Some lessons learned, so-far, from the HL-LHC experiments

Steve McMahon

RAL / Oxford

Vertex 2023 – 17th October 2023

Some lessons learned, so-far, from the HL-LHC experiments

and why is it taking so long

Steve McMahon

RAL / Oxford

Vertex 2023 – 17th October 2023

Lessons learned - Introduction

- This talk is about fabrication and progress towards the eventual delivery of silicon based charged particle tracking systems to the experiments. I cannot cover all areas in any detail, but I have picked a few areas of particular interest that might be of interest here.

- It is important to point out that the conclusions presented here are drawn from discussions with experts working on experiments of a very particular scale. These sub-detector collaborations are, in some cases, several hundred people from more than 100+ institutes from 25+ funding agencies. A lot of the lessons come from dealing with issues in communities of that scale.

Lessons learned - Introduction

- Before we start, it is instructive to ask who are these lessons for. Who is going to benefit? How can I use these lessons and when? And when one comes to want them will they be around. Are they applicable to the future? Did we learn anything from previous (past) experiments?

- Having said all of this I do hope everyone will find something they find useful. This week, I would also welcome hearing any other issues, particularly those that apply to small detectors and /or communities.

- This talk
    - does not deal with design.
    - does, not really, talk about technology choices sor technology problems.
- See other talks during this week

The people who helped – The consultants

- LHCb - C. Parkes (Univ. of Manchester)
- Sensors - Y. Unno (KEK) + M. Ullan (CNM-CSIC) + V. Fadeyev (SCIPP)
- RD53 - M. Garcia Sciveres (LBNL) + J. Christiansen(CERN-email)
- Strips - T. Affolder (SCIPP)
- Cooling - G. Viehhauser (Univ. Oxford) + L. Zwalinski (CERN)
- Strip ASICs - P. Keener (PENN)
- CMS tracker - D. Abbaneo (CERN)
- Electronics - H. Chen (BNL) + F. Vasey (CERN)
- ITk - S. McMahon (RAL – Univ. Oxford)

60-90 minute face-to-face interviews – based around a standard set of questions, sent by email in advance of the interview. Looking for common patterns between areas, looking for good advice. In particular looking for issues that appeared in both original LHC construction and the HL-LHC upgrade

The people who helped – The consultants

- ***I imagine that it is very unlikely that even one of the consultants would agree with all of the conclusions I present below, but I would like-to-think that most of them would agree with most of the conclusions.***

- The Phase-II constructions for ATLAS and CMS are only just starting. We can only comment on the project cycle up to the end of prototyping and preproduction.
- It will certainly be the case that the bulk production will have its own series of issues/challenges.

- ***My thanks you to the consultants. However, any errors or mistakes presented here are all mine.***

0. The Questions

- 1 What went right
- 2 What went , not so right
- 3 If you were to start again, how would you re-structure the management
- ~~4 The team and collaboration (benefits, draw backs)~~
- 5 The value of the TDR
- 6 The value of your experimental review process
- 7 The value of ongoing risk monitoring
- 8 New Concerns for the future

- The questions are only a starting-point and the interviewees were invited to take these issues wherever they want to

1. What went right

- Most of the experts were happy with where they are at this stage
- A lot of very valuable expertise was gained from the first (LHC) construction
- In many (but not all) project organisations scaled between projects (LHC → HL-LHC)
- Detector layout convergence – sub-system designs – final modularity - resilience
- Technology convergence was relatively smooth
- Good to assign groups to development/construction tasks early. Not as trivial as it might sound with frequent “over-subscription” of the more glamourous tasks.
- Important to get entire community to buy into big choices. The downside of this is that the process can be very time consuming.
- For risky technology choices, it is good to develop back-up plans from the outset.
- Development of tools/effort for ASIC verification went well.
- Development of C02 cooling systems – See later slide

2. What went not so well

- ASIC are always problematic – see additional sides for discussion
- Also slides on – procurement, vendor development, logistics and schedule.
- A global pandemic, a war in Europe, economic down-turn did not help.
- Hybrid development is always a challenge
    - Came up in most discussions with experts - a whole topic/talk in itself.
    - Problems are a mixture of electrical, mechanical, thermal, vendor …. “garbage collector”
- Putting excessive R&D time into sub-component development at risk of systems.
- Software & firmware requirements/resources (particularly firmware)
- LPGBT development – scope creep as everyone adds their particular requests
- In some cases, could/should have been more aggressive in mechanical design

3. Starting again, how to re-structure the management

- Most respondents were broadly happy with existing management structures.

