



UNIVERSITY OF IOANNINA

SCHOOL OF SCIENCES, PHYSICS DEPARTMENT

**Development of Trigger Systems and  
Algorithms for the CMS Experiment  
at the HL-LHC at CERN**

*Kosmas Adamidis*

Ioannina, 2024





ΠΑΝΕΠΙΣΤΗΜΙΟ ΙΩΑΝΝΙΝΩΝ

ΣΧΟΛΗ ΘΕΤΙΚΩΝ ΕΠΙΣΤΗΜΩΝ, ΤΜΗΜΑ ΦΥΣΙΚΗΣ

Ανάπτυξη συστημάτων και  
αλγορίθμων σκανδαλισμού για το  
πείραμα CMS στον επιταχυντή  
HL-LHC του CERN

*Κοσμάς Αδαμίδης*

Ιωάννινα, 2024



### **Three-member Advisory Committee**

- Konstantinos Foudas, Professor, Department of Physics, University of Ioannina, Greece (Thesis Supervisor).
- Ioannis Evangelou, Emeritus Professor, Department of Physics, University of Ioannina, Greece
- Nikolas Manthos, Emeritus Professor, Department of Physics, University of Ioannina, Greece

### **Seven-member Assessment Committee**

- Konstantinos Foudas, Professor, Department of Physics, University of Ioannina, Greece
- Ioannis Evangelou, Emeritus Professor, Department of Physics, University of Ioannina, Greece
- Nikolas Manthos, Emeritus Professor, Department of Physics, University of Ioannina, Greece
- Michail Bachtis, Assistant Professor, Department of Physics and Astronomy, University of Los Angeles, California, USA
- Vasilis Christofilakis, Assistant Professor, Department of Physics, University of Ioannina, Greece
- Ioannis Papadopoulos, Associate Professor, Department of Physics, University of Ioannina, Greece
- Yiorgos Tsiatouhas, Professor, Department of Computer Science & Engineering, University of Ioannina, Greece



εκ πάντων εν  
και εξ ενός τα πάντα

Ηράκλειτος





# Acknowledgements

The work described in this thesis would not be conducted without the assistance of my supervisor, Costa Fouda, to whom I am grateful.

I owe a special thank to Greg Iles, who gave me detailed explanations to many technical aspects of the Trigger system during the course of this work.

I am lucky to have worked around the people of CMS experiment. The interaction with many of them has contributed to the conclusion of this work. If I was to name some of them it would be Michalis Bachtis, Tom Williams, Ignacio Redondo and Alvaro Navarro.

During the times of this work the presence of Yannis Bestintzanos, both as a colleague and as a friend, has proven to be priceless.

The journey of the past years is concluded thanks to people who accompanied me in every aspect of life. I was blessed to have had Paris and Polydamas as friends and lab partners, as well as having Dimitris, Kostas and Stratos, whose support from distance was equal to having them next to me.

Lastly and most importantly, I want to thank my parents and family for being there during every little step towards the conclusion of this work.



# Abstract

The CMS experiment is being prepared for the Phase-2 era, when the High-Luminosity LHC (HL-LHC) will start its operations delivering more than 7 times the nominal LHC Luminosity. The aim of the CMS upgrade is both to maintain but also improve the detector performance, in order to extend the discovery reach of the detector. The amount and density of data produced by HL-LHC demands new detector systems, replacement of the majority of the on-detector electronics, and complete replacement of the Trigger system. The work of this thesis is part of the CMS Level-1 Trigger upgrade. The first part of the work was devoted to the development of an optical link protocol, operating at 25 Gbps and targeting FPGAs produced by Xilinx. The protocol will be the standard Phase-2 protocol, used by the Level-1 Trigger processors for their internal data exchange. A detailed description of its firmware implementation, as well as results of extended testing are presented. Moreover, the work here includes firmware developments that target the Barrel Muon Trigger Layer-1 subsystem, responsible for the generation of track segments of muons that cross the barrel region of CMS. An ATCA card based on a VU13P FPGA was designed to instrument this system. The design of certain modules of this card, as well as the development of its firmware infrastructure and the slice tests that demonstrated its performance using data from proton-proton collisions are presented in the second part of this thesis.



# Εκτεταμένη Περίληψη

## Καθιερωμένο Πρότυπο Στοιχειωδών Σωματιδίων

Οι ανακαλύψεις που επιτεύχθηκαν στον τομέα της σωματιδιακής φυσικής τους περασμένους αιώνες έχουν ως αποτέλεσμα την δημιουργία του Καθιερωμένου Προτύπου των στοιχειωδών σωματιδίων (Standard Model). Σε αυτό εμπεριέχεται το σύνολο της γνώσης μας σήμερα σχετικά με τους δομικούς λίθους της ύλης του σύμπαντος και των αλληλεπιδράσεων που πραγματοποιούνται μεταξύ τους. Αυτές αποτελούν στην ουσία τις δυνάμεις της φύσης, οι οποίες, όπως τις γνωρίζουμε σήμερα, είναι οι εξής: η ηλεκτρομαγνητική με φορέα το φωτόνιο, η ασθενής πυρηνική με φορέα τα  $W^\pm$  και  $Z$ , η ισχυρή πυρηνική με φορέα το γλουόνιο και η βαρυτική, ο φορέας της οποίας εικάζεται πως είναι το βαρυτόνιο. Τα αδιαίρετα σωματίδια που απαρτίζουν την ύλη γύρω μας χωρίζονται σε δύο βασικές κατηγορίες, τα λεπτόνια και τα κουάρκ. Για κάθε κατηγορία υπάρχουν έξι σωματίδια, τα οποία επίσης χωρίζονται σε τρεις γενιές. Έτσι, ως κουάρκ έχουμε ονομάσει τα: up ( $u$ ) και down ( $d$ ), charm ( $c$ ) και strange ( $s$ ), top ( $t$ ) και bottom ( $b$ ). Το ηλεκτρικό τους φορτίο συναντάται είτε στην τιμή  $2/3$  ή  $-1/3$  και η μάζα τους από μερικά MeV έως εκατοντάδες GeV. Τα κουάρκ διαθέτουν επίσης το φορτίο της ισχυρής δύναμης, το αναφερόμενο ως χρώμα, το οποίο μπορεί να είναι μπλε, πράσινο ή κόκκινο. Οι ενώσεις των κουάρκ φτιάχνουν νέα σωματίδια, τα αδρόνια και τα μεσόνια, τα οποία είναι άχρωμα. Η δεύτερη κατηγορία σωματιδίων, τα λεπτόνια, αποτελούνται από το ηλεκτρόνιο ( $e$ ) και το νεutrino του ηλεκτρονίου ( $\nu_e$ ), το μιονίο ( $\mu$ ) και το νεutrino του μιονίου ( $\nu_\mu$ ) και το τάο ( $\tau$ ) με το νεutrino του τάο ( $\nu_\tau$ ). Το ηλεκτρικό φορτίο του ηλεκτρονίου, μιονίου και τάο είναι  $-1$  και των αντίστοιχων νεutrίνων  $0$ . Η μάζα των πρώτων κυμαίνεται από εκατοντάδες KeV μέχρι μερικά GeV, ενώ αυτή των νεutrίνων είναι απειροελάχιστη. Σε κάθε ένα από αυτά τα σωματίδια αντιστοιχεί το αντισωματίδιό του, που έχει παρόμοιες ιδιότητες εκτός από το φορτίο, το οποίο έχει ίδια αλλά αντίθετη τιμή. Η τελευταία προσθήκη στο Καθιερωμένο Πρότυπο είναι αυτή του σωματιδίου Higgs ( $H$ ), το οποίο ανακαλύφθηκε στο CERN και ανακοινώθηκε το καλοκαίρι του 2012. Το Higgs είναι μποζόνιο με spin  $0$  και μάζα περίπου  $125$  GeV. Είναι το σωματίδιο φορέας του αντίστοιχου πεδίου Higgs, με το οποίο αλληλεπιδρούν τα σωματίδια ύλης και παίρνουν τη μάζα τους.

## Ο Μεγάλος Επιταχυντής Αδρονίων

Ο Μεγάλος Επιταχυντής Αδρονίων (LHC) αποτελεί τον μεγαλύτερο και ισχυρότερο επιταχυντή σωματιδίων σήμερα, ο οποίος λειτουργεί στον Ευρωπαϊκό Οργανισμό Πυρηνικών Ερευνών (CERN) στην Γενεύη της Ελβετίας. Πρόκειται για έναν κυκλικό

επιταχυντή με ακτίνα περίπου 27 χλμ, αποτελούμενο από δύο σωλήνες κενού μέσα στο οποίο κινούνται ομάδες πρωτονίων σε ταχύτητες πολύ κοντά σε αυτή του φωτός. Σε τέσσερα σημεία της ακτίνας του κύκλου οι δύο σωλήνες διασχίζουν ο ένας τον άλλο έτσι ώστε τα επιτευχθούν συγκρούσεις μεταξύ των σωματιδίων. Η μέγιστη ενέργεια του κέντρου μάζας των συγκρούσεων είναι 14 TeV και η ονομαστική τιμή της φωτεινότητας είναι  $L = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ . Ο βασικός σκοπός της κατασκευής του επιταχυντή ήταν η δημιουργία και ανίχνευση του σωματιδίου του Higgs. Ο σκοπός αυτός επιτεύχθηκε κατά της πρώτη φάση της λειτουργίας, από το 2010 έως τις αρχές του 2013. Η ανακάλυψη του σωματιδίου έγινε στα δύο μεγάλα πειράματα γενικού σκοπού, το ATLAS και το CMS. Από τότε, ο επιταχυντής συνεχίζει την λειτουργία του με σκοπό την περαιτέρω μελέτη του μηχανισμού Higgs, καθώς και την λεπτομερή μελέτη των διαφόρων αλληλεπιδράσεων του Καθιερωμένου Προτύπου. Επιπλέον, η αναζήτηση εστιάζει σε αναπάντητα ερωτήματα της Φυσικής στα οποία η ανακάλυψη οποιουδήποτε νέου σωματιδίου θα συνεισφέρει τα μέγιστα. Μερικά από αυτά είναι η προέλευση της σκοτεινής ύλης, η ασυμμετρία ύλης και αντιύλης στο σύμπαν, και άλλα.

Η επιτάχυνση πρωτονίων γίνεται σε ομάδες, οι οποίες διέρχονται από μια σειρά από μικρότερους επιταχυντές μέχρι να φτάσουν στην ενέργεια των 450 GeV. Έπειτα εισέρχονται στα δύο δαχτυλίδια του Μεγάλου Επιταχυντή Αδρονίων. Εκεί επιταχύνονται μέσω ηλεκτρικών πεδίων που λειτουργούν στην συχνότητα των 400.78 MHz και διατηρούν την καμπυλότητα τους μέσω ισχυρών υπεραγωγικών μαγνητών. Τα δύο δαχτυλίδια έχουν χωρητικότητα 3564 ομάδων το καθένα, αλλά ο τελικός αριθμός που υποστηρίζεται είναι κοντά στις 2808.

Η λειτουργία του επιταχυντή πραγματοποιείται σε φάσεις, οι οποίες περιλαμβάνουν κάποια χρόνια συγκρούσεων και κάποια χρόνια παύσης αυτών, έτσι ώστε τα μηχανήματα να συντηρηθούν ή/και να αναβαθμιστούν. Η πρώτη φάση ξεκίνησε με την πρώτη περίοδο συλλογής από το 2010 έως το 2013. Ακολούθησε παύση μέχρι το 2016, όταν και ξεκίνησε η δεύτερη περίοδος συλλογής μέχρι το 2018. Η δεύτερη παύση διήρκεσε σχεδόν 3 χρόνια ώσπου ξεκίνησε η τρίτη περίοδος συλλογής δεδομένων, την οποία και διανύουμε (2023). Με το τέλος της τρίτης περιόδου συλλογής δεδομένων τελειώνει και η πρώτη φάση λειτουργίας του επιταχυντή. Η δεύτερη του φάση προγραμματίζεται να ξεκινήσει το 2029, έπειτα από ένα κενό τριών χρόνων, όταν και θα επέλθει μεγάλη αναβάθμιση του επιταχυντή και των πειραμάτων του. Η νέα μηχανή, επονομαζόμενη ως Μεγάλος Επιταχυντής Αδρονίων Υψηλής Φωτεινότητας (HL-LHC), θα διατηρήσει τη μέγιστη ενέργεια λειτουργίας του υπάρχοντος αλλά θα είναι σε θέση να προκαλέσει μεγαλύτερο αριθμό συγκρούσεων ανά διάσχιση ομάδων. Η φωτεινότητα του νέου θα είναι σε θέση να φτάσει τα  $L = 7.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$  κατά την πλήρη λειτουργία της.


## Το πείραμα CMS

Το πείραμα CMS (Συμπαγής Σωληνοειδή Μιονίων) είναι ένα από τα δύο μεγάλα πειράματα που λειτουργούν στον LHC. Πρόκειται για έναν ανιχνευτή γενικού σκοπού, βασικός στόχος του οποίου ήταν η ανακάλυψη του σωματιδίου Higgs. Οι μετέπειτα έρευνές του συμπίπτουν με αυτές του LHC και εστιάζονται στη περαιτέρω μελέτη του Καθιερωμένου Προτύπου και την αναζήτηση στοιχείων που μπορεί περιέχουν ίχνη φυσικής πέρα από αυτό. Οι συγκρούσεις πρωτονίων λαμβάνουν χώρα στο κέντρο του ανιχνευτή κάθε 25 ns. Ο όγκος δεδομένων που παράγεται είναι αδύνατο να αποθηκευτεί άμεσα, επομένως

επεξεργάζεται επί τόπου από το λεγόμενο σύστημα Σκανδαλισμού, μέχρι να μειωθεί σε επίπεδα που επιτρέπουν την αποθήκευσή του σε δίσκους. Κατά το τέλος της πρώτης φάσης του LHC, μέρος των ανιχνευτικών συστημάτων του CMS, η πλειοψηφία των ηλεκτρονικών πάνω στον ανιχνευτή, καθώς και η ολότητα του συστήματος Σκανδαλισμού πρόκειται να αντικατασταθούν. Το αναβαθμισμένο πείραμα θα είναι σε θέση να ανταπεξέλθει στις αυξημένες απαιτήσεις του HL-LHC αλλά και να αυξήσει την αποτελεσματικότητά του.

Ο ανιχνευτής έχει κυλινδρικό σχήμα και χωρίζεται στην περιοχή του βαρελιού, και την περιοχή των δύο καπακιών. Η δομή του, ξεκινώντας από το κέντρο όπου λαμβάνουν χώρα οι συγκρούσεις και προχωρώντας προς τα έξω, έχει ως εξής: Πιο κοντά στην περιοχή συγκρούσεων βρίσκεται ο Tracker (Καταγραφέας Τροχιών). Το σύστημα της πρώτης φάσης πρόκειται να αντικατασταθεί από εντελώς νέο, το οποίο θα παρέχει περισσότερα και μικρότερα κανάλια ανάγνωσης. Η δομή του βασίζεται σε στρώσεις από αισθητήρες πυριτίου, στους οποίους φορτισμένα σωματίδια αφήνουν ηλεκτρικό ίχνος. Η πληροφορία αυτή χρησιμοποιείται στην ανακατασκευή των τροχιών φορτισμένων σωματιδίων. Έξω από την Tracker είναι τοποθετημένα τα καλοριμετρα του ανιχνευτή. Το κομμάτι του βαρελιού αποτελείται από δύο διαφορετικούς τέτοιους ανιχνευτές. Αρχικά, το Ηλεκτρομαγνητικό Καλορίμετρο (ECAL) είναι κατασκευασμένο από ράβδους βολφραμικού μολύβδου ( $PbWO_4$ ) και απορροφά ηλεκτρόνια και φωτόνια που παράγονται στις συγκρούσεις. Έπειτα βρίσκεται το Αδρονικό Καλορίμετρο (HCAL) το οποίο απορροφά τα αδρόνια παράγωγα των συγκρούσεων. Είναι κατασκευασμένο από επίπεδες πλάκες απορροφητικού ορείχαλκου εναλλασσόμενες από πλάκες πλαστικού σπινθηριστή. Τα δύο καλορίμετρα απορροφούν την ενέργεια των σωματιδίων που τα διαπερνούν και παρέχουν μετρήσεις σχετικά με την θέση και την ενέργεια των σωματιδίων αυτών. Οι ανιχνευτές του βαρελιού θα παραμείνουν ως έχει στην δεύτερη φάση που πειράματος αλλά το αντίστοιχο καλορίμετρο των καπακιών αποτελεί ένα εξ ολοκλήρου νέο σύστημα. Ονομάζεται Καλορίμετρο Υψηλής Λεπτομέρειας (HGCal) και αποτελείται από 52 στρώσεις ανιχνευτών πυριτίου ανάμεσα σε στρώσεις απορροφητικού υλικού.

Τα ανιχνευτικά συστήματα που αναφέρθηκαν έως τώρα στην περιοχή του βαρελιού είναι εγκατεστημένα μέσα στον όγκο ενός μεγάλου σωληνοειδούς πηνίου. Το πηνίο αυτό αποτελεί έναν ιδιαίτερα ισχυρό μαγνήτη, το πεδίο του οποίου προσεγγίζει τα 4 Tesla. Αποτελεί βασικό κομμάτι της λειτουργίας του πειράματος, καμπυλώνοντας τις τροχιές των φορτισμένων σωματιδίων. Με τον τρόπο αυτό επιτυγχάνεται ο έμμεσος υπολογισμός της ορμής των φορτισμένων σωματιδίων, καθώς και ο διαχωρισμός τους με αυτά που δεν κουβαλάνε φορτίο. Τα μόνα σωματίδια που φεύγουν έξω από τον μαγνήτη και μπορούν να ανιχνευτούν από τον CMS είναι τα μίονια. Για τον λόγο αυτό, οι ανιχνευτές μιονίων βρίσκονται έξω από το σωληνοειδές. Στην περιοχή του βαρελιού υπάρχουν τέσσερις στρώσεις από θαλάμους Drift Tube (DT) και στην περιοχή των καπακιών αντίστοιχες στρώσεις από θαλάμους Cathode Strip Chamber (CSC). Η λειτουργία και των δύο βασίζεται στην αλληλεπίδραση των μιονίων με το αέριο μέσα στους θαλάμους, η οποία παράγει ηλεκτρόνια τα οποία συλλέγονται, υποδηλώνοντας ότι ένα σωματίδιο άφησε το ίχνος του. Συνδυάζοντας ίχνη από τους σταθμούς των θαλάμων γίνεται η ανακατασκευή των τροχιών τους και ο προσδιορισμός της ορμής τους. Μια σχηματική απεικόνιση του ανιχνευτή CMS φαίνεται στην Εικόνα 1.



Σχήμα 1: Σχηματική απεικόνιση του ανιχνευτή CMS.

## Τεχνολογία του Συστήματος Σκανδαλισμού Επιπέδου-1 της δεύτερης φάσης του CMS

Το Επίπεδο-1 του Συστήματος Σκανδαλισμού του CMS είναι κατασκευασμένο από κάρτες ηλεκτρονικών σχεδιασμένες από τις διάφορες ομάδες του πειράματος. Οι κάρτες που προορίζονται για την δεύτερη φάση λειτουργίας βασίζονται στο εργαστηριακό πρότυπο ATCA (Advanced Telecommunications and Computing Architecture). Αυτό καθορίζει τις ακριβείς διαστάσεις των καρτών οι οποίες λειτουργούν μέσα σε ένα κιβώτιο, επίσης καθορισμένο. Το κιβώτιο αυτό παρέχει την τροφοδοσία κάθε κάρτας, βασικά σήματα διαδικτύου και άλλα σήματα που ορίζονται από τους χρήστες. Κάθε κάρτα, επομένως, πρέπει να είναι συμβατή με αυτές τις προδιαγραφές.

Το βασικό χαρακτηριστικό των καρτών του Σκανδαλισμού Επιπέδου-1 είναι η επεξεργασία και μεταφορά δεδομένων σε πολύ υψηλές ταχύτητες. Αυτό επιτυγχάνεται με την χρήση συστοιχίων επιτόπια προγραμματιζόμενων πυλών, τα λεγόμενα FPGA. Τα ολοκληρωμένα αυτά κυκλώματα αποτελούνται από τις βασικές δομές που χρησιμοποιούνται για την κατασκευή ψηφιακών κυκλωμάτων, όπως είναι οι λογικές πύλες, οι καταχωρητές, οι μνήμες, τα ρολόγια, και άλλα. Ένα FPGA εμπεριέχει μεγάλο αριθμό από τέτοια στοιχεία, συνδεδεμένα μεταξύ τους με επαναπρογραμματιζόμενες διασυνδέσεις. Με αυτό τον τρόπο ο χρήστης μπορεί να σχεδιάσει και να εφαρμόσει ψηφιακά κυκλώματα της επιλογής του. Ο σχεδιασμός τους πραγματοποιείται με την χρήση γλωσσών περιγραφής υλικού, όπως είναι η γλώσσα VHDL. Η περιγραφή των συνδέσεων και της λειτουργίας ενός κυκλώματος που αποτυπώνεται μέσω των γλωσσών αυτών μεταφράζεται από κατάλληλα λογισμικά στις τελικές συνδέσεις μέσα στο FPGA. Οι συνδέσεις αυτές καταλήγουν να εκτελούν τη λειτουργία που τους όρισε ο χρήστης, έχοντας το πλεονέκτημα του εύκολου επαναπρογραμματισμού. Τα FPGA που επρόκειτο να χρησιμοποιηθούν στο πείραμα



κατασκευάζονται από την εταιρία Xilinx.

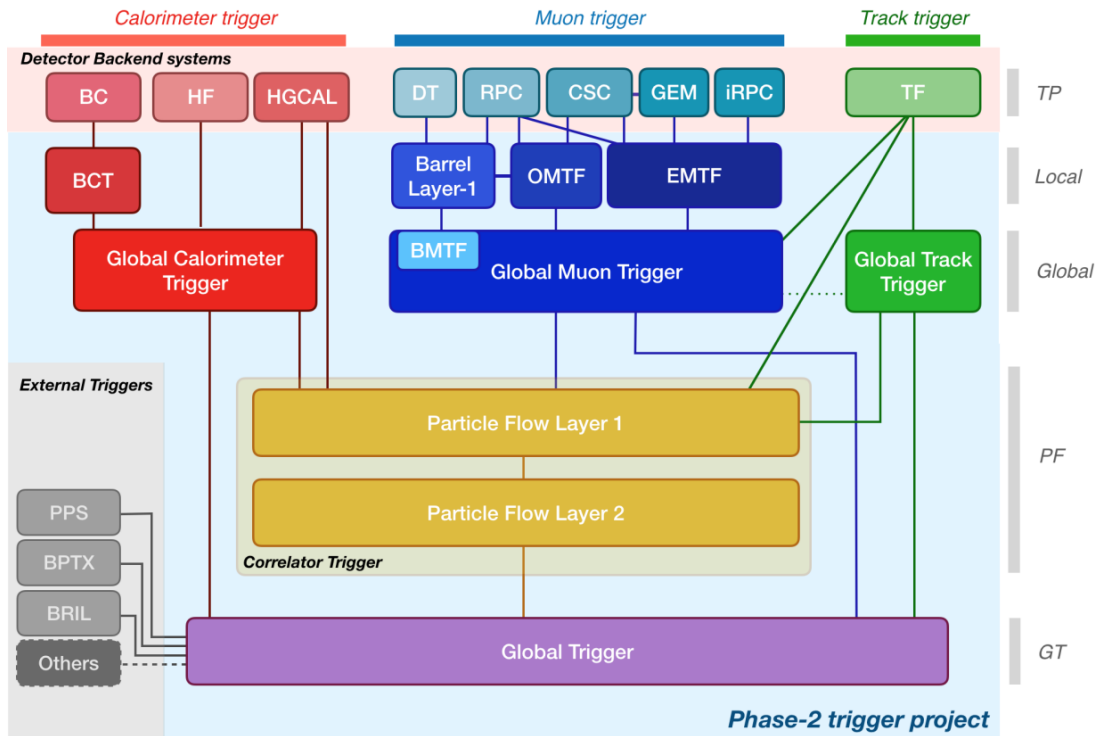
Το κομμάτι της μεταφοράς δεδομένων διεξάγεται με την χρήση οπτικών ινών, έτσι ώστε να επιτευχθούν οι μεγαλύτερες δυνατές ταχύτητες μετάδοσης. Η μεταφορά πραγματοποιείται από της συσκευές πομποδεκτών (transceivers) οι οποίες αποτελούν στοιχεία των FPGA. Οι πομποδέκτες αυτοί μετατρέπουν ψηφιακές λέξεις παράλληλης μορφής σε ακολουθίες σειριακών συμβόλων. Το σειριακό σήμα συνδέεται σε συσκευές μετάδοσης οι οποίες βρίσκονται πάνω στην κάρτα και μετατρέπουν το ηλεκτρικό σήμα σε οπτικό, ώστε να μεταδοθεί μέσω των οπτικών ινών. Δυο τέτοιες συσκευές που χρησιμοποιούνται στο πείραμα είναι τα λεγόμενα QSFP και τα αντίστοιχα Firefly.

## Το Σύστημα Σκανδαλισμού της δεύτερης φάσης του CMS

Το Σύστημα Σκανδαλισμού του CMS αποτελείται από το Επίπεδο-1 και τον Σκανδαλισμό Υψηλού Επιπέδου. Το πρώτο Επίπεδο λαμβάνει πληροφορία από κάθε σύγκρουση, την οποία επεξεργάζεται για να αποφασίσει εάν τα παράγωγα του κάθε γεγονότος εμπεριέχουν χρήσιμο υλικό. Η πληροφορία αυτή σε πολλές περιπτώσεις δεν στέλνεται στη μέγιστη ανάλυση της λόγω πρακτικών περιορισμών. Τα πλήρη δεδομένα όμως παραμένουν αποθηκευμένα σε μνήμες έως ότου το Επίπεδο-1 λάβει κάποια απόφαση. Εάν αποφασίσει πως το συγκεκριμένο γεγονός είναι χρήσιμο, στέλνει ένα σήμα αποδοχής (Accept) στις μνήμες ώστε να μεταφέρουν την πληροφορία στον Σκανδαλισμό Υψηλού Επιπέδου. Η χρονική διάρκεια μέσα στην οποία πρέπει να παρθεί απόφαση είναι 12.5  $\mu$ s και έχει σκοπό να μειώσει τον ρυθμό των δεδομένων από 40 MHz στα 750 KHz. Το σύστημα που μεταφέρει τα αποδεκτά γεγονότα ονομάζεται σύστημα Απόκτησης Δεδομένων (DAQ).

Το σχεδιάγραμμα του Σκανδαλισμού Επιπέδου-1 φαίνεται στο Σχήμα 2. Για κάθε ανιχνευτή του CMS υπάρχει και το αντίστοιχο υποσύστημα Σκανδαλισμού το οποίο επεξεργάζεται την αντίστοιχη πληροφορία. Κατά αυτό τον τρόπο υπάρχουν: το μονοπάτι για τα καλορίμετρα (Calorimeter Datapath), το μονοπάτι για τα Μιόνια (Muon Datapath) και το μονοπάτι για τον καταγραφέα τροχιών (Tracker Datapath). Οι έξοδοι αυτών των συστημάτων μεταφέρουν ανεξάρτητα σωματίδια (ηλεκτρόνια, φωτόνια, αδρόνια, μιόνια) και τροχιές στο λεγόμενο Particle Flow και στον Γενικό Σκανδαλιστή (Global Trigger). Το σύστημα Particle Flow τρέχει τον ομώνυμο αλγόριθμο ώστε να ανακατασκευάσει κάθε γεγονός και να αναγνωρίσει όλα τα σωματίδια του προέρχονται από την λεγόμενη αρχική κορυφή της σύγκρουσης. Η έξοδος του επίσης καταλήγει στον Γενικό Σκανδαλιστή, ο οποίος λαμβάνει την τελική απόφαση αποδοχής ή απόρριψης κάθε γεγονότος.

Τα δεδομένα που αποδέχεται το πρώτο Επίπεδο Σκανδαλισμού αποστέλλονται στο επόμενο επίπεδο επεξεργασίας μέσω της κάρτας DTH. Σε αυτήν στέλνει τα αποδεκτά δεδομένα το κάθε υποσύστημα Σκανδαλισμού. Τουλάχιστον μια τέτοια κάρτα είναι εγκατεστημένη σε κάθε κιβώτιο ATCA και σκοπός της είναι να λάβει τα δεδομένα και να τα αποστείλει στον Σκανδαλισμό Υψηλού Επιπέδου. Επιπλέον, είναι υπεύθυνη να διανείμει σημαντικά σήματα ελέγχου στα οποία βασίζεται η λειτουργία του πειράματος και όλων των καρτών του. Για παράδειγμα, το σήμα Αποδοχής μεταδίδεται μέσω αυτής, καθώς και το ρολόι του επιταχυντή, το λεγόμενο LHC clock.



Σχήμα 2: Σχηματικό διάγραμμα της αρχιτεκτονικής του Συστήματος Σκανδαλισμού Επιπέδου-1.

## Το Σύστημα Σκανδαλισμού Μιονίων του βαρελιού

Το σύστημα Σκανδαλισμού Μιονίων του βαρελιού είναι αυτό που ανακατασκευάζει μόνια τα οποία διασχίζουν το βαρέλι του CMS. Η διαδικασία ξεκινάει από την περιοχή των ανιχνευτών, από όπου οι κάρτες που είναι τοποθετημένες εκεί διαβάζουν και ψηφιοποιούν τον χρόνο της διέλευσης ενός μιονίου μέσα από τον κάθε θάλαμο DT. Οι κάρτες αυτές ονομάζονται OBDT και στέλνουν δεδομένα από τον ανιχνευτή μέσω μεγάλων οπτικών ινών στο δωμάτιο όπου βρίσκεται το Σύστημα Σκανδαλισμού Επιπέδου-1. Η κάρτα αυτή είναι ανθεκτική στην ακτινοβολία των συγκρούσεων, το ίδιο και το οπτικό πρωτόκολλο αποστολής δεδομένων. Το πρώτο στάδιο επεξεργασίας πραγματοποιείται από το σύστημα Σκανδαλισμού Μιονίων του Βαρελιού Επιπέδου-1 (BMTL1). Το σύστημα αυτό επεξεργάζεται δεδομένα από κάθε θάλαμο DT ξεχωριστά και ανεξάρτητα. Σκοπός του είναι να παράξει κομμάτια τροχιών, ή στάμπες, τα οποία υποδηλώνουν την τοποθεσία από όπου πέρασε ένα σωματίδιο στον θάλαμο και την γωνία στρέψης του. Το σύστημα BMTL1 τρέχει σε κάρτες που ονομάζονται BMTL1 ATCA. Είναι βασισμένες σε ένα ολοκληρωμένο VU13P FPGA, μεγάλο αριθμό εισόδων για να λαμβάνει οπτικές διασυνδέσεις από τον ανιχνευτή και μικρότερο αριθμό εξόδων οι οποίες στέλνουν στάμπες στο επόμενο σύστημα επεξεργασίας. Αυτό είναι το λεγόμενο GMT (Global Muon Trigger). Σκοπός αυτού του υποσυστήματος είναι να συγκεντρώσει μόνια από όλο τον ανιχνευτή και να τα στέλνει στα επόμενα υποσυστήματα του Σκανδαλισμού. Επιπλέον, τρέχει τον αλγόριθμο ο οποίος επεξεργάζεται τα κομμάτια τροχιών του βαρελιού και ανακατασκευάζει τροχιές υποψηφίων μιονίων. Το σύστημα αυτό τρέχει στις κάρτες X20, των οποίων η επεξεργαστική μονάδα είναι επίσης ένα VU13P FPGA.

