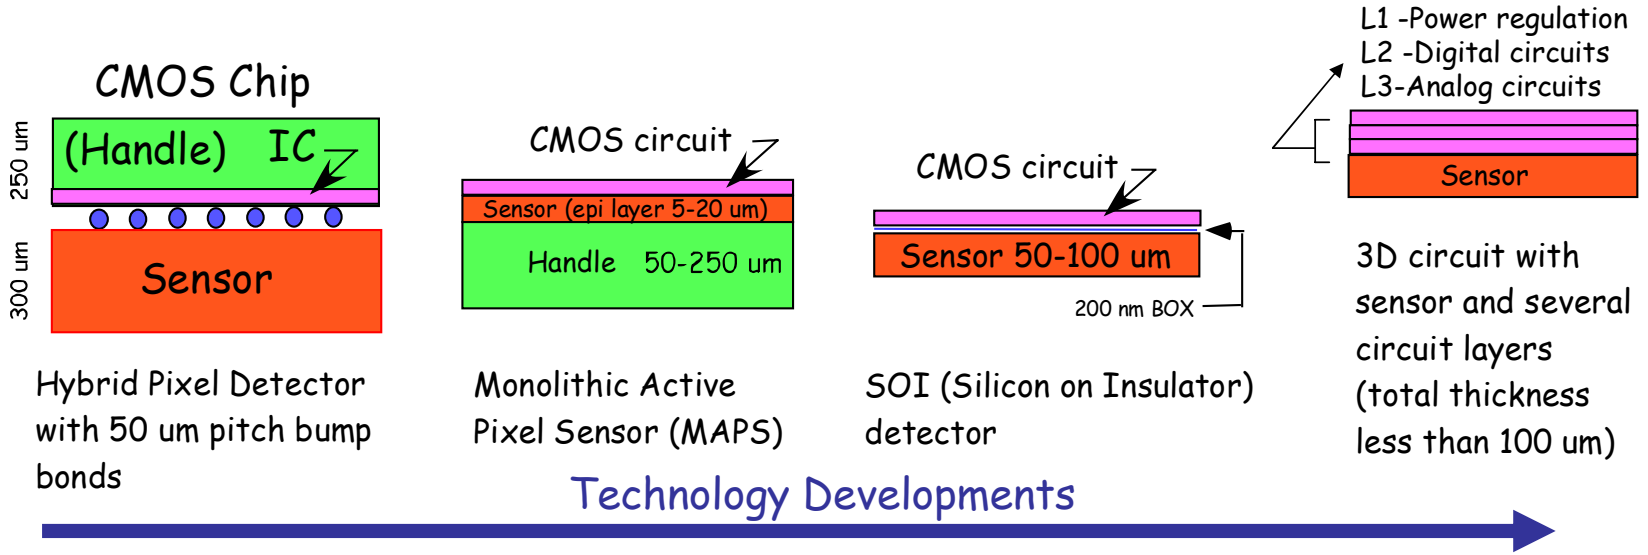

The RAPS04-3D structures

- ✓ Active Pixel Sensor 3T architecture with different photodiode area.

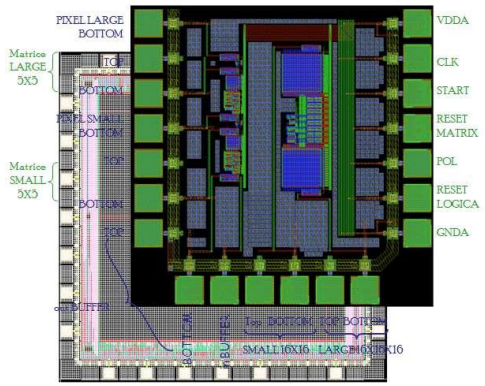

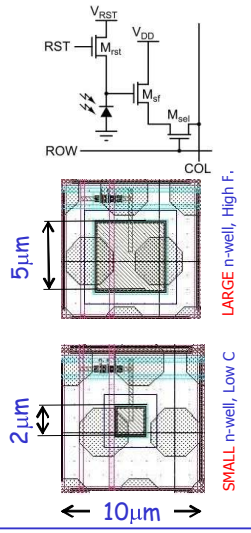

The RAPS04 chip structures

(Chartered) - GlobalFoundries/Tezzaron
3D-IC Integrated 2-tier stack
130nm CMOS

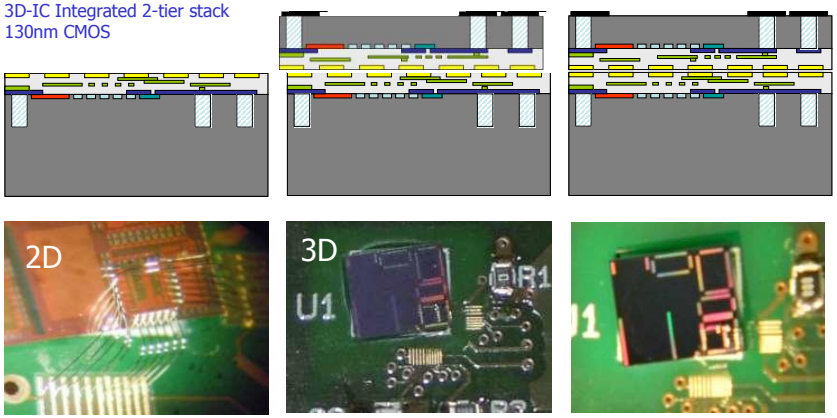

2D.

3D Not Aligned.

3D Aligned (Ziptronix/Tezzaron).

D Passeri et al 2012

3D devices

- Small but working!
  - from shared MPW project

62 MeV protons response

- ✓ RAPS04-3D 62MeV protons – Outer & Inner tier coincidence responses

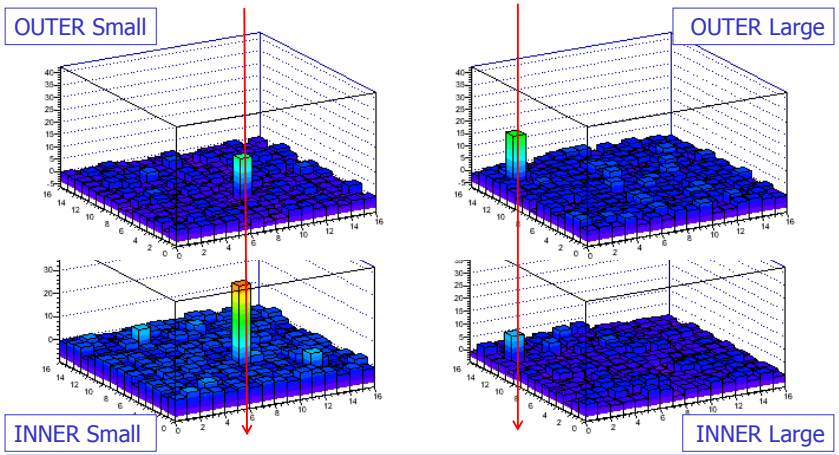

Further challenges

- ASIC design is not the only challenge
  - interconnects and assembly into modules increasingly important
    - generally reliant on commercial input & help
  - long term studies essential to find small faults or features
    - radiation effects, high rate, noise, power, assembly issues,…
  - building test & monitoring features into chips is vital

- Scope for new ideas
  - pixel cluster identification, p_T discrimination, …

- What is the right balance between performance, complexity and risk?

ATLAS Phase-0 upgrades (LS1)

- Insertable B-Layer
  - Installation of IBL in the pixel detector March 2014
  - FE-I4 Pixel Chip, 130 nm CMOS process
  - Will stay until Phase-II