- You need to remember “more often than not, we do not manage – we organise”
- Always a problem to recognise country/institutes agenda – expertise
- Countries (institutes) want/demand visibility in certain areas
- When people are “co-opted” into an activity area, supervision can be an issue

5. What is the value of the TDR

- “The TDR is a snapshot, nothing else”
- it does not represent, or very rarely, represnets the as built detector
- It is frequently too early in the project cycle. The timing is complex, but frequently revolves around releasing funding
- it is rarely used as a “how-to-build” reference
- However, it does serve some very particular useful functions
  - drawing communities together for a common endeavour
  - collects a bunch ideas in flux (frequently helping to down-select in the process),
  - It helps to start to develop major milestones
- Schedules, risks, WP organisations are often rudimentary/naive. It also establishes early funding and collaboration models.
- ***It is a real pity that there is not a similar exercise at the end of construction.***
  - *For the HL-LHC era this would seem to be more important than before.*

6. The value of your experimental review process

- Mixed response from the consultants
  - Different experiments review projects with different granularity
  - Reviews should not be confused with communication or rubber stamping
  - They can have an enormously beneficial value.
  - A lot of the value of a good review is getting the reviewees to organise their own thoughts, make plans more coherent, finalize & document ideas, connect otherwise disconnected ideas together into a system.
  - Also, important/helpful to develop and review early ideas around QA & QC procedures
  - Reviewers need real experts who can dedicate the appropriate amount of time to the review process, this including the preparation and execution.
  - Need written reports with fully conditional sign-off.
  - A fully active follow-up process is extremely important and "without it, the process is close to useless".

7. The value of ongoing Risk monitoring

- All projects, and countries/FA's maintain risk registers.
  - *It turns out that some of the peripheral (support) projects do not maintain registers*
- Most of these registers are variations around a fairly standard format
- Used for different reasons in different places, including release of contingencies.
- Very frequently viewed as a tick box exercise having limited/zero value
- Scepticism exists, and for very good reasons…we have limited success in
  - … capturing a complete set of risks
    - "we missed technical risks, risk with adhesives, movements of goods, ….."
  - … estimating the different probabilities of risk realization
  - … estimating the impact of risks being realized (finance & schedule)
- For these reasons Monte-Carlo estimates of probability of delivering the project within a given time tend to be over-optimistic

8. New Concerns for the future

- Complexities of next generation ASIC technologies
  - Designs are much-harder and costs are going to be “astronomical”
  - Fewer and fewer smaller academic institutes will be able to contribute
  - More complex developments will fall on national labs and larger universities
  - It is not assured that these institutes can find the right people with the needed skill sets
- Geopolitical issues around restrictions between countries cause concerns.

- Travel still not really back to pre-pandemic values – less face-to-face interaction
- Concern about travel in the future with carbon limitations

ASICs - I

- ASIC design and development (incl. testing, module & system level) is VERY demanding
- ASIC development is probably on most "top watch" lists for a significant fraction of the project development cycle for all projects.
- It is technically demanding and some group "over-reach" – far too ambitious in aspirations
- Concentrator chips seem to be particularly difficult and seemed to get "pushed to one side"
- Maintaining the required level of staff with the required expertise for the duration of the development is frequently very hard. The RD53 development (joint ATLAS & CMS) was a ten-year venture. This is much longer than many higher degrees, fellowships and institutional contracts. Despite this RD53 is in good shape.
- Specifications can/do creep and has caused problems
  - Physicists continue to have "great ideas" as the project goes along – continually moving goal posts
  - System architecture changing "underneath" the ASIC (example L0/L1→L0) can cause significant disruption
- Understanding process variation "peculiarities" (eg GF) changes production models
- Understanding details of the radiation tolerance and the impact on real estate has caused delays
- To get to the final design it is probably worth devoting more resource over a shorter period
- Poor coupling between developers and testers causes delays/misunderstandings.

ASICs - II

- When problems have occurred, mitigating actions have been taken
  - eg CHIPS initiative at CERN
- It took time to fully understand and exploit the full suite and power of verification tools. The verification tools demand a particular kind of expertise, not immediately available in numbers in HEP.
- Complete verification is a time-consuming process (~1 yr for RD53 chip)
- Some vendor specific “peculiarities” were well known to some parts of the ASIC community but not others, the GF – TID effects are an example.