## Υλικολογισμικό υποδομής

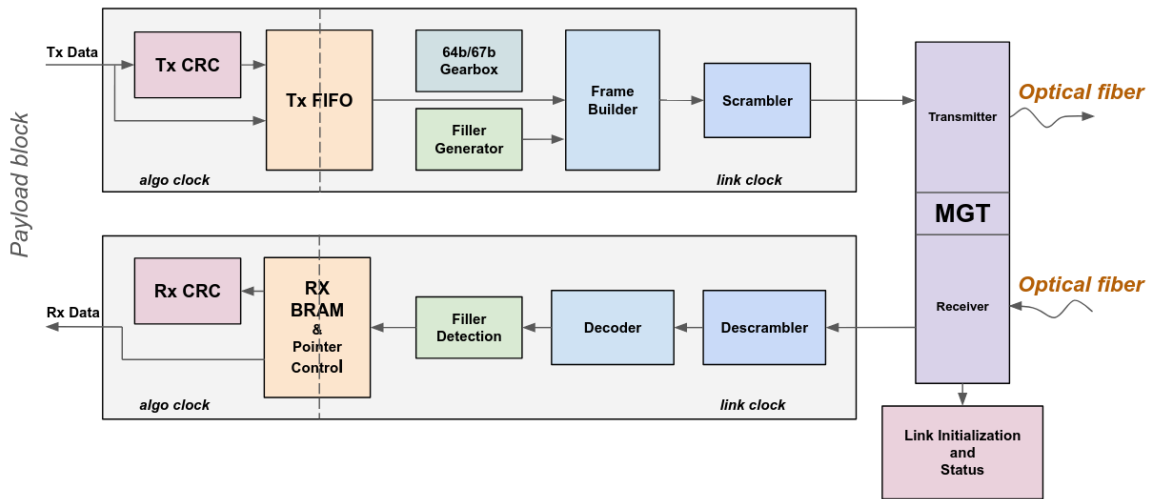
Το υλικολογισμικό υποδομής αποτελεί κομμάτια κώδικα VHDL τα οποία είναι κοινά στους επεξεργαστές Σκανδαλισμού και παρέχουν τις απαραίτητες λειτουργίες ώστε μία κάρτα να είναι σε θέση να λειτουργεί στο πείραμα. Ένα από τα υλικολογισμικά τέτοιου τύπου είναι το λεγόμενο EMP Framework. Τα βασικά κυκλώματα τα οποία περιέχει είναι τα εξής. Το λεγόμενο TTC block το οποίο λαμβάνει και διανέμει όλα τα βασικά σήματα του πειράματος, όπως το σήμα Αποδοχής και το LHC clock. Το IPbus block, το οποίο υλοποιεί το ομώνυμο πρωτόκολλο που χρησιμοποιείται για τον έλεγχο των καρτών μέσω λογισμικού υψηλού επιπέδου. Το Datapath block που υλοποιεί τα πρωτόκολλα επικοινωνίας οπτικών ινών μέσα στο FPGA και παρέχει μνήμες όπου δεδομένα μπορούν να γραφούν και να διαβαστούν για δοκιμαστικούς σκοπούς. Τέλος, το λεγόμενο Algorithm block είναι αυτό μέσα στο οποίο υλοποιούνται οι αλγόριθμοι του κάθε συστήματος. Ως τέτοιο, η ανάπτυξη κώδικα σε αυτό εκτελείται από τις ομάδες που χρησιμοποιούν την κάθε κάρτα.

Υπάρχουν δύο κατηγορίες οπτικών πρωτοκόλλων που χρησιμοποιούνται στο πείραμα. Αυτά που επικοινωνούν με τον ανιχνευτή και λειτουργούν σύγχρονα με το ρολόι του LHC και αυτά με τα οποία αποστέλλεται πληροφορία μεταξύ των επεξεργαστών Σκανδαλισμού και λειτουργούν ασύγχρονα σε σχέση με το ρολόι του LHC. Το σύγχρονο πρωτόκολλο που θα χρησιμοποιηθεί στην δεύτερη φάση του πειράματος λέγεται IpGBT. Στη διαδρομή από την ανιχνευτή προς τον σύστημα Σκανδαλισμού η ταχύτητα μεταφοράς του είναι είτε 5.12 Gbps ή 10.24 Gbps. Στην αντίθετη διαδρομή είναι 2.56 Gbps. Τα δεδομένα που αποστέλλονται με αυτό κωδικοποιούνται με τρόπο ώστε κατά την αποκωδικοποίησή τους να είναι εφικτή η διόρθωση λαθών που έχουν προκύψει κατά τη διάρκεια της μεταφοράς.

Ένα ασύγχρονο πρωτόκολλο που αναπτύχθηκε για χρήση στο EMP Framework είναι το λεγόμενο Hermes. Το Hermes λειτουργεί σε ταχύτητες των 16 Gbps ή των 25 Gbps. Έχει μικρή καθυστέρηση και είναι εύκολο προς στη χρήση. Αποστέλλει δύο τύπους λέξεων των 64 bit, τα Δεδομένα (Data) και τις λέξεις ελέγχου (Control words). Για να διαχωρίσει τις δύο μεταξύ τους χρησιμοποιεί το σύστημα κωδικοποίησης 64b/67b, το οποίο για κάθε λέξη των 64 bit στέλνει και μία επικεφαλίδα των 3 bit. Η ασύγχρονη λειτουργία επιτυγχάνεται χρησιμοποιώντας δύο ξεχωριστές θέσεις μνήμης FIFO, μία στον πομπό και μία στον δέκτη. Από τη μία τους μεριά αυτές λειτουργούν με το ρολόι του LHC και από την άλλη με το ρολόι του πομποδέκτη. Για να εξακριβώσει ότι δεν έχουν υπάρξει αλλοιώσεις των δεδομένων κατά την μεταφορά, το Hermes εφαρμόζει την τεχνική των CRC (Cyclic Redundancy Check), με την οποία είναι σε θέση να καταλάβει εάν έχουν υπάρξει λάθη κατά την μεταφορά. Το σχηματικό διάγραμμα του πρωτοκόλλου φαίνεται στο Σχήμα 3.

## Το οπτικό πρωτόκολλο CSP

Το CSP (CMS Standard trigger link Protocol) είναι το οπτικό πρωτόκολλο που θα χρησιμοποιηθεί στην δεύτερη φάση λειτουργίας του πειράματος από όλους του επεξεργαστές Σκανδαλισμού. Ορίζει τους κανόνες αποστολής και λήψης δεδομένων μέσα από οπτικές ίνες στην ταχύτητα των 25.78125 Gbps. Η ανάπτυξή του είναι το αποτέλεσμα της ένωσης δύο πρωτοκόλλων που αναπτύχθηκαν σαν προτάσεις από δύο διαφορετι-

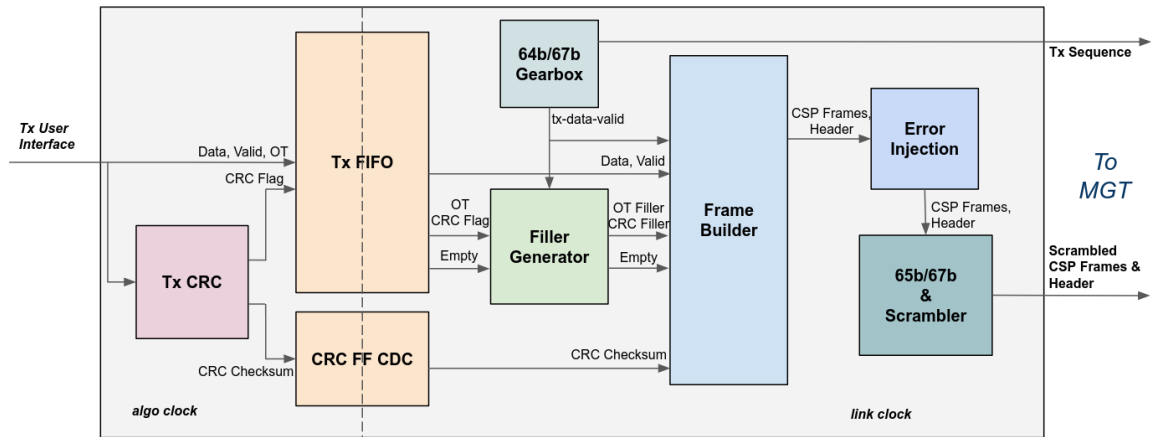


Σχήμα 3: Σχηματικό διάγραμμα του πρωτοκόλλου Hermes.

κές ομάδες του πειράματος. Το ένα από αυτά είναι το Hermes, όπως περιγράφηκε στο προηγούμενο κεφάλαιο. Η τελική εκδοχή του πρωτοκόλλου Σκανδαλισμού, του CSP, υιοθετεί τα σημαντικότερα χαρακτηριστικά των δύο σε ένα. Η κατάληξη σε ένα κοινό πρωτόκολλο για το σύστημα ήταν απαραίτητη, μιας και όλες οι κάρτες του πειράματος θα πρέπει να είναι σε θέση να επικοινωνήσουν με όλες τις υπόλοιπες.

Το CSP, όπως και το Hermes, διαχειρίζεται δύο είδη λέξεων, τα Δεδομένα και τις λέξεις Ελέγχου. Επιπλέον, είναι βασισμένο στην ασύγχρονη αρχιτεκτονική η οποία υλοποιείται με την χρήση FIFO. Η αρχιτεκτονική αυτή διαχωρίζει το διαθέσιμο εύρος ζώνης σε δύο μέρη: το μέρος των Δεδομένων και το Πληρωτικό μέρος. Το μέρος των Δεδομένων χρησιμοποιείται αποκλειστικά για την αποστολή λέξεων που περιέχουν πληροφορία του ανιχνευτή. Το Πληρωτικό εύρος ζώνης αξιοποιείται για την αποστολή πληροφορίας ελέγχου, όπως είναι οι λέξεις ελέγχου λαθών CRC, η ετικέτα τροχιάς και διάφορες πληροφορίες της ταυτότητας της εκάστοτε διασύνδεσης. Η αποστολή των παραπάνω πληροφοριών πραγματοποιείται μέσω της χρήσης διαφορετικών λέξεων Ελέγχου. Αυτές διαχωρίζονται μεταξύ τους από ένα πεδίο τεσσάρων bit. Το πεδίο αυτό, όπως και η Ετικέτα των 3 bit που αποστέλλεται με κάθε λέξη, κωδικοποιούνται σε κώδικες διόρθωσης λαθών. Αυτό συμβαίνει ώστε να αποφευχθεί ο λαθεμένος διαχωρισμός των λέξεων του πληρωτικού εύρους ζώνης με λέξεις που ανήκουν στο εύρος ζώνης των δεδομένων.


Το κύκλωμα του CSP αποτελείται από τα μονοπάτια του αποστολέα και του παραλήπτη. Το σχηματικό διάγραμμα του αποστολέα φαίνεται στο Σχήμα 4. Πριν τα δεδομένα εισέλθουν στο πεδίο ρολογιού του πομπού, διέρχονται από το κύκλωμα του CRC ώστε να παραχθεί η αντίστοιχη λέξη ελέγχου λαθών. Έπειτα, μέσω της μνήμης FIFO, μεταβαίνουν στο πεδίο του πομπού, όπου ο χτίζονται οι λέξεις CSP, εφαρμόζοντας τους κανόνες του πρωτοκόλλου. Κατά την έξοδό τους από το μονοπάτι του αποστολέα γράφονται στις αντίστοιχες πόρτες της συσκευής πομποδέκτη ώστε να μετατραπούν σε σειριακό σήμα και να αποσταλούν.



Σχήμα 4: Σχηματικό διάγραμμα του κυκλώματος αποστολέα του CSP.


## Η κάρτα επεξεργασίας BMTL1

Η κάρτα επεξεργασίας BMTL1 ATCA επρόκειτο να χρησιμοποιηθεί στο υποσύστημα BMTL1. Έχει σχεδιαστεί και κατασκευαστεί κατά κύριο λόγο από την ομάδα του Πανεπιστημίου Ιωαννίνων και η πρώτη κάρτα που παράχθηκε παραλήφθηκε τον Απρίλιο του 2022. Η κάρτα φαίνεται στο Σχήμα 5. Αφού ολοκληρώθηκαν οι πρώτες δοκιμές, οι οποίες επιβεβαίωσαν ότι όλα τα κυκλώματα υλικού πάνω της λειτουργούν ορθά, ακολούθησαν δοκιμές με την χρήση του VU13P FPGA, το οποίο αποτελεί την κεντρική μονάδα επεξεργασίας της.



Σχήμα 5: Η κάρτα BMTL1 ATCA με οπτικές διασυνδέσεις.

Βασικό κομμάτι των δοκιμών αυτών αποτέλεσε η λήψη των λεγόμενων διαγραμμάτων ματιού, τα οποία επαληθεύουν την λειτουργία τόσο των ρολογιών αναφοράς που υπάρχουν στην κάρτα, όσο και των οπτικών διασυνδέσεων όταν η μεταφορά γίνεται μέσα από οπτικές ίνες. Τα αποτελέσματα των ματιών για τα 40 κανάλια που λειτουργούν στην ταχύτητα των 25 Gbps φαίνονται στο Σχήμα 6. Από το άνοιγμα της μπλε περιοχής των διαγραμμάτων συμπεραίνουμε πως το ηλεκτρικό σήμα των σειριακών δεδομένων που εισέρχονται στο αναλογικό μέρος του πομποδέκτη είναι εξαιρετικά ποιοτικό.



**Σχήμα 6:** Διαγράμματα ματιού των οπτικών διασυνδέσεων που λειτουργούν στην ταχύτητα των 25 Gbps.

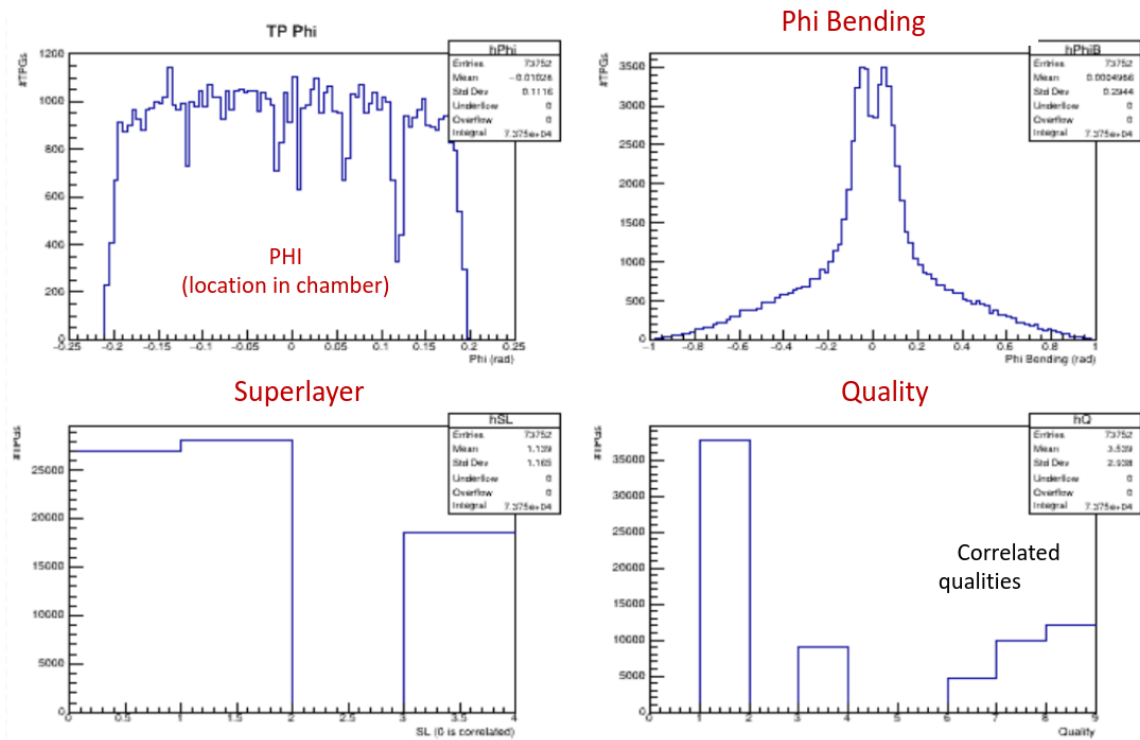
Με την ολοκλήρωση των λειτουργικών δοκιμών της κάρτας, και αφού αποφάνθη πως λειτουργεί όπως ήταν αναμενόμενο, αναπτύχθηκε το υλικολογισμικό υποδομής της κάρτας. Αυτό βασίζεται στο ήδη υπάρχον EMP Framework, στο οποίο πραγματοποιήθηκαν όλες οι απαραίτητες τροποποιήσεις ώστε να μπορεί να λειτουργεί στην κάρτα BMTL1. Μετέπειτα, εισήχθη στο υλικολογισμικό της κάρτας ο αλγόριθμος Analytical Method. Ο αλγόριθμος αυτός είναι υπεύθυνος για την επεξεργασία των δεδομένων που έρχονται από τους θαλάμους Drift Tube των ανιχνευτών μιονίων του βαρελίου. Το αποτέλεσμα της επεξεργασίας τους έχει την μορφή ιχνών από σωματίδια που διασχίζουν τους θαλάμους. Αποστέλλεται με την χρήση οπτικών διασυνδέσεων των 25 Gbps στο επόμενο στάδιο επεξεργασίας, που είναι το λεγόμενο Global Muon Trigger. Στο υποσύστημα αυτό τα δεδομένα χρησιμοποιούνται για την ανακατασκευή των τροχιών μιονίων.

## **Δοκιμές με υλικό του συστήματος Σκανδαλισμού Μιονίων του Βαρελίου**

Το σύστημα Σκανδαλισμού Μιονίων του Βαρελίου της δεύτερης φάσης του πειράματος δοκιμάστηκε στην περιοχή ενσωμάτωσης των Drift Tube, καθώς και υπό συνθήκες

αληθινών συγκρούσεων. Αρχικά, δοκιμές έγιναν σε όλη την αλυσίδα υλικού, που αποτελείται από αληθινό Drift Tube (DT) chamber, μία κάρτα OBDT, την κάρτα BMTL1 ATCA και την κάρτα OCEAN του συστήματος GMT. Ο θάλαμος DT, εγκατεστημένος στην επιφάνεια του πειράματος, παρήγαγε δεδομένα από κοσμικά μίονια που διέσχισαν την επιφάνειά του. Τα δεδομένα αυτά αποστέλονταν στην κάρτα OBDT, η οποία τα μετέτρεπε σε πληροφορία TDC και τα έστειλε στην κάρτα BMTL1. Εκεί, επεξεργάζονταν από τον αλγόριθμο Analytical Method με σκοπό την παραγωγή σταμπών μιονίων, τα οποία και αποστέλονταν στην κάρτα OCEAN με την χρήση οπτικών ινών. Η κάρτα αυτή κατέγραφε τα δεδομένα που ελάμβανε, με σκοπό την μετέπειτα επεξεργασία τους.

Έπειτα από την επιτυχή εκτέλεση των παραπάνω δοκιμών στην επιφάνεια του πειράματος, το σύστημα κατέβηκε κάτω, στην αίθουσα ελέγχου του CMS. Εκεί εγκαταστήθηκε ξανά και συνδέθηκε με 13 OBDT κάρτες, οι οποίες βρίσκονται σε έναν από τους τομείς μιονίων του πειράματος. Αυτές χρησιμοποιήθηκαν για την αποστολή δεδομένων που παρήχθησαν από αληθινές συγκρούσεις πρωτονίων στο κέντρο του ανιχνευτή. Τα δεδομένα αυτά λήφθηκαν από την κάρτα BMTL1, η οποία τα επεξεργάστηκε και έστειλε τα αποτελέσματα στην OCEAN για την καταγραφή τους. Διαγράμματα από τα αποθηκευμένα αυτά δεδομένα φαίνονται στο Σχήμα 7. Το διάγραμμα πάνω αριστερά απεικονίζει την θέση των μιονίων, καθώς διέσχισαν έναν από τους θαλάμους του τομέα. Παρατηρούμε ότι η κατανομή τους στην επιφάνεια του θαλάμου είναι επίπεδη, όπως θα περίμενε κανείς λόγω της τυχαιότητας στην θέση που παράγονται τα μίονια. Το πάνω δεξιά διάγραμμα απεικονίζει την γωνία στρέψης των μιονίων σε σχέση με το κέντρο του θαλάμου. Παρατηρούμε ότι η κατανομή είναι συμμετρική γύρω από την κεντρική τιμή, το οποίο είναι αναμενόμενο. Επίσης, οι δύο κορυφές αποτελούν τα δύο φορτία των μιονίων, καθώς αυτά στρίβουν σε αντίθετες κατευθύνσεις λόγω του μαγνητικού πεδίου του ανιχνευτή.



Σχήμα 7: Διαγράμματα από ίχνη μιονίων που παρήχθησαν από αληθινές συγκρούσεις μιονίων.



# Contents

## Abstract

<b>Εκτεταμένη Περίληψη</b>	<b>i</b>
<b>1 Introduction</b>	<b>6</b>
1.1 The Standard Model of Particle Physics . . . . .	6
1.2 Discovery beyond Standard Model . . . . .	8
<b>2 Large Hadron Collider</b>	<b>9</b>
2.1 Injection chain and bunch structure . . . . .	10
2.2 Luminosity . . . . .	11
2.3 Experiments . . . . .	13
2.4 High Luminosity LHC . . . . .	15
<b>3 The CMS Detector</b>	<b>17</b>
3.1 Introduction . . . . .	17
3.1.1 CMS coordinate system . . . . .	19
3.2 Superconducting Solenoid . . . . .	19
3.3 Tracker System . . . . .	20
3.3.1 Phase-1 Tracker . . . . .	21
3.3.2 Phase-2 Tracker . . . . .	21
3.4 Calorimeters . . . . .	24
3.4.1 Electromagnetic Calorimeter Barrel . . . . .	24
3.4.2 Hadronic Calorimeter Barrel . . . . .	25
3.4.3 High Granularity Calorimeter . . . . .	26
3.5 Muon System . . . . .	28
3.5.1 Drift Tube system . . . . .	29
3.5.2 Cathode Strip Chambers . . . . .	30
3.5.3 Resistive Plate Chamber system . . . . .	31
3.5.4 Improved RPC . . . . .	32
3.5.5 Gas Electron Multiplier detectors . . . . .	33
3.6 MIP Timing Detector . . . . .	33

<b>4</b>	<b>Technology of Phase-2 Level-1 Trigger and Data Acquisition systems</b>	<b>34</b>
4.1	ATCA Standard . . . . .	34
4.2	Field Programmable Gate Arrays . . . . .	36
4.2.1	Architecture . . . . .	36
4.2.2	Hardware Description Languages . . . . .	37
4.2.3	FPGAs at CMS . . . . .	38
4.3	Multi Gigabit Transceivers . . . . .	39
4.3.1	IBERT tool . . . . .	42
4.3.2	Optical Transceiver modules . . . . .	43
<b>5</b>	<b>CMS Phase-2 Trigger and Data Acquisition System</b>	<b>44</b>
5.1	Level-1 Trigger . . . . .	45
5.1.1	Calorimeter Trigger . . . . .	45
5.1.2	Muon Trigger . . . . .	47
5.1.3	Track Trigger and Global Track Trigger . . . . .	50
5.1.4	Correlator Trigger . . . . .	51
5.1.5	Global Trigger . . . . .	52
5.1.6	Scouting System . . . . .	53
5.1.7	Hardware . . . . .	53
5.2	Data Acquisition and High Level Trigger . . . . .	57
5.2.1	DAQ system Overview . . . . .	57
5.2.2	DTH ATCA board . . . . .	59
5.2.3	Data to Surface and Event Builder . . . . .	60
5.2.4	High Level Trigger . . . . .	61
<b>6</b>	<b>Barrel Muon Trigger</b>	<b>62</b>
6.1	Overview . . . . .	62
6.2	Front-End Electronics . . . . .	63
6.2.1	Drift Tube Front-End . . . . .	63
6.2.2	RPC Front-End . . . . .	65
6.3	Barrel Muon Trigger Layer-1 . . . . .	65
6.3.1	Analytical Method algorithm . . . . .	66
6.3.2	Hardware . . . . .	67
6.4	Global Muon Trigger . . . . .	68
6.4.1	Kalman Muon Track Finder algorithm . . . . .	68
6.4.2	Hardware . . . . .	70
6.5	Architecture . . . . .	70
<b>7</b>	<b>Firmware Infrastructure for Phase-2</b>	<b>74</b>
7.1	Introduction . . . . .	74
7.2	EMP Firmware Framework . . . . .	75

7.2.1	TTC block . . . . .	75
7.2.2	IPbus Block . . . . .	76
7.2.3	Readout Block . . . . .	77
7.2.4	Datapath Block . . . . .	77
7.2.5	Algorithm Block . . . . .	79
7.2.6	Declaration File . . . . .	79
7.2.7	Constraint Files . . . . .	80
7.2.8	EMP Software . . . . .	81
7.3	Front-End Optical Protocols . . . . .	81
7.3.1	GBT . . . . .	82
7.3.2	lpGBT . . . . .	82
7.4	Hermes: A Back-End Optical Link Protocol . . . . .	83
7.4.1	Encoding Layer . . . . .	84
7.4.2	Asynchronous Architecture . . . . .	84
7.4.3	Framing Layer . . . . .	85
7.4.4	Data transmission modes . . . . .	86
7.4.5	Cyclic redundancy check . . . . .	87
7.4.6	Link Alignment . . . . .	88
7.4.7	Protection Mechanism . . . . .	89
<b>8</b>	<b>The CSP Trigger Link protocol</b>	<b>90</b>
8.1	Protocol syntax . . . . .	90
8.1.1	Index Number . . . . .	91
8.1.2	Control Word Payload . . . . .	91
8.1.3	Data Integrity . . . . .	92
8.1.4	Scrambling logic . . . . .	94
8.2	Transmitter Datapath . . . . .	94
8.2.1	Transmission modes - User Interface . . . . .	95
8.2.2	CRC generation . . . . .	95
8.2.3	Data Domain Crossing . . . . .	95
8.2.4	Synchronous Gearbox . . . . .	96
8.2.5	Deterministic Filler Generation . . . . .	96
8.2.6	Frame Builder . . . . .	97
8.2.7	Error Injection . . . . .	97
8.2.8	Scrambler - 65b/67b implementation . . . . .	97
8.3	Receiver Datapath . . . . .	98
8.3.1	Descrambler . . . . .	99
8.3.2	Hamming Decoder . . . . .	99
8.3.3	Link Initialization and Alignment . . . . .	100
8.3.4	Filler detection and BRAM control . . . . .	101

8.3.5	Index Correction Mechanism . . . . .	102
8.3.6	Domain Crossing . . . . .	102
8.3.7	CRC Check . . . . .	102
8.3.8	User Interface . . . . .	103
8.4	Control and Status registers . . . . .	103
8.5	Software Interface and Configuration . . . . .	104
8.6	Firmware Performance . . . . .	105
8.6.1	Final CSP Testing . . . . .	106
<b>9</b>	<b>The BMTL1 ATCA Trigger Processor</b>	<b>108</b>
9.1	Introduction . . . . .	108
9.2	Hardware Design . . . . .	108
9.2.1	High Speed Serial Links . . . . .	110
9.2.2	Clocking Network . . . . .	110
9.3	Hardware Testing . . . . .	112
9.3.1	Optical links . . . . .	113
9.3.2	ZYNQ to FPGA interface . . . . .	116
9.3.3	Zone 2 connections . . . . .	116
9.4	EMP Framework at BMTL1 board . . . . .	118
9.4.1	BMTL1 In and Out ports . . . . .	119
9.4.2	FPGA to ZYNQ IPbus interface . . . . .	121
9.4.3	Declaration Files . . . . .	122
9.4.4	TCDS2 interface . . . . .	123
9.4.5	QPLL based lpGBT . . . . .	124
9.4.6	lpGBT In CSP Out Hybrid links . . . . .	125
9.4.7	Constraint files . . . . .	127
9.4.8	Building the project . . . . .	127
9.5	Analytical Method algorithm Integration . . . . .	130
9.5.1	Single-Chamber Integration . . . . .	130
9.5.2	Sector Implementation . . . . .	133
<b>10</b>	<b>Barrel Muon Trigger Slice Tests</b>	<b>136</b>
10.1	SX5 Single Chamber Slice Test . . . . .	136
10.1.1	Setup description . . . . .	136
10.1.2	Validation of optical interfaces . . . . .	138
10.1.3	Configuring the slice test . . . . .	139
10.1.4	Results . . . . .	139
10.2	USC Multi Chamber Slice Test . . . . .	140
10.2.1	The USC setup . . . . .	141
10.2.2	Configuring the slice test . . . . .	142

10.2.3 Results . . . . .	143
<b>11 Conclusions</b>	<b>147</b>
<b>A CSP Control and Status Registers</b>	<b>148</b>
A.1 Channel Control Registers . . . . .	148
A.2 Common Control Registers . . . . .	151
A.3 Channel Status Registers . . . . .	152
A.4 Common Status Registers . . . . .	155

# Chapter 1

## Introduction

The urge of discovery can be considered as a fundamental characteristic of human beings. It can be traced back to millennia and is one of the main drive forces that leads our evolution through space and time. The discovery domains can contain anything from new lands and unknown territories of our planet, to laws and society systems that assist the existence of humans in large groups. One domain where discovery is of crucial importance is the study of the universe and its physical phenomena, as conducted by the field of Physics. Physics tries to give a scientific explanation to the structure and behaviour of the world around us, from the very small (micro) up to the very large (macro). The first theories, as we know, that relate to the microcosm have been expressed about 2,500 years ago, when Democritus introduced the concept of indivisible matter that he called *atom*. Since then, and especially during the last 200 years, physicists are constantly in search of the "truly indivisible" building blocks of the universe, as well as the interactions between them. This process has led to the creation of the Standard Model of Particle Physics, the model that describes our best knowledge of the elementary particles and their fundamental forces, up to day.

### 1.1 The Standard Model of Particle Physics

The Standard Model (SM) is being developed since the early 1970s, as the result of discoveries of all the more elementary particles. Today, it represents the core of our understanding in the particle physics field, describing the properties of elementary particles and their interactions. SM provides a consistent theoretical description for three out of the four known fundamental forces of the universe: electromagnetic, weak and strong. The fourth force is gravity, which has not been incorporated within the SM context yet. However, the effects of gravity are far weaker than those of the other three interactions and can be neglected. The predictions of SM at energy scales that are accessible by our technology today are remarkably successful, even though it lacks of explanation for a plethora of phenomena, such as the neutrino masses, the dark matter origin, the hierarchy problem, etc.

The Standard Model of Elementary Particles is shown in Figure 1.1. It contains two main categories of particles: matter particles and force carriers. The matter particles are further subdivided into two basic types: leptons and quarks. Particles and carriers (which are also particles) are considered as indivisible, or elementary,

since their internal structure cannot be accessed by the energy scales achieved by today's largest collider. Leptons and quarks are fermions with spin equal to  $1/2$ . As shown in the Figure below, they are arranged in three *families*, or generations. There are six types (flavors) of quarks, pairs of two from every generation: up ( $u$ ) and down ( $d$ ), charm ( $c$ ) and strange ( $s$ ), top ( $t$ ) and bottom ( $b$ ). The electric charge of  $u$ ,  $c$  and  $t$  is  $2/3$  and that of  $d$ ,  $s$  and  $b$  is  $-1/3$ . Their mass varies from about 2.2 MeV (natural unit system  $c=1$ ) the mass of up quark to about 173.1 GeV the mass of top quark. An additional quantum property of quarks is the colour (*green, blue, red*), which is the charge of the strong interaction. Baryons and mesons are formed by quarks arranged in colorless combinations. Similar to quarks, there are six leptons that are categorized in generations. They consist of the electron ( $e$ ) and the electron neutrino ( $\nu_e$ ), muon ( $\mu$ ) and the muon neutrino ( $\nu_\mu$ ) and tau ( $\tau$ ) and the tau neutrino ( $\nu_\tau$ ). The electric charge of electron, muon and tau is  $-1$  and that of their neutrinos is 0. The mass of the former three ranges from 511 of KeV to 1.77 GeV, while that of neutrinos is very, very small. For each one of the leptons and quarks there is a corresponding anti-lepton and anti-quark, respectively, that has identical mass and opposite charge sign.




Figure 1.1: The Standard Model of elementary particle physics.

The SM theory views the interactions between particles as exchange of the force mediators, which are bosons with spin 1. In particular, each fundamental force has its own boson or bosons. The electromagnetic force is mediated by the massless and

electrically neutral photon ( $\gamma$ ). The weak force is mediated by the electrically charged  $W^\pm$  and the neutral  $Z^0$  bosons, with masses around 80 GeV and 91 GeV respectively. Finally, there are eight mediator particles for the strong force, the gluons, that are massless and electrically neutral. Just like quarks, they carry color charge that allows them to interact with themselves. Among the three forces, electromagnetic is the only long-range force, exhibiting infinite range. On the other hand, the weak force and the strong nuclear force have short-range of around  $10^{-18}$  m the former and around  $10^{-15}$  m the latter. The latest insertion to the SM is the Higgs boson, the discovery of which was announced at 2012 [1]. The Higgs boson has spin 0 and mass around 125 GeV. Its existence was predicted about 50 years ago, as it represents an essential component of the Standard Model. The Higgs field, mediator of which is the Higgs particle, is the field to which the  $W$ ,  $Z$  bosons and the fermions are interacting to obtain their masses.

## 1.2 Discovery beyond Standard Model

The Standard Model predicts successfully the interaction of particles in our accelerators so far. However, there are many aspects not yet understood, as well as problems that this theory cannot provide an answer. Some notable examples are the origin of dark matter, the matter-antimatter asymmetry we observe in the universe, the hierarchy problem and the Higgs's mass, and others. These kinds of questions cannot be answered by the SM theory. In addition, new particles that could provide us with insights regarding physics beyond the Standard Model have not been discovered yet by today's experiments.