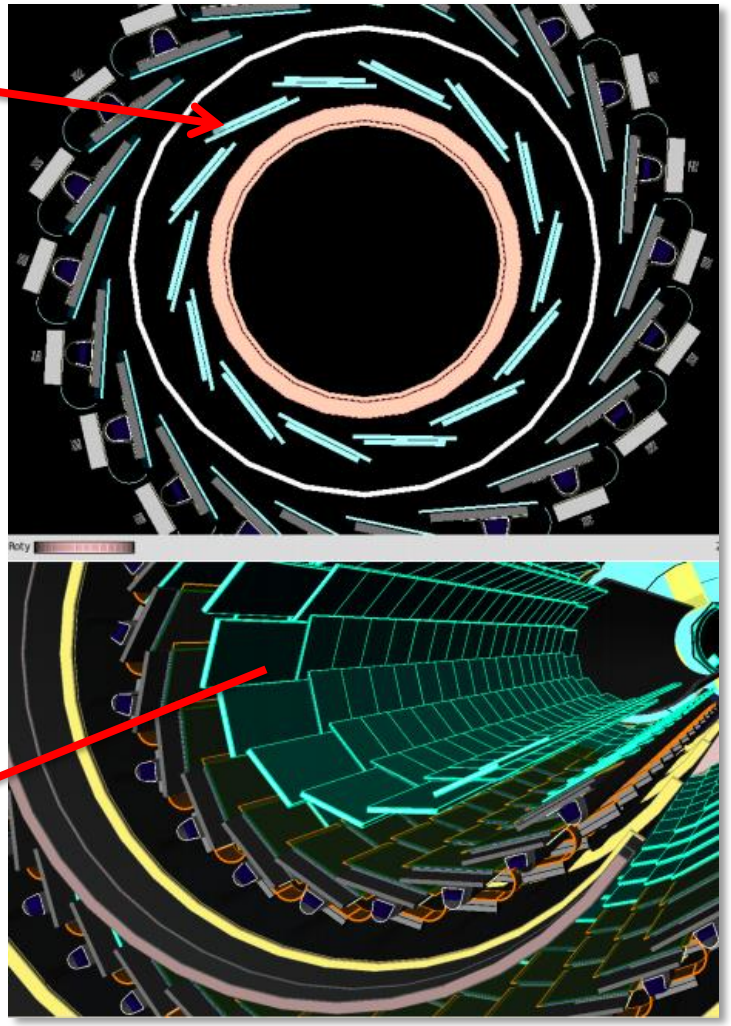

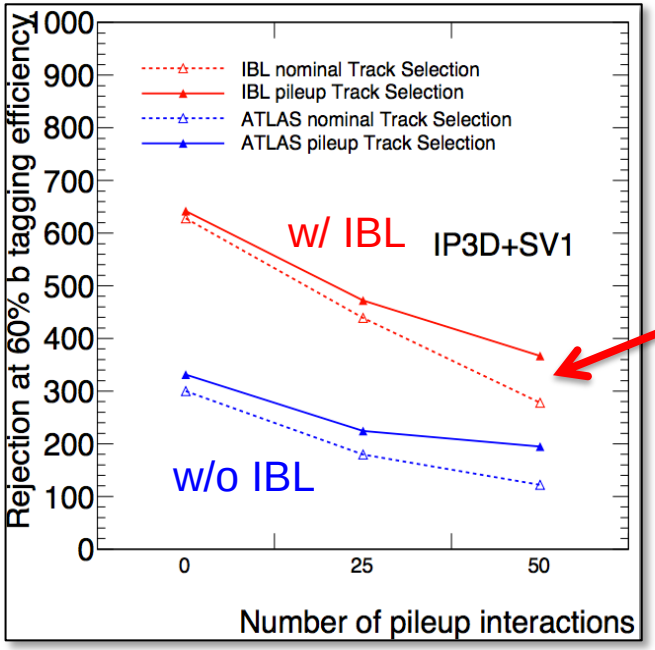

b-tagging rejection vs pile-up

ATLAS FE-I4 pixel ASIC

- 19 x 20 mm² → ~6 times size of FE-I3
- Pixel size 50 x 250 µm → FE-I3: 50 x 400 µm
- 26.880 pixels
- Organized in 336 rows and 40 double columns
- Readout organized in four pixel regions, hits buffered at pixel level until LV1-trigger → cope with high occupancies

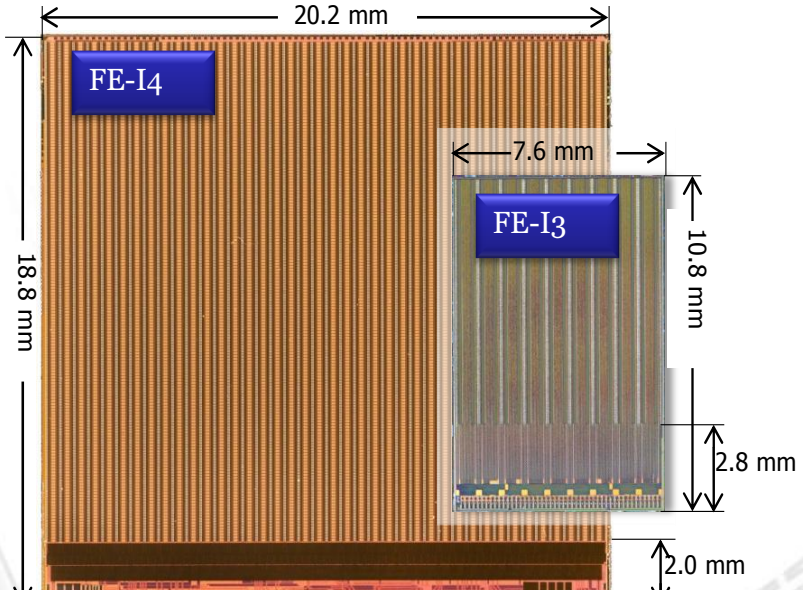

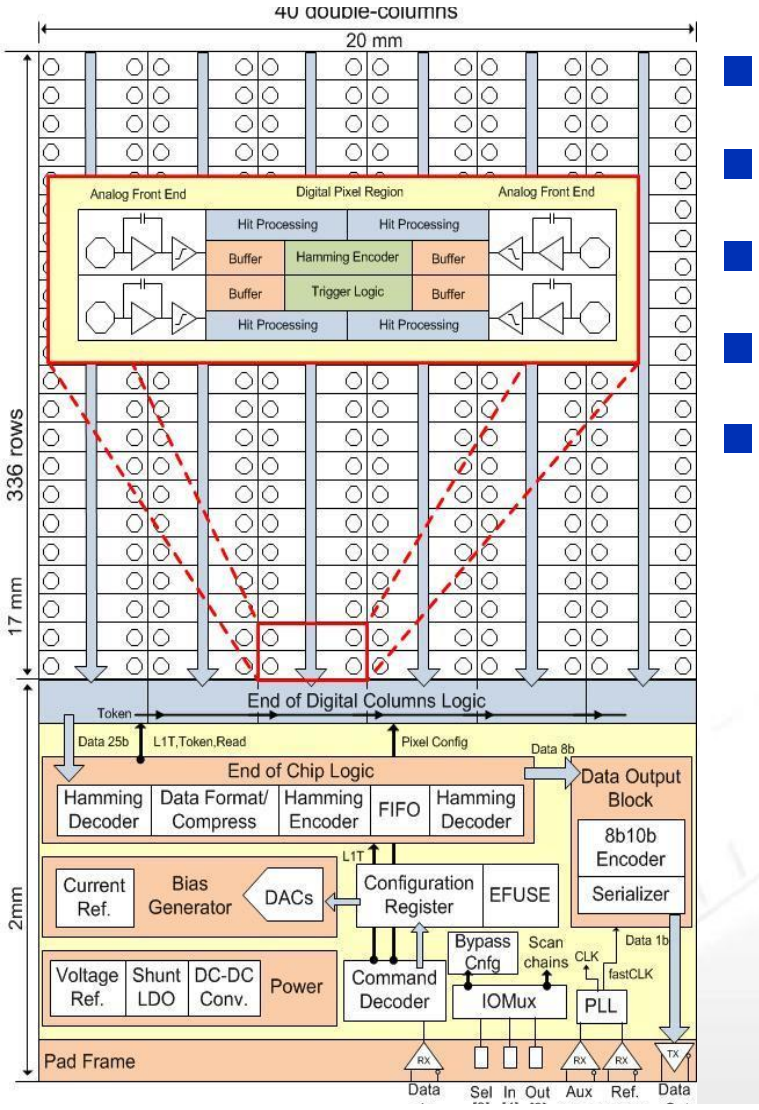

D Braga, M Raymond et al 2014

CMS Tracker ASIC evolution

- 1999: APV25 0.25µm
  - 7 mm x 8mm (128 chan)

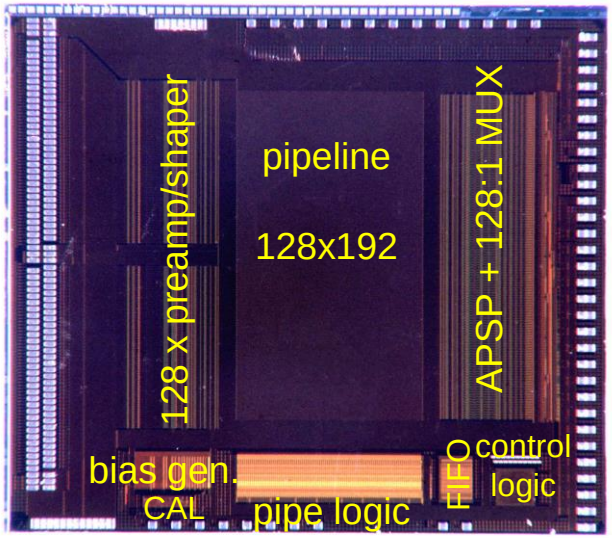

programmable settings (now standard)
analogue data
~4 µs latency
wire-bondable
pulse-shaping choice

2011: CBC 0.13µm

7mm x 4mm (128 chan)

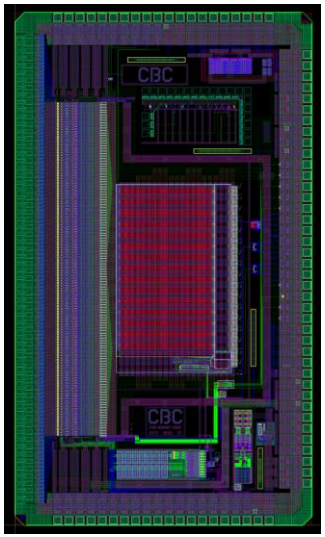

binary data,
6.4 µs latency
wire-bondable

2013: CBC2 0.13µm

11mm x 5mm (254 chan)