Sensors for large areas (>100m^2)

- One of those issues that appear as issues in both HL and HL-LHC.
- HPK are a massive player in the construction of CMS and ATLAS trackers and HGCAL. Largest contract for CERN in Phase-II. They are an outstanding vendor with an impeccable heritage. HEP have a long history of working with HPK.
- But relying on only one vendor represents a significant risk of single point failure.
- Development of additional vendors in the years ahead of purchasing phase only really identified Infinion. After 5+ years of development with both ATLAS and CMS they pulled out at the start of the MS phase for "commercial reasons ". Similar things happened in the run up to LHC production (eg e2v). ATLAS and CMS worked well to recover from the situation.
- When contracts were established HPK could not meet required (integrated over all orders) production rate. Loss of schedule.
- All other sensor vendors failed at the CERN market survey stage
- Why so few? Development cycles are long with multiple cycles. Lessons include: niche vendors need their hands held. HPK learning curve was much shorter. Cultural issues!

Procurement

- The complexity of the procurement process has been underestimated.
- In many cases the physicist are going into this for the first time
- This is particularly true for procurement through CERN where the rules are strict, cumbersome and the overheads to the project are significant.
- Training is given but every particular procurement process has its own pitfalls, each of which invariably cost time and often money. Not unusual for procurements to last > 12m. Not unusual to need to stop, restructure and relaunch procurements
- Remember we build these detectors every 10+ years, and we are frequently low volume. Some vendors refuse to engage for this level of business.

Vendors in development & production

- Groups typically work closely with one or two vendors during the R&D process.
  - These is a very natural tendency to want to work with same vendors through procurement
  - This can become a real problem in procurement – need to explore full field
  - Not uncommon to be asked to relax specifications in Tender process to allow more vendors into the process
  - Can end up with un-familiar vendors – this can develop into problems

- Particle physics is frequently a very small perturbation on the flow all but the smallest business.
- Changing commercial environment following pandemic, vendors risk adverse, long lead times… vendors wanting to re-negotiate contracts.

- It can happen that business refuses to engage in production. It can happen that one can engage with a company during the R&D phase that refuse to submit returns to invitations to tender.

Schedule : development & maintenance

- A whole talk/session by itself
- The effort and scope of the task, to meet need by dates, are frequently underestimated. Ideally the schedule management has its own dedicated, full time, effort.
- A continuous struggle to maintain deadlines, float and robustness.
- The dicipline of schedule management, revision, update and change control is a continuous process for a project on the scale of an LHC experiment . Ideally the schedule should be agile and robust to changes in external circumstances. Frequenetly it is not.

Some high-level schedule generalities

- Discovering individual technical problems early is extremely beneficial. But there are a hierarchy of problems. Some problems effect every item and are easy to identify (possibly not fix), then comes the problems that happen 1 in 10 and then 1 in 100. Get to the infrequent problems as fast as possible.

- It is important to get through component problems and into system problems as quickly as possible. Don’t extend the prototyping of individual components for too long, get to primitive systems quickly and make infrastructure as realistic as possible.

Logistics

- Production of a large detectors is typically a highly distributed process with many institutes (or clusters) in multiple countries on several continents.
- The efficient, timely movement of parts between these sites is not trivial.
- The scope of the logistics work is frequently underestimated, and the required resource seems to be missing from everyone’s plans.
- Issues include, legal restrictions in several countries, contracts, shipping, incl. managing dual use electronics, trying to assure minimum or zero VAT & duties, time-lags in some countries (time to country = 5 days, time to institute inside countries = 5 days).
- Ensuring goods are traceable in the database and available just in time deliveries
- Chasing goods that go missing