A theory that tries to address these issues is Supersymmetry (SUSY). SUSY is a framework that extends the SM by adding new particles, the super-symmetric particles. It foresees that for every SM fermion exists a twin boson particle, and also, for every SM boson there is a SUSY fermion. This way, the theory predicts a number of new particles that can fill the gap of many unresolved areas. For example, the lightest stable and electrically neutral supersymmetric particle is a promising dark matter candidate. However, many of the particles predicted by SUSY should exist in a mass range that is already accessible by colliders. The fact that none of them has been discovered so far raises questions about the validity of this theory.

The largest particle factory today is the Large Hadron Collider (LHC). The discovery of the Higgs boson was done by the experiments operating at LHC. However, no other clue of new physics has been observed so far. For this reason, an extension of its operation has been decided, leading to an upgrade of the current machine. The new collider, called HL-LHC, will not increase the center-of-mass energy, but is going to deliver much higher number of collisions, as described in the next Chapter.

Targeting towards the long term future of high energy physics and the possibility for searches at high energy scales, the construction of a new collider is under consideration, which is foreseen to take place at CERN. The new machine, called Future Circular Collider (FCC), will extend the center-of-mass energy to 100 TeV (7 times the LHC), in search of new particles beyond Standard Model. Its construction is set to start towards the end of LHC operations (2041) and will have two phases. A positron-electron ( $e^+e^-$ ) machine in 2045 and a hadron-hadron machine in 2065.




# Chapter 2

## Large Hadron Collider

The Large Hadron Collider (LHC) is the most powerful particle accelerator and collider that has ever been built. It is located at the European Organization for Nuclear Physics (CERN) at Geneva, Switzerland. The construction of LHC started at 2001 inside the 26.7 Km long underground tunnel that already existed and hosted the Large Electron-Positron (LEP) accelerator [2]. LHC accelerates bunches of protons in opposite directions inside two vacuum circular pipes. The maximum energy a particle can reach inside each ring is 7 TeV, resulting to maximum center-of-mass collision energy at 14 TeV. The collisions take place at four Interaction Points (IP) where the two beams cross each other. The highest luminosity is delivered in two of these IPs and the nominal LHC value is  $L = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ . LHC also performs special runs where heavy (Pb) ions are accelerated to energy of 2.8 TeV.

The main goal of constructing the LHC was the production and discovery of the theory-predicted Higgs boson, as well as the search for signatures from theories beyond the Standard Model. Its basic goal was achieved during the first run of operations and the results of the Higgs discovery were announced on the 4th of July 2012 [1]. Since then, LHC continues its operation with a rich physics program which, apart from the further study of the properties of Higgs production, involves searches for physics beyond standard model (BSM). These include indications of the nature of dark matter, answers to why matter dominated over antimatter on the early universe, signs of super symmetric particles, or any other anomaly that cannot be explained by Standard Model.

LHC started its inaugural beam tests at 2008, achieving beam energy of 1.18 TeV. The first physics program started at 2010 and lasted up to early 2013 (Run 1), delivering proton-proton collisions with center of mass energy at 7 TeV. The two years followed by the so called Long Shutdown 1 (LS1). During this period, LHC was shut down in order to get upgraded to achieve collisions at the nominal energy of 14 TeV. The next run took place the years from 2015 to 2018 (Run 2) with operating collisions energy at 13 TeV. At that time the delivered luminosity reached the value of  $L = 2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ , surpassing by two times the nominal luminosity of the machine. After the next Long Shutdown (LS2) the years 2018-2022, LHC restarted operations at 2022 and will continue until 2025, when Run 3 will end along with the first Phase of LHC's operation (Phase-1). The total integrated luminosity delivered by the LHC up to 2023, as recorded by CMS experiment, versus time (in years) for Run 1 & 2 & 3, is shown in Figure 2.1.



**Figure 2.1:** The total integrated luminosity delivered by the LHC and recorded by the CMS experiment during Run 1, Run 2 and the beginning of Run 3.

## 2.1 Injection chain and bunch structure

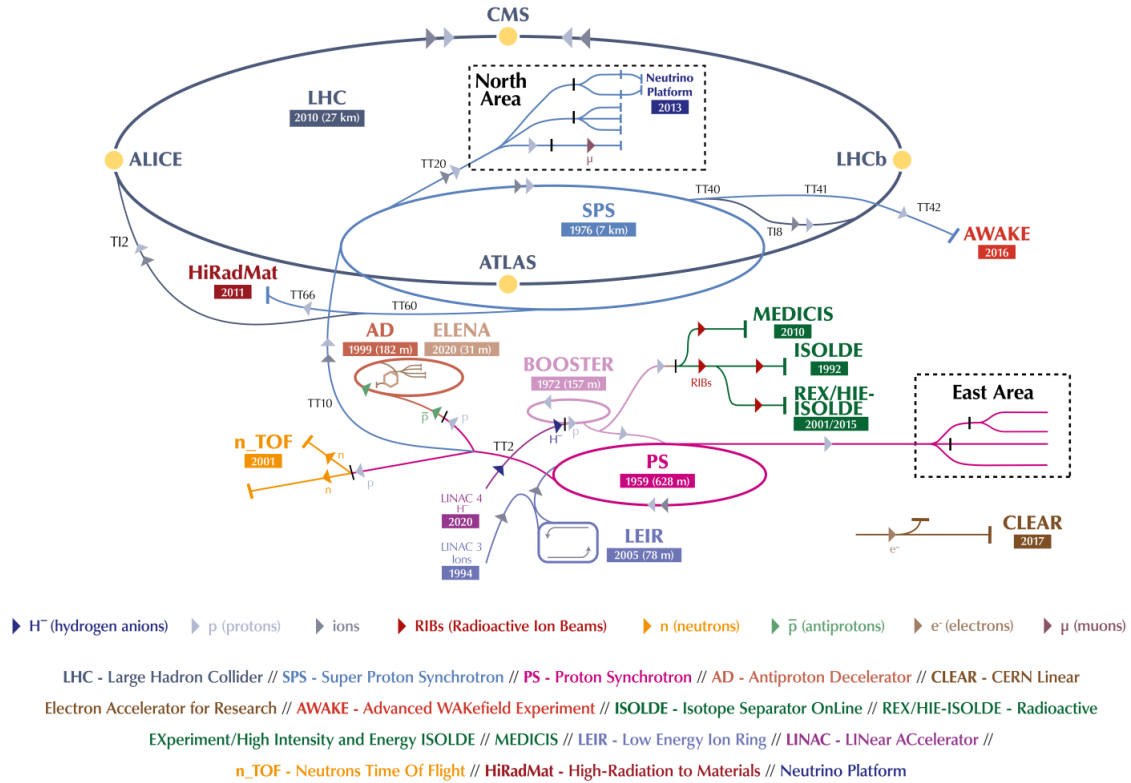
The proton bunches accelerated at LHC are produced and prepared by CERN’s accelerator complex. As shown in Figure 2.2, a number of smaller machines exists that gradually accelerate proton beams up to the energy of 450 GeV, when bunches are able to be injected inside the two LHC pipes. These accelerators are older experiments and today, apart from accelerating protons for LHC, they are also used to deliver beams of lower energy to other experiments operating at CERN.

The source of proton beams at the LHC complex used to be Linear Accelerator 2 (LINAC2) but after 2020 it has been replaced by LINAC4 [3]. LINAC4 is 76 meters long and accelerates negative hydrogen ions (consisting of one additional electron) to 160 MeV energy (LINAC2 used to boost protons up to 50 MeV). LINAC4 includes transfer and measurement lines where ions are stripped off the electrons and the remaining protons are injected to the Proton Synchrotron Booster (PSB) [4]. Inside the four superimposed synchrotron rings of PSB, protons reach the energy of 2 GeV and are then injected to Proton Synchrotron (PS). PS accelerates trains that include 72 proton bunches. A nominal bunch consists of about  $1.15 \times 10^{11}$  protons. These trains are accelerated up to 25 GeV energy and are then injected inside CERN’s second largest machine, the Super Proton Synchrotron (SPS). The number of trains that can be injected to SPS are 3 or 4. These batches are accelerated to 450 GeV energy before they are ready to be injected inside the LHC machine.

LHC captures and accelerates the SPS proton bunches using 400.8 MHz superconducting Radio Frequency (RF) cavities. In total, 16 such cavities instrument the LHC RF system, 8 for each of the two rings. To restrain the beams inside the circular vacuum tubes, LHC uses superconducting magnets. Their operation temperature is below 2 K and the produced magnetic field reaches 8 T.

The proton bunches existing inside the rings every moment form an LHC *orbit*. The frequency of an orbit is defined as the time it takes one bunch to complete a full circle. This frequency is 11.2455 KHz. The available number of bunches in every orbit is 3564, with a frequency of 40.078 MHz (or 25 ns period), defined by the frequency of the RF systems. However, not all bunch places are filled and their structure is defined by the previous accelerator machines and the injection characteristics. The

## The CERN accelerator complex *Complexe des accélérateurs du CERN*




**Figure 2.2:** Schematic view of the accelerator complex at CERN. The protons used by LHC originate from LINAC4 and are gradually accelerated by the Booster (PSB), the PS and SPS. They are injected to LHC at 450 GeV energy. Beams from the accelerator complex are also used by other experiments operating at CERN.

structure of a standard LHC orbit consists of 2808 bunches and can be seen in Figure 2.3. This structure defines as the start of an orbit the first bunch inside of it, called Bunch Crossing 0 (BC0). Furthermore, a gap with 119 missing bunches exists at the end of every orbit.

## 2.2 Luminosity

The two LHC rings cross each other at four interaction points. Prior to the crossing points, the beams are squeezed by superconducting quadrupole magnets. This operation increases their intensity as much as possible in order to increase the possibility of achieving hard proton-proton collisions. The search of rare physics events requires high number of collisions per bunch crossing and one of the most important parameters to measure the ability of a collider to produce such events is Luminosity. The number of events per second  $N_{event}$  that are generated in LHC collisions is given by:



**Figure 2.3:** *The bunch structure of a nominal LHC fill. The total number bunches fitting in an orbit is 3564. A typical orbit consists of about 2808 proton bunches, arranged as shown in the image.*

$$N_{event} = L\sigma_{event}$$

where  $\sigma_{event}$  is the cross section of the interaction in study and  $L$  the machine luminosity. The machine luminosity depends only on the beam parameters and for a Gaussian beam distribution it can be written as:

$$L = \frac{N_b^2 n_b f_{rev} \gamma_r}{4\pi \epsilon_n \beta^*} F$$

where  $N_b$  is the number of particles per bunch,  $n_b$  the number of bunches per beam,  $f_{rev}$  the revolution frequency,  $\gamma_r$  the relativistic gamma factor,  $\epsilon_n$  the normalized transverse beam emittance,  $\beta^*$  the beta function at the collision point and  $F$  the geometric luminosity reduction factor due to the crossing angle at the interaction point:

$$F = \frac{1}{\sqrt{1 + \left(\frac{\theta_c \sigma_z}{2\sigma_*}\right)^2}}$$

where  $\theta_c$  is the full crossing angle at the IP,  $\sigma_z$  the RMS bunch length and  $\sigma_*$  the transverse RMS beam size at the IP. Hence, in order to achieve high number of collisions, the LHC beam energy and beam intensity must be as high as possible.

The four IPs of LHC serve proton-proton and heavy ion collisions to corresponding experiments. At two of them, luminosity reaches its peak nominal value of  $L = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ . ATLAS and CMS are the experiments located in these points. The other two locations operate lower luminosity experiments. ALICE aims at a peak luminosity of  $L = 10^{32} \text{ cm}^{-2}\text{s}^{-1}$  and it mainly targets heavy ion runs. At the fourth IP, where the LHCb experiment is installed, the peak luminosity reaches  $L = 10^{27} \text{ cm}^{-2}\text{s}^{-1}$ .

The instantaneous luminosity is useful to quantify the density of a beam at one instant of time. Another very useful quantity is integrated luminosity, which depicts the accumulation of luminosity over a specific period of time:

$$L = \int_0^T L(t) dt$$

## 2.3 Experiments

- **ALICE** (A Large Ion Collider Experiment) [5]:

is a general-purpose, heavy ion detector located at one of LHC's four interaction points. Its goal is to exploit the heavy nuclei runs that LHC produces (Pb-Pb collisions) in order to study the strong interaction sector of the Standard Model (QCD). When heavy nuclei collide inside ALICE, the extreme energy density and temperature produced forces protons and neutrons to "melt" into their elementary constituents, quarks and gluons, forming the quark-gluon plasma. This primordial state of matter, where quarks and gluons are freed, is considered as dominant in the universe during the first millionths of a second after the Big Bang. The hot reaction zone rapidly expands and cools and then the ordinary matter particles are formed. Therefore, recreating this primordial state of matter in the laboratory and tracking precisely its evolution, helps in addressing questions about how matter is organized, the color confinement mechanism of QCD etc. ALICE's overall dimensions are  $16 \times 16 \times 26 m^3$  and it weights 10000 t. It was build by physicists and engineers from over 100 institutes from 30 countries.

- **ATLAS** (A Toroidal LHC ApparatuS) [6]:

is one of the two high-luminosity general purpose detectors of LHC designed to exploit the full discovery potential of  $pp$  collisions at LHC. The physics program of ATLAS experiment is quite extend, ranging from precise measurements of Standard Model parameters to searches for new physics phenomena. After the completion of one of the most important goals for the ATLAS and CMS experiments, the discovery of the Higgs boson on 2012 [35], further studies are ongoing to investigate its properties, interactions with other SM particles, the production and decay mechanisms etc. Besides Higgs physics, the electroweak sector of the SM (e.g accurate measurement W-mass) and flavor physics (e.g t-quark properties, B and D-mesons) are also thoroughly investigated. Additionally, searches for new particles involve those predicted from theories Beyond Standard Model to address open physics questions such as the matter-antimatter asymmetry, the existence of extra dimensions, the nature of dark matter, and others. ATLAS detector is 46 m long, 25 m high, 25 m wide and 7000 t heavy and more than 5500 scientists are members of the ATLAS experiment (as of 2023).

- **CMS** (Compact Muon Solenoid) [7]:

is the second large general-purpose detector that operates at LHC and investigates phenomena produced by  $pp$  collisions and also heavy ion runs. Its physics program is similar to that of ALTAS, that is a mix of measuring and further studying Standard Model parameters and mechanisms on the one hand, and on the other studies for physics beyond Standard Model. The key differences

between CMS and ATLAS are found in the overall conceptual design of the detector systems, that is the tracking, calorimeter and muon systems, as well as the magnetic fields and the triggering systems of the two. CMS detector has a length of 21.6 m, diameter of 14.6 m and weights 12500 t. Further description of the CMS experiment is given in the following chapters.

- **FASER** (ForwArd Search ExpeRiment) [8]:

is a new and rather small experimental facility installed during LS2 (2019-2021) targeting to operate during LHC Run 3 (2022-2025). FASER is located 480 m downstream of the ATLAS experiment and is aligned with its collision point axis. The purpose of the location is to take advantage of light, weakly-interacting particles that are produced from  $pp$  collisions in the forward direction at ATLAS's interaction point and escape through the holes of the beam line. The targeted particles are long-lived particles that travel hundreds of meters before decaying into known states of matter. FASER aims to the detection of particles that will decay inside its detector systems, about 5 meters long.

- **LHCb** (Large Hadron Collider beauty) [9]:

is the fourth experiment that is located in one of LHC's lower luminosity interaction points. LHCb's primary goal is to provide insight regarding the matter-antimatter asymmetry that we observe in the universe today. To achieve this it focuses on studying heavy flavor physics at LHC, and more specifically it searches for indirect evidence of new physics in CP violation and rare decays of beauty and charm hadrons. A new source of CP violation beyond standard model could explain the amount of matter in universe. For this reason, LHCb takes advantage of the high number of B mesons produced at LHC but in using different method than that of closed detectors, such as CMS and LHC. It consists of a series of sub-detectors designed to mainly detect particles in the forward region. In total, the detector is 21 m long, 10 m high and 13m wide weighting 5600 t.

- **LHCf** (Large Hadron Collider forward) [10]:

is an experiment that aims to collect data from an environment that simulates cosmic ray particles but inside a laboratory. The reason is to provide calibration that is needed in Monte Carlo of the hadron interaction models that are used for the description of ultra-high-energy cosmic ray collisions with the earth's atmosphere. Such information cannot be extracted from any other source in the Earth and LHC is the appropriate machine for this. LHCf consist of two small detectors that are placed symmetrically and 140 m away from ATLAS's two sides, exactly at zero degrees from the collision axis. High energy particles that are thrown away by forward collisions at ALTAS simulate high energy cosmic ray particles. Each of the two detectors weighs only 40 kilograms and measures 30 cm long by 80 cm high and 10 cm wide.

- **MoEDAL** (Monopole and Exotics Detector at the LHC) [11]:

is solely dedicated to extend the search for exotic particles at LHC. More specifically, MoEDAL's prime motivation is to directly search for magnetic monopoles

and other ionizing Stable (or pseudo-stable) Massive Particles (SMPs), predicted by theories Beyond the Standard Model (BSM). The MoEDAL detector is deployed around the intersection region of LHCb experiment and acts as giant camera sensitive only to new physics, as well as a trap for potential BSM particles.

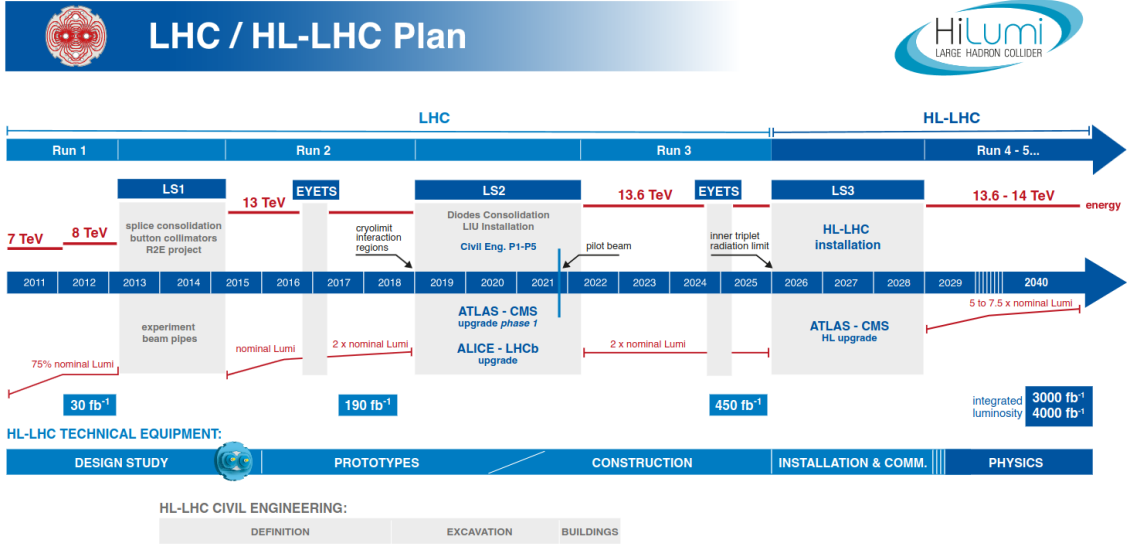
- **TOTEM** (TOTAl cross section, Elastic scattering and diffraction Measurement at the LHC) [12]:

aims the precise measurement of the total proton-proton cross section as well as to explore the structure of proton. It does so by measuring protons emerging from the collision point in the region very close to the particles beam (forward region), which is inaccessible by other LHC experiments. It is located at the two sides of CMS, where two tracking telescopes, T1 and T2, are installed on each side in the pseudorapidity region  $3.1 < \eta < 6.5$ . Furthermore, Roman Pot detector stations are placed at distances of  $\pm 147$  m and  $\pm 220$  m from CMS's interaction point.

## 2.4 High Luminosity LHC

The Large Hadron Collider started to operate in 2008, marking a new era for High Energy Physics. The production and discovery of Higgs boson at the first Run of the machine marked a major milestone of both the theoretical and experimental physics world. This achievement, however, would not be possible without the cooperation and expertise of many scientific and engineering disciplines of people and institutes around the globe. At the time of writing this thesis (2023), LHC is in the phase of Run 3, as shown in Figure 2.4. This run will last until the end of 2025 and will be the last run of the first phase of LHC's operation, the Phase-1. During this period, thousands of physicists have worked for LHC's experiments and contributed to the record and analysis of the huge amount of collision's data produced by the machine. This effort will not stop at the end of Phase-1 but will continue to push the boundaries of scientific progress with the second phase of LHC, the Phase-2.

LHC is and will remain the highest energy accelerator in the world for at least the next two decades. However, in order to maintain and extend its discovery potential, a major upgrade has been decided to be implemented after the end of Run 3. Result of this upgrade will be the so called High Luminosity LHC (HL-LHC), or HiLumi-LHC [13]. The aim of the upgrade is double. On the one hand, after operating for one and a half decades, LHC requires upgrades in order to maintain its operability by another decade or more. On the other hand, most important goal is the increase of instantaneous luminosity and thus, increase in the amount of collided particles per bunch crossing. More specifically, the new machine is expected to deliver more than five times higher luminosity and a tenfold increase of the integrated luminosity with respect to LHC's nominal design values. Figure 2.4 illustrates the overall schedule of LHC and HL-LHC since its first run and up to the end of Phase-2. According to the current plan, Run 3 will stop at the end of 2025 and be followed by Long Shutdown 3 (LS3), lasting for the next three years. During LS3, the decommission of LHC and commission of HiLumi-LHC will take place, as well as the commission of the detectors. The first Run of the Phase-2 era, Run 4, will begin the year 2029.



**Figure 2.4:** *The LHC and HL-LHC timeline plan, as of 2023.*

In total, Phase-2 is scheduled to operate three periods of collision runs, Run 4, Run 5 and Run 6, and have two periods of stops, LS4 and LS5, extending the lifetime of the accelerator up to the year 2041.

Since the maximum energy of the accelerator depends on the circumference of the machine and the field strength of the magnets, which will not change, the key point of the upgrade is the increased number of proton-proton collisions per bunch crossing. However, all new magnetic circuits have been designed with a 8-10 percent margin with respect to the peak 7 TeV energy value, providing the possibility of reaching 'ultimate' beam energy at 7.5 TeV. The aim for peak luminosity at the upgraded machine is  $L = 7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ , targeting a maximum of 200 collisions per bunch crossing inside the CMS and ATLAS detectors. This level of luminosity should be able to deliver 300 to 400  $\text{fb}^{-1}$  per year, meaning that the ultimate performance is able to deliver up to 4000  $\text{fb}^{-1}$  of total integrated luminosity.

The Phase-2 upgrades do not involve only the LHC machine, but also the accelerator complex and the two detectors that operate at LHC's interaction points. The upgrade of the accelerator chain is subject of the LHC Injectors Upgrade (LIU) project of CERN. Its purpose is the implementation of all necessary improvements of the injection complex that are required for the increased luminosity of HL-LHC. The activities of LIU have already started during LS2. LINAC4, which will also be used at Phase-2, is such an example, along with upgrades to the SPB, PS and SPS. Furthermore, ATLAS and CMS experiments are already working on the design and production of their upgraded detector systems. A description of the upgraded version of CMS experiment is given in the next Chapter.



# Chapter 3


## The CMS Detector

### 3.1 Introduction

Compact Muon Solenoid (CMS) is one of the two general-purpose experiments operating at LHC [7]. It is installed at Interaction Point 5 (P5) about 100 meters below surface. Proton-proton collisions take place at the center of the detector at LHC's nominal luminosity, resulting to thousands of secondary particles being generated and distributed all around the interaction point. The number of inelastic collisions achieved per bunch crossing (every 25 ns) during Run 3 is around 60 (pile up). For every such event, the CMS detector systems identify the position and momentum of every secondary particle in order to re-create the event and run the corresponding physics program. The initial purpose of CMS was the discovery of Higgs boson. This milestone was achieved during the first run of LHC and was announced at July 4 2012, alongside with the discovery at the ATLAS experiment. Since then, CMS focuses on a broad range of physical phenomena, with the most notable of them being the following. Further study of the Higgs mechanism and its interaction with matter, precise measurements of the Standard Model, searches for new physics, involving new particles that could provide insights regarding the current unanswered questions of physics.

A drawing of the CMS detector is illustrated in Figure 3.1. Key element of the design is the superconducting solenoid magnet that operates at 4 T. The choice of using such a strong magnetic field is driven by the desire to measure the momentum of muons at high precision, using the bending of energetic charged particles inside a strong magnetic field. Starting from the innermost region and moving outwards, first detector system is the Tracker, placed as close as possible to the interaction point. The next layer is covered by the Electromagnetic Calorimeter (ECAL), absorbing and measuring photons and electrons. The Hadronic Calorimeter is installed exactly behind ECAL to absorb and measure hadron particles. In the next layer is where the solenoid magnet is located, in a way that all sub-detectors mention above are installed inside its volume. With this configuration, the only traceable secondary particles escaping the solenoid are muons. The detection of muons is performed using four muon stations and by reconstructing the track segments they leave when passing through the detectors. The information produced on every event is processed online by the CMS Trigger system. The collision rate of 40 MHz produce a huge amount of

events that are impossible to be stored directly. For this reason, the Trigger system performs an online event selection and reduction in order to only store potentially interesting events. This action is performed in two sequential stages, the Level-1 Trigger and the High Level Trigger. The accepted events are read-out in their full resolution by the Data and Acquisition (DAQ) system and are stored in disks for detailed offline analysis.



**Figure 3.1:** *Schematic view of the CMS detector.*

CMS is installed 100 meters below the ground at the interaction point 5 of LHC (P5). The detector subsystems were fabricated at CERN and in academic institutes and factories all around the world. They were transferred at the surface area of P5 in the constructions building called SX5. At SX5, CMS was pre-assembled in slices that were moved underground through a large hole that connects SX5 with the experimental cavern. The experimental area room is called UXC (Underground eXperiment Cavern). It is the place where high energy collisions occur and hence, the radiation levels are above normal. For this reason, all electronic equipment installed at UXC is radiation hard. The information that is produced at the detectors of CMS is captured, amplified and digitized by on-detector electronics. This region of the experiment is called front-end. From the front-end, information is transmitted to a service room that is located next to UXC, called USC (Underground Service Cavern). Between the two there is a 7-meter concrete wall to bring radiation down to normal levels. Thus, electronics installed at USC can be standard, commercial and not radiation hard components. From the front-end information is transmitted using 90-meter long optical fibers to the so called back-end. The back-end boards are used both to receive data from the sub-detectors but also to configure the front-end electronics by transmitting fast and slow commands, clock, and other important signals. In most cases, back-end boards that interface the detector are the Layer-1


boards of the Level-1 Trigger system (see next chapter).

### 3.1.1 CMS coordinate system

The coordinate system used by CMS is defined by the interaction point at the center of the detector. As illustrated in the left of Figure 3.2, y-axis points vertically upwards and x-axis points at the center of LHC. The z-axis points along the beam direction and towards the Jura mountains. The x-y plane, also called transverse view, is where the azimuthal angle  $\phi$  is measured. The polar angle  $\theta$  is measured from the z-axis. Apart from the Cartesian and cylindrical coordinates, CMS uses a quantity that remains invariant to Lorentz boosts along the beam axis (z-axis.) This coordinate is called pseudorapidity,  $\eta$ , and is defined as:

$$\eta = -\ln \tan\left(\frac{\theta}{2}\right)$$

Thus, to determine the exact location of a secondary particle that hit a specific detector sub-system, the coordinates that need to be extracted are  $\eta$  and  $\phi$ . The relation of pseudorapidity with the polar angle  $\theta$  can be seen at the right of Figure 3.2. Lastly, the momentum and energy of detected particles is measured on the transverse plane and are denoted as  $p_T$  and  $E_T$ , respectively. An also very important measurement at CMS is that of energy imbalance over the total energy measured in the geometry of the detector and is called missing  $E_T$ , or  $E_T^{miss}$




**Figure 3.2:** Left: The CMS detector coordinate system. Right: The relation between the pseudorapidity,  $\eta$ , with the polar angle  $\theta$ .

## 3.2 Superconducting Solenoid

The Solenoid magnet of CMS, illustrated in Figure 3.3, is the largest of its kind that has been used in high energy experiments. Its nominal magnetic field is designed to reach 4 Tesla. The usage of such a strong field has been chosen in order to maximize the physics performance of CMS. More specifically, one of the CMS design goals was a muon system that measures muon momenta with very high resolution (1 % at 200 GeV). On this scope, the 4 T field benefits substantially the muon tracking both inside and outside the solenoid. Furthermore, the resolution of the trajectory measurements of all energetic charged particles heavily benefits from this choice.

The magnet of CMS detector is the central element of the experiment. Due to its size and structure, it is used as the principal support for the sub-detector systems. The four muon stations are located outside the solenoid and inside are placed the barrel part of the Hadronic and Electromagnetic Calorimeter and the Tracker. It has a diameter of 6 m and length of 12.5 m and is capable of storing 2.6  $GJ$  of energy at full current. The flux of the solenoid's magnetic field is returned through a 10,000 tons iron return yoke, installed in layers outside the solenoid. In order to produce such high magnetic field, superconducting elements are used to cope with the extreme current flux. The superconducting coil is made of a 4-layer winding made from stabilised reinforced NbTi conductor. The nominal current of the coil is 19.14 kA.



**Figure 3.3:** *Left: Artistic view of the CMS solenoid magnet. Right: The magnet during the construction of CMS at UXC. The red iron layers outside the coil are the return yokes.*

The magnet was initially deployed and tested at the surface area of CMS (SX5) during 2006. It then moved 90 meters underground to be installed at its final position inside the experimental cavern. First time it cooled down and ramped up to its nominal field was in November 2008. Its operation will continue as is for the Phase-2 era of CMS.

### 3.3 Tracker System


The Tracker system of a detector is the system that reconstructs tracks of charged particles, or, in other words, provides the optical view of how the secondary particles spread in space around the interaction point. This operation is mandatory as it enables the separation between charged and neutral particles, as well as the reconstruction of the primary vertex of the hard interaction in every bunch crossing. For this reason, the Tracking system has to be installed as close as possible to the beam pipe, has high granularity and low sampling time.

At CMS, the Tracker system is designed to provide high precision and efficient measurement of the trajectories of particles produced by proton-proton collisions of LHC. This means that every 25 ns the Tracker must identify trajectories of more

than 1000 particles, products of the inelastic  $pp$  collisions. To achieve this and also allocate them to the correct bunch crossing, the detector technology features high granularity and fast time response. This task, however, is quite challenging since the location of the Tracker imposes many limitations and difficulties. First of all, the amount of material used in the detector must be as low as possible in order to reduce phenomena such as multiple scattering, bremsstrahlung, photon conversion and nuclear interactions. Furthermore, the region around the interaction is where particle flux is the most intense, causing continuous radiation damage to the tracking system during the years of operation of the detector. To overcome these challenges, CMS tracker is made entirely from silicon sensors.

### 3.3.1 Phase-1 Tracker

The Tracker system of CMS for Phase-1, shown in Figure 3.4, is 5.8 m wide and has a diameter of 2.5 m. The innermost part is a Pixel detector, arranged in three layers at the barrel and three disks per endcap. It is made of 66 million pixel cells that have a size of  $100 \times 150 \mu\text{m}^2$ . Outside the Pixel, a Silicon strip tracker has been installed, consisting of ten layers in the barrel and twelve disks in the endcap regions. It consists of 9.3 million strip sensors. The pseudorapidity area coverage is  $-2.5 < \eta < 2.5$ .



**Figure 3.4:** *The CMS Phase-1 Tracker.*


### 3.3.2 Phase-2 Tracker

At the end of Run 3 the current Tracker system will reach the end of its lifetime and be completely replaced by a new tracking system. The Tracker system for Phase-2 [14] has to be able to cope with the increased HL-LHC particle flux. Comparing to Phase-1 Tracker, it will provide increased forward acceptance and radiation hardness, as well as higher granularity. A very important new feature will be the delivery of tracks to the Level-1 Trigger system, contrary to the current architecture where tracks are only used by the High Level Trigger. Phase-2 Tracker will be composed by the Inner Tracker and the Outer Tracker.

## Inner Tracker

The Inner tracker is designed not only to cope with the radiation levels and pile up of HL-LHC (up to  $4000 \text{ fb}^{-1}$ ) but also to improve the tracking and vertexing capabilities of the current detector. The silicon sensors (thin planar n-in-p) will have a thickness of  $100\text{-}150 \mu\text{m}$  and be segmented into pixels. Their final size is not yet decided, but will be either  $25 \times 100 \mu\text{m}^2$  or  $50 \times 50 \mu\text{m}^2$ , providing even higher detector resolution (factor of 6 reduced area with respect to Phase-1). The pseudorapidity coverage will also be extended up to  $|\eta| \sim 4$ .