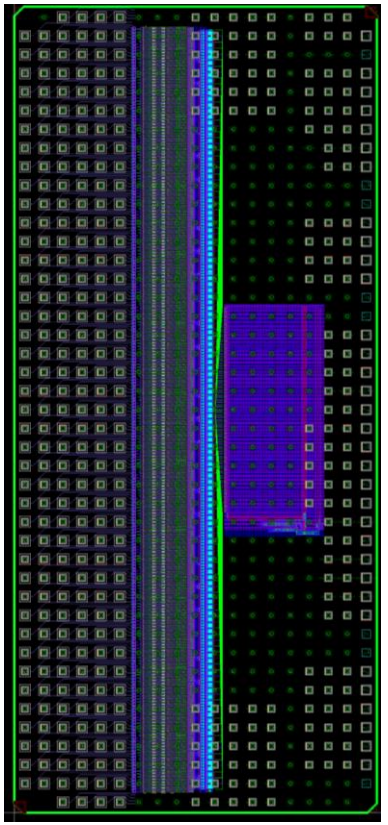

bump-bondable,
cluster & correlation logic

CMS Phase II Outer Tracker design

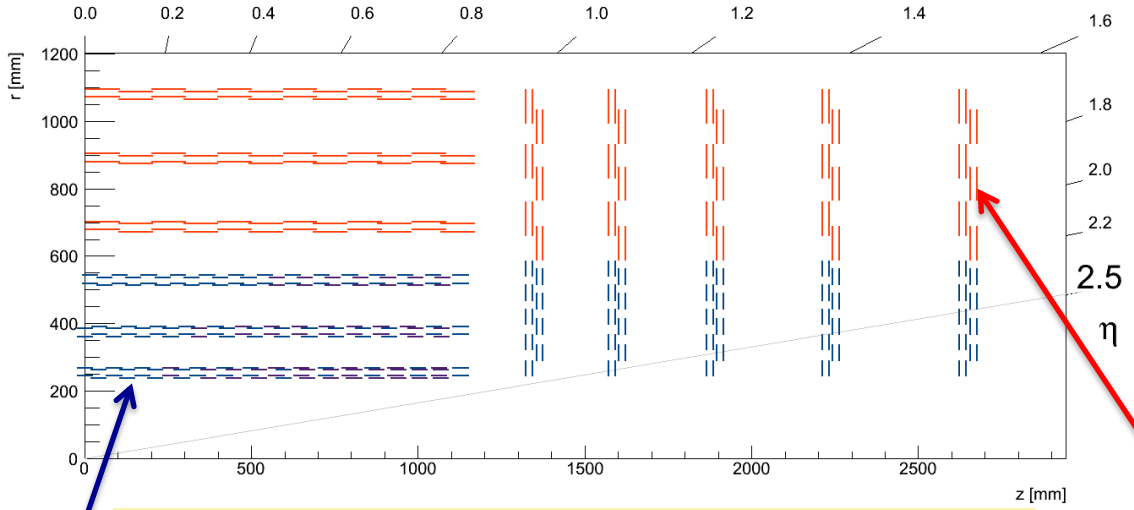

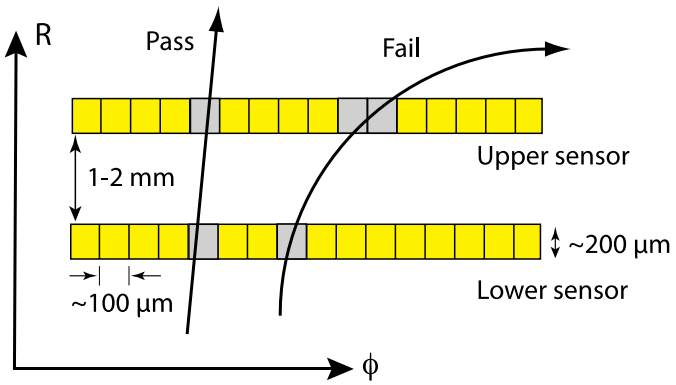

- ~15000 modules transmitting
  - p_T-stubs to L1 trigger @ 40 MHz
  - full hit data to HLT @ 0.5-1 MHz