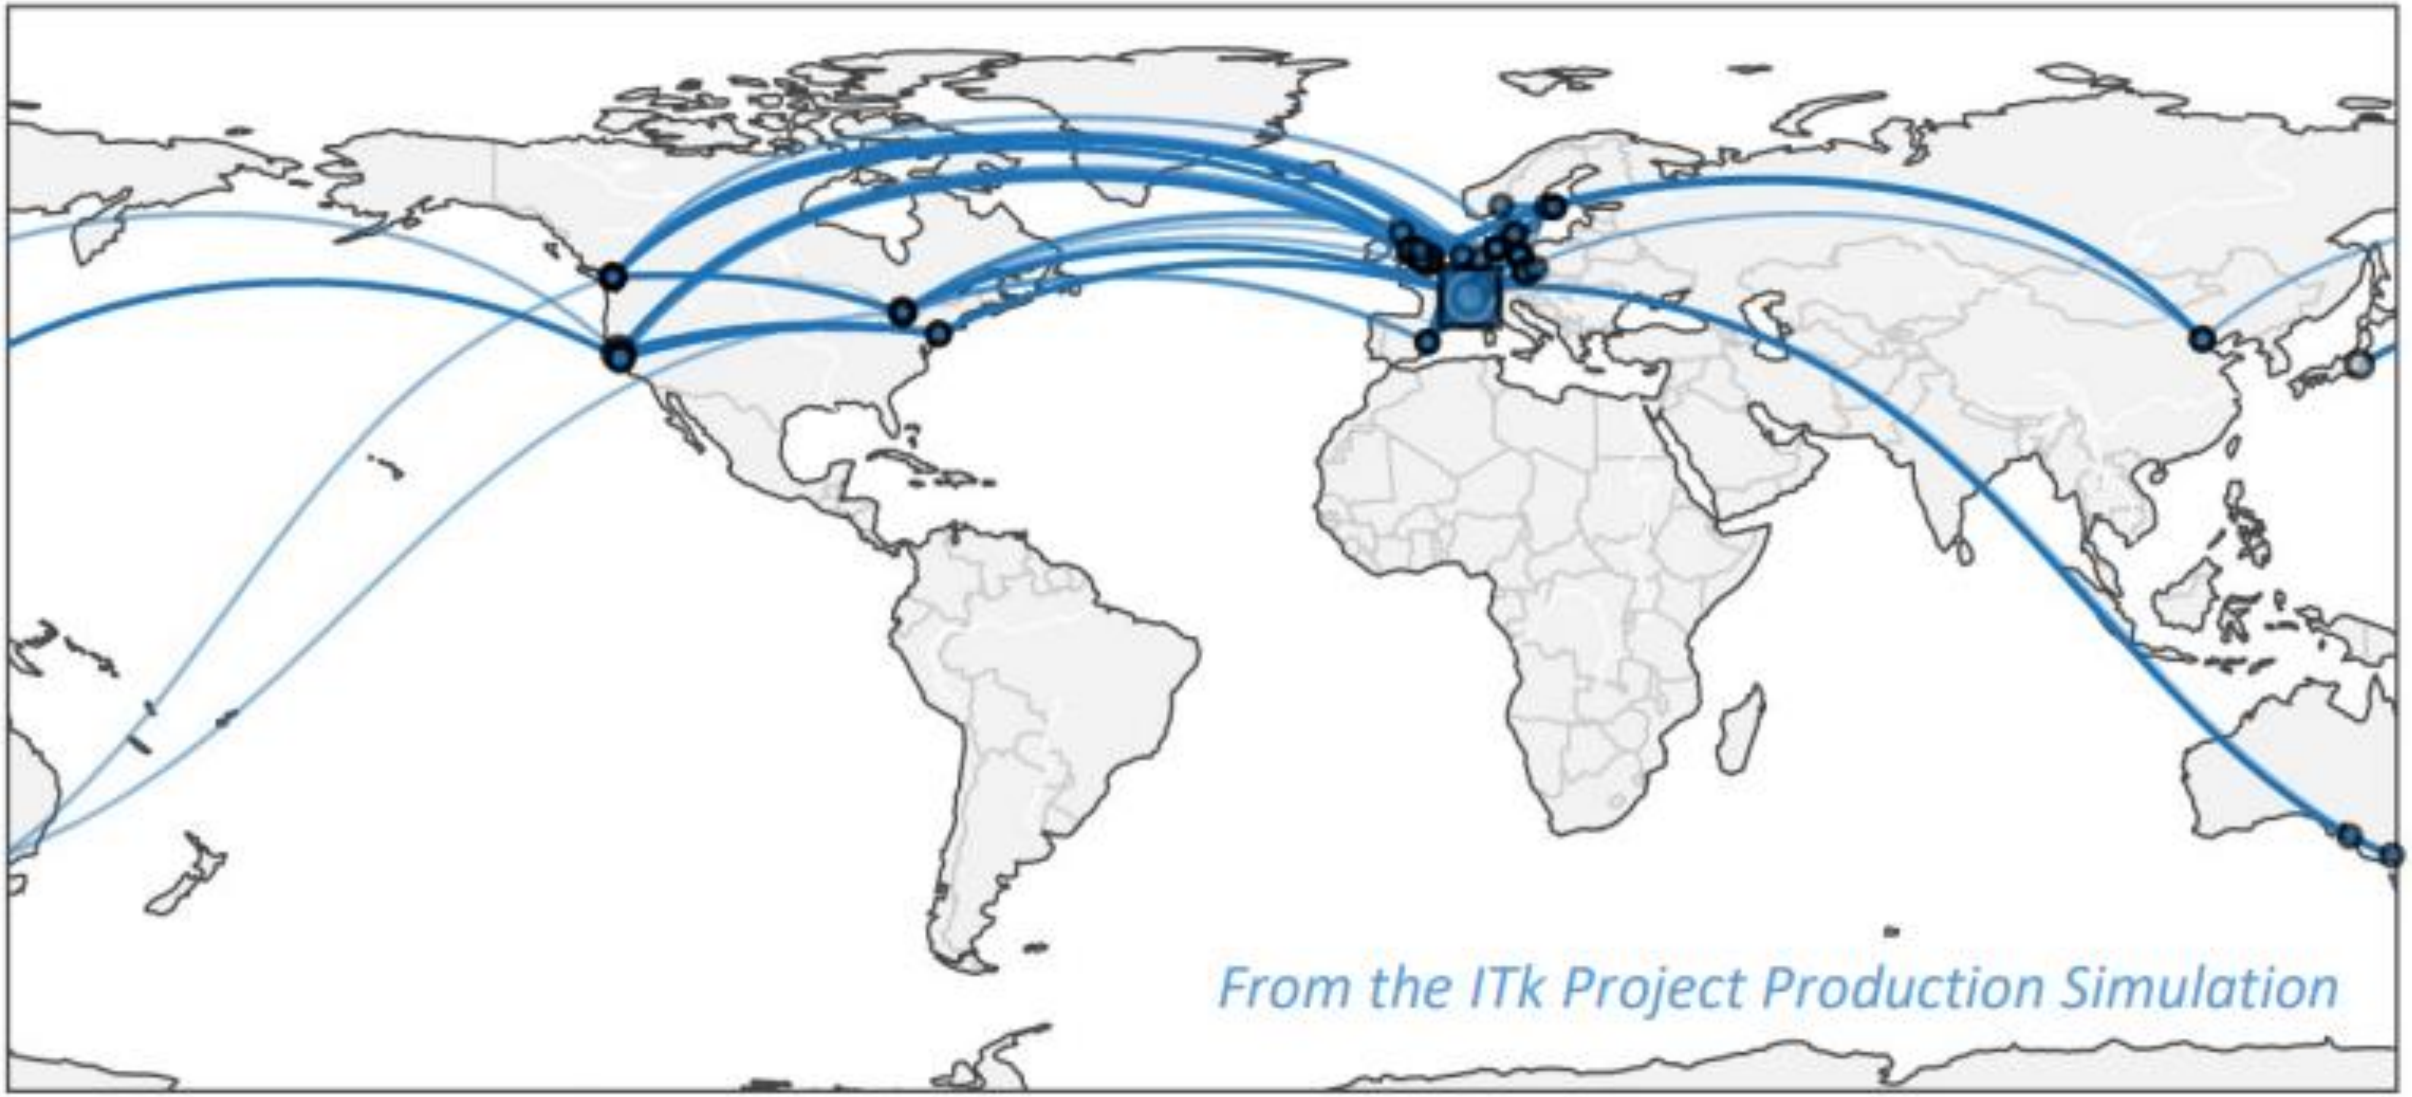
From the ITk Project Production Simulation