A perspective view of one quarter of the Inner tracker is shown in Figure 3.5. The detector is divided in three parts. The barrel part, called Tracker Barrel Pixel Detector (TBPX) contains four layers of sensor modules arranged in *ladders* so that they overlap in  $r - \phi$ . Further along in both sides of  $z$  is the Tracker Forward Pixel Detector (TFPX), made of eight small double-discs per side. Finally, Tracker Endcap Pixel Detector (TEPX) is the last sub-detector and in both sides it consists of four large double-discs. The total active area covered by Inner tracker is  $4.9 \text{ m}^2$ . The detector is designed in a way to allow the replacement of large parts over an Extended Technical Stop.

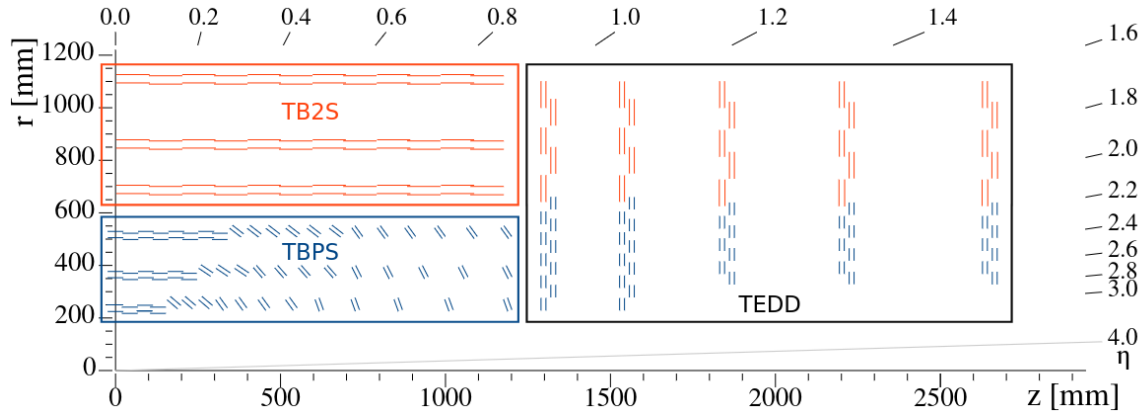


**Figure 3.5:** View of one quarter of the CMS Phase-2 Inner Tracker detector.

The basic unit that Inner Tracker is called pixel module. It consists of a pixel sensor, several Pixel ReadOut Chips (PROC's), a flex circuit and a mechanical support. Apart from the sensor's operation, other functions involve shipping the data out to back-end electronics, provide clock, trigger and control signals, as well as power distribution. The pixel module will be constructed in several types in order to match the requirements of each sub-detector.

## Outer Tracker


A representation of one quarter of the Outer Tracker can be seen in Figure 3.6. In total, it provides a coverage between  $r = 21$  cm and  $r = 112$  cm in the x-y plane and covers the region of  $|z| < 270$  cm. Outer Tracker is divided in three sub-systems depending on the location and the tracker module used in each case. These are: the Tracker Barrel with PS modules (TBPS), the Tracker Barrel with 2S modules (TB2S) and the Tracker Endcap Double-Discs (TEDD). The barrel region of the Outer Tracker consists of six cylindrical layers extending up to  $|z| < 120$  cm. The endcap region will be instrumented with five double-disc layers, covering  $120 < z < 270$  cm. The arrangement of the layers provides that at least six module layers are crossed by particles coming from the luminous region  $|z| < 70$  at  $|\eta| < 2.4$ .



**Figure 3.6:** *The r-z view of one quarter of the CMS Phase-2 Outer Tracker.*

The basic unit will be the so called  $p_T$  module. It comes in two flavors, the PS (Pixel-Strip) and 2S (2-Strip) modules, shown on the right and left of Figure 3.7, respectively. The modules are instrumented with sensors on both of their sides. The 2S module consists of strips of size  $5 \text{ cm} \times 90 \mu\text{m}$  and covers total active area of  $2 \times 90 \text{ cm}^2$ . The PS module is half the 2S, covering an area of  $2 \times 45 \text{ cm}^2$ . It consists of both strips with size  $2.4 \text{ cm} \times 100 \mu\text{m}$  and macro-pixels with size  $1.5 \text{ mm} \times 100 \mu\text{m}$ . The modules are mounted in a service board called *hybrid*, where the two single-sided sensors are closely-spaced in both sides. The hybrid provides all mechanical aspects necessary for the construction of the detector, as well as the circuitry needed for the read out and control of the sensors. The CMS Binary Chips (CBCs) are used for reading out the strips. Eight of them exist on 2S hybrids and one Concentrator Integrated Circuit serves as the interface between all CBCs and the readout link. The PS front-end hybrid uses eight Short Strip ASICs (SSAs) to read out the strip sensor and one CIC, similarly to the 2S module. Transferring data between the detector front-end chips and the back-end processors is performed using the low-power Gigabit Transceiver (lpGBT) chip [15]. The communication is performed optically with a specific link module called Versatile Link Plus (VL+) [16].

The most impressive feature of the new Tracker is that it will deliver tracking information of the Outer Tracker to the Level-1 Trigger system. The current implementation only transmits track data to the High Level Trigger. Further description of the Track Trigger is given in section 5.1.3.



**Figure 3.7:** *CMS Phase-2 Outer Tracker hybrids. Left: The 2S module. Right: The PS module, having half the size of 2S.*

## 3.4 Calorimeters

The calorimeter detectors at the barrel region of CMS are installed right after the Tracker system and are defined within the volume of the Solenoid magnet. The barrel systems are supplemented with the corresponding endcap detectors at the two sides of CMS. The Phase-1 architecture is based on two calorimeters. Moving from the interaction point and outwards the first detector system is the Electromagnetic Calorimeter (ECAL), followed by the Hadronic Calorimeter (HCAL). Both of them are divided in two parts of similar architecture. There is the ECAL Barrel (EB) and Endcap ECAL (EE), as well as the HCAL Barrel (HB) and HCAL Endcap (HE). Both barrel detectors will remain as is for Phase-2, but the endcap calorimeters will be completely replaced by a new detector.

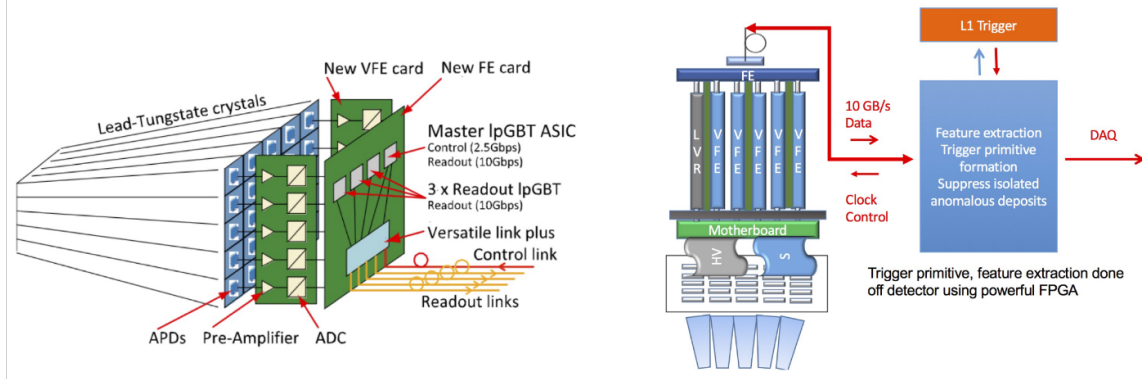
### 3.4.1 Electromagnetic Calorimeter Barrel

The Electromagnetic Calorimeter (ECAL) [7] is a homogeneous crystal calorimeter and provides high precision measurements of photons, electrons and positrons. In the barrel region, it consists of 61,200 lead tungstate crystals ( $\text{PbWO}_4$ ) and covers the pseudorapidity range  $|\eta| < 1.479$ . An image of ECAL can be seen at the right of Figure 3.9. When electromagnetic particles pass through the crystals they decay and decelerate, converting their kinetic energy into radiation. The choice of  $\text{PbWO}_4$  fits to the requirements of LHC. Their high density ( $8.28 \text{ g/cm}^3$ ), short radiation length ( $X_o = 0.89 \text{ cm}$ ) and small Moliere radius ( $r_M = 2.2 \text{ cm}$ ) create a compact calorimeter with high granularity. The length of the crystals is 23 cm, which corresponds to total radiation length  $25.8 X_o$ . On the front they cover an area of  $22 \times 22 \text{ mm}^2$  and on the back  $26 \times 26 \text{ mm}^2$ . The light produced in the crystals is detected using Avalanche Photodiodes (APDs). Two are used per crystal. The performance of both the crystals and the APD highly depends on the temperature. For this reason, a cooling system based on water is used in order to maintain stable temperature at  $18 \text{ }^\circ\text{C}$  ( $9 \text{ }^\circ\text{C}$  during Phase-2).

The descriptions so far refer to the legacy ECAL Barrel detector, which will remain as is for HL-LHC as well. In order to cope with Phase-2 requirements, all front-end electronics and the data processing and transmission methods will be upgraded [17].



The function of APD's pulse amplification, shaping and digitization will be performed by the new Very Front End (VFE) card. This new card will provide better timing resolution and noise filtering. The digital signal from 5 VFE boards is passed to the Front End (FE) card. FE card performs the control of VFE chips, buffering and transmission of data to the back end electronics. The transmission is executed by the lpGBT ASIC using the VL+ optical link.



**Figure 3.8:** Schematic of the front-end electronics architecture for Phase-2. Left: The connections between ECAL crystals, the APDs, the new VFE card and the new FE card. Right: The architecture of data path between the new FE card and the Level-1 Trigger and DAQ systems.


Contrary to Phase-1, where ECAL trigger primitives (TPs) are formed by FE boards on the detector, in Phase-2 the generation of TPs will be performed off-detector, at CMS counting room. Main advantage of this topology is the usage of commercial electronics, since radiation levels at USC are normal. At the Level-1 Trigger processors, analysis of incoming data will be performed, including rejection of spikes and basic clustering of localized energy. A pre-processed set of TPs will be transmitted to the next layer of the Level-1 Trigger system. The subsystem responsible for this operation, called Barrel Calorimeter Trigger, is discussed at section 5.1.1.

### 3.4.2 Hadronic Calorimeter Barrel

The Hadronic Calorimeter (HCAL) [7] at the barrel region of CMS (HB) is placed right after the ECAL in the available space inside the volume of the solenoid magnet. More specifically, it is installed at radii  $1.77 < R < 2.95$  covering pseudorapidity range  $|\eta| < 1.3$ . Due to lack of available space, the amount of absorber material at HB is not the optimum that is needed to absorb the hadronic shower. For this reason, a complementary outer calorimeter (HO) has been installed right outside the solenoid magnet.

The HB, shown at the left of Figure 3.9, is a sampling calorimeter made of alternating layers of flat brass absorber plates and tiles of plastic scintillator. It is divided into 36 identical azimuthal wedges forming two half-barrels, HB+ and HB- (18 wedges each). The absorber consists of 14 layers of brass-plates sandwiched between two steel plates. The first steel plate is 40 mm thick, the next brass plates 50.5 mm thick, six brass plates have 56.6 mm thickens and the last steel plate is 75

mm thick. The total absorber thickness at  $90^\circ$  is 5.82 interaction lengths ( $\lambda_I$ ) and at  $|\eta| < 1.3$  is  $10.6 \lambda_I$ . The active material used at HB is scintillator tiles and wavelength shifting fiber. The number of tiles installed is about 70,000, grouped into scintillator tray units depending on their  $\phi$  position. The first scintillator layer is installed in front of the first steel plate. It is made of Bicron BC408 material and has thickness of 9 mm. The next 15 layers are made of Kuraray SCSN81 and are 3.7 mm thick. The last layer has the same Kuraray material but is 9 mm thick in order to correct for late developing showers leaking out of the HB. At the longitudinal view HB is divided in 16  $\eta$  sectors, resulting in a segmentation  $(\Delta\eta, \Delta\phi) = (0.087, 0.087)$ . The scintillator light is collected by a 0.94 mm thick green doubled-cladded wavelength-shifting (WLS) fiber.




**Figure 3.9:** *Left: The Hadronic Calorimeter at the surface assembly area of CMS. Right: The Electromagnetic Calorimeter installed inside the HCAL.*

Studies have shown that the total degradation of HB (scintillator and absorber units) until the end of HL-LHC will be lower than that of Endcap HCAL (made out of the same material) at the end of Phase-1 [17]. However, the Hybrid PhotoDiodes (HPD) that are currently used to read-out the light from WLS fibers need to be replaced. The new read-out chips will be Silicon PhotoMultipliers (SiMPs), installed during LS2.

### 3.4.3 High Granularity Calorimeter

The endcap calorimeter detectors of CMS during Phase-1 are homogeneous calorimeters based on  $\text{PbWO}_4$  (ECAL) and plastic scintillator (HCAL). They were designed for an integrated luminosity of  $500 \text{ fb}^1$ . Beyond that, their performance degradation is so high that sets them unable to be used during the HL-LHC era. As a result, both calorimeters on the two endcap sides of CMS are going to be replaced the by High Granularity Calorimeter (HGCal) [18]. HGCal is a sampling calorimeter aiming to provide unprecedented transverse and longitudinal segmentation at pseudorapidity range  $1.5 < \eta < 3.0$ . It is divided in two compartments, the electromagnetic (CE-E) and the hadronic (CE-H). They are both made of layers of sampling sensors, followed by absorber material. Silicon (Si) sensors will be used at CE-E and the absorber material will be copper and lead. The CE-H will use both silicon sensors and scintillator

tiles at larger radii. Its absorber material is going to be steel. The layout of one half of HGICAL is shown in Figure 3.10.

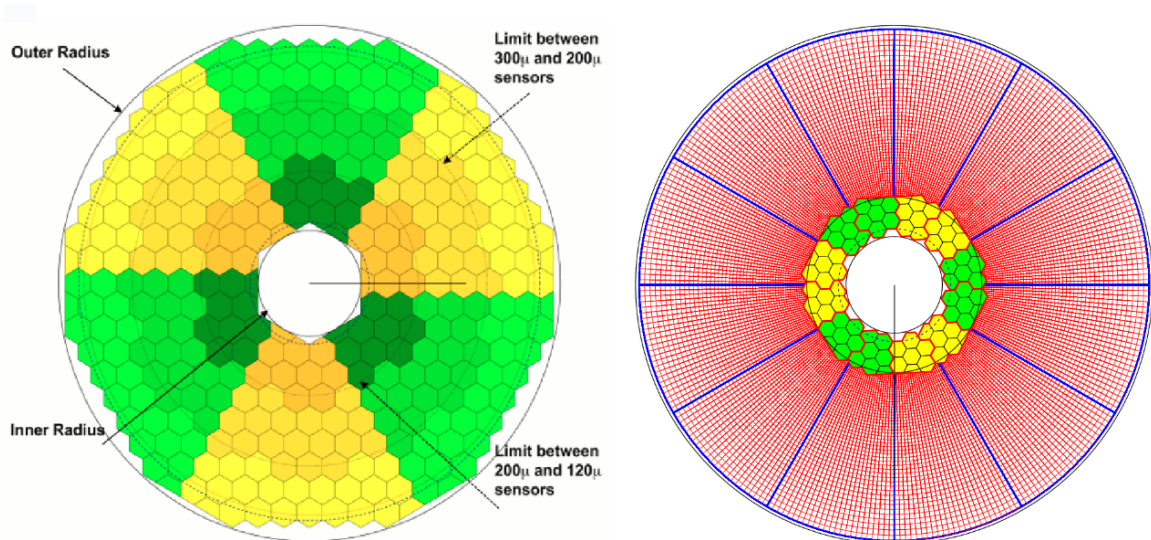


**Figure 3.10:** *Layout of one half of the HGICAL detector.*

CE-E consists of 28 sampling Si layers with a total thickness of 34 cm. Its depth is approximately  $26 X_0$  and  $1.7 \lambda$ . The so called silicon modules use silicon as the active detector element sandwiched between a 1.4 mm thick copper lead material (WCu) and the printed circuit board that hosts the front-end electronics. The silicon sensors come from 8 inch wafers and have a hexagonal shape, in order to exploit the maximum wafer area. There are three different module flavors depending on the radii each of them will be installed. At radii close to the beam axis ( $r < 70$  cm) sensors with thickness of  $120 \mu\text{m}$  and size of  $0.5 \text{ cm}^2$  will be used. This choice arises from the much higher fluence experienced in the forward region and the sensitivity requirements due to higher number of particle showers. The silicon modules in this area consist of 400 individual cells in total. At higher than 70 cm radii the thickness increases to 200 and  $300 \mu\text{m}$  and the cell size to about  $1 \text{ cm}^2$ , creating silicon modules of about 200 individual cells. The layout of hexagonal modules shaping Layer-9 of CE-E can be seen on the left of Figure 3.11.

The hadronic compartment of HGICAL is composed by both silicon and scintillator sensors. Steel will be used as its absorber material. The region of  $|\eta| < 2.4$  is exposed to high particle fluence and will be covered exclusively by silicon sensors. In region of  $|\eta| > 2.4$  where the integrated dose and fluence are low-enough, scintillator sensors can be used as the active material. In total, CE-H will be made of 24 layers. The layout of Layer 22 is shown in left of Figure 3.11. Small scintillator tiles are arranged in an  $r, \phi$  grid with cells of increasing size. In the region near the beam axis the size of cells starts from  $4 \text{ cm}^2$  to reach that of  $32 \text{ cm}^2$  at the outer edge. The produced light is read out by Silicon Photo-multipliers (SiMP) chips. The absorber material will be stainless steel.

The read-out of both the silicon and SiMP information is performed by the ded-



**Figure 3.11:** *Left: Drawing of Layer-9 of CE-E, made of HGCal hexagonal sensors. Right: Layout of layer 24 of CE-H, made of both silicon and scintillator sensors.*

icated radiation hard ASIC called HGCROC (HGCal Read Out Chip). It comes in two variants that serve each sensor separately, with the only difference being an adaptation to the first amplifier stage of the SiMP version with respect to the Si version. The operation of HGCROC is to measure and digitize the energy deposition of the sensor cells, as well as to record the time of arrival of the pulses. The digital information follows two different paths, the DAQ and the trigger path. On the DAQ path, raw information is written in circular buffers waiting for the Level-1 accept signal. Upon reception, it is sent to the ECON-D (DAQ) ASIC which then transmits it via the lpGBT chip and VL+ to the DAQ readout. For the trigger path, HGCROC computes sums of energy for neighboring cells and transmits them to back-end processors. Data transmission is performed by the ECON-T (Trigger) and via lpGBT and VL+ as well. The back-end system that uses this information to form the HGCal trigger primitives and reconstruct electromagnetic and hadronic objects is described in section 5.1.1.

### 3.5 Muon System

Many signatures of interesting processes that originate both from the Standard Model and the Higgs mechanism, and also from new physics, decay eventually to high energy muons. For this reason, the design of the CMS experiment was based on the efficient measurement of muon particles. The detector systems are separated into those located at the barrel region and those at the two endcaps. In addition, a specific detector system that specializes in the precise measurement of time is installed in both regions. All detector systems are based on gaseous ionizing chambers that are read out by fast front-end electronics [7].

Studies have shown that chambers of all muon detectors are able to operate during HL-LHC [19]. The majority of the front-end electronics, however, will have to be replaced by new ones, capable of handling the increased radiation dosage, higher data

rates, longer trigger latency, etc. Moreover, new detectors are going to be installed at the very forward region of the endcaps to provide coverage at higher pseudorapidity range and also to provide increased measurement efficiency. Figure 3.12 illustrates the r-z view of one quarter of the CMS detector. The muon subsystems are represented by the colored layers, or stations, located outside the magnet coil.

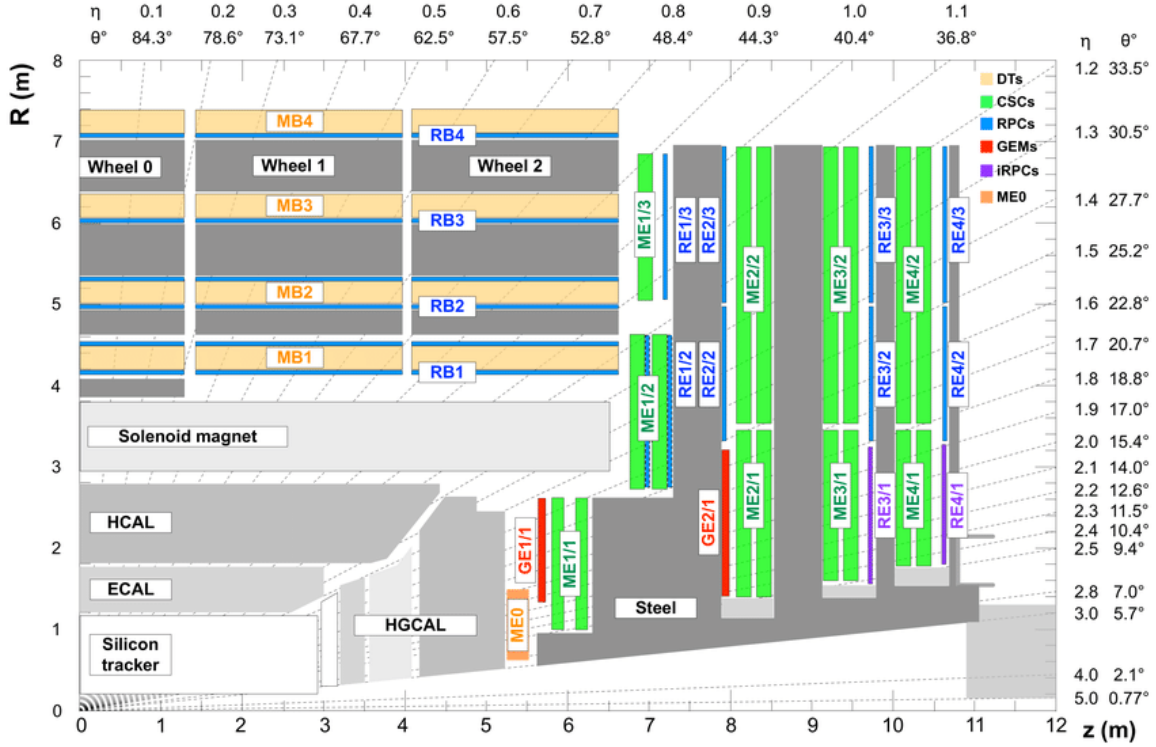



Figure 3.12: A quadrant of the muon system.

### 3.5.1 Drift Tube system

The Drift Tube (DT) system is responsible for the detection of muons crossing the barrel region of CMS. It is located outside the Solenoid magnet, hence, the only traceable secondary particles crossing the DT system are muons. The system consists of layers of stations interchanged with the iron return-yoke. The yoke is used to provide homogeneity to the magnetic field produced by the Solenoid coil. On the longitudinal view the DT system is divided in five wheels, wheel -2, -1, 0, 1 and 2. On the transverse view every wheel is made of 12 Sectors, hence, the total number of Sectors is 60. The sectors can also be grouped in 12 wedges with each wedge consisting of 5 Sectors in the z-axis. A schematic of the DT system showing one wheel in the r- $\phi$  view is illustrated in Figure 3.13. In every Sector there are 4 DT stations, forming concentric circles around the beam axis. At the top and bottom Sectors the number of stations is 5 since the outer layer consists of two instead of one. Thus, the total number of DT stations, or chambers, is 250.

The basic unit of the system is a drift tube chamber cell, depicted at the right of Figure 3.14. The size of a cell is 42 x 13 mm<sup>2</sup> on the transverse view. At the center of each cell there is a gold-plated stainless steel anode wire and on the sides of the cell




**Figure 3.13:** *Layout of the CMS barrel muon DT chambers and sectors in one of the 5 wheels.*

walls are the cathodes. They extend to a length of 2 meters. High voltage current flows in both of them (+3600 V on the anode and -1200 V on the cathode), producing high density electric field which is assisted by field forming strip electrodes on the top and bottom cell walls. The ionizing gas used is a mixture of 85 % Ar and 15 % CO<sub>2</sub>, providing a maximum drift velocity of about 54 μm/ns and a corresponding maximum drift time of about 390 ns. The DT cells are grouped in four half-staggered layers which form one Super Layer (SL). The crossing of a muon inside a SL can be reconstructed and the position and direction of a particle crossing the SL can be obtained. In every DT station there are 2 SLs with their anode wires oriented parallel to the beam axis in order to measure the  $r$ - $\phi$  view. In addition, 1 SL, placed vertically to the beam axis, exists in the 3 inner stations to provide the  $\theta$  measurement. The structure of a DT Chamber with the SLs and the DT cells is illustrated on the left of Figure 3.14. In total, the barrel muon system consists of 172,000 DT cells covering  $|\eta| < 1.2$ .

The current muon detector of the barrel is expected to provide very good performance during the whole operation of HL-LHC. However, the front end electronics have to be replaced with new ones that are able to cope with the increase in particle flux and data rate. A detailed description of the Phase-2 Barrel Muon system is given in Chapter 6.

### 3.5.2 Cathode Strip Chambers

Muon detection at the endcap regions of CMS is performed by Cathode Strip Chamber (CSC) detectors. High particle flux of the forward region combined with the non-




**Figure 3.14:** *Left: Structure of a DT Chamber consisting of three Super Layers. Each SL is made of four layers of DT cells, two oriented parallel to  $z$ -axis and one vertically to it. Right: View of a Drift Tube cell.*

uniform magnetic field set the usage of CSC the most suitable choice. The system is made of trapezoidal chambers installed in four disk layers at each endcap side and perpendicular to the beam axis. There are four layers, as shown in Figure 3.12 (illustrated in green), and in total 540 cathode strip chambers. In each side, the first layer consists of three rings and 108 chambers and the remaining three layers have two rings and 54 chambers. The overall size of the chambers depends on the radial location. Their length varies from 160 cm to 340 cm. The top and bottom widths (trapezoidal shape) are between the ranges 61-153 cm and 31-90 cm, respectively. The thickness of the chambers is fixed at 25 cm except from those of the first ring in the first station where the thickness is 15 cm. Each chamber covers azimuthal angle of  $10^\circ$  or  $20^\circ$ . The total pseudorapidity coverage of the CSC system is  $0.9 < \eta < 2.4$ , having an overlap with the DT system at  $0.9 < \eta < 1.2$ . The second station of the CSC detector (ME2) is illustrated in Figure 3.15.

The CSCs are multiwire proportional chambers made of 6 anode wire planes, running azimuthally to define the track's radial coordinate, interleaved among 7 cathode panels. The panels are milled with strips that run radially with constant azimuthal width  $\delta\phi$ . The gaps created between the cathode panels are filled with a mixture of gas which is made of 40% Ar + 50% CO<sub>2</sub> + 10% CF<sub>4</sub>. The precise position of a muon is defined by interpolating charges induced on strips, produced by the gas ionization.

### 3.5.3 Resistive Plate Chamber system

The Resistive Plate Chambers (RPC) are gaseous detectors that have been installed at CMS to provide additional, more precise time measurement to the muon system [7]. They are capable of tagging events with resolution of 2 ns, setting them able to identify unambiguously the Bunch Crossing number of the ionizing particle, i.e. muons that are crossing the detectors. This functionality is very important to the muon system where tagging particles to the correct BX is of high importance. RPCs can also operate in an environment where high particle rate is present, making them ideal to be used at the forward region of the CMS endcaps as well. Moreover, apart from excellent time resolution, they also provide a good enough spatial resolution.



**Figure 3.15:** *The ME2 station of CSCs.*

The RPC system is installed both at the barrel and endcap region covering pseudorapidity range of  $|\eta| < 1.9$ . At the barrel region their placement is similar to that of the DT stations. They form 6 coaxial cylinders around the collision axis and are attached either in the inner or outer side of DT chambers, or at both sides. In the first 2 stations there are 2 RPC chambers in each station placed in both sides of the DT chambers. In the 2 outer stations there are again two layers of RPCs but at the inner side of the corresponding DT station. The total number of the barrel RPC chambers is 480. At the two endcaps RPC chambers have trapezoidal shapes and are arranged in four concentric disks around the beam axis. There are in total 576 chambers at both endcap sides. The arrangement of the RPC layers at the barrel and endcap region can be seen in blue color at Figure 3.12.

### 3.5.4 Improved RPC

The current pseudorapidity coverage of the legacy RPC system does not cover the very forward region of CMS, a region very challenging in terms of backgrounds and momentum resolution. To overcome this absence, the CMS collaboration decided to install two additional layers of detectors that will cover the region  $1.8 < \eta < 2.4$ . The new system will consist of improved RPC chambers (iRPC) installed at disks RE3/1 and RE4/1, as depicted with purple at Figure 3.12. The new chambers are going to have thicker electrodes and gas gap than the previous ones (1.4 mm) resulting to reduced numbers of electrons generated. These can be absorbed faster by the electrodes leading to increased sensitivity and rate capability of the detector. The Front-End Electronics Board is also going to be upgraded to cope with the low charge signals while keeping the efficiency as high as in the usual RPC. The installation of prototype chambers has already began during 2022/2023 and the full commissioning is foreseen to be completed by 2024.



### 3.5.5 Gas Electron Multiplier detectors

The installation of a new detector system based on Gas Electron Multiplier (GEM) lies within the efforts of CMS to extend the efficient muon detection at the very forward region near the beampipe. This region is going to suffer the most after the high luminosity upgrade and during the Phase-2 era. The detector systems at this region will have to cope with higher number of collisions per bunch crossing that will generate greater radiation dose and higher number of event rate. At the two endcap sides, the new GEM detectors will cover pseudorapidity range of  $1.5 < \eta < 2.8$ , extending and improving the measured muon's polar bending angle. The location of the foreseen installation of the GEM detector is illustrated at Figure 3.12 with red and orange colors. They are denoted as GE1/1, GE2/1 and ME0.

## 3.6 MIP Timing Detector

MTD (stands for Minimum Ionizing Particle (MIP) Timing Detector) is a new detector that will be commissioned at CMS for Phase-2 [20]. This detector will be able to measure the production time of MIPs in high precision. Specifically, MTD's resolution is going to be 30-40 ps at the beginning of HL-LHC and about 50-60 ps during the end of HL-LHC, due to degradation by radiation damage. This information can be used to efficiently assign charged tracks to the correct interaction vertices, even in scenarios of 200 pile up collisions. The pseudorapidity coverage provided by MTD will be  $|\eta| < 3.0$ . The detector system will be a thin layer installed between the Tracker and the calorimeters, divided in the barrel and endcap regions.

# Chapter 4


## Technology of Phase-2 Level-1 Trigger and Data Acquisition systems

### 4.1 ATCA Standard

The CMS Level-1 Trigger (L1T) and Data Acquisition (DAQ) systems for Phase-2, described in Chapter 5, are being built using custom designed electronic boards. This choice arises from the very specific tasks and requirements for data processing and management imposed by HL-LHC. The two systems are planned to be installed in the underground service cavern, USC, located next to the CMS detector room, UXC. Inside USC, the L1T and DAQ boards will be installed inside crates that are placed in rows of electronic racks. The radiation levels at USC are normal and thus commercial electronics can be used. For the Phase-2 systems, the L1T and DAQ have adopted the Advanced Telecommunications and Computing Architecture (ATCA) as the industry standard to be used. The ATCA specification defines the physical and electrical characteristics of the shelf, or chassis, that hosts easily swappable ATCA blades. The ATCA also provides a high density and flexible backplane that allows interconnections between the blades inside a crate.

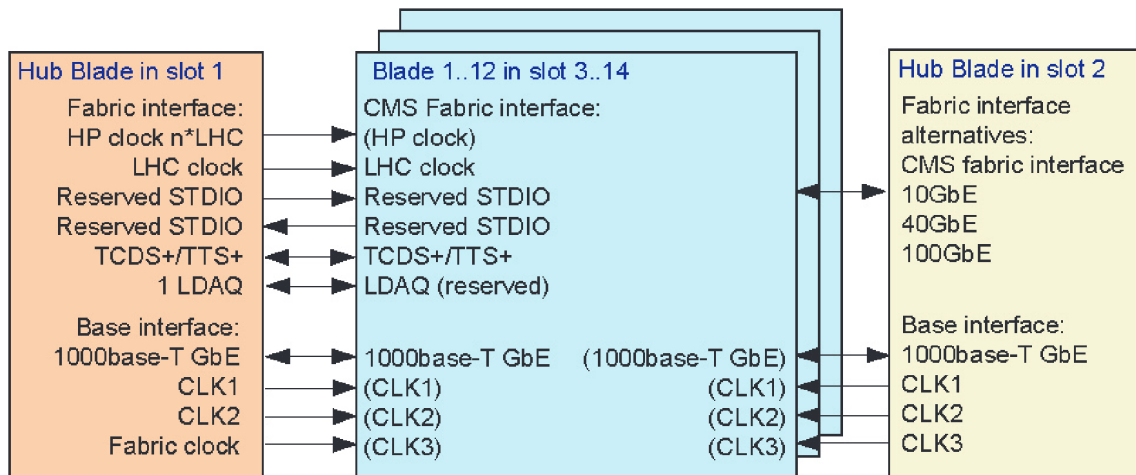
A drawing of an ATCA front board along with a Rear Transmission Module (RTM) is illustrated in Figure 4.1. The left side of a blade is called front panel and can be used for custom connections depending of the board's functionalities. The backplane is located between the front board and the RTM and consists of the Zone 1, used for power and management connections, and Zone 2. The data transport connectors are located in Zone 2. Zone 3 is used to connect signals between the front board and the RTM and as that, its usage in L1T blades is optional.