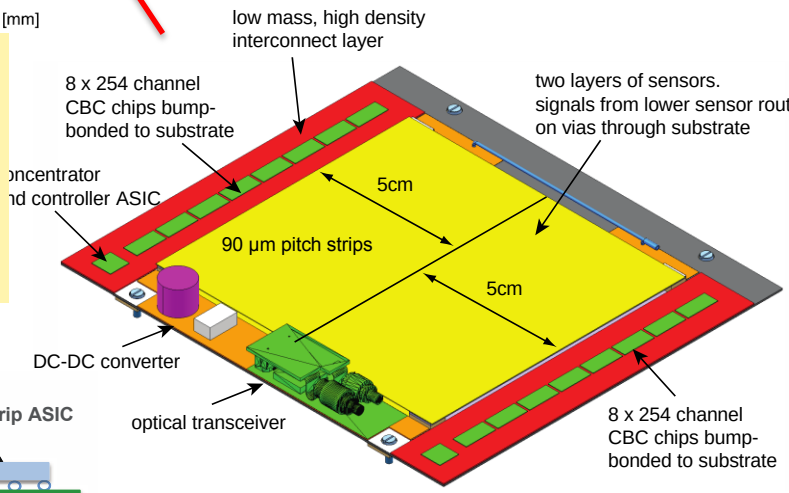

~7100 PS-modules

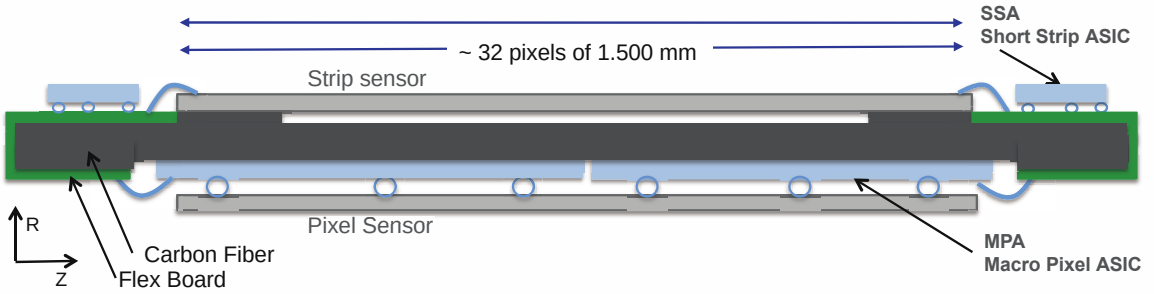

~8400 2S-modules

Stub-finding logic

reject wide clusters

identify stubs of p_T>2 GeV tracks

mask noisy channels

fully functional logic

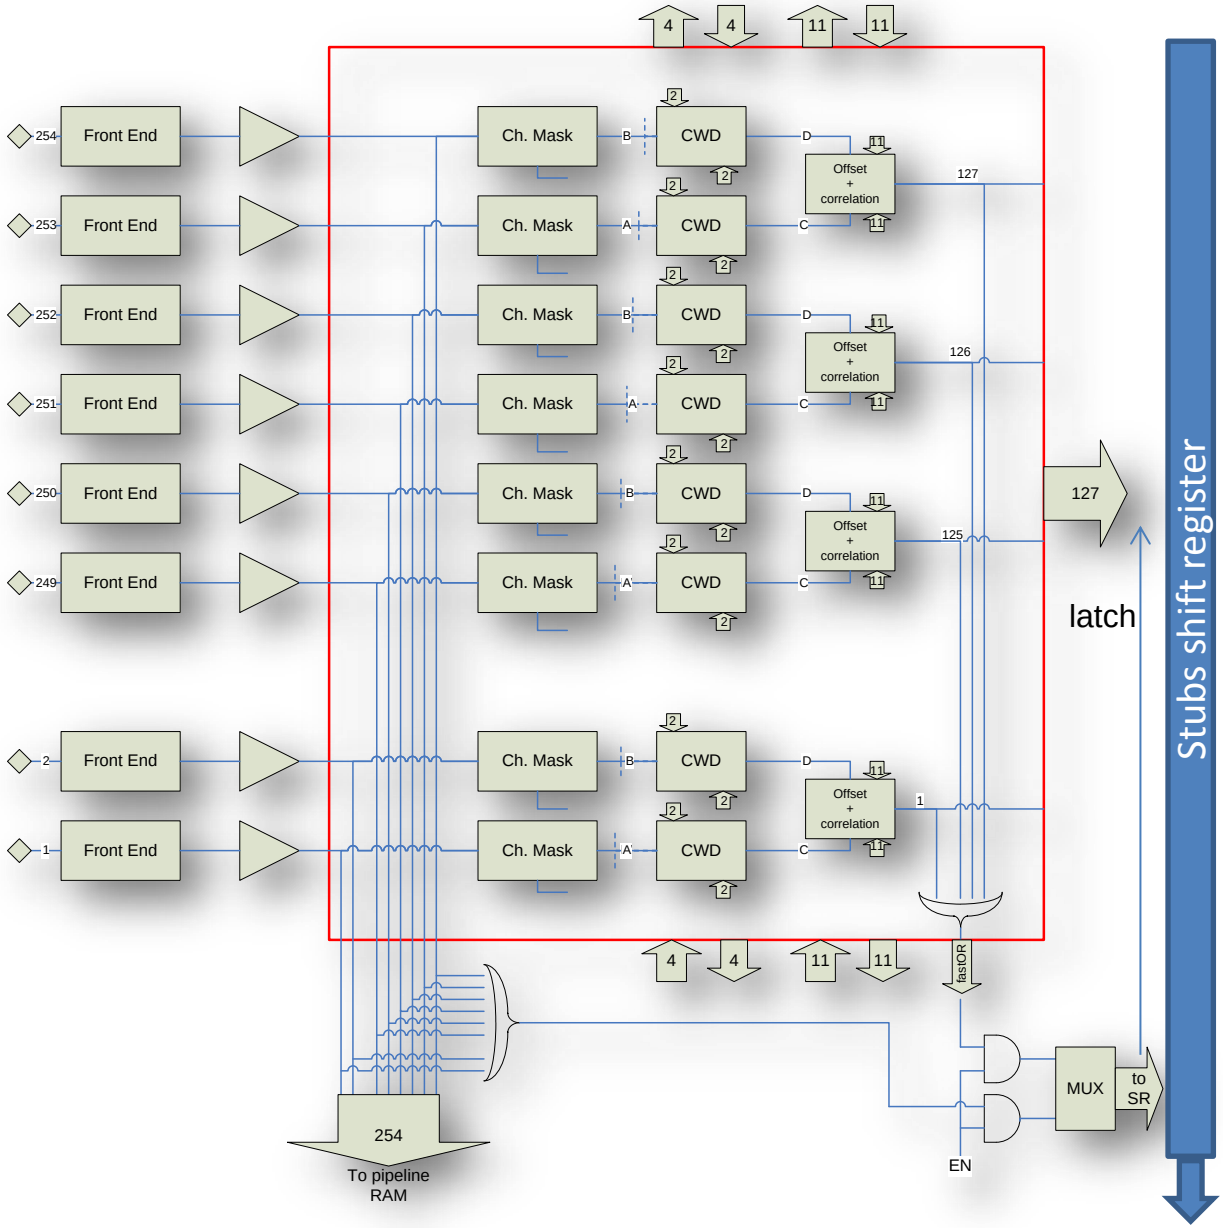

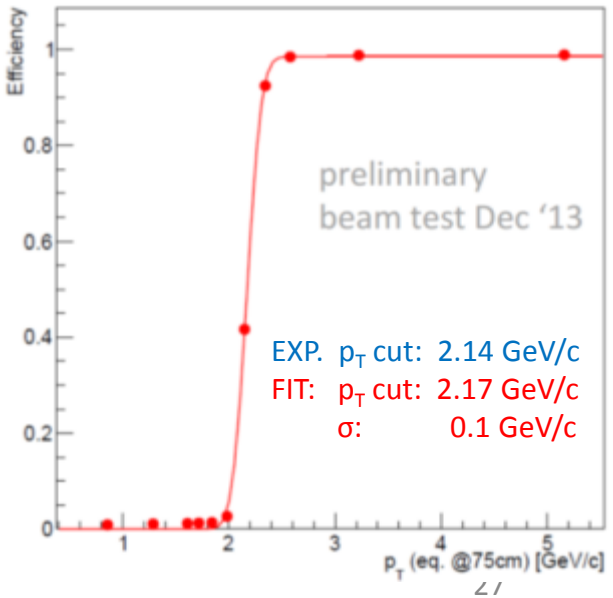

PS MODULE BLOCK DIAGRAM

D Ceresa et al 2014

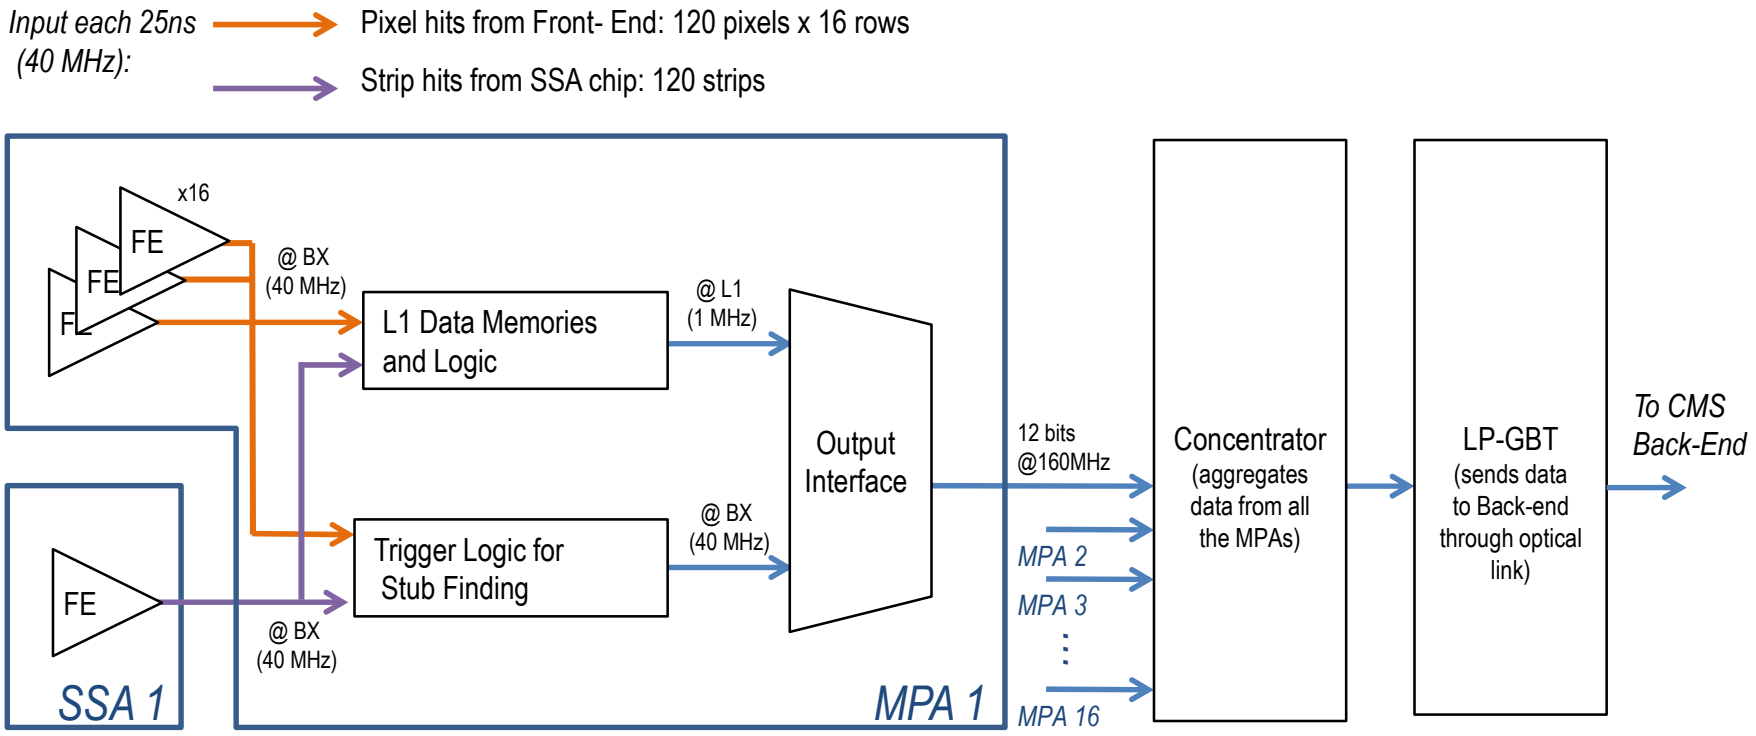

L1 Data Memories stores the events for the duration of the L1 latency. Upon arrival of L1 signal, the event is processed by the L1 Data Logic and sent to the Output Interface.

Trigger Logic elaborates inputs synchronously with BX and looks for the coincidences between pixel and strip hits in order to generate stubs.

Output Interface organizes data from the two data paths and transmits them through the 12 bits bus towards the concentrator at 160 MHz

THE MPA-LIGHT DEMONSTRATOR

- A reduced size MPA design
  - 16 x 3 pixels (instead of 120 x 16)
  - Size of single pixel (as final): 100 x 1446 um
  - Bump-bond pad size: 100 mm
  - Pitch 200 mm horizontal, 200 mm vertical
  - Wirebond pads for hybrid connections
  - Pixel floorplanning similar to the final MPA
  - Scalable to the final design
- Purpose
  - Prototype & qualify the analog FE circuitry
  - Facilitate the development of the sensor
  - Understand and solve the numerous technical aspects of the Module Assembly
  - Digital Logic Verification and Low Voltage Supply Testing

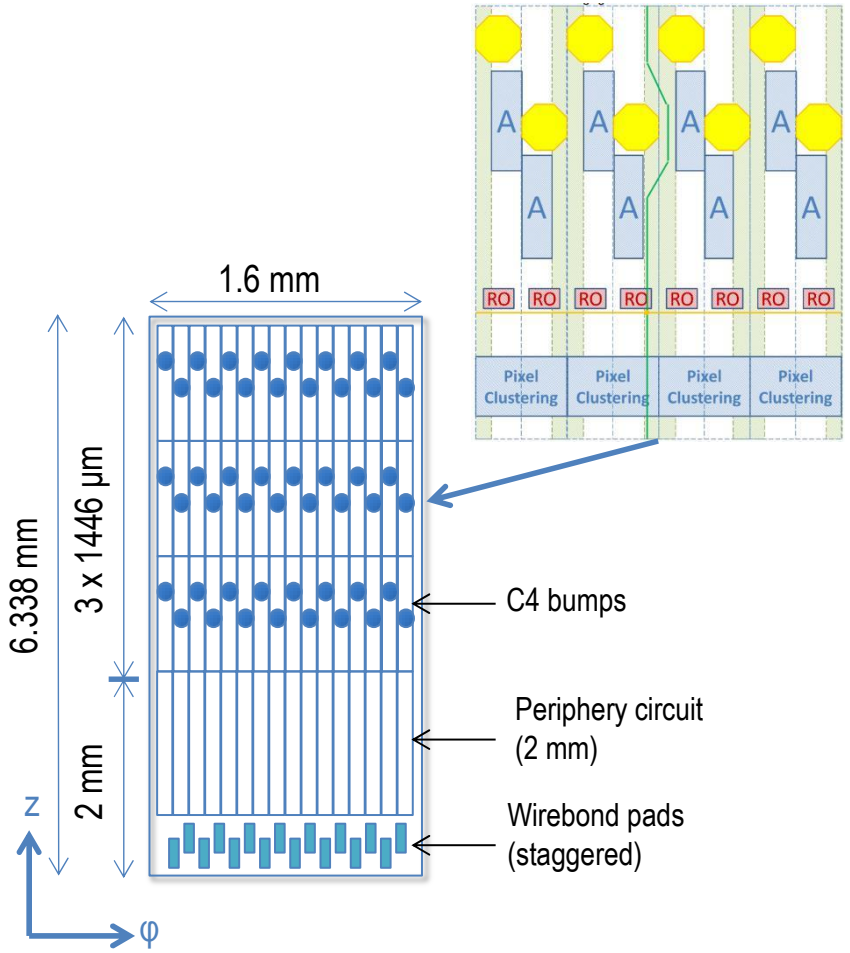

D Ceresa et al 2014

Optical links

- First large scale applications at LHC – and vital to success
  - now an established – but evolving vigorously – technology
  - profited from serendipitous commercial progress (telecomms & internet)
- Almost noiseless data transmission, immune to electrical interference, low power
  - major advantage in dense, congested collider experiments with high channel count
  - customisation required to meet mass and radiation constraints
    - commercial requirements also different for telecomms c.f. HEP data
  - connectors, fibre and wavelength to be chosen
  - not yet truly standard component set – will there be?
  - on-detector and off-detector components are different
    - radiation hard electronics and sensors needed, so lag in performance

The need for digital transport is growing exponentially

Information is of little use if you have to keep it to yourself

- Humans have a desire to interact (Cell phones, YouTube, …)
- Requires huge transport capacities (especially for real time app's)

Computers also want to talk:

- 1 Flop triggers ~1 Byte/s of transport
- Coupled with exponential growth in computing power