Production flow – HCCStar chips

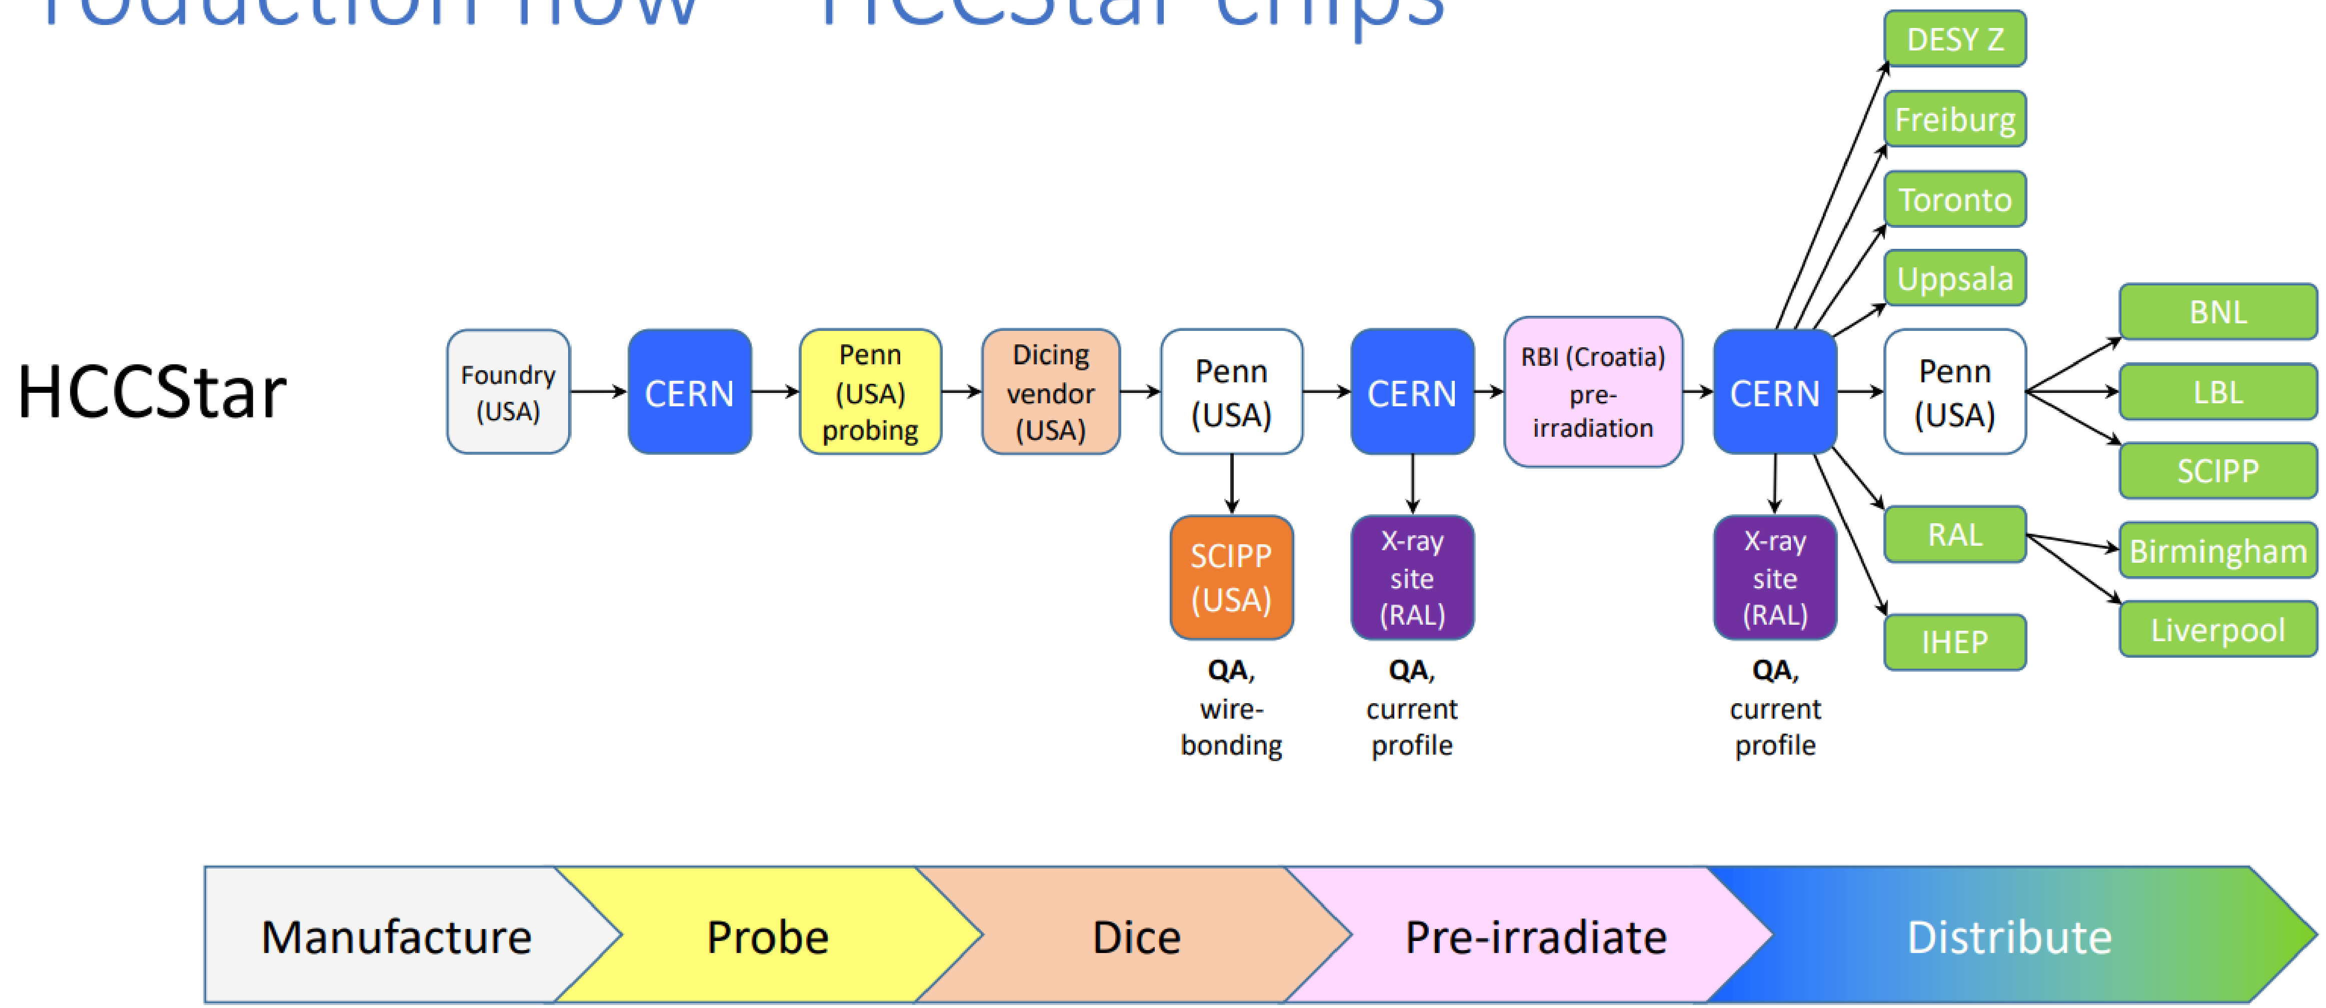

Production flow – ABCStar chips

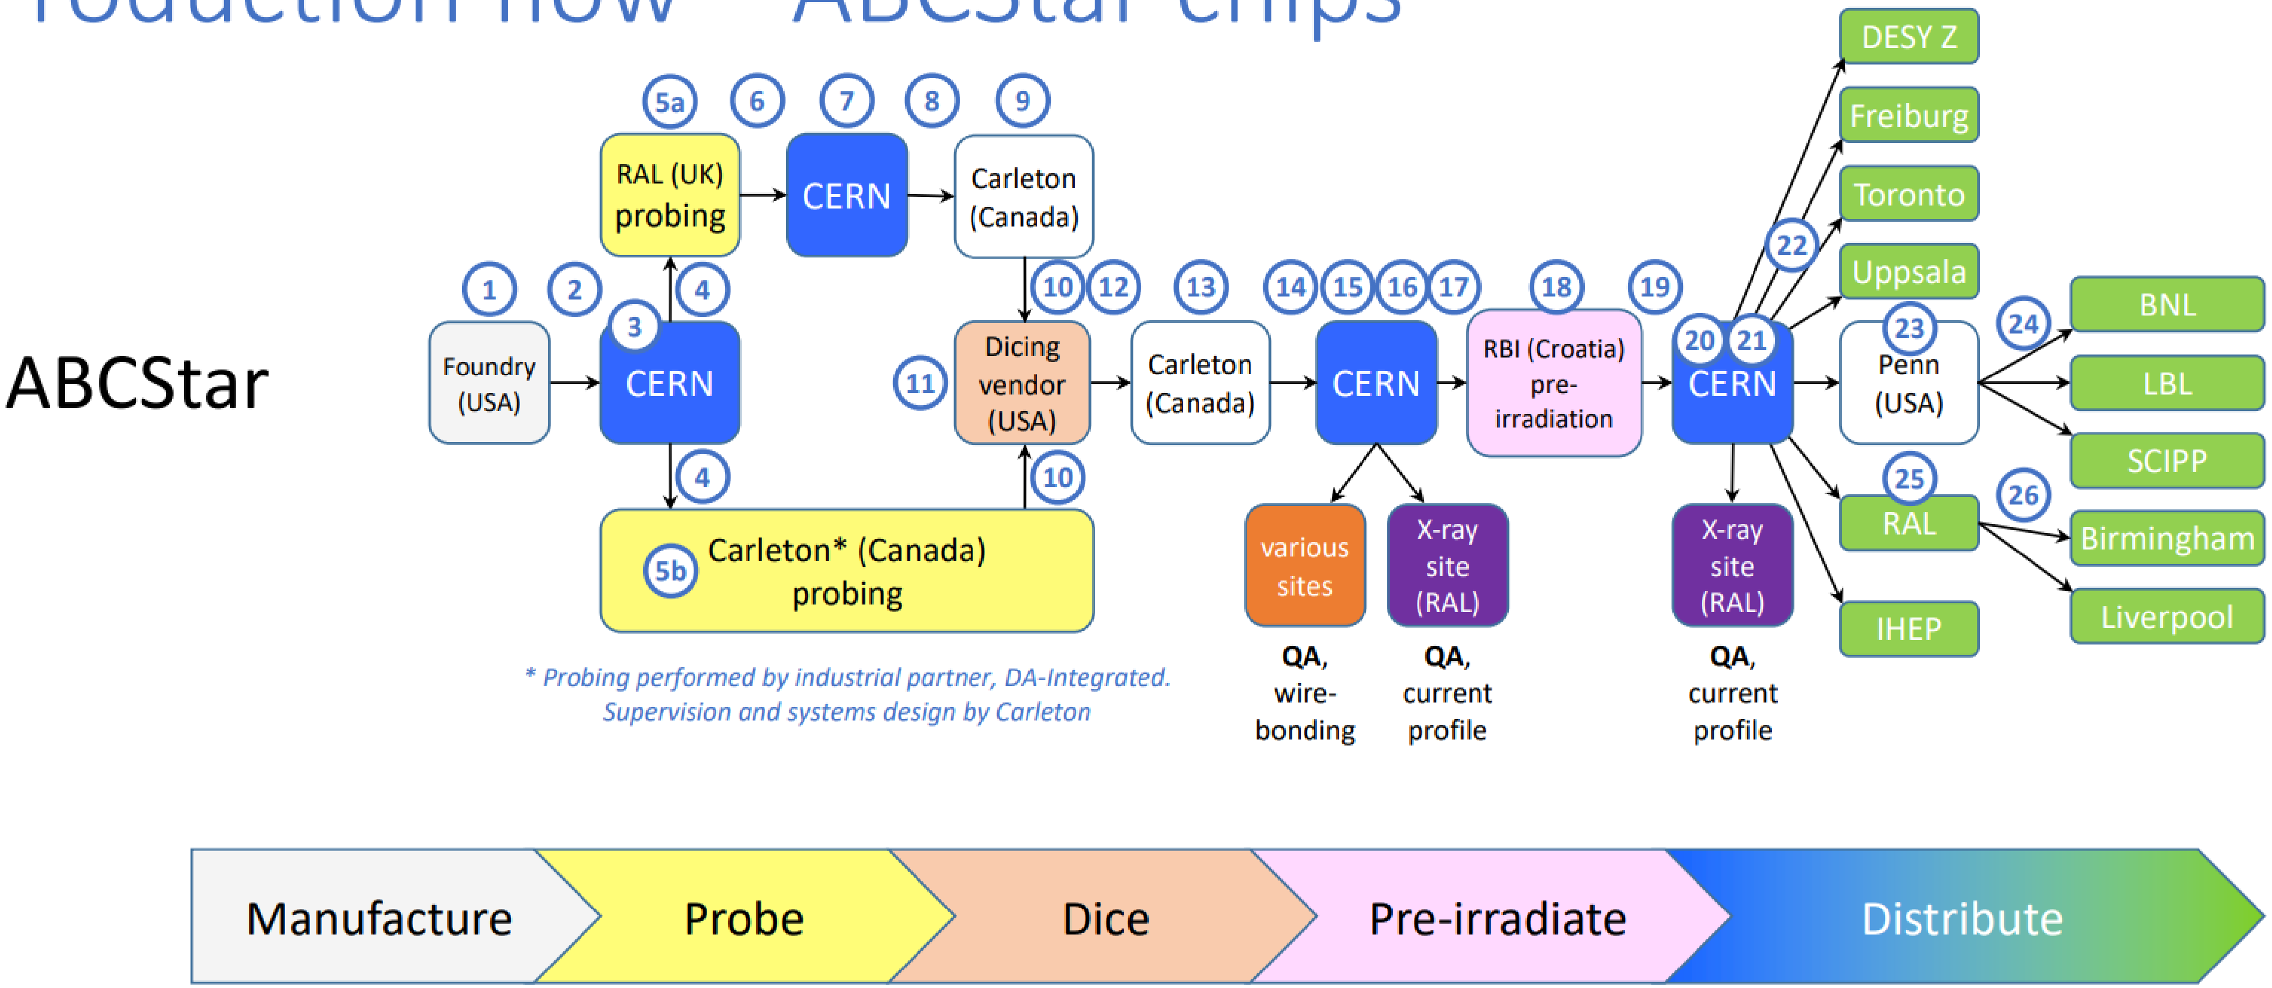

Cooling - Extra

- Transition from LHC to HL-LHC (via Phase-I upgrades) was a success for cooling effort.
- We now have the right teams, with appropriate expertise over multiple experiments. The synergies between the experiments/detectors have been exploited in an appropriate and efficient way.
- Good requirements capture and staged R&D approach to production systems has worked well (Baby DEMO, DEMO…)
- Detector side testing currently lags behind larger systems.
- Feel that TDR was of no value, too early.
- System Risk Analysis less important than Safety Analysis

Conclusions

- The construction of the trackers for the LHC experiments are large, expensive (~120MCHF core) and technically demanding projects. The collaborations that build them are enormous and highly distributed.
- We only build these detectors very infrequently, and in some senses this is not our core business. We end up learning a few things over and over.
- It is not clear when we will make the next set of detectors on this scale, but I think we can improve, we can and should take a long hard look at how we execute these projects before we do the same again. But it is not clear who we pass this message on to.