There are in total 14 available blade slots inside an ATCA crate. Slots 1 and 2 are intended for hub blades. In a 'Dual Start' configuration, four bidirectional general purpose signal pairs are connected from each hub board to every blade node in the crate. Furthermore, a base interface provides a 1000base-T Ethernet, three clock pairs and a 100base-T Ethernet for management. These signals, shown in Figure 4.2, are distributed through the Zone 2 connector. At CMS, one hub slot must be facilitated with the DAQ blade, called DAQ and TCDS2 Hub, or DTH (see section 5.2.2). For the L1T node boards, only the Zone 2 J23 connector is needed, while



**Figure 4.1:** Drawing of the side view of an ATCA blade attached with an RTM. The figure highlights the backplane connectors of the ATCA standard.

DTH requires connectors J20 through J24. For power consumption considerations, the CMS guideline foresees up to 10 L1T ATCA blades to be installed in the 12 remaining slots. The custom backplane connections between the hub and the main boards are also illustrated in Figure 4.2. They include the main LHC clock running at 40.0x MHz, High Precision (HP) clock multiple of the LHC clock and trigger timing and control signals (TCDS2, described in section 5.2.1). The remaining interfaces are reserved.



**Figure 4.2:** Data interface between the hub blades and the processor boards inside an ATCA crate, as defined by CMS.

Managing of the crate takes place outside the card cage. It is executed by a commercial shelf manager which is responsible for managing the fans, monitoring the power distribution, temperatures, etc. Management of the blades themselves is performed through the Intelligent Platform Management Interface (IPMI) bus and

carried out by the IPM Controller (IPMC). IPMC is installed on every blade inside the crate. It offers functions such as hotswap, sensor monitoring, etc. It also includes Ethernet connection that provides distant control of the power of the board. For Phase-2, CERN produces its own CERN-IPMC mezzanine card that can be used by CMS and ATLAS experiments [21].

## 4.2 Field Programmable Gate Arrays

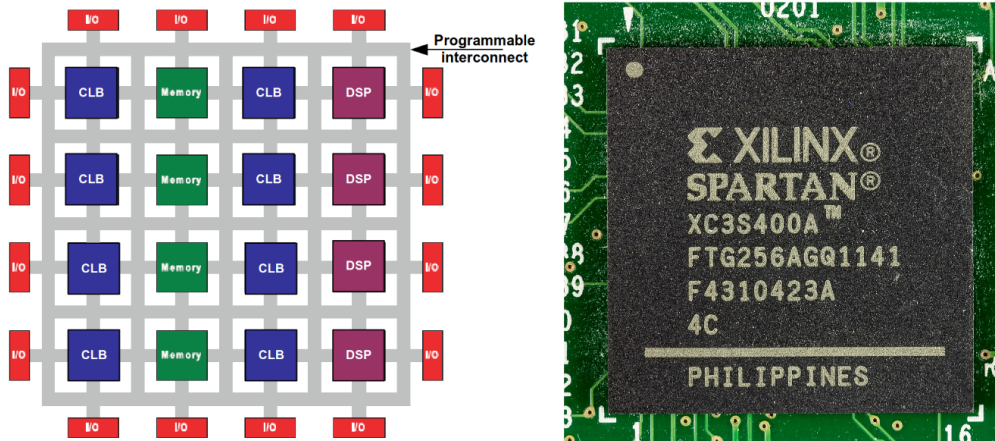
Field Programmable Gate Arrays (FPGA) are widely used at CMS since the first years of the experiment's operation. At L1 Trigger, they constitute the basic processing unit. The ATCA blades that are being designed to build the system are all based on FPGA processors to transfer data and execute the algorithmic logic.

FPGAs are user programmable integrated chips, introduced around the early 1980s. They consist of an array of programmable logic elements such as look-up tables, flip-flops and memory elements. All these electronic pieces are interconnected with a matrix of wires, which is also programmable. As a result, by enabling the usage of specific elements connected by corresponding routes, the user can implement any digital circuit that can fit in a specific device. During the first years of the FPGA evolution, the number of logic gates inside a chip was at the order of tens of thousands. The latest FPGAs as of 2023 consist of up to tens of millions. Furthermore, FPGAs are becoming even more sophisticated by hosting additional digital circuits, such as Digital Signal Processors (DSP), clocking resources, blocks of Random Access Memory (RAM), high speed serial transceivers, and others.

### 4.2.1 Architecture

The main building block of an FPGA is the Configurable Logic Block (CLB). As shown on the left of Figure 4.3, copies of CLBs are repeated inside the chip and their number can range from thousands to millions, depending on the size of the FPGA. A CLB is typically made of a set of Look Up Tables (LUT), Flip-Flops, Carry logic and multiplexers. LUTs are essentially memory chips that match input values to output values. When a logical function is to be performed, instead of executing the actual calculation, LUTs have stored the result and retrieve it to the output pins. LUTs can also be used to store user data that can be accessed any time by a logical circuit. Flip-flops are the primitive storage circuits that can store single bit information for multiple clock cycles. A multiplexer has a number of input signals and each time it outputs only the one selected by the user. High speed carry logic is provided to perform fast arithmetic functions, such as addition and multiplication.

Apart from CLBs, FPGAs include other types of resources, providing the ability to implement digital circuits that can serve any kind of application. One example are tiles of Block Random Access Memory (BRAM) that can be used to store large amounts of data that are easily accessed by the user logic. Moreover, Digital Signal Processing (DSP) slices are being included in FPGAs to aid applications that require fast and efficient arithmetic functions to be implemented. All resources described so far are capable of processing large amount of data very fast. For this reason, the ability to drive data in and out the FPGA in high speeds is mandatory. To achieve



**Figure 4.3:** *Left: Block design of fundamental building blocks and programmable interconnects of an FPGA. Right: Photo of a Spartan FPGA produced by Xilinx.*

this, FPGAs provide a significant amount of user Input/Output (I/O) pins to be used either for serial or parallel data transfer. In addition, transceiver devices are being used to handle serial data interfaces in rates that exceed that of gigabits per second (described in the next section).

Another device that has been introduced during the last twenty years is the System on Chip (SoC). A SoC is an integrated circuit that includes the components of an entire system into it. In the case of FPGAs, a SoC connects the FPGA programmable logic (PL) with a processing system (PS) in one single die. The two are interconnected with dedicated high performance busses that facilitate data transactions between the two. The PS can be a commercial processor that hosts embedded operating system (OS). From inside the OS, high level applications can interface the programmable logic to create high-performance applications.

## 4.2.2 Hardware Description Languages

The behaviour of an FPGA is defined by Hardware Description Languages (HDL). An HDL is used to describe digital circuits using written computer language. The typical structure of code is divided in two parts. One where inputs and outputs of the system are declared, and the main body, where the behaviour of the digital circuit is described. This representation is also referred to as Register Transfer Level (RTL). The structural form of HDL languages allows the abstraction of a complex circuit into smaller components that are easier to understand and handle. Two of the most widely used HDL languages are Verilog and VHDL (Very high-speed integration circuit HDL).

The transformation of an RTL into the configuration file that will program the FPGA is performed by an Electronic Design Automation (EDA) tool. The procedure starts with analyzing written language and producing the netlist of the digital circuit. This process is called Synthesis. The netlist contains a gate-level representation of the elements that are used and the connections they share. Next stages require that the tool has knowledge of the specific resources of the targeted FPGA device. Thus, the compilation tool is provided by the company that produces the corresponding

FPGA chip. These steps involve the procedures of placement and routing. During placement, the compilation tool places the design onto the resources of the target device. The routing stage searches for the most efficient connections, or routes, that should be used to connect together the different elements of the circuit in order to implement it without placement or timing violations. The last stage is the one that generates the so called bitfile. The bitfile is the configuration file that programs the FPGA with the digital circuit designed by the user.

### 4.2.3 FPGAs at CMS

Most of the FPGA devices used at CMS are produced by Xilinx. The work of this thesis is conducted exclusively in devices manufactured by Xilinx, as well as the usage of the corresponding EDA software, the Vivado Design Suite [22]. More specifically, the Phase-2 activities of the experiment have been involving the Ultrascale [23] and Ultrascale Plus [24] families, that were already introduced in the market since 2015. These generations have utilized many technological innovations that allowed their production in 20 nm, 16 nm and 14 nm, enabling applications of ultra-high performance and massive bandwidth.

The initial stages of research and development at CMS were being conducted using Ultrascale FPGAs and small parts of the Ultrascale Plus generation. During the course of the years, it has become clear that the needs of the Level-1 Trigger (described in the next chapter) can only be met by utilizing one of the biggest parts of the Ultrascale Plus family, the VU13P FPGA. A list of FPGAs that have been used during the years, and will be used at the final system, is shown in Figure 4.4. The list also includes a couple of ZYNQ devices that are mentioned in the next chapters.

Device Name	KU040	KU15P	ZU4CG	ZU19EG	VU9P	VU13P
<b>CLB Flip-Flops (K)</b>	485	1,045	176	1,045	2,364	3,456
<b>CLB LUTs (K)</b>	242	523	88	523	1,182	1,728
<b>Total Block RAM (Mb)</b>	21.1	34.6	4.5	34.6	75.9	94.5
<b>UltraRAM (Mb)</b>	-	36	-	36	270	360
<b>DSP slices</b>	1,920	1,968	728	1,968	6,840	12,288
<b>16.3 Gbps Transceivers</b>	20	44	16	44	-	-
<b>32.77 Gbps Transceivers</b>	-	32	-	28	120	128

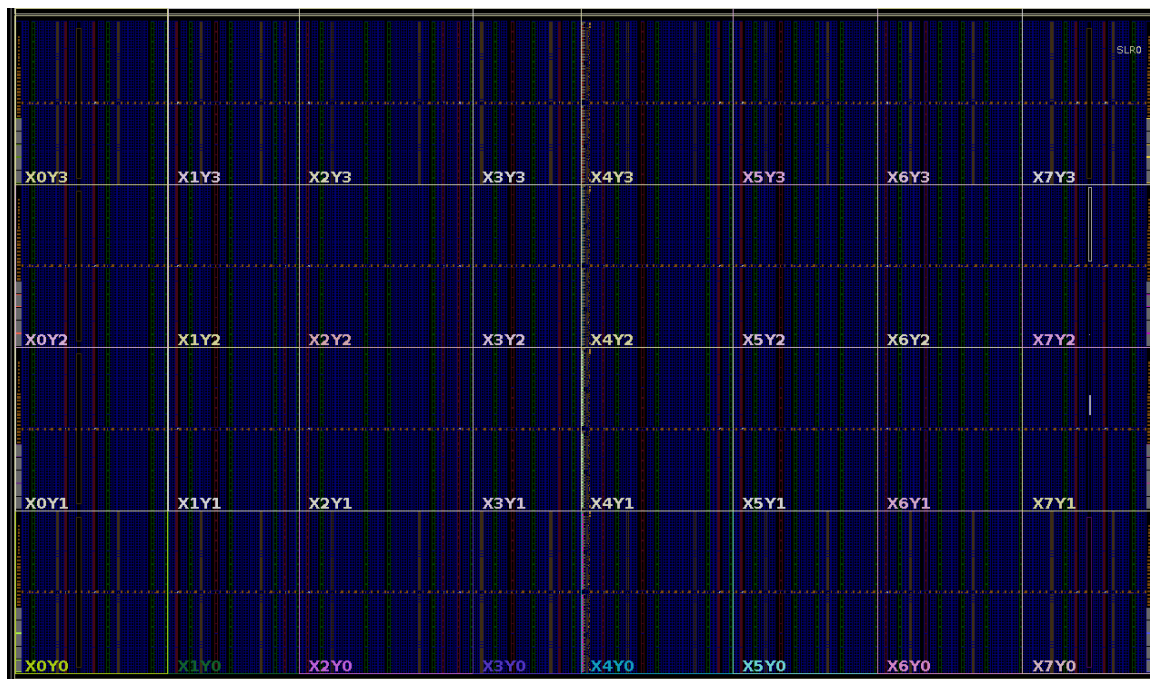
**Figure 4.4:** Product table of FPGAs produced by Xilinx and are used by the CMS Level-1 Trigger for Phase-2.

### Xilinx VU13P FPGA

The VU13P FPGA is going to be the standard device used by the systems of the Level-1 Trigger. As such, it has been used as baseline for the R&D work presented in this thesis. It is one of the largest chips of the Ultrascale Plus family, and by far the largest of other FPGAs that were used until now at CMS for Phase-2. The VU13P chip is made of four SLRs (Super Logic Region). Each SLR is a single FPGA die slice that interconnects with the others into one single package. Routing signals

between SLRs takes place through SLLs (Super Long Lines) and most of the times their efficient usage is conducted by the Vivado software.

The floorplanning of one SLR of the VU13P chip is shown in Figure 4.5. This drawing is a representative image of the chip resources, its regions and Banks. The Banks at the left and right borders of the chip contain the high-speed transceivers that are highly used to transfer data in and out of the device. Their operation is described in the next section.



**Figure 4.5:** *Floorplanning of SLR 0 of a VU13P FPGA, extracted from the Vivado software.*

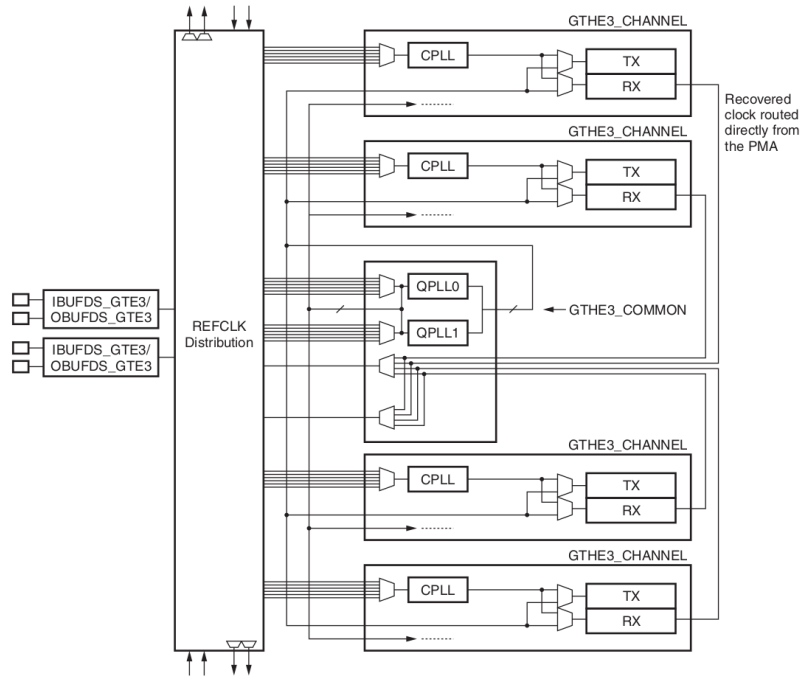
### 4.3 Multi Gigabit Transceivers

One very important advantage of the FPGA, apart from the processing power and parallelism they offer, is high speed I/O serial interfaces. These transceivers increase the total amount of bandwidth the device can receive and transmit without requiring high count of I/O pins. Furthermore, they can provide efficient connectivity through serial backplanes, optical interfaces, or between chips placed on the same board. Data transmission via optical links is highly used at the CMS L1 Trigger and Data Acquisition systems. Collision information is transmitted from the detector and between ATCA boards through optical fiber cables. Moreover, crucial timing and control signals are distributed back to the front-end electronics also via high speed optical links.

The Xilinx FPGAs listed in Figure 4.4 include a high number of serial transceivers that operate at line rates up to 32.77 Gbps. The Xilinx high speed transceiver devices are called Multi Gigabit Transceivers (MGT). Depending on characteristics such as the internal bus width they process and the line rate they can achieve, MGTs are

categorized in different types. The GTH transceivers are supported by both the Ultrascale and Ultrascale Plus families and are capable of transmitting data to line rates up to 16.3 Gbps. The GTY type reaches line rate up to 32.77 Gbps and is supported only by Ultrascale Plus devices.

The MGTs inside Xilinx FPGAs are grouped as Quads. Each Quad contains four transceiver channels, two LC tank-based Quad Phase-Locked Loops (Quad PLL or QPLL) and four ring-based Channel PLLs (CPLL). There are two dedicated clock pins to provide the reference clock required for the operation of each channel. These clock sources can be routed to the QPLLs and be used by all four channels of the quad, or routed to each CPLL independently. The block design of the structure of a Quad is illustrated in Figure 4.6.

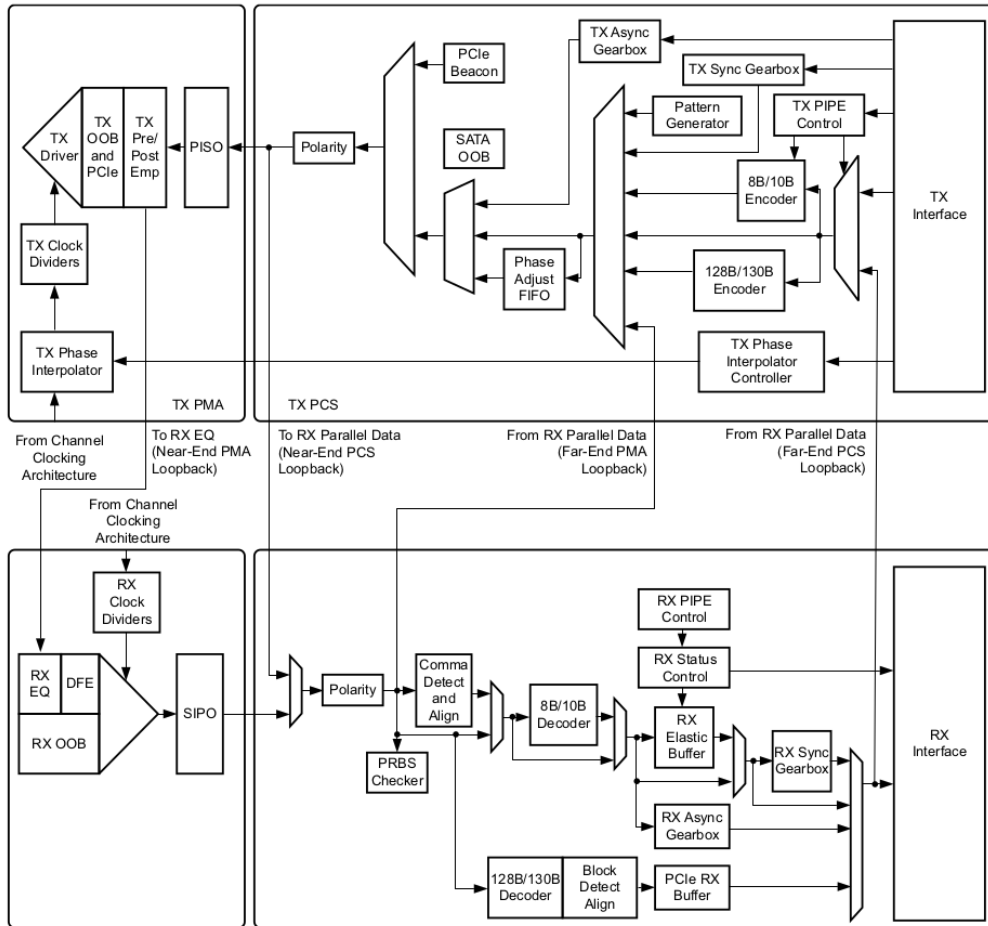


**Figure 4.6:** Block design of the MGT Quad structure.

An MGT channel consists of one transmitter block (TX), one receiver block (RX) and one CPLL. The reference clock can be provided either from the QPLL or the channel's CPLL. The schematic block diagram of a GTY channel is shown in Figure 4.7. Both TX and RX contain one Physical Coding Sublayer (PCS) and one Physical Medium Attachment (PMA). Parallel data are interfaced between the user logic and the MGT through the PCS. On each side, this layer contains optional blocks to assist the implementation of most frequently used commercial transmission protocols. It also contains necessary blocks to facilitate the data flow inside the transceiver. Such blocks are Encoder and Decoder for the 8b/10b scheme, 64b/66b and 64b/67b support, Pseudo-Random Binary Sequence (PRBS) Generator and Checker, First In First Out (FIFO) blocks to facilitate clock domain crossings between internal clocks, and others.

In the PCS layer, data are in parallel form and can have width of either 4, 6 or 8 bytes. The first block on the TX PMA layer is a Parallel-In Serial-Out (PISO) to convert data to the serial stream that is going to be transmitted. Other optional





**Figure 4.7:** MGT Channel Block design showing the PCS and PMA of the TX and RX functional blocks.

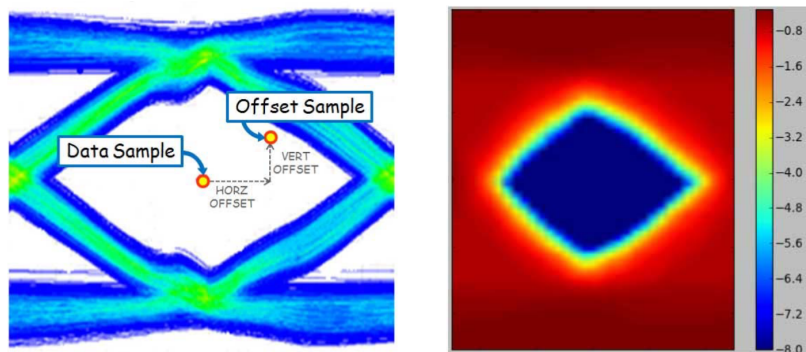
blocks existing on the TX PMA are the Pre and Post Emphasis, the TX Out-Of-Band signaling and beacon signaling for PCI express designs. Transmission of the serial signal outside the chip is performed by the TX Driver.

On the receiving side the serial stream at the RX PMA is captured by the RX analog front end. It is a high-speed current-mode input differential buffer that includes configurable RX termination voltage and calibration resistors. The signal is then fed into the RX Equalizer. Xilinx Transceivers provide two adaptive filter blocks to reduce the losses imposed during the transmission of the signal. These are a Decision Feedback Equalizer (DFE), recommended for high channel losses, and a Low Power Mode (LPM) equalizer for lower channel losses and less power consumption. Once equalized, the serial data stream is directed to the Clock Data Recovery circuit (CDR). CDR is responsible of determining the phase of the digital signal and extracting the clock of the stream, known as *recovered clock*. Once the RX clock is recovered, the serial stream can enter the Serial-In Parallel-Out circuit. The resulting parallel data bus is routed to the RX PCS. Inside the RX PCS data flow through decoder blocks and return to the form the transmitter has sent them. They can then be used inside the RX interface by user applications.

### 4.3.1 IBERT tool

Xilinx provides a tool in the form of IP (Intellectual Property) core that facilitates validation testing of the MGT transceivers, called IBERT (Integrated Bit Error Ratio Tester) IP core. The IBERT tool connects the MGT ports and dynamic reconfiguration port attributes to a graphical interface where the user can easily access them. This way, the user can validate the operation of both the MGT reference clocks and that of the serial link. Initially, the reception of a reference clock to the MGT is realized through the *QPLL Locked*, that is exposed to the user. To evaluate the serial link operation, the core uses the pattern generator and checker blocks that existing in the PCS. The transmitted data are in the form of PRBS (Pseudo-Random Binary Sequence) sequences. When patterns are created by the same polynomial, the link status is *Locked* and the receiver can extract the number of potential bit flips that occurred during transmission. If for any reason the interface is broken, the link status will be *Unlocked*. This way, the operation of a link can be validated by running an IBERT project for hours or days and counting the number of errors detected. A statistical estimation of the validity of data transferred can be calculated this way.


A method of measuring the signal integrity of a transmission line is through eye scans. The eye scan, or eye diagram, provides a statistical eye view of the accumulation of errors area over an extended sampling period. The Xilinx GTY transceivers include the necessary circuitry required to perform an eye scan. Sampling of the serial stream takes place right after the equalizer block. There are two samplers that operate in parallel, shown on the left of Figure 4.8. The first one, called *Data Sampler*, samples data on the time domain at the location determined by the CDR block and on the voltage domain at differential zero. The second one, called *Offset Sampler*, samples in programmable horizontal (time) and vertical (voltage) offset positions. On every offset position, the Bit Error Ratio (BER) is calculated as the ratio of the error count to the data count. Plotting the result to a 2-dimensional phase space results to the eye diagram, as shown on the right of Figure 4.8. An eye diagram can be made of thousands or up to  $10^{14}$  samples and the more open the blue area is (np-errors area), the better the quality of the transmission line is.



**Figure 4.8:** *Left: Data Sampler and Offset Sampler. Right: Eye-diagram 2D plot.*

### 4.3.2 Optical Transceiver modules

High speed transceivers can be used in multiple use cases, such as on-board chip-to-chip data transfer or communication between boards through a backplane. The prominent way of transmitting data at high line rates and larger distances is through optical fibers. Data transmission from the CMS detector to the service room is performed by optical links, as well as the communication between the ATCA boards of Level-1 Trigger and Data Acquisition systems. On the physical layer, the transition of the signal from electrical to optical and vice versa is performed using optical transceiver modules attached on the ATCA blades. Figure 4.9 illustrates two such modules that are used in CMS, a Quad Short Form-factor Pluggable (QSFP) on the left and a Samtec Firefly on the right.



**Figure 4.9:** *Optical Transceiver Modules. Left: A Quad Short Form-factor Pluggable (QSFP). Right: Samtec Fireflies.*

The QSFP is the most commonly used transceiver type today. The original module, SPF, represents a compact, hot-swappable and single channel transceiver. As requirements for data bandwidth increase, new variants are being introduced to accommodate these needs. The SFP, SFP+ and SFP28 have all the same form factor but are capable of supporting different maximum line rates, 1 Gbps, 10 Gbps and 28 Gbps respectively. The QSFP is a similar module but has different form factor as it includes four transceiver channels in a single module. As such, different types also include the QSFP+ and QSFP28. They are widely used for Ethernet applications.

The Samtec Firefly is a compact transceiver module. The transceiver plugs in a small footprint connector that can be placed near the FPGA chip, reducing the high-speed traces on the Printed Circuit Board. There are variants of the Firefly module that support line rates of 14, 16, 25 and 28 Gbps. Firefly modules have two designs, the x4 and x12. The x4 part contains 4 bi-directional transceiver channels. The x12 type can contain either 12 transmitters or 12 receivers and is not supported at 28 Gbps line rates.

# Chapter 5

## CMS Phase-2 Trigger and Data Acquisition System

The proton-proton collisions rate of 40 MHz at the center of CMS produces a vast amount of data coming from the detector sub-systems. The total data stream produced every second reaches the order of tens of Petabytes, making them practically impossible to be stored directly in disk. Furthermore, the majority of the p-p collisions lead to known and well studied interactions of the Standard Model. For these reasons, a reduction of the event rate must take place before data can be efficiently stored. This reduction is executed by the CMS Trigger system. The CMS Trigger is an online event selection system that interfaces with the detector in real time, reconstructs every collision event and decides whether it contains meaningful information or whether it should be discarded. The event selection is performed by the trigger menu of algorithms, which defines the physics channels of interest. During Phase-2, CMS aims to perform more precise measurements of the Standard Model coefficients, further study the Higgs boson sector and also search for traces of particles predicted by theories beyond Standard Model.

The event rate reduction is performed in two stages. The first stage, called Level-1 Trigger (L1T), executes the first data processing by processing data of every single event, every 25 ns. It is made out of custom designed trigger processors that are based on sophisticated FPGAs interconnected with high speed optical links. L1T reduces the event rate to 750 KHz with a latency of 12.5  $\mu$ s. Until the decision of whether to accept or discard an event is taken, the raw information of the event is stored in buffers on the detector. Upon reception of an accept signal, full resolution of the corresponding event is read out by the Data and Acquisition system (DAQ) and is transmitted to the High Level Trigger (HLT). HLT, the second stage of the online analysis, is made of commercial CPUs and GPUs where more complex, high level algorithms are possible to run. By running a more detailed processing of the events it further reduces the rate down to 7.5 KHz. The selected events are then permanently stored in disks and are available for offline analysis by physicists around the world.

The upgraded Trigger system of CMS has to be able to cope with the unprecedented number of collisions produced by HL-LHC. At its ultimate configuration, the luminosity at the interaction point of CMS will reach  $L = 7.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ , re-

sulting to about 200 aggregate proton-proton collisions per event. To manage the new processing requirements, the current Trigger system of Phase-1 will be decommissioned and completely replaced by the new Trigger system that is currently under development. This chapter describes the Trigger and Data Acquisition, as of 2020, the time when the Technical Design reports of projects were published. While the core of the projects will remain the same, it is possible that small modifications might occur since the projects are still in the stage of research and development.

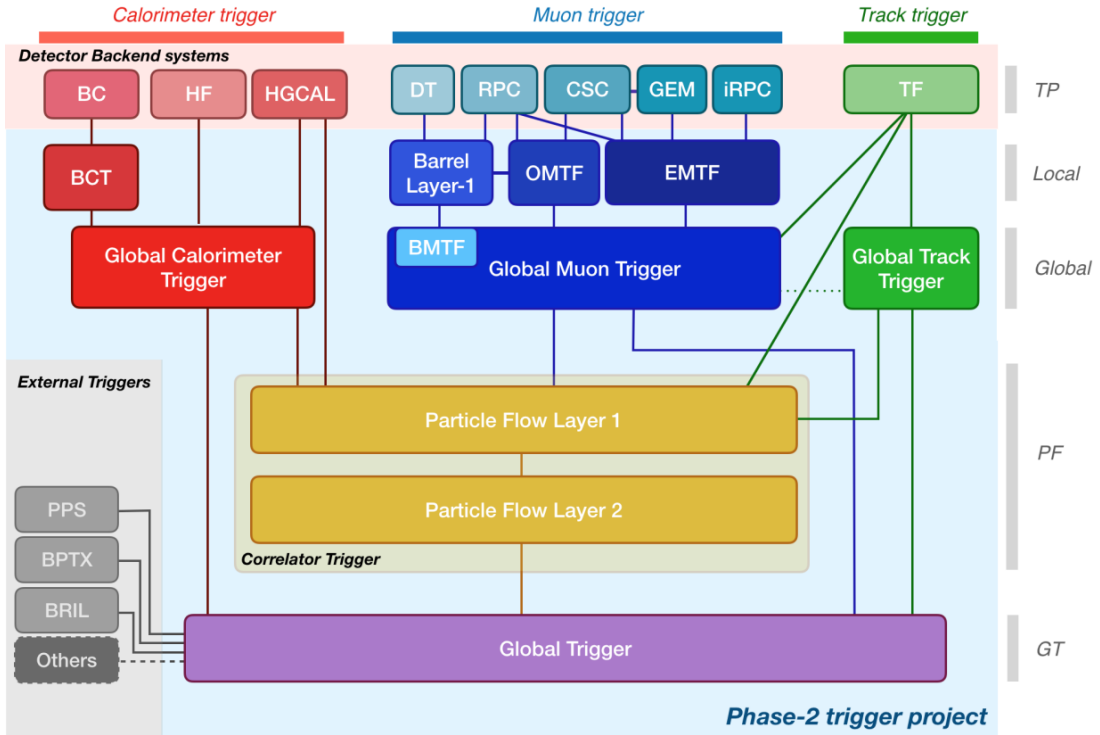
## 5.1 Level-1 Trigger

The Level-1 Trigger system is the system that performs the first online analysis of the data produced on the CMS detector by the products of proton-proton collisions. The system is based on custom designed boards that follow the ATCA standard and use powerful FPGAs as processing units. It is installed in racks at the underground service room (USC), located next to the cavern room (UXC) where the CMS detector is located.