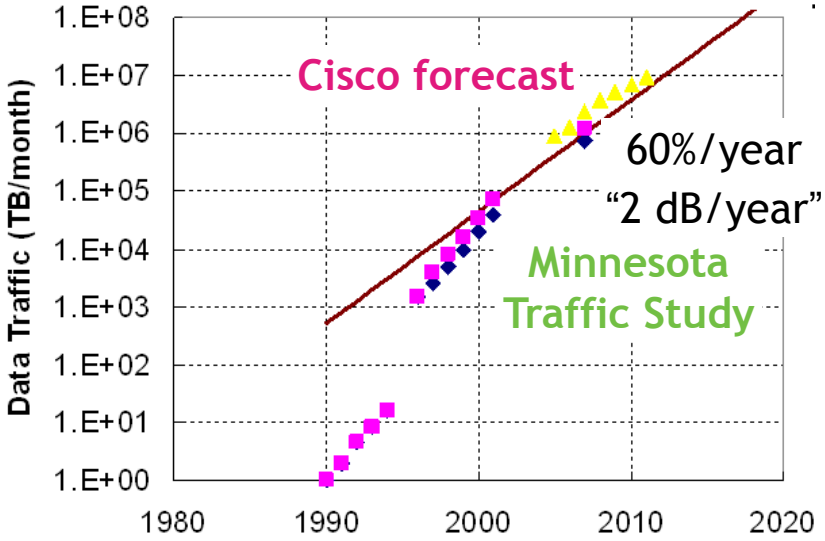

tv ORF.at on demand

YouTube

Telepresence

Alcatel·Lucent

Fiber-optic transmission systems to provide high capacity - Basics

- Guided, isolated from ext. interferences
- **Very low attenuation**
  - 0.3dB/km @ 1310nm
  - 0.2dB/km @ 1550nm (down to ~0.16dB/km)

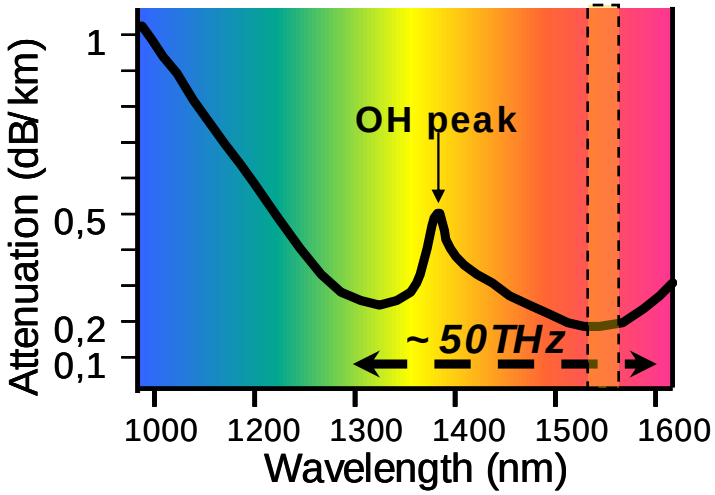

Huge available bandwidth ➔ high capacities ?

- Virtually 50THz
- In practice, operate w/ 4-5THz bandwidth all-optical Erbium Doped Fiber Amplifiers

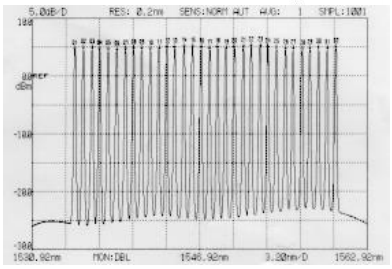

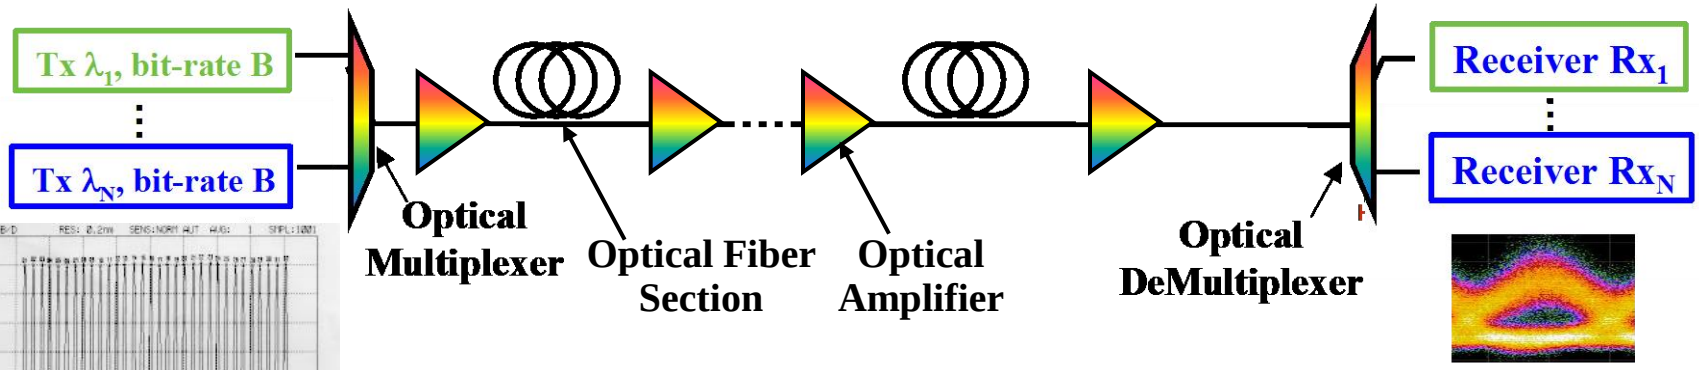

Wavelength Division Multiplexing
Typical bandwidth: 1529-1565nm
50-100GHz channel spacing

➔ **Capacity = sum of channel rates**

M Ritter IBM 2010

Optical Interconnect Modeling

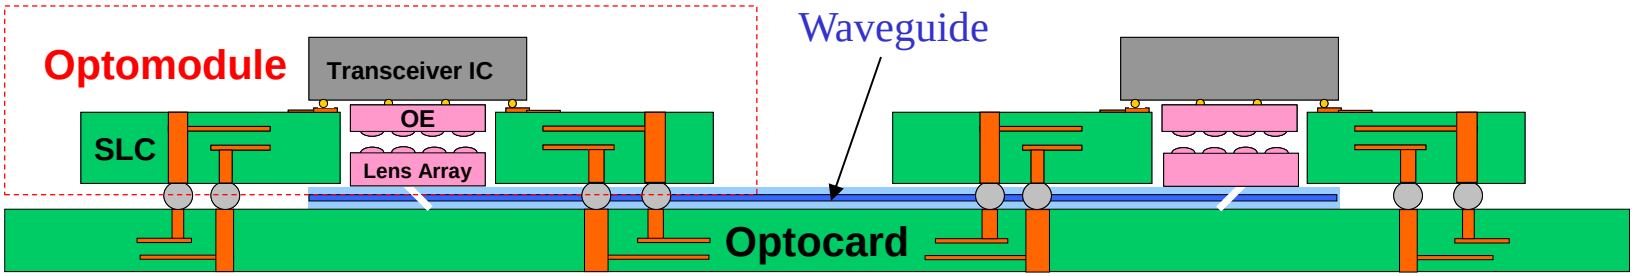

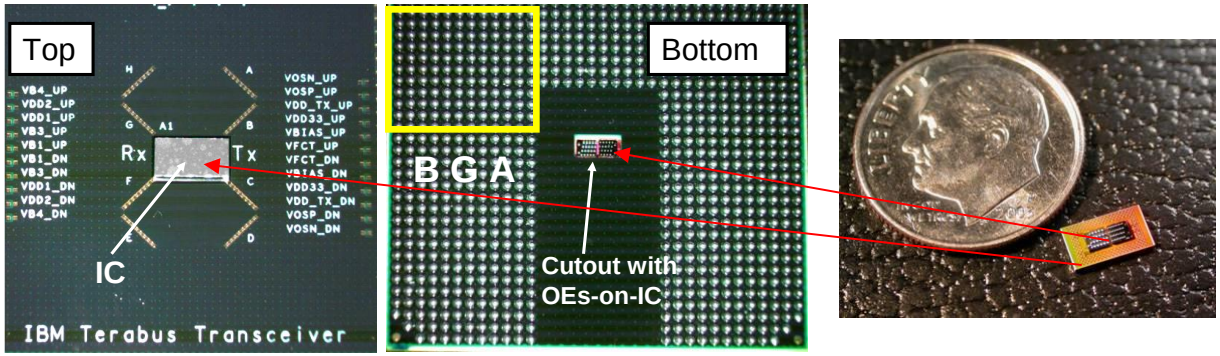

Terabus
"Optomodule"
Kash et. al.

- **Lowest-power links use optics as "analog repeaters" of signal with no clock recovery**
  - E-O-E modeling required: jitter adds up over two electrical links, one optical link

Complex assemblies to be customised by industry for radiation, low mass environment

Optical link challenges

- Two kinds of requirements in HEP (plus computing/DAQ)
  - data (and control) communications from (and to) detector
    - customised systems and components required for tracking and calorimetry
    - wide range of components and choices, which evolve
    - modulators: transmission or reflective
    - laser transmitters: VCSELs, Fabry-Perot edge emitters
    - wavelength: long distance links use longer wavelengths eg ~1300-1600 nm
    - analogue data was used but future systems all binary?
      - A to D conversion on-detector
  - data transmission in trigger systems
    - no radiation requirements but high speed and reliability required
    - commercial devices, steadily increasing BW with challenges in developing processing boards to match

FPGAs and digital processing

- Another huge change in landscape over last 1-2 decades
  - mid-1990s: custom digital ASIC considered for off-detector data processing for CMS silicon tracker
  - mid-2000s: completed production of 500 boards, each with 8 “large” FPGAs
    - commercial state-of-the-art, and most costly
    - custom optical receivers
  - mid-2010s: production of trigger processing board, with 1 large FPGA and 72 + 72 optical links
    - ~several 100x data throughput cf Tracker
    - vastly more capacity for data processing
    - but also more difficult to program

P Alfke (Xilinx Corp)