The block diagram of the Phase-2 Level-1 Trigger is illustrated in 5.1 [25]. L1T utilizes a pipelined architecture, processing in layers the information that comes directly from the on-detector electronics. The first data processing takes place in three different datapaths: the Calorimeter trigger, the Muon Trigger and the Track trigger. The first layer of each datapath receives digitized information from the front-end electronics through dedicated optical links at up to 10 Gbps and using the lpGBT protocol. At this stage, raw or minimum processed data are fed into algorithms that perform the Trigger Primitives Generation (TPG). Depending on the nature of the Trigger Primitives (TP), they usually contain information such as local energy deposition, particle hits, position in  $\phi$  and/or  $\eta$ , quality of the measurement, bunch crossing that corresponds to that TP, and others. At the next processing layer, TPs are used by each subsystem for the reconstruction of particles. For every bunch crossing, particle objects are sorted and eventually transmitted to the Correlator Trigger datapath that runs the Particle Flow algorithm to reconstruct the whole event. The Global Trigger is the final processing block that receives inputs from all previous datapaths. It runs the menu of algorithms to calculate the trigger decision. The result, called Level-1 Accept signal, is transmitted back to the trigger and front-end subsystems through the Trigger and Control Distribution System for phase-2 (TCDS2). Its reception initiates the transmission of the detector information to the HLT, which is performed by the Data and Acquisition system.

### 5.1.1 Calorimeter Trigger

The Calorimeter Trigger processes information from the calorimeter detectors of CMS, which are the ECAL and HCAL at the barrel region, and HF and HGAL at the endcap regions. The main calorimeter trigger objects are electrons, photons, jets, hadronic taus and various energy sums, provided for the pseudorapidity coverage of  $|\eta| < 5$ . A block diagram of the calorimeter trigger is shown at 5.2. The Global Calorimeter Trigger (GCT) receives TPs from the back-end electronics of HGAL, HF and the Barrel (or Regional) Calorimeter Trigger (BCT). BCT is the back-end system

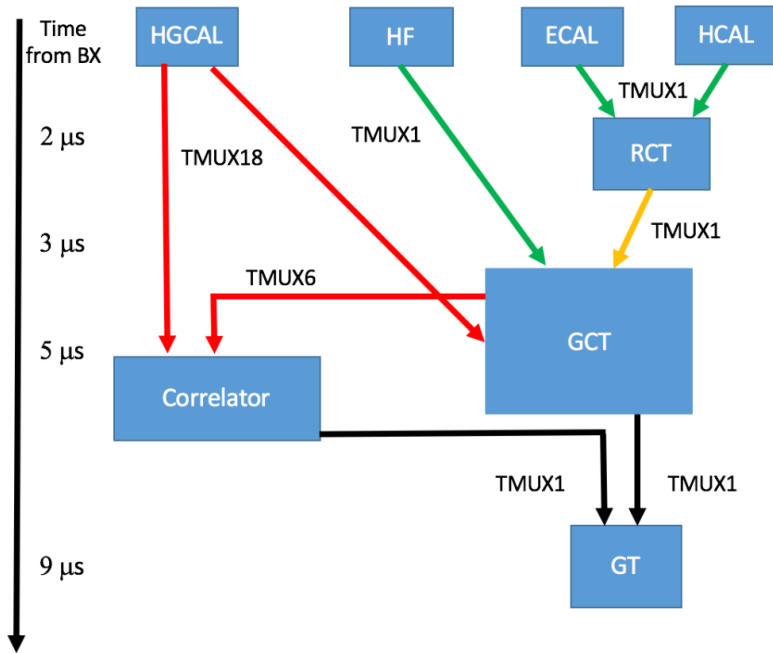


**Figure 5.1:** Functional block diagram of the CMS Level-1 Trigger system for Phase-2.

that generates TPs of the ECAL and HCAL detectors. TPs for the HGCAL are generated by its own back-end system, which is not considered part of the Calorimeter Trigger.

Data from every single ECAL crystal will be transmitted by the new Very Front-End electronics to the BCT FPGAs. The sampling frequency for Phase-2 is 160 MHz, four times higher than Phase-1. The coverage of the 61,200 crystals of the barrel ECAL requires the usage of 108 BCP boards. These boards will perform the readout of the ECAL system, the clustering of crystals to 3x5 towers and the generation of ECAL TPs. The same system also receives data from the 2304 towers of the HCAL. The sampling takes place at 40 MHz. TPs from both calorimeters are transmitted from the BCT system to the GCT boards in Time Multiplexed (TMUX) fashion with period 1.

The TP generation of the endcap calorimeter, HGCAL, is performed in a separate system and the result is transmitted to central L1T. The magnitude and complexity of the HGCAL detector is such that this separation is required. Each of the two parts of HGCAL is composed of 28 sensitive layers made of silicon for the electromagnetic section, and 22 layers of plastic scintillator for the hadronic section. The total number of trigger cells is about 60 million. They are read-out by HGCROC (High-Granularity Calorimeter ReadOut Chip) chips and pre-processed and transmitted by the ECON-T (Endcap CONcentrator Trigger) ASIC. The transmission to the back-end boards is realised through almost 5000 optical links. For the completion of the TP generation, a two stage back-end system is designed that consists 48 boards in the first stage and 24 boards in the second (as of 2019). The information of HGCAL is sent to GCT and



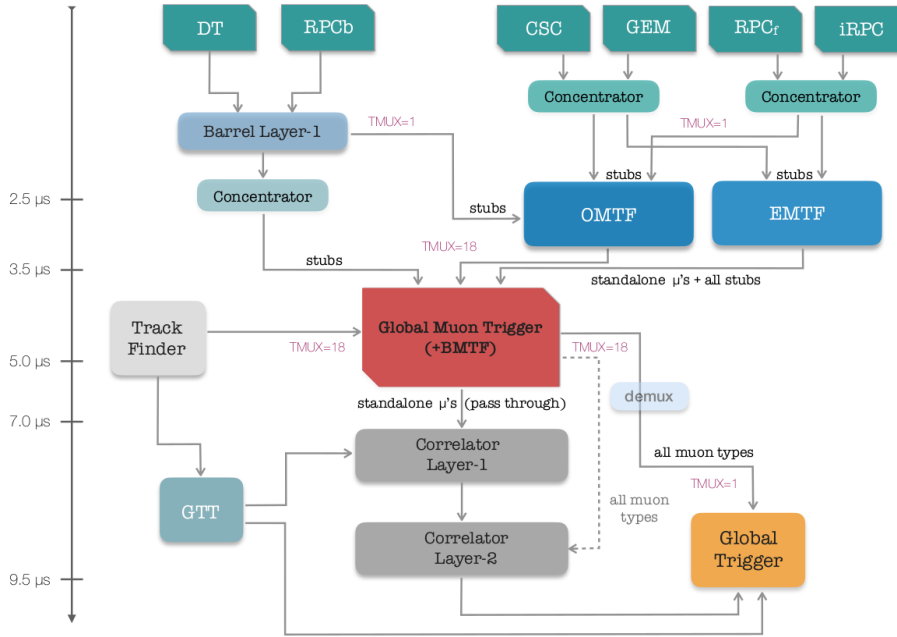
**Figure 5.2:** Block diagram of the Phase-2 Calorimeter trigger.

the Global Trigger (CT) simultaneously. In the two endcaps, TPs are also generated for the Hadronic Forward calorimeter, containing information of energy and timing of the forward hadronic calorimeter towers.

The GCT boards receive TP data from the BCT, the HF and the HGCAL. The main tasks of this subsystems are two. One to prepare and send the information of the BCT and HF to the Correlator Trigger (CT). The second is to combine information from all calorimeter subsystems (BCT, HF and HGCAL), calculate calorimeter-based trigger quantities and send them to the global trigger (GT). The number of GCT boards required to cover the calorimeter region is three.

### 5.1.2 Muon Trigger

The Muon Trigger datapath is responsible for the muon identification and reconstruction at the CMS detector. In order to accommodate the geographical characteristics of the muon detectors, standalone muon reconstruction is performed in three subsystems: the Barrel Muon Track Finder (BMTF) at the barrel region, the Endcap Muon Track Finder (EMTF) at the endcap region and the Overlap Muon Track Finder at the overlap region. EMTF and OMTF transmit their output to the Global Muon Trigger (GMT) system. The BMTF algorithm is decided to be placed inside the GMT system, since the utilization footprint of the algorithm is low enough to easily fit in the FPGA processors of these boards. In addition, GMT receives tracks from the Global Track Trigger. With this information, it is able to combine standalone muons with tracks to produce track-matched muons, and also generate global muons. The GMT output is transmitted both to the Correlator trigger and to the Global trigger. The overall architecture of the Muon Trigger is illustrated in Figure 5.3.



**Figure 5.3:** Block diagram of the Phase-2 Muon Trigger system.

### Endcap Muon Trigger

The endcap muon detector is made of four CSC disk layers with a total of 540 chambers. The pseudorapidity coverage provided is  $0.9 < \eta < 2.4$ . The CSC chambers are organized in stations called ME1, ME2, ME3 and ME4. There are six layers in every chamber with each layer containing cathode strips that run radially and wires placed almost orthogonal to the strips. Trigger primitives at the endcap, also called local charged tracks, are build by the front-end electronics. Their generation requires hits that create straight line patterns to at least four layers. They are produced by the CSC Trigger Motherboards (TMBs) and are transmitted to the back-end track finder systems through the Muon Port Cards (MPC). Out of the whole CSC range, the region  $1.2 < \eta < 2.4$  is sent to the EMTF.

At the region of  $0.9 < \eta < 1.9$  four stations of RPC detectors are also installed in the muon endcaps. This coverage is extended and reaches  $|\eta| < 2.4$  by the new improved RPC chambers. RPC hits are transmitted to the EMTF system using optical links. A concentration of the RPC and iRPC is possible to take place before data are received to EMTF boards. They are used to assist the timing resolution of the generated TPs by combining them with the CSC hits.

The EMTF subsystem processes the TPs from the endcap chambers to reconstruct standalone muons and assign them their track parameters. The processing is divided in six sectors divided in the  $\phi$  view. Having one back-end board processing one sector, 12 boards are needed in total to facilitate the EMTF. The produced muons as well as all the non-zero trigger primitives are transmitted to the GMT in a Time Multiplexed architecture with period of 18.



## Overlap Muon Trigger

In order to assist the efficient muon detection at the region where CSC and DT chambers overlap, the OMTF system is used. It covers the pseudorapidity range of  $0.9 < \eta < 1.2$  and receives trigger primitives that belong to this region both from the barrel and the endcaps. OMTF performs track matching, reconstructs muons and assigns them their track parameters. It does so by splitting the processing in three sectors in azimuth on each side of the barrel. That structure requires six trigger processor boards to be used by the system. The output of OMTF, that consists of standalone muons, along with all non-zero trigger primitives is transmitted time-multiplexed to the GMT. The period will be of 18.

## Barrel Muon Trigger

The barrel muon detector is made of 60 Sectors where DT and RPC chambers are installed. In each Sector, there are four DT stations and four RPC stations. The DT chambers consist of three superlayers of DT cells, two of which measure the phi view and one the theta view. When a muon crosses a DT cell, a signal is generated and measured by the on-detector electronics. This signal is digitized and the resulting TDC hit is transmitted to the back-end electronics using optical fibers. The trigger primitive generation on the barrel region is performed on the back-end by a system called Barrel Muon Trigger Layer-1 (BMTL1). Processing of TDC hits for every chamber results to the production of stubs, or track segments, that include information of the position in  $\phi$ , the bending of the stub, bunch crossing number, quality of the TP, etc.

To cover the totality of the barrel muon detector, BMTL1 requires 60 trigger processor boards in the scenario that every board processes one Sector. However, the scenario where one board processes two Sectors is also possible. In the latter case, an additional processing layer might be used before data are received to GMT. Muon reconstruction for the barrel is performed in the GMT boards. Thus, the generated TPs are transmitted to GMT over optical fibers in a time-multiplexed architecture with period 18.

The barrel muon trigger is discussed in depth at Chapter 6.


## Global Muon Trigger

The Global Muon Trigger (GMT) system collects reconstructed muons from the overlap and endcap muon track finders and trigger primitives from the barrel region. In addition, track data are sent to GMT from the L1 track finder. Barrel muons are reconstructed at the GMT boards using a Kalman filter algorithm, described in section 6.4.1. The Kalman output includes standalone muons with the possibility to also detect displaced and slow muons. Next operation of GMT is to detect and cancel duplicate muons that are found on the borders of the previous subsystems. These actions are all completed before the reception of L1 tracks. The track information is used by GMT to match standalone muons with tracks and produce the so-called tracker muons. This computation provides a significant increase to the efficiency of the measured muons and would not be possible without the track information.

The GMT system receives data from all subsystems in a time-multiplexed period of 18. The number of boards needed to build the system is 18 as well. Optical links are received from BMTL1, OMTF, EMTF and GTT. The output of the GMT is sent to the Correlator and the Global Trigger systems.

### 5.1.3 Track Trigger and Global Track Trigger

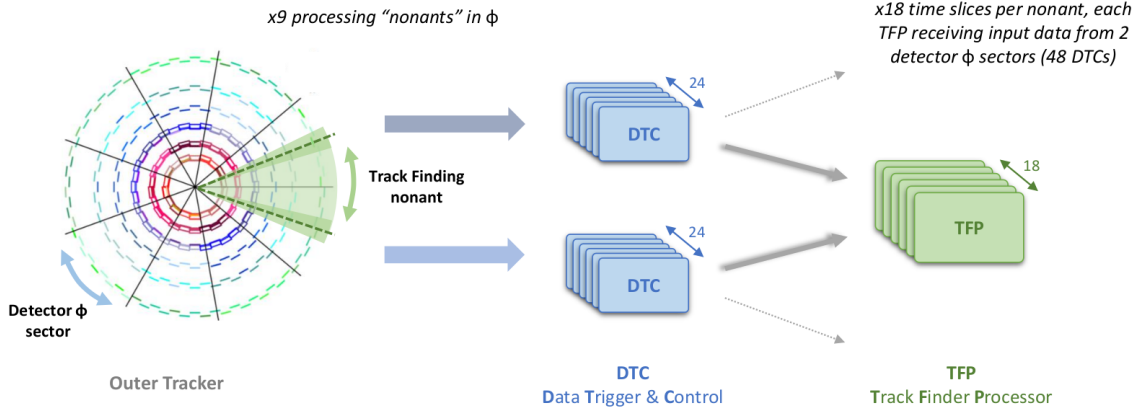
The inclusion of track information to the Level-1 trigger is a new feature that opens a wide range of data processing capabilities at L1T, that otherwise could only be performed to the High Level Trigger. The new Tracker of CMS is designed so that hit information from the Outer Tracker ( $|\eta| < 2.4$ ) is delivered to a new back-end system, called Track Finder. This delivery, however, cannot be achieved if raw hit information was to be sent at 40 MHz. The number of particles produced inside the detector volume at pileup scenarios of 200 reaches that of thousands. For this reason, the Outer Tracker  $P_T$  modules (discussed at section 3.3.2) are designed to perform a first processing to apply a momentum cut by identifying its bending due to the strong magnetic field. As illustrated in Figure 5.4, stubs from the two sensor sides are correlated and a threshold of around 2 GeV is applied for the production of the selected hit pairs, called stubs. The applied threshold value is programmable and can be tuned independently for every module, depending on its location. The amount of data remaining after this reduction is feasible to be transmitted to the back-end system at 40 MHz. In parallel, full resolution information is stored in pipeline buffers and is read-out upon reception of L1 Accept signals.



**Figure 5.4:** Channel correlation at the two sides of a  $P_T$  module. Green channels show hits that pass the 2 GeV threshold.

The back-end system that receives the Outer Tracker stubs and produces particle tracks is called Track Trigger. The architecture of the system is illustrated at Figure 5.5. Processing of the detector is divided in nine  $\phi$  sectors, called nonands, with a margin of overlap foreseen. Stubs are transmitted from the front-end ASICs through optical fibers to the boards of the first back-end layer. These boards are called Data Trigger and Control (DTC) and are based on FPGA processors. The number of DTC boards required to process one nonand sector is 24, making 216 the total DTC boards needed. One of their main functions is to receive information from the  $P_T$  modules, perform pre-processing of the stubs and distribute them to the next processing layer. In addition, DTCs are in charge of interfacing the detector to send timing, trigger and control data to the front-end electronics and perform the detector calibration.

Readout of selected events to the Data Acquisition system is also executed by DTC boards.



**Figure 5.5:** *Track Trigger architecture*

The second layer of processing is made of the Track Finder Processor (TFP) boards. The interface between the two layers is time-multiplexed (TM) with period of 18 and each TFP receives data from 2 neighboring sectors. On every TM period, one TFP processes information of one  $\phi$  sector. Hence, to cover the whole detector in every time slice the total number of boards required is 162. The algorithm that runs inside the TFP FPGAs uses stubs to reconstruct particle tracks. As illustrated in 3.6, Outer Tracker consists of 6 barrel layers and 5 disks. Reconstruction starts by forming the first seeds, called tracklets, by combining stubs from different combinations of two adjacent layers or disks. An additional requirement for the tracklets is to originate from the beamspot. On the next stage, tracklets are projected in both sides to find more stubs in a narrow window of other disks and layers. To form a track, at least four stubs are required. Once tracks are found, a Kalman filter algorithm is used to identify the best track candidates and assign them their track parameters. Furthermore, in order to enhance the search for long-lived particles the usage of *triplets* as the seed of the track reconstruction without the constraint of passing through the beamspot is also used.

The Track Trigger system transmits tracks to the Global Track Trigger (GTT) subsystem. The interface between the two is performed in a time-multiplexed architecture and the period is 6. In every time slice, two GTT boards are required for the processing of a whole event, making the total number of GTT boards being 12. GTT boards are responsible for the reconstruction of the primary vertex (or vertices) and its transmission to the Correlator Trigger Layer-1 subsystem. In addition, GTT reconstructs track stand-alone objects, such as track jets, track  $E_T^{miss}$  and single isolated tracks. Those are transmitted to the Global Trigger boards, as well as to the GMT subsystem. Apart from those, there are many possibilities for additional GTT algorithms that are being studied.

#### 5.1.4 Correlator Trigger

The Correlator Trigger subsystem (CT) exists at the center of the L1 Trigger design, receiving trigger objects from all previous subsystems. Namely, inputs to CT are the

Track Finder, GTT, HGAL, GCT and GMT. The CT uses all this information to run, for the first time in L1 Trigger, the Particle Flow [26] and PileUp Per Particle Identification algorithms (PUPPI) [27]. Using these algorithms already in L1 Trigger is a great advantage of the upgraded system. Its implementation is feasible due to a combination of reasons, such as the existence of track information to L1, the increased latency of  $12.5 \mu\text{s}$  and latest developments of the technologies used at L1.

The Particle Flow (PF) algorithm is being used by HLT and offline reconstruction at CMS already since the first Run of LHC. By using information from all detector subsystems, PF identifies and reconstructs every single particle produced in every event. Tracker information is mandatory for this process. First step is to match stand-alone muons with tracker tracks and exclude the best matching tracks from further processing. Next, tracks from charged particles are combined to energy clusters to define charged hadrons and electrons. This process is executed independently in the endcap and barrel regions and the improved granularity provided by the upgraded detectors enhances significantly the result. Clusters not linked to tracks are characterized either as photons or neutral hadrons. Once PF objects are computed, CT executes the PUPPI algorithm to determine which particles originate from the primary vertex and which are considered as pileup.

The CT subsystem is implemented in two layers, the CT Layer-1 (CTL1) and CT Layer-2 (CTL2). The implementation of PF and PUPPI algorithms, as described in the paragraph above, takes place in the boards of the first layer. Since PF algorithm runs locally, separate boards process different regions of the detector. To cover the total detector, 18 CTL1 boards are going to be used. The PF objects produced by CTL1 are transmitted to the boards of CTL2. At layer 2, the particle-flow objects are used for the reconstruction of physics objects, that are jets, hadronic taus, electrons, photons and global energy sums. This is performed in 30 boards that will comprise the layer-2 of the correlator trigger. The output is then transmitted to the last L1 subsystem, the Global Trigger.

### 5.1.5 Global Trigger

The Global Trigger is responsible for running the CMS menu of algorithms and produce the L1 Accept signal. It receives sorted lists of physics objects from the Global Calorimeter Trigger, the Global Muon Trigger, the Global Track Trigger and the Correlator Trigger Layer-2. Received data are buffered, de-multiplexed and fed to the global trigger algorithms. There are  $O(1000)$  algorithms that run at GT, each one describing a specific event signature of the trigger menu. All algorithms evaluate in parallel the input data and a trigger decision is made in case any of those is fired. This results to the production of the L1A signal alongside with the type of trigger that fired. Some additional functionalities of GT that exist to aid the overall trigger performance are the ability to prescale algorithms, enable individual algorithms to fire only in specific bunch crossings of the orbit and the monitoring of the rate of each algorithm.

The architecture of the GT system is split in algorithm boards and in the final-OR board. Each of the algorithm boards receives the same copy of the total GT inputs and runs a subset of the algorithms menu. The number of boards required to fit the whole menu is twelve. Their output is transmitted to the GT Final-OR where the

final-or of the trigger is computed. Other functions that may be performed in the Final-OR board include the prescaling and monitoring of the algorithms rate. The accept decision as well as the trigger type are sent to the TCDS2, which delivers it to all previous L1 subsystems as well as to the front-end pipeline buffers. In addition, GT interfaces the DAQ and HLT systems to transmit readout records with information regarding the triggers that have fired.

### 5.1.6 Scouting System

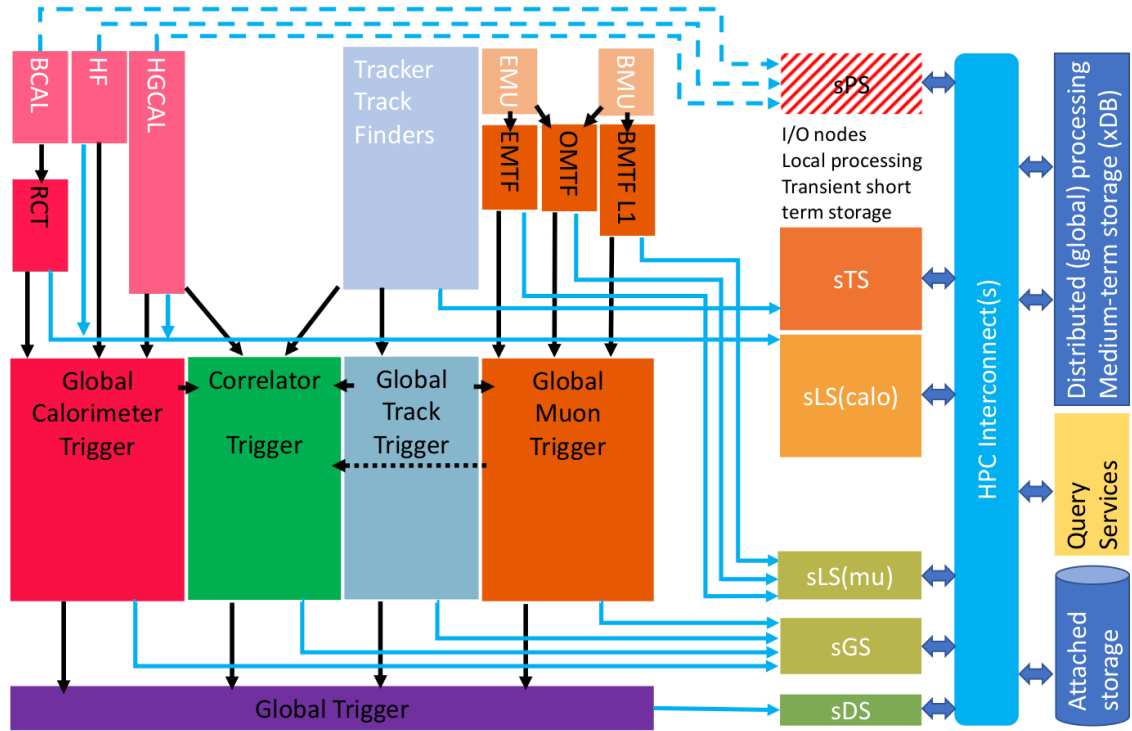
The Scouting System [28] is a new feature of CMS data taking. It is part of the Phase-2 upgrade project but it is already being demonstrated during Run 3 by capturing part of Level-1 data using commercial boards. The idea of the Scouting System is to collect Level-1 data at the production frequency of 40 MHz. In other words, provide a semi-online analysis in the totality of the produced events. This possibility enables the study of specific physics channels that require lower signal rates and/or thresholds that are otherwise not provisioned by the GT menu. In cases where potential signals are found, dedicated algorithms can be implemented afterwards in GT to aid further study. Furthermore, monitoring of the functionality of different L1 Trigger subsystems becomes possible by capturing their inputs and outputs.

The Scouting System is a completely independent system that runs in parallel with L1 Trigger. The design collects data from spare output links of different trigger components at 40 MHz. The block design is illustrated in Figure 5.6. Blue arrows represent links that transmit data to the different scouting systems. The connection deployment will be done in stages. First stage connects the output of the GT Final-OR and some inputs of the GT to the scouting Decision System (sDS). Also, the output of the GCT, GMT, GTT and CT will be sent to the scouting Global System (sGS). This first stage of data capturing requires modest bandwidth and will provide great diagnostic functionality for the GT, as well as interesting physics capabilities. The second stage involves the scouting Local System (sLS) and the scouting Track System (sTS). Outputs from the local muon and regional barrel and endcap calorimeter triggers are captured by sLS and tracks from the track finders by the sTS. Last stage would be the connection of primitives themselves coming from the endcap and barrel calorimeter to sPS. This stage involves, however, throughput capability of about one order of magnitude larger than the others combined.

### 5.1.7 Hardware

As already mentioned, the Level-1 Trigger system runs on custom designed electronic boards. More specifically, trigger processors are based on FPGA devices that provide high speed data transfer and processing. The design and production of L1T boards is carried out by collaborations of institutes all around the world. To accommodate the hardware requirements of every trigger subsystem, there are four hardware production lines. Some of them design general purpose boards that will be used in multiple subsystems, while others are targeting more specific use cases.

The ATCA standard, described in section 4.1, is the chosen form factor for all L1T boards. In the subsections that follow, a description is given of prototype boards produced by every hardware line, as presented at the latest L1T Technical Design Report



**Figure 5.6:** Block design of the Phase-2 Level-1 Trigger including the 40 MHz scouting system.


in 2020 [25]. At the time of writing, the hardware development has matured and current designs have reached the final production versions of the boards. However, not all boards have been produced and tested and/or not all designs are finalized. Hence, the description below refers to the prototype boards whose operation is demonstrated and verified.

### APd1 ATCA board

The Advanced Processor (AP) consortium consists of several CMS institutes that collaborate for the production of general purpose ATCA hardware, software and firmware. The main board family, called APx, provides a large processor based on the Ultrascale Plus architecture of Xilinx. The board adopts a modular design that splits many of the onboard circuitry on different mezzanine card types.

Figure 5.7 depicts the first APx prototype board, the APd1. The FPGA used is a XCVU9P and the optical connectivity includes 76 transceivers that operate to line rates up to 28 Gbps. They are arranged in 19 Samtec Firefly optical modules [29]. In addition, 24 channels are connected to the Rear Transmission Module through the backplane connector Zone 3. The main hardware control is performed by the Embedded Linux Mezzanine (ELM) and the IPMC. Once the card is successfully power on, a Linux operation system runs on a XCZU04CG ZYNQ on the ELM and is used for many functions, such as initialization and programming the FPGA, configuration of support devices, such as the Firefly modules and clock circuits. To communicate between the ELM and the FPGA, an AXI (Advanced eXtensible Interface) bridge is implemented. The Ethernet connection from the ATCA Hub Slot 2 is performed by

a 10 GbE Ethernet PHY.




**Figure 5.7:** *The APd1 ATCA board featuring a VU9P FPGA.*

### Serenity ATCA board

The Serenity collaboration is also established to produce general purpose ATCA hardware for Phase-2, as well as the corresponding firmware and software. The ATCA platform consists of two basic elements. One is the base board which includes all necessary electronic components and services. Second is the processing FPGA chip that comes attached on the so called daughter-cards. The daughter-card communicate with the base board through an interposer Printed Circuit Board (PCB). This way, different flavors of FPGAs can be easily changed to suit the applications of different subsystems. The prototype Serenity version, shown in Figure 5.8, is able to host two such daughter-cards into one single ATCA blade.

Serenity boards use the CERN IPMC [21] that provides the IPMC functionality required by the ATCA standard. High level control of the FPGA and the on-board circuits is performed by a COM Express computer-on-module. The Linux operating system hosted in the COM Express communicates with the two daughter-card sites through PCI Express lanes. This way, control and monitor of the firmware registers is executed using high-level software commands. In addition, a third, small FPGA is used (Artix-7) to carry out the protocol and voltage conversions required to interface the JTAG and I2C chains. On every FPGA site there are 72 optical channels served by Samtec Firefly transceivers [29]. There are also 64 differential pairs routed through the PCB between the two daughter-card sites. Moreover, 2 channels are served by a Quad Small Form-factor Pluggable (QSFP) module. The variety of FPGAs that exist on daughter-card PCBs include Ultrascale Plus devices, such as KU15P, VU7P and VU9P.



**Figure 5.8:** *The Serenity ATCA board hosting two daughter-cards.*


### **Ocean ATCA board**

Ocean is a general-purpose prototype ATCA board. Its main difference with the other hardware lanes is that it utilizes a large System-On-Chip (SoC) that combines the FPGA logic with a multi-core processor in a single chip. This way, one device is used for the algorithmic logic and the control of both the firmware and the peripheral devices of the board, reducing significantly the complexity of the PCB design. The SoC used by Ocean is a Xilinx ZYNQ Ultrascale Plus (ZU19EG). A Linux operating system runs on the processor of the ZYNQ which is interconnected with the FPGA logic through several AXI lanes. These can be used for applications such as configuration, monitoring and real-time scouting. The optical connectivity supports 28 transceivers capable of speeds up to 28 Gbps, and 44 supporting line rates up to 16 Gbps. The 72 in total optical channels are instrumented with Samtec Firefly modules. Furthermore, Ocean hosts the CERN-IPMC and a small Artix-7 FPGA. A picture of the Ocean board can be seen in Figure 5.9.

### **Barrel Layer-1 demonstrator**


The Barrel Layer-1 board is the only hardware platform that aims to design and manufacture a ATCA blade to be used in a specific trigger subsystem. The first prototype can be seen in Figure 5.10. Its purpose was the development and evaluation of trigger hardware technologies and as such, it does not follow a specific PCB form factor. The processing unit is a Kintex Ultrascale KU040 FPGA. Optical connections include 12 channels routed to a Samtec Firefly connector and 4 routed to a QSFP module. All 16 links can reach speed up to 16 Gbps. The board also hosts a small ZYNQ SoC that is connected with the clocking network and the power modules.





**Figure 5.9:** *The Ocean ATCA board featuring a ZU19EG ZYNQ.*

It is used to provide basic control and monitoring over these circuits. The ATCA prototype that was produced after this demonstrator is described in Chapter 9.