The view in 1999

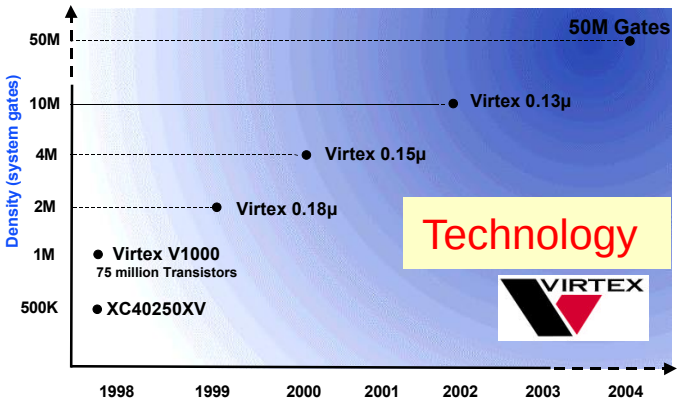

current Virtex-7: 16-28nm

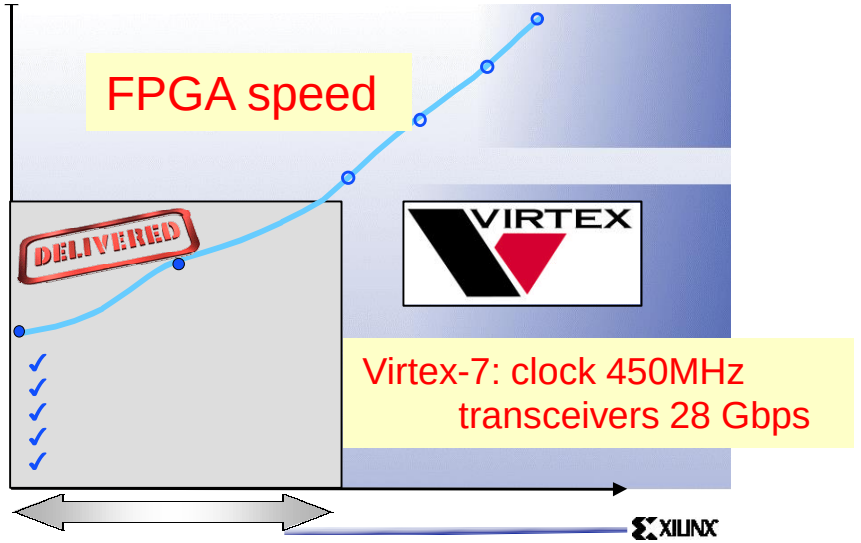

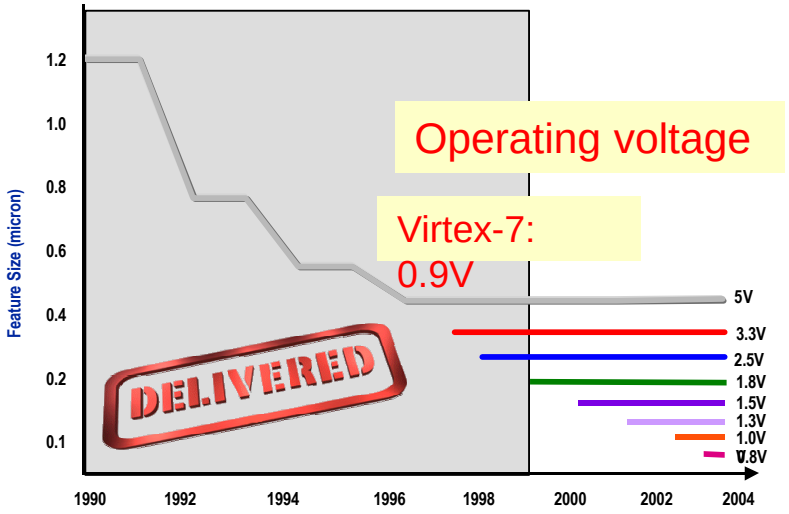

	1965	1980	1995	2010(?)
Max Clock Rate (MHz)	1	10	100	1000
Min IC Geometries (µ)	-	5	0.5	0.05
# of IC Metal Layers	1	2	3	10
PC Board Trace Width (µ)	2000	500	100	25
# of PC-Board Layers	1-2	2-4	4-8	8-16

- Every 5 years: System speed doubles, IC geometry shrinks 50%
- Every 7-8 years: PC-board minimum trace width shrinks 50%

www.xilinx.com

Current Xilinx products

- Increasingly powerful devices (and other similar manufacturers)
  - the challenge is how to exploit them

	Spartan-6	Artix-7	Kintex-7	Virtex-7	Kintex UltraScale	Virtex UltraScale
Logic Cells	147,443	215,360	477, 760	1,954,560	1,160,880	4,407,480
BlockRAM	4.8Mb	13Mb	34Mb	68Mb	76Mb	132.9Mb
DSP Slices	180	740	1,920	3,600	5,520	2,880
Transceiver Count	8	16	32	96	64	120
Transceiver Speed	3.2 Gb/s	6.6 Gb/s	12.5 Gb/s	28.05 Gb/s	16.3 Gb/s	32.75 Gb/s
Total Transceiver Bandwidth (full duplex)	50 Gb/s	211 Gb/s	800 Gb/s	2,784 Gb/s	2,086 Gb/s	5,886 Gb/s
Memory Interface (DDR3)	800	1,066	1,866	1,866	2,400	2,400
I/O Pins	576	500	500	1,200	832	1,456
I/O Voltage	1.2V – 3.3V	1.2V – 3.3V	1.2V – 3.3V	1.2V – 3.3V	1.0 – 3.3V	1.0 – 3.3V

FED 2006

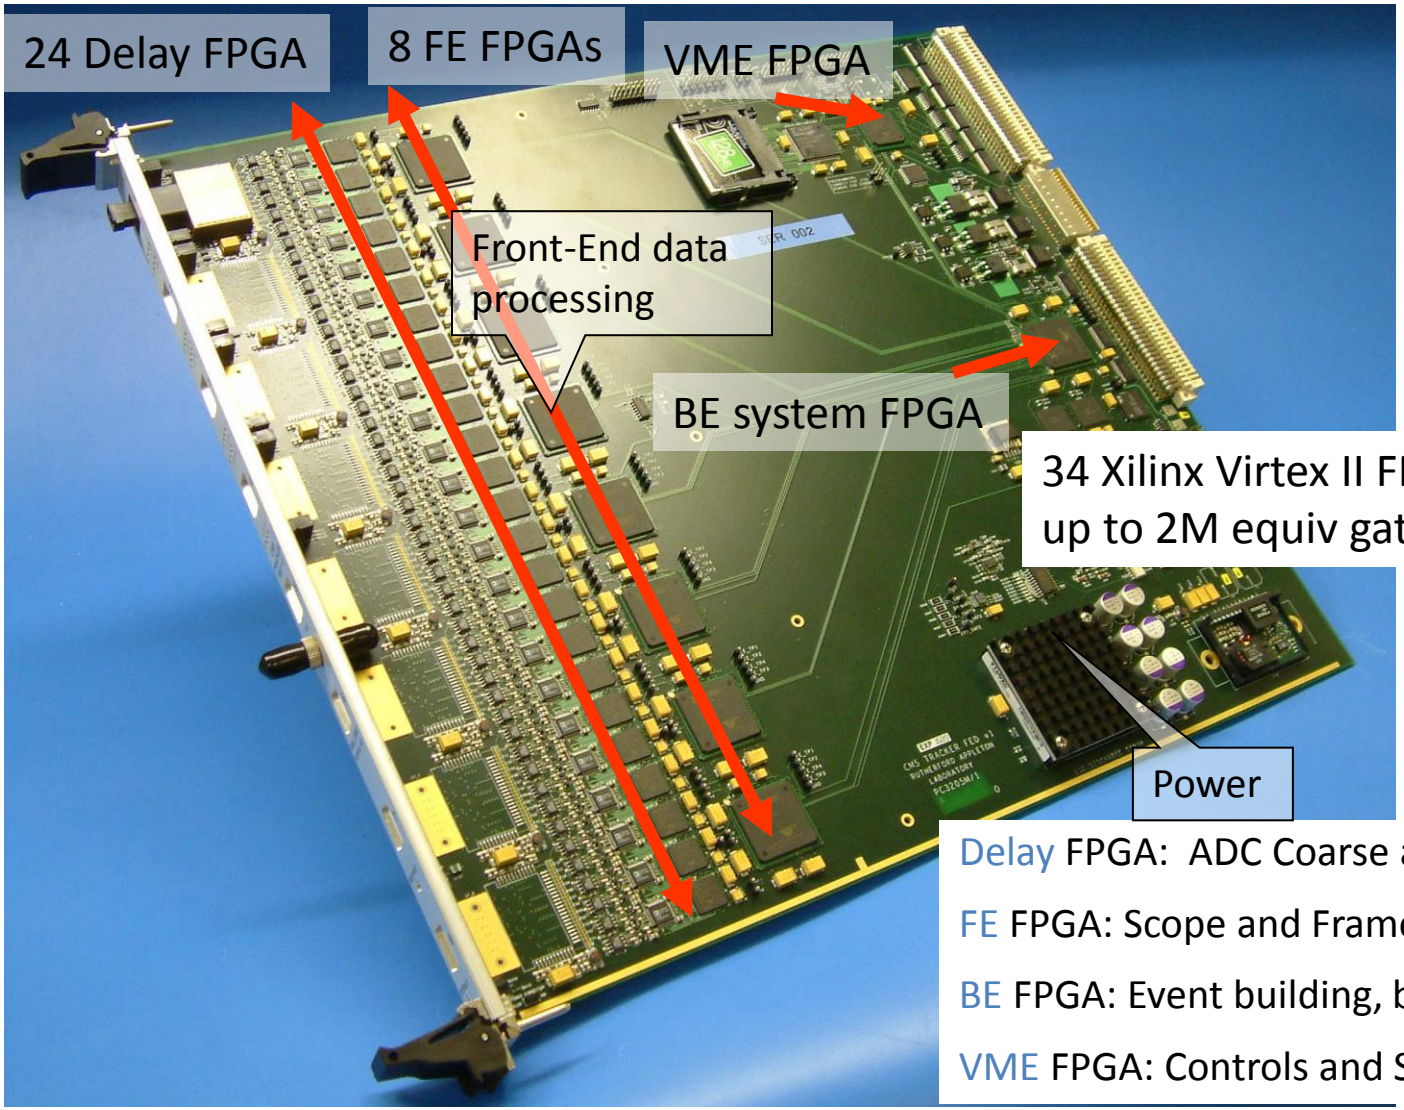

- CMS Tracker off-detector module
- Total input data rate ≈ 96 x 40 MHz x 10 bits ≈ 3.8 Gbps

34 Xilinx Virtex II FPGAs
up to 2M equiv gates each

Delay FPGA: ADC Coarse and Fine Clock Skewing.