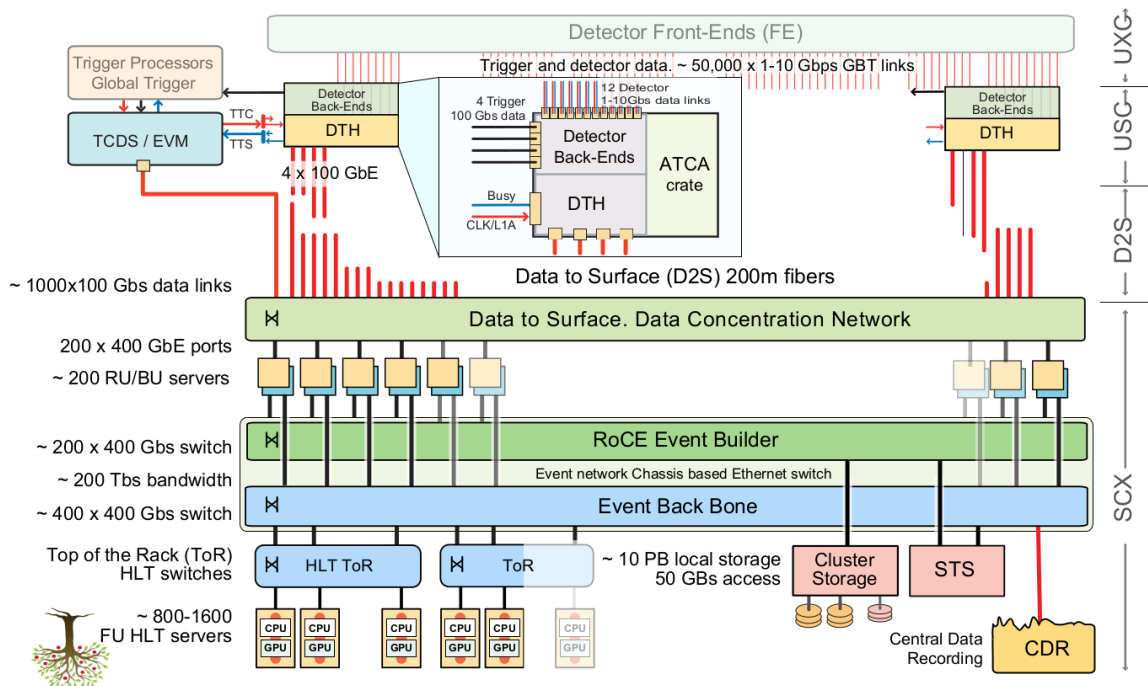
**Figure 5.10:** *The Barrel Layer-1 demonstrator board featuring a KU040 FPGA.*

## 5.2 Data Acquisition and High Level Trigger

### 5.2.1 DAQ system Overview

The collision rate at HL-LHC will be 40 MHz and the luminosity at its ultimate configuration is going to reach  $L = 7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ . On every bunch crossing, the proton collisions produce thousands of particles that are detected by the CMS detector systems. This information is read by the front-end sensors, digitized and transmitted

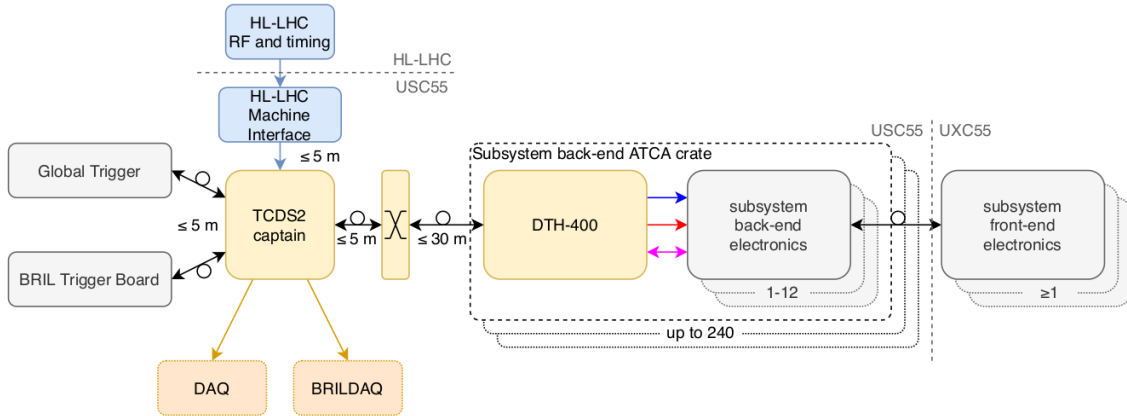
to the back-end electronics in order to be processed by the Level-1 Trigger system. The L1T performs the first stage of online data analysis in order to reduce the event rate. During the ultimate luminosity phase, the output rate of L1T is going to be 750 KHz and the latency of the accept or reject decision  $12.5 \mu\text{s}$ . Due to bandwidth and processing power requirements, data transmitted to L1T do not always provide the high resolution information that is produced. However, full resolution data are stored inside front-end pipeline buffers. The size of these buffers is equal to the latency of the L1 accept decision. This decision is calculated by the Global Trigger and is distributed back to the front-end buffers through the Trigger Timing and Control Distribution System (TCDS2 at Phase-2). Upon reception of the L1 Accept, front-end electronics deliver the information from the corresponding event to the readout links of the Data and Acquisition (DAQ) system. In some detector systems this procedure takes place directly at the front-end electronics and in others at the back-end TPG systems. In the latter case, the full resolution information is received from the detector. Data from the selected events are transmitted by DAQ to the High Level Trigger (HLT). HLT runs in a farm of commercial CPUs and GPUs, performs more sophisticated event reconstruction and data processing and further reduces the event rate. The output, or storage, rate of HLT at the ultimate scenario will be 7.5 KHz.



**Figure 5.11:** Block design of the structure of the Data Acquisition system.

The block diagram of the DAQ architecture for Phase-2 is depicted at Figure 5.11 [30]. The detector, as well as front-end electronics are installed in the UXC cavern. They are interfaced by the detector back-end board, that are located in the service room of CMS, UCS. Communication between the two is bi-directional and performed by optical links using the lpGBT radiation hard protocol. The datapath from the front-end to the back-end, called uplink, transmits digitized data from the detector systems to the Level-1 TPGs. Information such as status of the front-end buffers, error counters, etc is also propagated through this link. The opposite direction

datapath is called downlink and is utilized for the distribution of the Trigger, Timing and Control (TTC) signals. These are the master LHC clock, the L1 Accept signal and various fast control signals, such as resets. The status of the front-end pipeline buffers is also transmitted through the uplink datapath. This information is used by the Trigger Throttling System (TTS) which controls the accept signal generation in order to prevent buffer overflow and as a consequence, front-end data loss. Both TTC and TTS are parts of the TCDS2 system. The TCDS2 architecture is showing in Figure 5.12.



**Figure 5.12:** *Structure of the TCDS2 system.*


The DAQ system will be deployed using a custom designed ATCA board, called DAQ and TCDS2 Hub (DTH). There will be two flavors of the DTH board to facilitate two specific functionalities, the DTH-400 and the DTH-800. As mentioned in section 4.1, all L1 processor boards will be installed in ATCA crates. It is foreseen that the total number of ATCA crates will be about 130 (as of 2020, DAQ TDR numbers). Independently of whether or not a trigger subsystem interfaces the detector front-end, TTC information has to be distributed to all ATCA blades that operate in CMS. Hence, at least one DTH blade is going to be installed at slot 1 (hub slot) of every ATCA crate. The TTC signals are then distributed to the backplane bus and reach the processors chips of all cards installed inside the crate. The transmission is performed by high speed lanes and through a custom protocol.

Apart from accommodating the TCDS2 system, DTH is also responsible for the central readout of the selected events. The desired information is kept in buffers at the different sub-detector boards, waiting for the L1 Accept signal. The transmission is performed using a custom bi-directional protocol, called SlinkRocket, and through optical fiber connections of the front panel of the ATCA blades. Data from every accepted event must be concentrated, combined into event data structures and transmitted to the surface data center (data to surface, D2S). This transmission is done using standard commercial protocol that provides the possibility of re-transmission in case of losses.

## 5.2.2 DTH ATCA board

The two flavors of DTH boards, also called D2S boards, will be used to instrument the DAQ system. First flavor, DTH-400, is the basic board as it provides both the

TCDS2 and readout functionalities that DAQ should perform. It must be installed in the first slot of every CMS ATCA crate. Its block design is illustrated at the left side of Figure 5.13. The TCDS2 functionality is carried out by a Xilinx VU35P FPGA that receives the optical link from the TCDS2 master and distributes it to the appropriate backplane bus of the ATCA crate. The readout is performed by a second VU35P FPGA. To accommodate the receiving optical links, this FPGA connects 24 high speed transceivers to 6x4 Firefly optical modules. On the other side, transmission to the surface is performed by 5 QSFP modules, 100 Gbps of bandwidth each. The maximum average throughput of this board is 400 Gbps. The prototype board is shown on the right of Figure 5.13.



**Figure 5.13:** *Left: Block design of the DTH-400 board. Right: The prototype DTH-400 ATCA board.*

The second D2S board flavor is called DTH-800. This board instead of facilitating the TCDS2 functionality it contains a second readout FPGA, thus providing an additional 400 Gbps readout bandwidth to the back-end ATCA shelf. It is designed for projects that require more throughput than that provided by DTH-400. Since it doesn't contain the TCDS2 block, it is considered as a standard ATCA board and receives the TTC stream through the backplane connector. In that sense, it cannot replace DTH-400 but be used complementary to it. The block design of DTH-800 is shown on the right of Figure 5.13.

### 5.2.3 Data to Surface and Event Builder

As illustrated in the DAQ block design of Figure 5.11, transmission of the readout data from the DTH boards to the High Level Trigger is performed by the Data to Surface (D2S) system. Events that have been accepted are packed in readout packets and sent over the SLinkRocket protocol to the D2S boards. They then have to be delivered to the Event Builder system which is located on the surface of P5. The transmission will be performed using a commercial and reliable protocol, the TCP/IP. This choice offers the possibility of using a standard switched network at the reception side. For the D2S boards to ensure reliable transmission, large on-chip high-bandwidth memory must be used to avoid data loss in cases of congestion in the receiving side.

Before the readout data are directed to HLT, they have to be processed by the Event Building (EB) system. Between the D2S boards and the EB processors, data are handled by an InfiniBand switched fabric. The main purpose of the event builder is to collect all packets that originate from accepted events and merge data that correspond to the same L1 accept into records. These records must be buffered until a HLT node is available for processing. Event Building is implemented in software and it is divided in three units: the Event Manager (EVM), the Read-out Unit (RU) and the Builder-Unit (BU).

#### 5.2.4 High Level Trigger

The High Level Trigger (HLT) performs the second stage of online processing of the collision events that are accepted by the Level-1 Trigger. The ultimate input rate of HLT is foreseen to be 750 KHz and its output, after rate reduction, 7.5 KHz. HLT runs in a farm of commercial processors on the surface of the CMS experiment. It receives accepted events from the Event Builder and performs event reconstruction and selection depending on the HLT physics menu. Data processing is similar to that of Level-1 Trigger but using full resolution of the detector data. Furthermore, the available latency permits the usage of commercial CPUs and GPUs. This means that the version of the reconstruction algorithms that run at HLT can be much more sophisticated than the corresponding versions at L1T. The infrastructure that runs on HLT is the CMS Software (CMSSW). This is the same framework used for offline reconstruction. Data processing at HLT can run in parallel by distributing events to processor nodes. The average processing time of an event is 200-300 ms. Events selected by HLT are stored locally in order to be assembled into larger data-set files which are then transferred to central computing network (Tier-0). Events are then reconstructed offline and permanently stored to be used for analysis.

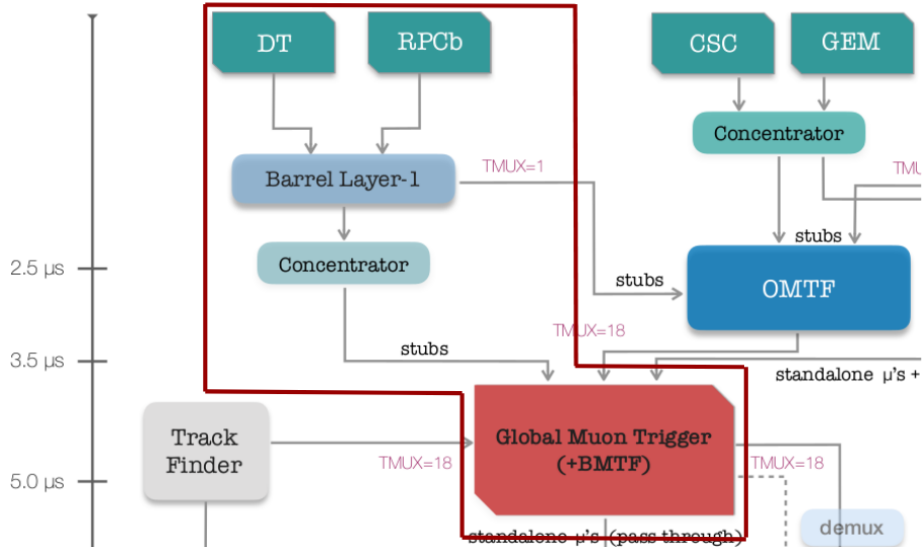
# Chapter 6

## Barrel Muon Trigger

### 6.1 Overview

The Barrel Muon Trigger (BMT) subsystem is responsible for the reconstruction of muon particles in the barrel region of the CMS detector for Phase-2, at pseudorapidity range of  $|\eta| < 0.9$ . As described in section 3.5.1, the barrel muon detector consists of four layers of Drift Tube (DT) chambers installed outside the solenoid magnet. When muons pass through a DT station, they ionize the gas inside the chamber, producing a measurable amount of electrical current. This signal is amplified, digitized and transmitted to the back-end processors in the form of TDC (Time to Digital Converter) hits. The transmission takes place in long optical fibers connecting the DT front-end electronics with the BMT Layer-1 (BM TL1) boards. The BM TL1 subsystem generates Trigger Primitives (TP) for every DT chamber independently. TPs, also called stubs or track-segments, contain information of the position and direction of a muon candidate inside each station. In addition, timing information provided by the RPC detectors is combined to generate the so called *super-primitives*. The next processing layer is the Global Muon Trigger (GMT). Super-primitives are transmitted to GMT and are used for the reconstruction of muon tracks inside the barrel. The momentum of the reconstructed muon candidates can be extracted from the track curvature inside the magnetic field. Eventually, muon objects are transmitted from GMT to the Correlator Trigger and Global Trigger subsystems. A block diagram of the BMT subsystem is highlighted in red at Figure 6.1.

The barrel muon system is made of DT chambers grouped in Sectors. There are four chamber stations in every Sector stacked in layers one behind the other, having three layers of the return iron-yoke placed between them. The DT Sectors are grouped in five disks in the  $z$ -axis, called Wheels. Starting from the negative direction and moving to the positive side of  $z$ , the Wheels are named as: Wh-2, Wh-1, Wh0, Wh+1 and Wh+2, where Wh0 is located at  $z = 0$ . Each wheel contains twelve Sectors. Thus, the total number of DT Sectors is 60. The stations inside a Sector are labeled as MB1, MB2, MB3 and MB4. MB1 is the first chamber installed right after the solenoid magnet and MB4 the outer one. Each station is made of one chamber, except from MB4 of the top and bottom Sectors that contain two chambers due to their size. Thus, the total number of Chambers is 250. The transverse view of one DT wheel is depicted at Figure 3.13.



**Figure 6.1:** *Barrel Muon Trigger Architecture.*

Complementary to the DT system, RPC detectors are installed in the barrel region to assist the trigger primitive generation with their finer time resolution. Similar to the DT architecture, there are six layers of RPC chambers installed around the beam axis, forming concentric cycles. Their on-detector location is at the sides of the DT chambers. At MB1 and MB2, four RPC stations are installed in both sides and at MB3 and MB4 two are installed in one side, as shown in Figure 3.12.


## 6.2 Front-End Electronics

### 6.2.1 Drift Tube Front-End

The basic unit of the muon barrel detector is a drift cell, as described in 3.5.1. Drift cells are grouped together in four layers of half-staggered cells to form a superlayer (SL) (Figure 3.14). Every DT chamber consists of two superlayers (SL1, SL3) oriented parallel to the  $z$ -axis (beamline), providing a measurement on the transverse ( $r$ - $\phi$ ) plane. Stations except for MB4 are also instrumented with a third superlayer (SL2), the wires of which are placed vertically to the  $z$ -axis. This way it offers a measurement on the  $r$ - $\theta$  view. When muons pass through the cells they ionize the Ar/CO<sub>2</sub> gas, producing an amount of electric charge that drifts to the anode wire with a maximum time of about 390 ns. The resulting current is captured, amplified and digitized by the front-end DT boards.

The Phase-1 Read Out Boards (ROB) and all on-detector electronics are going to be replaced for Phase-2 by the new OBDT boards (On detector Board for Drift Tubes) [31]. Contrary to the current architecture, where DT TPs are generated on-detector, the upgraded front-end system will only perform time digitization of muon hits. The first prototype version of the OBDT board can be seen in Figure 6.2. It is based on a radiation hard FPGA, the Microsemi Polarfire MPF300, that is responsible for the time digitization of the chamber signals. The total number of input channels 240. Time digitization is performed inside the FPGA logic by sampling each channel with

a 600 MHz Double Data Rate (DDR) deserializer. The deserializer output ends up to a parallel bus of 30 bits for every bunch crossing, clocked to the 40.078 MHz clock. This 30-bit bus encodes the hit time as a transition of 0s to 1s, with a resolution of 30 bins for every 25 ns, or resolution of 0.83 ns. The final TDC hit encodes the bit position of the first binary 1 that is found in the 30-bit bus, thus requiring a field of 5-bits per hit at the final OBDT payload.



**Figure 6.2:** *The first version of the OBDT board.*

The OBDT data interface is facilitated by a SFP+ and a QSFP+ module. The SFP+ is dedicated for slow control and synchronization. It establishes a bi-directional link that implements a front-end protocol, either the Phase-1 GBT or the Phase-2 lpGBT (described in 7.3). Through it, TTC and slow control information reaches the OBDT, which in turn distributes it to the front-end boards. It includes the LHC 40.078 MHz clock, Bunch Counter and Orbit Counter resets, and other configuration and monitoring signals. The QSFP+ transceiver delivers hit data asynchronously to the back-end through optical links and by using a front-end optical protocols, as well. The data format of one hit, as generated by OBDT, can be seen in Figure 6.3. It contains the wire number of the hit cell, the bunch crossing it occurred and the corresponding TDC time with respect to the start of the LHC orbit.

24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
OBDT Channel Number										Bunch Crossing										TDC Value				

**Figure 6.3:** *Data format of a single TDC hit.*

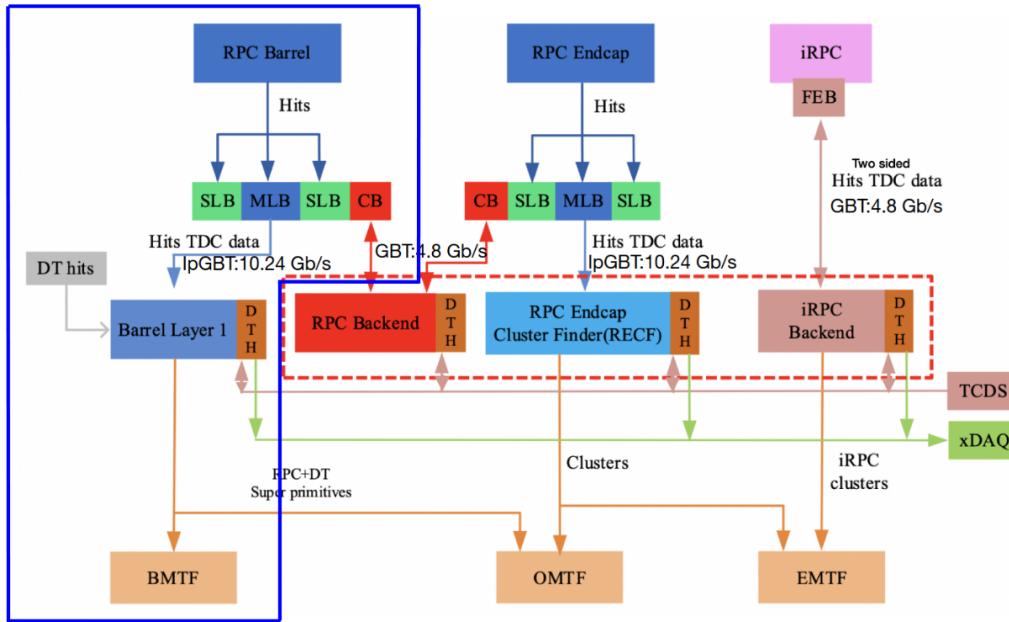
At the time of writing this thesis (2023), two versions of the OBDT board have been produced. The basic architecture and functionality of the two remains the same. An important difference between the two is the optical transceiver they utilize. The first prototype is designed to use the Phase-1 GBT ASIC and the second revision is based on the Phase-2 lpGBT ASIC. In both cases, the transceiver ASIC defines the protocol that is used by both the SFP+ and the QSFP+ optical interfaces. Another important note is that the second revision, similarly to the final architecture, is produced in two flavors. The differentiation occurs by the location each OBDT flavor is



going to be installed, which can be either an  $r\text{-}\phi$  SL or an  $r\text{-}\theta$  SL. As a consequence, the former board flavor is called OBDT-phi and the latter OBDT-theta.

## 6.2.2 RPC Front-End

The architecture of the RPC and iRPC system for Phase-2 is illustrated in the block diagram of Figure 6.4. The part that belongs to the barrel muon trigger is enclosed inside the blue lines. Muon hits from the six RPC stations of the barrel are sent to the Master Link Board (MLB). At the FPGA of the MLB boards (*Kintex-7*) the hit data are processed and converted into RPC TDC hits. The TDC hits are then transmitted for further processing to the BMTL1 boards at the CMS counting room.



**Figure 6.4:** Block diagram of the RPC Phase-2 system. The Barrel part is highlighted in the blue lines.

## 6.3 Barrel Muon Trigger Layer-1

The first processing layer of the Phase-2 Barrel Muon Trigger is the BMT Layer-1, or BMTL1. This subsystem will be installed at the CMS service room (USC), that is located next to the experiment cavern (UXS). Long optical fibers will transmit muon hit data directly from the on-detector OBDT and MLB boards. The objective of BMTL1 is to generate Trigger Primitives (TP), also called stubs or track segments, for every DT chamber. In addition, RPC information is combined for the production of the super-primitives (SP). The SPs are generated independently for every chamber of the barrel. The main information they carry is position on the chamber, the slope, or bending angle, of the hit and the bunch crossing number. The generation is performed in the BMTL1 ATCA boards that host a powerful FPGA device. DT


hit data are processed by the Analytical Method algorithm and the produced SPs are transmitted to the GMT subsystem.

### 6.3.1 Analytical Method algorithm

The Analytical Method algorithm generates TPs by processing hit data originating from the DT barrel muon system [32]. The algorithm inputs are the wire numbers from all DT cells of one DT chamber, together with their corresponding hit time. The hit time arrives in the form of TDC hits, calculated with respect to the start of the LHC orbit. This information, combined with the known value of drift velocity ( $54 \mu\text{m}/\text{ns}$ ), is analytically processed to reconstruct the position of a hit and its bending angle. Furthermore, the Bunch Crossing (BX) number of the collision that produced the muons is extracted.

The first processing stages of the algorithm are performed on super-layer basis. The initial step, called grouping, is to find cells containing geometrical hit patterns that are possible to originate from a muon trajectory. Only patterns that correspond to straight lines are considered, as shown in Figure 6.5. At these regions, 10 cells at a time are grouped, labeled as candidates and propagated to the next step.

The next step is called fitting, as it implements analytical expressions to determine the collision time and the parameters of the track segment. Since the anode wire that collects the ionized electrons exists in the center of the cells, the hit time is agnostic with respect to the position of the muon (left or right). Thus, the calculations assume all potential combinations of lateralities. Candidates are considered cases with either 3 or 4 hit cells. In each case, analytical calculations are performed for every laterality combination and those that result to physical solutions are considered as valid. The physical information obtained by this method include the position of the hit candidate, its bending angle and the BX number that generated the hit.



**Figure 6.5:** *Left: Groups of 10 cells where combinations of hits are searched for. Right: All cell layouts compatible with a muon straight line inside a SL.*

Once candidates of both  $r$ - $\phi$  superlayers are calculated, they are propagated to the Correlation step. Hit candidates from SL1 and SL3 are correlated if their BX number is within a window of  $\pm 25$  ns. For the correlated candidates the TP parameters are calculated again. The position and BX time are now the arithmetic mean of the previous two. The bending angle is calculated by taking into account the difference in the position of the two hits, since the lever arm is larger due to the distance of the two superlayers. If no match is found, the single SL primitives are kept. Apart from correlated and uncorrelated TPs, there are also the confirmed TPs when two individual candidates are found within a window of  $\pm 1$  cm. This information is

tagged to each TP using an integer number called quality. The values and meaning of each quality are shown in Table 6.1.

Quality	Description	Type
1	3 hit segment	uncorrelated
2	3+2 hit segment	confirmed
3	4 hit segment	uncorrelated
4	4+3 hit segment	confirmed
5	non existing label	
6	3+3 hit segment	correlated
7	4+3 hit segment	correlated
8	4+4 hit segment	correlated

**Table 6.1:** *Qualities of the DT trigger primitive.*

In order to avoid potential duplicates or fake primitives being produced, cleaning filters are also placed in different stages of the algorithmic logic. In the end, TPs are being translated to the CMS global sector coordinates. The output of the Analytical Method algorithm can be up to four TPs per chamber for every Bunch Crossing. The size of a TP is 64 bits and its format is depicted in Table 6.2. On every bunch crossing, stubs from every chamber of the DT system are transmitted to the Global Muon Trigger.

Field	Valid	Index	Wheel	Ch	SL	Q	RPC	$\phi_B$	$\phi$	BX	Fine t0
Bits	1	2	3	2	2	4	3	13	17	12	5

**Table 6.2:** *Data format of a DT Trigger Primitive.*

### 6.3.2 Hardware

The BMTL1 subsystem will be instrumented with a custom ATCA board, designed specifically for BMT Layer-1. This board is the ATCA successor of the demonstrator board for Barrel Layer-1, which is described in section 5.1.7. The main processing unit is a powerful VU13P FPGA. Around it are routed high speed optical transceivers facilitated by Samtec Firefly connectors. There are in total 20 Firefly modules: 10 of them are x4 bi-directional up to 25 Gbps, 7 are x12 receives up to 16 Gbps and 3 are x12 transmitters up to 16 Gbps as well. The asymmetry in the number of inputs and outputs exists due to the larger number of receiving fibers, coming from the detector front-end, with respect to the number of output links. Furthermore, the receiving channels will operate at either 5.12 or 10.24 Gbps, following the standard of the lpGBT protocol, while the outputs will operate at 25 Gbps using the CMS Standard trigger link Protocol.

An image of the board can be seen at the left of Figure 6.7. Apart from the FPGA, the board also features a ZYNQ System-on-Chip. It is used to provide basic configuration and monitoring of both the FPGA logic and the different peripherals attached

on the board. These can be the Firefly modules, chips of the clocking network, the power modules, and others. A Linux operating system runs on the processor system of the ZYNQ, enabling high level control as well as remote access through Ethernet connection. The BMTL1 ATCA board is further described in Chapter 9.

## 6.4 Global Muon Trigger

The Global Muon Trigger (GMT) subsystem collects muon information from the three regions of the CMS detector. Specifically, the endcap ( $1.2 < \eta < 2.4$ ) and overlap ( $0.9 < \eta < 1.2$ ) regions transmit standalone muon objects, as produced by EMTF and OMTF track finders respectively. The barrel region ( $|\eta| < 0.9$ ) transmits TPs in the form of DT track segments. Furthermore, GMT receives track data from the Track Finder subsystem.

There are two main tasks that are implemented at GMT. First is to reconstruct standalone muons of the barrel region. This procedure is performed using track segments that are generated at BMTL1 and transmitted to GMT over optical links. The algorithm that performs muon reconstruction is called Kalman Filter Track Finder (KMTF). It has been decided that KMTF algorithm can run in the FPGAs of the GMT processors, due to its low utilization footprint and availability in latency. Moreover, this decision simplifies the overall architecture of the muon trigger. The second main task of GMT is to match tracks from the Track Finder with standalone muons. The products of this operation are called tracker muons and their efficiency improves significantly with respect to that of standalone muons (muons reconstructed only by hits at muon detectors). Lastly, GMT performs ghost and duplicate cleaning at the overlapping regions and transmits sorted list of muon objects to the Correlator Trigger and the Global Trigger.

### 6.4.1 Kalman Muon Track Finder algorithm

Muon reconstruction at the barrel region of CMS is performed using a Kalman Filter based algorithm. Kalman Filters are widely used in dynamic systems to estimate future states of a discrete linear set. For this reason, they have also been adapted to hadron collider experiments for reconstruction and fitting of particle tracks [33]. At CMS, Kalman Filter based algorithms are being used already since the beginning of Phase-1 by offline reconstructions algorithms. At L1 Trigger, the KMTF algorithm for the barrel region, as described in this section, has been operational since the start of Run 3 [34].

KMTF uses state vectors to define the track parameters of every DT station. The vector is defined as  $x_n = (k, \phi, \phi_b)$ , where  $k = q/p_T$  and  $q \pm 1$  is the muon charge. The first seed of a track is the outer available station and the track parameters and their uncertainties are propagated inwards. The energy loss at this region of the detector is taken into account. The relation between the old station and the new state is given by  $x_{n+1} = Fx_n$ , or:

$$\begin{pmatrix} k \\ \phi \\ \phi_b \end{pmatrix}_{(n+1)} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & b \\ c & 0 & d \end{pmatrix} \begin{pmatrix} k \\ \phi \\ \phi_b \end{pmatrix}_n$$


where the propagation matrix  $F$ , consisting of  $(a, b, c, d)$ , is produced by the geometry of the detector and simulation. The state uncertainties are expressed by a covariance matrix  $P$  and they are propagated to the next station by the transformation  $P_{n+1} = FP_nF^T + Q$ , where  $Q$  is another covariance matrix expressing the multiple scattering of muons crossing the return iron yoke. Once the state is propagated to the next layer, the closest stub is selected and the track parameters are updated based on the values and uncertainties of the measurement. These values are represented by a vector  $z_k = (\phi, \phi_B)$  and the corresponding uncertainty by a  $2 \times 2$  covariance matrix  $R$ , which is different depending on the positions resolution of the detector.

The state update is a matrix manipulation process that involves matrix inversion. The result produces a Kalman gain matrix,  $G$ , that is used in order to update the state with the following formula:

$$x_n = x_n + Gr_n$$

where  $r_n$  is the residual between the propagated state and the measurement.


The values of the Kalman Gain are pre-calculated in order to reduce calculations in the FPGA logic and save resources. They depend on the curvature and different combinations of station hits that produce the muon track. Thus, the KMTF algorithm propagates the track from station to station, starting from MB4 and going to MB1, and on every station updates it using the pre-calculated Kalman Gain. The procedure is illustrated in Figure 6.6. Once it reaches MB1, the measurement is stored and is called unconstrained. Next, the track is propagated to the center of CMS. At this stage the material of the magnet, calorimeters and tracker are taken into consideration for the calculation. A Kalman update is performed again and the result provides a beamspot constrained measurement.



**Figure 6.6:** Graphical illustration of a slice of the CMS in the transverse view. The figure highlights the steps followed by the KMTF algorithm to reconstruct a muon track.

## 6.4.2 Hardware

The GMT group designs and produces its own hardware for the Phase-2 upgrade. The ATCA board that will facilitate the subsystem is called X2O, and is shown on the right of Figure 6.7. The same board will also instrument the EMTF and OMTF subsystems.



**Figure 6.7:** *Left: The BMTL1 ATCA board. Right: The X2O board used by GMT.*

The X2O is a modular ATCA platform. It is divided in three separate modules, each of which executes different tasks. On the right of the board is the power and control module. It includes the ATCA backplane connectors, the IPMC card and a host ZYNQ device used to control the board. Furthermore, it delivers the ATCA power to the processing module. The processing module, also called Octopus, hosts a Xilinx VU13P FPGA. The dimensions of Octopus permit the usage of two such modules in one ATCA platform. A key difference of this board with the other hardware lines of the Phase-2 L1 Trigger is that it does not use Samtec Firefly transceivers. Instead, it uses the Ethernet standard QSFP modules, attached in the corresponding Optical module at the left of the board. The connection of the high speed signals from the FPGA to the QSFPs is done through copper cables that are attached very close to the FPGA chip. This way, the high speed PCB traces are kept as short as possible, resulting to minimum signal integrity losses of the high speed optical connections.

## 6.5 Architecture

This section describes the architecture of the Barrel Muon Trigger, as evolved during the time of writing this thesis (fall 2023). The Phase-2 L1 Trigger TDR [25] defines a system with 60 BMTL1 boards and 18 GMT board, in which every BMTL1 board processes data from one DT Sector. The document also foresees the case of having 30 BMTL1 boards instead, with each board processing two Sectors instead of one.

Sector	Number of OBDTs
1, 2, 3, 5, 6, 7, 8, 10, 12	14
4	16
9, 11	12

**Table 6.3:** *The number of OBDTs for every Sector of a wheel of the DT system. The numbers do not change in any of the 5 wheels.*

This scenario also includes an extra layer, called Concentrator, to manage the link interfaces between BMTL1 and GMT. With the additional knowledge gained during the years after TDR, it has been decided to implement an architecture similar to the latter.

The 60 DT Sectors are going to be instrumented by OBDT boards. The OBDT inputs are wires from the DT cells of a super-layer. Each board can digitize information from up to 240 such channels. In the cases of MB1 and MB2, one OBDT board can fit the DT cells of one super-layer. Thus, 3 OBDTs are needed per chamber to instrument SL1, SL2 and SL3. The super-layers of MB3 and MB4 are larger than the other two, resulting to more DT cells per SL. For this reason, MB3 will use 4 OBDTs instead of 3. Lastly, MB4 does not contain a  $\theta$  SL but in most Sectors its size requires the usage of 2 OBDTs per SL. There are cases where 2 are needed in total for small MB4 chambers and others where 6 are needed for the largest ones. Eventually, the number of boards that will be used at every Sector can be seen in Table 6.3. The total number of OBDTs that will be installed for Phase-2 is 830.