FE FPGA: Scope and Frame Finding modes.

BE FPGA: Event building, buffering and formatting.

VME FPGA: Controls and Slow Readout path.

FED Overview

Input ~ 3 GBytes/sec

after Zero Suppression : Output: ~ 200 MBytes/sec

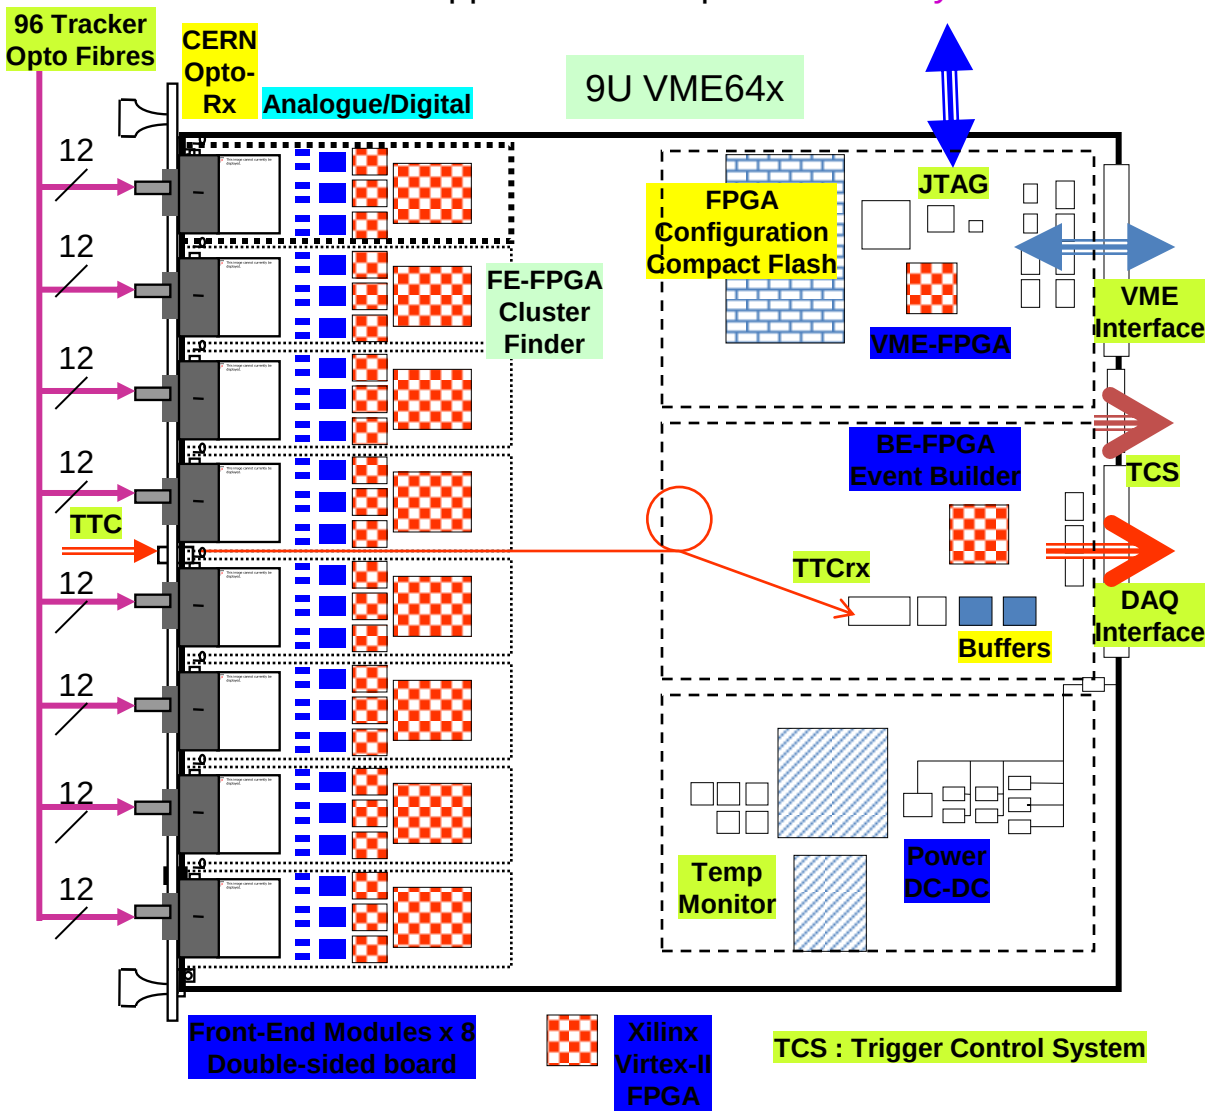

Modularity

9U VME64x Form Factor

Modularity matches Opto Links

25,000 Si strips / FED

440 FEDs in Total.

8 x Front-End "modules"

OptoRx/Digitisation/Cluster Finding

Back-End module / Event Builder

VME module / Configuration

Power module

Other Interfaces:

TTC : Clk / L1 / BX

DAQ : Fast Readout Link

TCS : Busy & Throttle

VME : Control & Monitoring

JTAG : Test & Configuration

MP7 (Virtex-7 XC7VX690T)

future generations will improve, but don't yet know precisely how

purpose-built µTCA card for CMS upgraded L1 calorimeter trigger

TM performance & calo algorithms demonstrated in recent integration tests

- **72 input/72 output** optical links

-all links operate at 12.5 Gbps
(10 Gbps in CMS)

- total bandwidth **> 0.9 Tbps**

tested, currently in production

NB good cooling required!

G Iles et al

MP7 2013

- IN
- 72 x 12.5 Gbps = 0.9 Tbps
- OUT
- 72 x 12.5 Gbps = 0.9 Tbps
- flexible processing
- NB smaller form factor than 9U FED
- Power similar

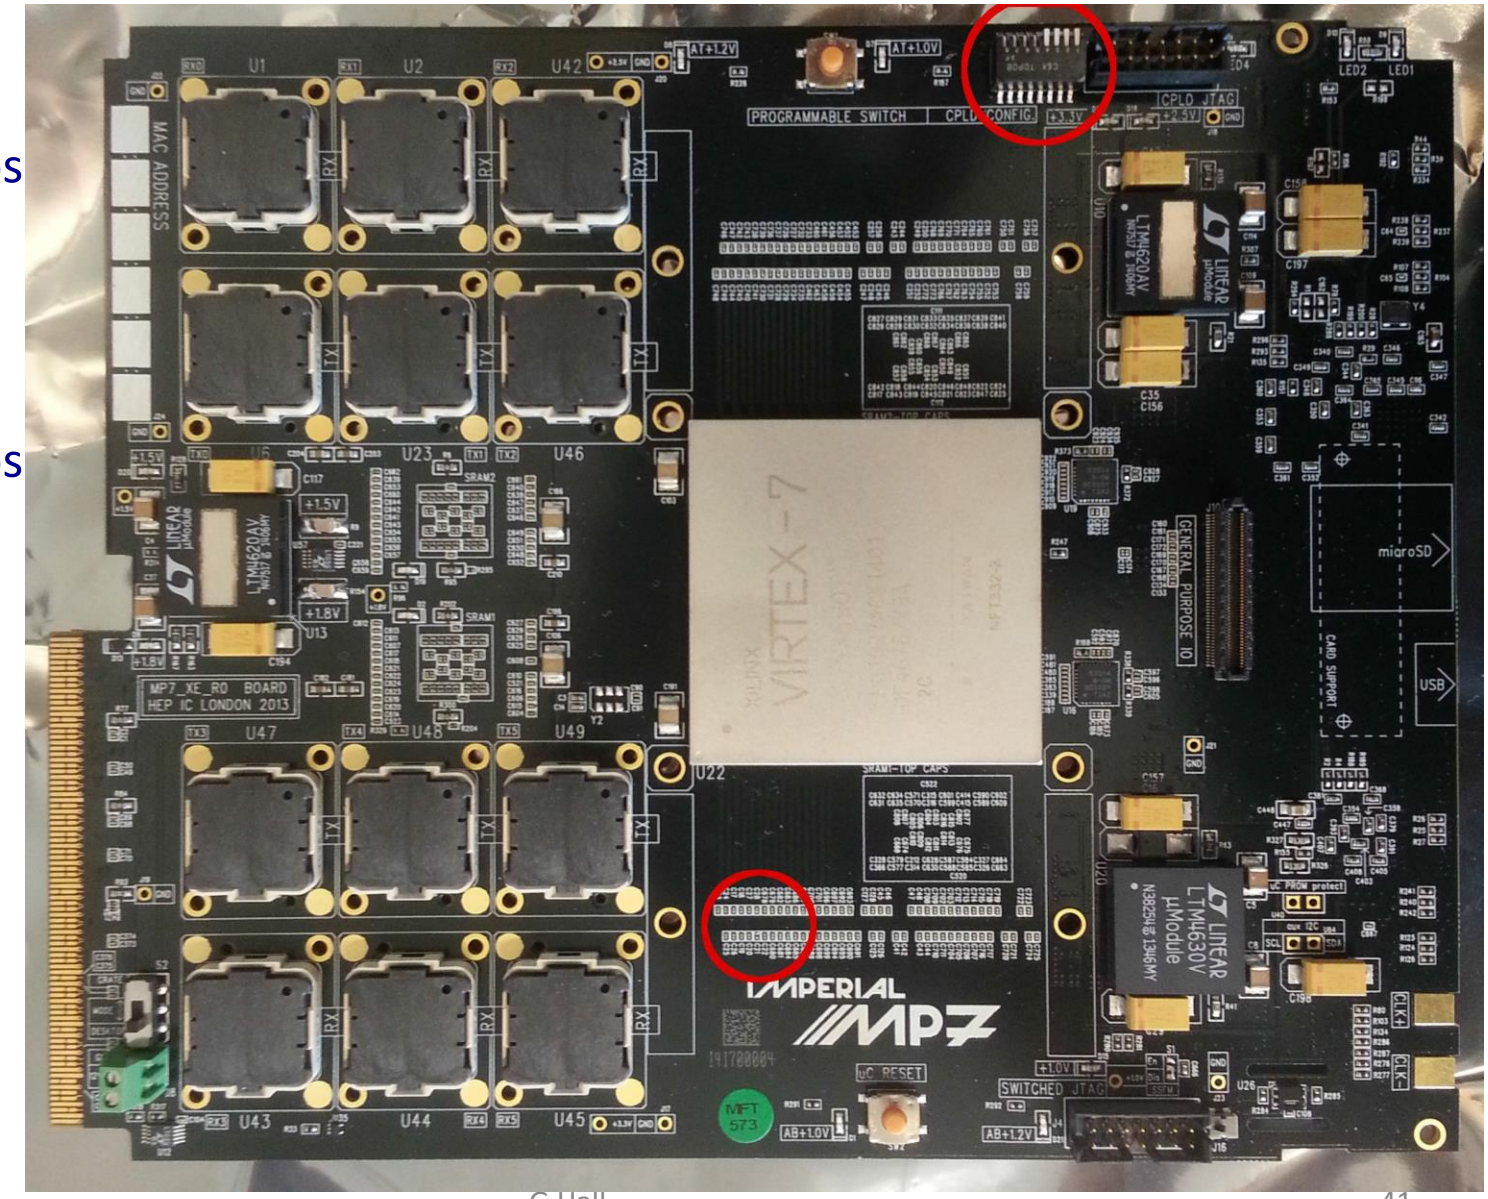

All problems solved?

- Programming FPGAs is not quite like developing software
  - HEP does not use FPGAs like many other users
    - eg video processing, for large scale use in products
  - typically we often require fast, efficient algorithms
    - adapted to changing experimental conditions
  - simple conceptual problems may not be easy to implement
    - we also know that software code can be inefficient and optimisation is essential for future data processing
    - but there is a large community of software developers

A simple example of Routing Congestion

- (G Iles) Created simple design to find routing limit in modern FPGA
  - 30x36 2x2 tower clusters ("electrons") with 10bit energy
  - 432 Gb/s (without 8B/10B)
    - Approximately ¾ of CMS
  - Sum 16 clusters to create "pseudojets"
  - No other firmware (e.g. no sort, no transceivers, no DAQ, etc)
  - XC7VX485T – Place & Route fails even though LUT usage only at 29%

Bare minimum "physics" algorithm

but number of LUTs is not the whole story…

A bigger FPGA may not solve all the problems…

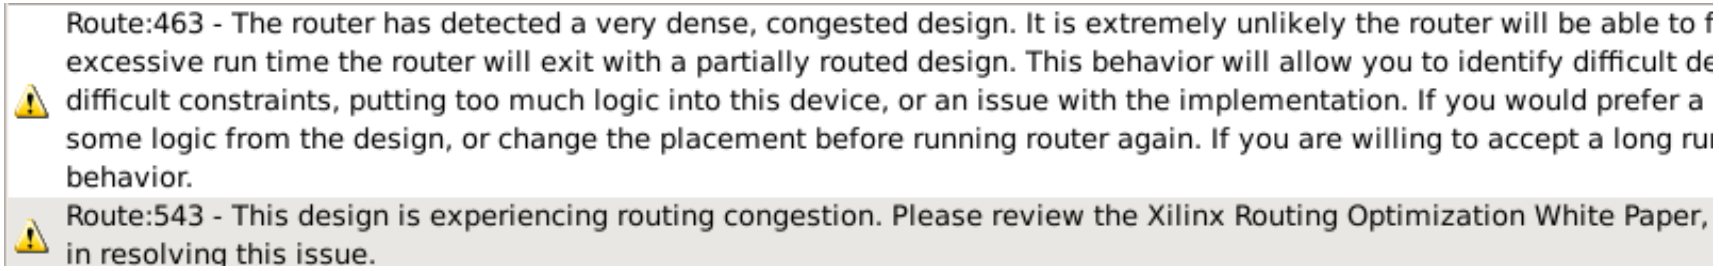

Summary

- ASIC design is an example
  - driven so far mainly by trackers (largest application)
- Some of the questions to be asked
  - which manufacturing process? do we need smallest feature size? what radiation tolerance is required? What performance is needed? What are the constraints from power, space, etc? How long can the ASIC development take? Can more performance arise from adapting existing designs?
- Similar questions can be posed of all the technologies described. Various decisions must be made, including
  - Power v performance
  - Risk v flexibility
  - Development time v level of perfection
  - Redundancy v reliability
  - …

So - what is the answer to the original question?

New challenges for signal processing

- "New" means evolution
  - creeping changes, whose dramatic scale only become obvious in comparison with the past
  - our systems are of increasingly high complexity
    - no time for discussion of how to operate or program them but many years were needed, to construct and debug them, and provide tools to use them
- Which stage are we at?

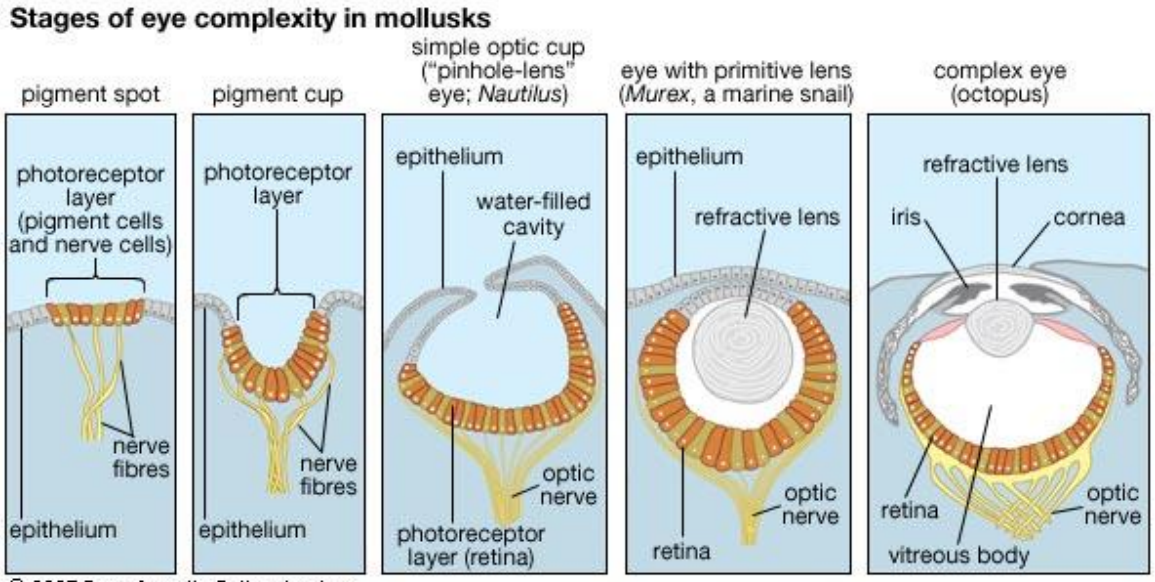

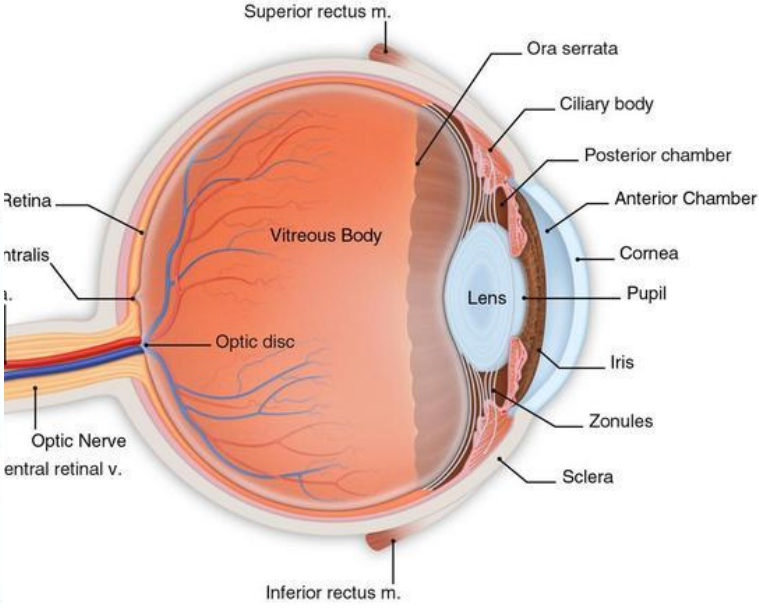

from Steve Jones. *Almost Like a Whale.*

A parable of evolution

- To make soap powder, liquid is blown through a nozzle. As it streams out, the pressure drops and a cloud of particles forms…
- Thirty years ago, the spray came through a simple pipe that narrowed from one end to the other… it had problems with irregularities in size of grains, liquid or blockages…
- The problem was too hard to allow even the finest engineers to explore with mathematics and design… they tried another approach… **evolution**
- Take a nozzle that works quite well and make copies, each changed at random. Test them for how well they make powder. Then impose a struggle for existence by insisting that not all can survive.
- The few able to do a superior job are allowed to reproduce and are copied – but again not perfectly. As generations pass, there emerges a new and efficient pipe of complex and unexpected shape.
- The nozzle has become an intricate duct, longer than before, with many constrictions and chambers. The liquid follows a complex path before it sprays.

New challenges for signal processing

- What next?
  - how to manage complexity?
  - potential for all system elements to increase further in complexity
    - this is the trend in consumer products
    - but who now maintains their own car, or writes their own software?
    - increasing reliance on third parties, or sophisticated tools
  - HEP can choose its level of control
    - and depends on the young members of its community to do the work and have new good ideas

BACKUP MATERIAL