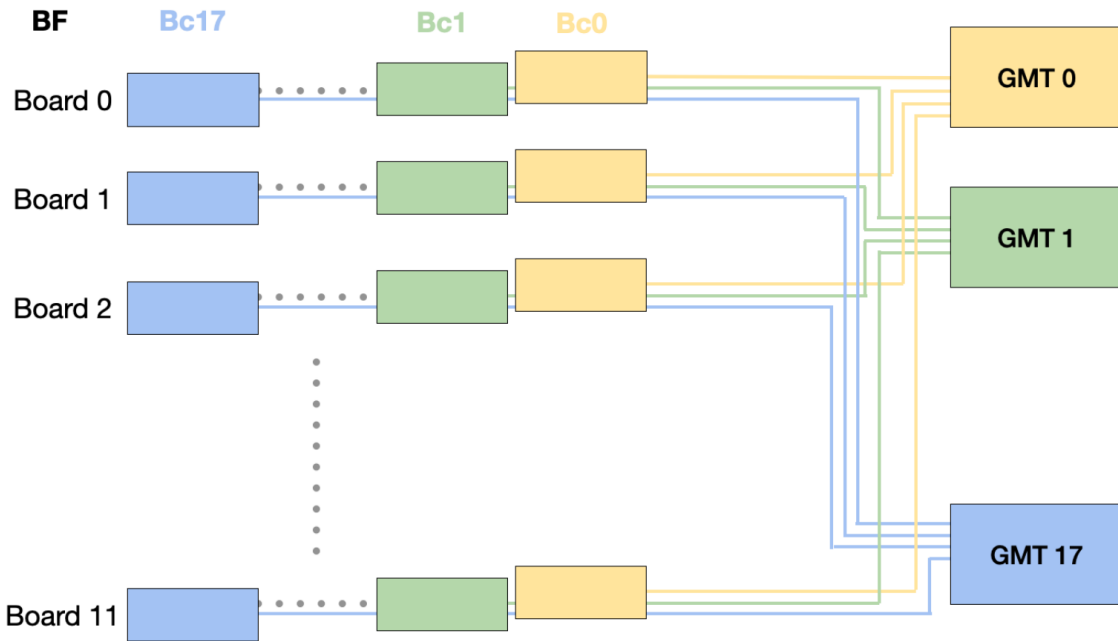
The FPGA of the OBDT boards convert the signal from every DT cell to TDC hits that are transmitted to BMTL1. The interface uses long optical fibers that start from UXC and travel through a tunnel inside the wall that exists between the detector and service room, ending up to the BMTL1 boards at USC. The current estimate is that the payload of one OBDT can be served by one optical fiber. Thus, the number of links that are sent to one BMTL1 board from DT Sectors is going to be either 26, or 28, or 30, depending on the Sectors. These links will operate the lpGBT protocol. As described in section 7.3, lpGBT payload in FEC12 mode (that will be used by OBDT) is 192 bits per BX. Considering that the hit size of one cell is 25 bits (Table 6.3) the maximum number of TDC hits per SL per BX is 7. In case the bandwidth of one optical fiber is not enough, the output link number per OBDT will get doubled.

TDC hits that originate from the RPC detectors will be transmitted by the Master Link Boards. The number of links required per Sector is 5. Hence, 10 RPC fibers will be connected in every BMTL1 board, resulting to a total of 36, 38, or 40 receiving links. In case the OBDT links get doubled, the maximum number of input links to one BMTL1 board will be 70.

The BMTL1 subsystem will consist of two sub-layers. The first is the core BMTL1 layer that receives detector data from all 60 Sectors. This layer will be instrumented with 30 boards and every board will process two barrel muon Sectors. Processing of the DT hits is performed with the Analytical Method algorithm. The algorithm that processes the RPC hits has not been implemented yet and hence, it is not know in which of the two sub-layers it will run. The second sub-layer is now called Barrel Filter (BF) (instead of Concentrator) and will host 12 additional BMTL1 boards. Its

exact functionality is not yet officially defined. However, it is foreseen to implement processes such as ghost/noise suppression and sorting of the DT TPs and handling of data transmission to OMTF and GMT. Furthermore, it is possible that algorithm pieces such as the theta matching, RPC clustering and matching with DT TPs are going to be implemented at the BF layer.

The link interface between the BMTL1 and BF boards is not yet completely defined. It is known, however, that one BF board will receive TP data from 9 BMTL1 boards. The interface between BF and GMT will take place in a time-multiplexed (TM) architecture with a period of 18, same as the number of X2O boards. This means that every BF card will output 18 links, with each one of them be connected to one GMT board. These links will operate at 25 Gbps line rate. A schematic illustration of the connection between the two systems can be seen in Figure 6.8.



**Figure 6.8:** A schematic illustration of the interface architecture between BF and GMT.

The time-multiplexed architecture defines that every GMT board processes information from the whole detector. Hence, it receives input links from every BF board. If we take as example data from BC0, they are all transmitted at GMT board 0. Data from the next BX are going to be transmitted to GMT board 1, while board 0 continues to process the previous BX. This procedure is repeated until all GMT boards are processing a specific BX. Then, data from BC18 will be directed again to GMT board 0. The result of this architecture is that every GMT board has 18 x 25 ns time to complete its data processing. Moreover, a direct feature is the optical link bandwidth allocation. For example link 0 from BF board 0 transmits data from BC0 to GMT board 0. Data from BC1, will be transmitted from link 1 of BF board 0 to GMT board 1. Thus, link 0 is able to continue transmitting data from BC0 to GMT board 0. In reality, the available bandwidth ends up being that of 1 BC times 18. At 25 Gbps the number of 64-bit frames that can be transmitted in one BX is 9. As a result, the number of frames per link in a TM 18 architecture is  $9 \times 18 = 162$ .



Every BF board will transmit TPs that are produced from 5 muon Sectors. The output of the AM algorithm for every chamber is 4 TPs, or frames, on the  $\phi$  view and 4 on the  $\theta$  view. Thus, the maximum number of TPs generated for one Sector is 28 ( $4 \times 4 \phi\text{-SL} + 4 \times 3 \theta\text{-SL}$ ) and that for 5 Sectors is 140. Hence, 140 out of the available 162 frames per link will be used to transmit the TP data. The remaining 22 frames can be used to transmit any other information is decided. However, it is expected that in most cases the generated number of primitives will be much less than 140 per BX. In this case, and by implementing a technique called zero-suppression, it is possible to transmit all TPs in fewer number of frames. The remaining bandwidth can be used to transmit data from plus or minus one BX, that can be used for searches of long-lives particles, for example.

# Chapter 7

## Firmware Infrastructure for Phase-2

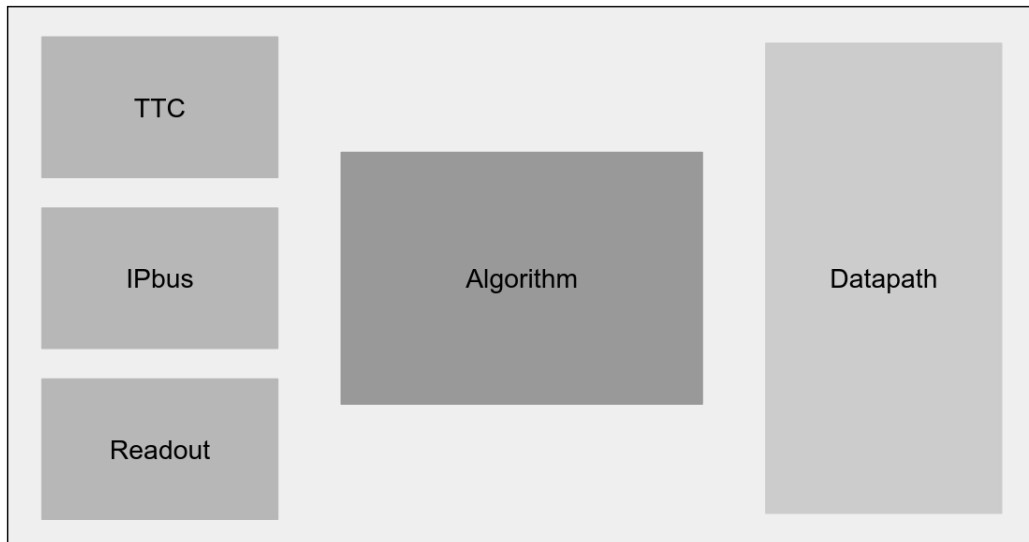
### 7.1 Introduction

This section presents firmware infrastructure that will be used by ATCA processors of the Phase-2 Level-1 Trigger system. By the term *firmware* we refer to digital circuits that are written using the VHDL (or Verilog) language and run on FPGA devices. As described in section 4.2.3, the FPGAs used at L1 Trigger are produced by Xilinx. The term *infrastructure firmware* refers to digital circuitry and applications that a L1T processor is required to include. They facilitate tasks such as data transmission and reception, control and monitoring via software, delivery of important trigger signals, and others. Generally, firmware of L1T processors includes the infrastructure logic and the algorithmic logic. Thus, the infrastructure logic creates the appropriate environment so that the algorithmic logic can be implemented, by providing all necessary signals.

The groups which develop hardware for L1T at Phase-2, in parallel of producing ATCA boards, they also develop the corresponding firmware and software tools that handle the operation of the FPGA and the board overall. Most of the ATCA board flavors use their own framework that implements more or less similar functionalities, but uses different coding approaches. Large part of the work of this thesis involves firmware infrastructure developments, most notable of which is the Hermes protocol and its successor, the CSP protocol. The firmware framework around which both of them are developed is the EMP-Framework, which is the one used by the Serenity boards. This framework has also been integrated in the BMTL1 ATCA board, as described in section 9.4. For this reason, a presentation of its functionalities is given in this chapter. Furthermore, this chapter includes a section about the Hermes protocol, as well as a description of the GBT and lpGBT front-end optical protocols. The Hermes and CSP protocols have also been adapted to the firmware framework of the Ocean and X2O boards. However, description of this work is not included in the chapters of this theses.

## 7.2 EMP Firmware Framework

The EMP Framework (stands for Extensible, Modular firmware framework for Phase-2 upgrades) is a firmware framework that targets L1T processors for Phase-2. It implements the necessary circuitry that is required for the implementation of CMS trigger algorithms inside the FPGAs. A simple block diagram of the framework can be seen in Figure 7.1. The main components of the firmware are the TTC block, the Datapath block, the IPbus block, the Readout block, and the Algorithm (or Payload) block.



**Figure 7.1:** *Block diagram with the basic firmware components of the EMP Framework.*

The original hardware platform that uses the EMP Framework is Serenity. As described in section 5.1.7, Serenity can host different types of FPGAs using daughter-cards that connect to the main PCB through an interposer connector. These types are Xilinx FPGAs of both the Ultrascale and Ultrascale Plus families. For this reason, EMP Framework is inherently modular and able to support all daughter-card variants. A description of the basic firmware components is given below.

### 7.2.1 TTC block

The TTC block handles all Timing Trigger and Control signals of the framework, which are part of the TCDS2 stream (Timing and Control Distribution System for Phase-2). They provide synchronization between all subsystems of the CMS experiment and must be distributed to every single processing node that operates in the L1 Trigger. They include the main 40.078 MHz LHC clock, a bunch crossing zero tag, the Level-1 Accept, back-pressure, and others. Reception of the TCDS2 stream to the FPGA is performed using a high speed link and through the backplane Zone 2 connector. The stream is delivered to all blades inside a crate by the DTH board, as discussed in section 5.2.2.

The TCDS2 link operates at 10.24 Gbps using a custom protocol. The interface firmware is implemented inside the TTC block, where the information is decoded and

the corresponding signals are routed to the other firmware blocks. This protocol is synchronous to the LHC clock and there are two methods that it can operate. The simple one uses the LHC clock as it is received from the backplane connector both as the main firmware clock, and also to synchronize the MGT reference clock. The other, that is not yet used, uses the recovered clock of the received data. This method extracts the LHC clock from the serial stream, routes it outside the FPGA to the input of a jitter cleaner chip and then returns it back to the FPGA. It can then be used as the main LHC clock and also used to synchronize the corresponding MGT reference clocks.

The LHC clock is the main clock that drives collisions data in and out the FPGA, and also drives the algorithms that are used to process them. Since the FPGA devices are able to operate in much higher frequencies, the LHC clock is multiplied in a number of different frequencies, such as 160 MHz, 240 MHz and 360 MHz. This way, the number of operations and data transactions that take place in one bunch crossing time (25 ns) can be significantly multiplied. Otherwise, the very strict latency of the L1 Trigger would not be possible to be met.

Apart from the LHC clock, the CMS systems get synchronized using two signals: the orbit tag and bunch crossing zero (BC0) tags, that are also delivered through the TCDS2 stream. The BC0 signal is asserted at the first clock cycle of the LHC orbit. As described in section 2.1, the LHC orbit frequency is 11.2455 KHz and that of the proton bunches 40.078 MHz, resulting to 3564 available bunches per orbit. Hence, the TTC block contains a bunch counter that locks to BC0 and starts counting from 0 to 3563. A similar counter logic can be implemented for the orbit number, even though for Phase-2 it is not applicable yet. The bunch counter, however, is very important when two systems communicate data with each other. The pipeline architecture of the L1 Trigger can only operate if all systems are aware of which event they process.

The above description regarding the LHC information reception to the trigger processors refers to the case where the boards are inserted inside an ATCA crate together with a DTH board. However, most of the initial development stages take place in the lab and on a test bench environment. For this reason, other means of receiving the LHC signals are also implemented in the TTC block. The simple one is dedicated to single board tests, where both the 40.078 MHz clock and the BC0 can be generated on-board, operating in free-run mode. In cases where synchronization between multiple boards is needed but outside an ATCA crate, for R&D purposes, it is foreseen that these signals can be delivered to the boards from other external sources. The TTC block supports these methods as well, and can be configured to operate using any of them.

## 7.2.2 IPbus Block

The IPbus block implements the IPbus transaction protocol. It is a custom packet-based protocol that is used to access firmware registers using software [35]. IPbus was developed during the early years of CMS operation to provide a hardware control mechanism for the  $\mu$ TCA and ATCA standards. Since then, its usage has been adopted by the majority of Phase-1 hardware and will also be the basic control protocol for a large number of Phase-2 boards.

The protocol implements 32-bit data transactions of a 32-bit memory addressing

between a client and a host. The client is usually a PC that runs the IPbus software to initiate transaction to the host, which is typically an FPGA that implements the corresponding firmware. The IPbus information includes a Packet Header followed by IPbus requests. The requests can be: a register read or write of either a specific address space or incrementing address space, read-modify-write bits for specific subset of a 32-bit register, read-modify-write sum to add values to a register. For these transactions the software client sends a request and waits for the IPbus device to respond back. One packet can direct more than one transactions at a time.

The typical IPbus implementation during Phase-1 uses the network application layer to wrap IPbus packets. The chosen transport protocol is UDP (User Datagram Protocol) due to its simplicity. Loss of packet is guaranteed by the implementation of a reliability mechanism running at the IPbus Control Hub. The UDP packet is then wrapped to an IP (Internet Protocol) and an Ethernet Packet in order to be directed through the Ethernet physical layer to and from the FPGA device. The Phase-2 hardware platforms include both the client and the host onto the same PCB. As a result, the network transport layer is replaced either by the on-board peripheral bus PCI Express, or by the chip-to-chip AXI interface. In both cases, IPbus registers are mapped using direct memory addressing from the firmware logic directly to the memory of the CPU that runs the client. A description of an AXI interface for the BMTL1 ATCA card is given in section 9.4.2.

The building blocks of IPbus are the IPbus firmware, the Control Hub and the uHAL. The IPbus firmware is modular and mainly includes a decoder block and a bus master control to interface the slave registers in the FPGA logic. Control Hub is a software application that mediates transaction requests coming from multiple user processes to one or more devices and then passes the response back to the original sender. It also implements a reliability mechanism for UDP based transactions. The uHAL (micro Hardware Access Library) is a Python/C++ library used by the end-users to write their software applications.

### 7.2.3 Readout Block

The Readout block performs the readout of a trigger processor to the DTH board. The interface takes place via optical fibers and by using the Slink Rocket protocol. The users have to implement a buffer logic that holds information of multiple events until a Level-1 Accept signal is received. If not, the data from a particular event are discarded.

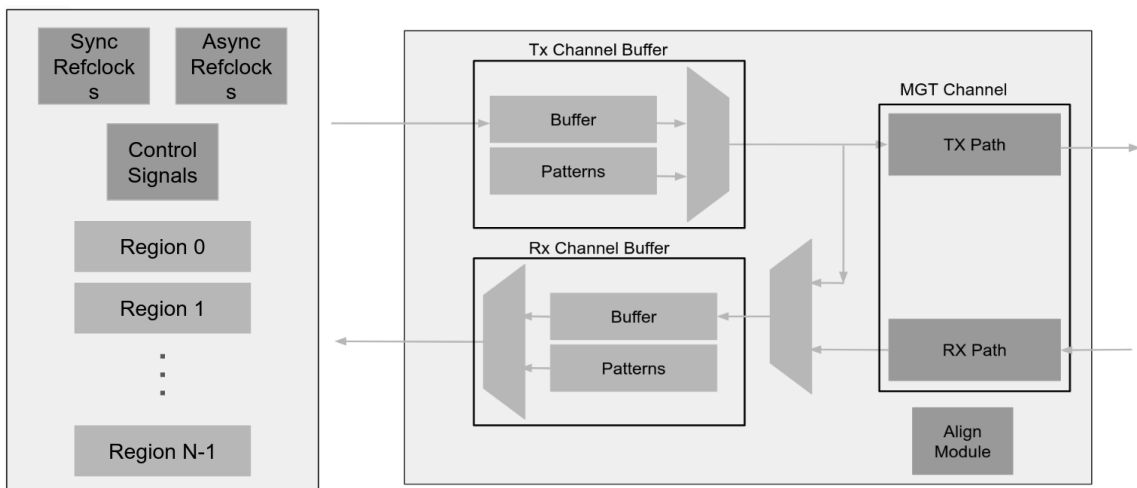
### 7.2.4 Datapath Block

The Datapath block handles data transmission and reception of the framework for a particular FPGA device. The interface is performed through serial optical links using the Xilinx MGT transceivers (described in section 4.3). Depending on whether a board communicates with the front-end or with other back-end processors, it should utilize the corresponding link protocol. The standard Phase-2 protocols for each communication respectively are the lpGBT and the CSP.

Implementation of the two, or any other protocol that is used, is accommodated by the Datapath block. It provides the appropriate firmware environment, as it includes,

among others, the necessary IPbus interfaces, TTC signals, clocks and resets. The clock sources provided to the Datapath are: reference clocks for every supported MGT QPLL, the LHC clock and multiple versions of it, and a free-running general purpose clock. There are two types of reference clock sources used by the MGTs. One of them is called synchronous, as it drives the interface of protocols that operate synchronously with respect to the LHC, such as the lpGBT and the DTH protocols. These clocks are produced by clock synthesizers on the board that use the LHC 40.078 MHz clock as reference. The back-end protocols can operate using asynchronous reference clock sources, which are synthesized in free-running mode. Other signals provided by Datapath to every Quad are a dedicated ipbus control and monitor bus and TTC signals, such as the bunch crossing zero tag.

As described in section 4.7, the MGT transceivers that exist in a Xilinx FPGA are grouped in Quads. These Quads are placed around the FPGA resources, at the left and right borders of the chip. Following the same philosophy, the firmware implementation of the optical protocols maintains a Quad structure and their code is contained in a specific block inside Datapath, called *Region*. As illustrated on the left of Figure 7.2, identical copies of the Region block are implemented. Their number,  $N$ , is equal to the number of MGT Banks in the chip. Hence, in a VU13P FPGA, which contains 128 MGT channels grouped in 32 Banks, there are 32 Regions.



**Figure 7.2:** *Left: Block diagram of the basic components inside the Datapath block. The Region block is implemented  $N-1$  times, where  $N$  is the total number of Regions supported by the FPGA type. Right: Block diagram of the Region block.*

A block diagram of the Region block is shown on the right of Figure 7.2. For illustration purposes the diagram depicts one MGT channel. However, the same components are implemented four times to form a Quad. The firmware components of optical protocols are contained in the MGT channel block. The kind of protocol that is going to be implemented is defined by the user, as described two sections below. To be used complementary with the MGT links, the Region block contains a Tx and Rx channel buffer for every channel, as well as an Alignment module.

## Channel Buffer

The channel buffers are implemented to assist the development stages of firmware parts of the framework. They are also very useful debugging tools. Every MTG channel in the EMP Framework includes a set of one Tx and one Rx channel buffer. Each buffer instantiates a 72 Kbit BRAM Tile ( $72 \times 1024$ ) that can be filled with 1024 64-bit words, leaving 8 bits for other meta-data signals. The buffers, shown on the right of Figure 7.2, are user configurable. They can either inject data or capture data. Furthermore, their direction can point either the MGT Tx or Rx path or the Payload path. For example, one use case is the transmission of a known pattern of data to the MGT interface. The MGT can be configured to loop data back in the receiving side and the Rx channel buffer can capture the received data and compare them with the transmitting ones. One other scenario is playing data to the algorithm block from the Rx channel buffer and capture the output of the algorithm to the Tx channel buffer. The payload of the buffers can be either a predefined pattern or a user defined stream of 64-bit words. This functionality makes the channel buffers a very powerful tool in the development stages of both a link protocol and the algorithmic logic.

## Align Module

The align module is another important piece of the Region block. It regulates latency of every back-end link by modifying the read pointer of the RX BRAM of the CSP protocol. This way, the exact frame and bunch crossing at which the words of each channels are routed to the framework can be controlled. The result is that all channels can be bonded together to the same bunch crossing. More details about the link alignment process are given in section 7.4.6.

### 7.2.5 Algorithm Block

The algorithm block is the central component of the framework. Data interface to and from every MGT channel is routed to the algorithm block, since the algorithms inside it are the main producers and consumers. Furthermore, it receives all versions of the LHC clocks and resets, the bunch counter, orbit counter and the L1 Accept signal. Moreover, an interface to the SLink Rocket protocol exists inside the algorithm block. The number of readout channels per project is defined by the user, as well as the data buffering and handling to the Slink.

The standard algorithm (or Payload) block of the EMP Framework is a null component, since algorithm implementation is conducted by the Trigger subsystems. The algorithm block is actually the only component of the Framework developed by the user. All functionalities around it are in the response of the infrastructure firmware.

### 7.2.6 Declaration File

The philosophy of the EMP Framework is to provide to the user all firmware pieces that are required for the development of trigger algorithms. To fit all different scenarios depending on the application, the optical interfaces of the framework are defined inside a VHDL component, called *project declaration file*. An example can be seen

in Figure 7.3. The declaration file creates the *REGION\_CONF* array, every row of which corresponds to one Region block. In every region, the user can select the transmitting and receiving link protocol (first and fifth columns) and whether Tx and Rx channel buffers are going to be implemented (second and fourth columns). The third column is not applicable by the current version of the framework. The example shown instantiates the first five regions of the project. In region 0, the *gty25* setting will implement the CSP protocol running at 25 Gbps on a GTY MGT, both on the receiver and the transmitter sides. Region 1 will implement the transmitter of a GTY channel running CSP at 16 Gbps, while in region 2 will implement the receiver of a GTH channel running at 16 Gbps. Regions 3 and 4 will implement the front-end *gbt* and *lpGBT* protocols, respectively. All five regions will implement channel buffers on both the Tx and Rx sides. The remaining regions will implement a *dummy* region, which does not include MGTs or channel buffers.

```
constant REGION_CONF : region_conf_array_t := {
  0 => (gty25, buf, no_fmt, buf, gty25),

  1 => (gty16, buf, no_fmt, buf, no_mgt),

  2 => (no_mgt, buf, no_fmt, buf, gth16),

  3 => (gbt, buf, no_fmt, buf, gbt),

  4 => (lpGBT, buf, no_fmt, buf, lpGBT),

  others => kDummyRegion
};
```

**Figure 7.3:** Example of a declaration file for an EMP project.

## 7.2.7 Constraint Files

The EMP framework contains a number of constraint files that facilitate specific operations on the framework. First of all, there are the board specific constraints that connect the ports of the firmware to corresponding pins of the FPGA chip. Constraining of the MGT channels is performed by a different file that runs a set of scripts written in the Tcl programming language. The *mgt-constrains* file locates every region where a protocol is instantiated and it then determines the name of corresponding MGT Banks. Once it has this information it executes the Vivado command that constrains the channels of that specific Quad to the correct FPGA location. Moreover, there are constraint files that declare the clocks of the framework and set those that originate from the same source in asynchronous groups.

The floorplaning of the chips is manipulated by the creation of the so called *pblocks*. The pblocks are areas of the chip that are defined and are used during the Implementation process to place specific firmware circuits inside them. By default, the EMP framework creates pblocks to define the area of every region and the area where the payload will be placed. The region pblocks exist in the area around the MGT Quads, which are located on the sides of the chip. The middle area is where the payload pblock exists.



## 7.2.8 EMP Software

The firmware of the EMP Framework is controlled using the corresponding software tool, called *emp-butler*. Emp-butler contains a set of libraries and applications that allow the user to control and monitor a project. The commands run on a Linux terminal and include every possible action that is supported by the framework. The first command a user should run, is the *reset*. This command, apart from sending a general reset to the FPGA, it also configures the source of the LHC clock, the BC0, and monitors the status of the bunch crossing counters. As mentioned in the TTC block section, there are different ways of sourcing the LHC signals. As a consequence, the reset command expects an additional option, that can be one of the: internal, legacy, tcds2 and external. The internal option uses on-board generated signals and the tcds2 can be used only inside a crate with a DTH board. The legacy and external options execute specific configurations of the Serenity boards. An example of the output of the *reset internal* is show in Figure 7.4.

```
[root@bmt11-2 fpga-ctrl]# empbutler -c connections.xml do serenity reset internal
21-11-23 13:47:04.651942 [281473289334784] WARNING - Address overlaps observed - report file w
21-11-23 13:47:04.657688 [281473289334784] NOTICE - mmap client with URI "ipbusmmap-2.0:///dev
Resetting device 'serenity'
Changing clock and TTC source to: Internal
Clock 40 locked after 0 ms
TTC BC0 locked
Global BC0 locked

STATUS
  Clock source: Internal
  TTC source: Internal
  Internal BC0: True

  Clock locked: True
  Frequency: 40.000 MHz

  BC0 locked: True
  Dist locked: True

  Bunch count: 0x000093
  Orbit count: 0x0015fd
  Event count: 0x000000
  BC0s recvd: 0 valid, 0 missing, 0 invalid
[root@bmt11-2 fpga-ctrl]#
```

**Figure 7.4:** Output of the *reset internal* command using the *emp-butler* software.

## 7.3 Front-End Optical Protocols

The information generated at the CMS detector from every LHC bunch crossing is digitized by the on-detector electronics and transmitted to the CMS service room, USC. The communication is performed using optical fiber cables, the size of which is 90 meters during Phase-1. Since the radiation dose at the detector cavern (UXC) is higher than normal, the electronics placed at UXC should be able to cope at this environment. Similarly, the optical protocol that is used for data transfer should be radiation hard and able to correct errors that might occur during operation.

The task described above is facilitated by the GigaBit Transceiver (GBT) project during Phase-1 and will be continued by the low power GBT (lpGBT) project at Phase-2. The underlying purpose of both is to build the electronic and protocol components that are required to transmit data to and from the detector. On the detector side, the GBTX and lpGBT chips are custom designed transceiver ASICs. They

are essentially serializer-de-serializer devices that implement the GBT and lpGBT protocols respectively. The optical on-detector module is also custom designed and operated by the Versatile Link at Phase-1 and Versatile Link Plus at Phase-2. On the other side of the links, reception and transmission of detector data will be facilitated during Phase-2 by specific Firefly modules, called CERN-B and the MGT transceivers of Xilinx FPGAs.

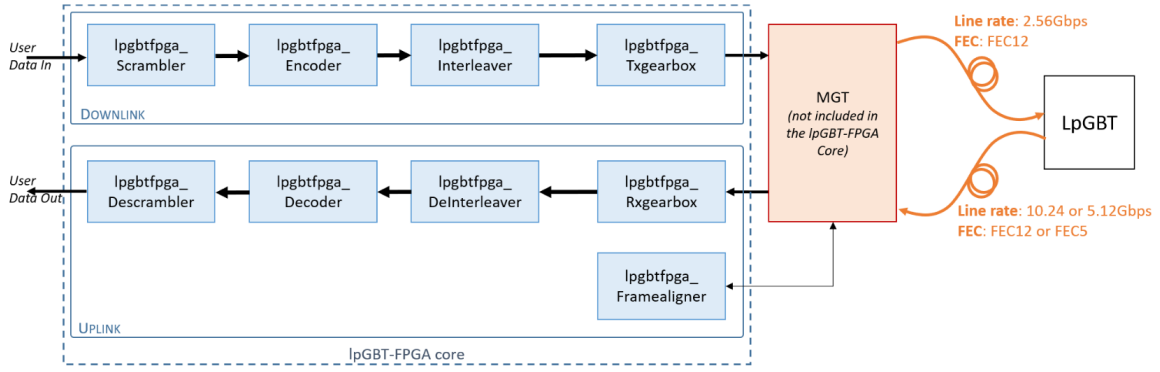
The nature of data transmitted through this bi-directional link can be separated in three categories: timing and control signals (TTC), detector data and slow control. The TTC fraction of the data deliver signals such as the L1 Accept, orbit and bunch counter resets and back-pressure. The LHC clock is distributed to the detector using the serial data stream. Both GBT and lpGBT protocols run synchronously with respect to the machine clock and hence, the recovered clock of the link can be cleaned and be used as the LHC clock. The detector data are transferred from the detector to the back-end trigger processors or are directed to the Data Acquisition system, depending on the detector sub-system. The slow control contains commands transmitted from the back-end boards to the detector in order to configure and monitor the front-end electronics.

### 7.3.1 GBT

The GBT protocol is being used during Phase-1 to transmit data between the detector and back-end electronics. The line rate of the link is 4.8 Gbps on both directions (uplink and downlink). The GBT parallel frame is a 120-bit word driven by the LHC 40.078 MHz clock. It is divided in fields, as defined by the GBT protocol. The 32 Least Significant Bits are used to transmit a Forward Error Correction (FEC) code that is applied to the payload bits. The FEC is based on Reed-Solomon (RS) (15,11) codes that are able to correct up to 16 consecutive bit errors. Next 80 bits are devoted to the user payload. Next 4 bits are called Slow Control bits. Two of them are used for External Control, targeting the control of other front-end chips. The other two, called Internal Control, are used to control the GBTX ASIC itself. The remaining four bits form a Header that is used by the receiver to align the serial stream to the correct word boundaries. As a result, the actual user bandwidth is 3.2 Gbps.

### 7.3.2 lpGBT

The lpGBT protocol is the Phase-2 evolution of GBT [15]. A block diagram of the FPGA firmware is shown in Figure 7.5. The core firmware blocks of the protocol are provided by the lpGBT group and are agnostic to the FPGA device. They are divided in the Downlink, path from back-end to the detector, and Uplink, path from the detector to the back-end boards. On the one side, they both connect to the user logic that handles the payload transmitted through the link. On the other side the two paths connect to an MTG block that is responsible for the serialization/de-serialization of the lpGBT frames and the serial transmission of the stream. The MGT block is developed by the end user of the FPGA, but an example for a KU040 FPGA is provided along with the core firmware blocks. The serial transmission at CMS takes place using optical fibers that connect to the lpGBT ASIC on the detector through the VL+ module.



**Figure 7.5:** Block diagram of the lpGBT-FPGA firmware.

The lpGBT link is bi-directional, as the GBT is, but asymmetric. The downlink line rate is fixed at 2.56 Gbps and the uplink line rate is configured to be either 5.12 Gbps or 10.24 Gbps. Furthermore, the uplink frame supports two FEC implementations, FEC5, consisting of RS (31,29) codes, and FEC12, consisting of RS (15,13) codes. The lpGBT frame contains the same bit fields as the GBT frame, but the FEC and user payload field size varies depending on the lpGBT implementation. For the downlink, the size of the frame is fixed at 64 bits. It includes a 24-bit FEC code, correcting up to 12 consecutive errors, 4 bits for Internal and External Slow Control, 32 bits for user data and a 4-bit header. The size of the bit fields for every possible configuration of the uplink are given in Table 7.1.

Config	5.12G/FEC5	5.12G/FEC12	10.24G/FEC5	10.24G/FEC12
<b>Header</b>	2	2	2	2
<b>SC</b>	4	4	4	4
<b>Payload</b>	112	98	230	202
<b>FEC</b>	10	24	20	48

**Table 7.1:** Size of the different lpGBT Uplink frame bit fields.

## 7.4 Hermes: A Back-End Optical Link Protocol

One of the main parts of the work for this thesis involved the development of an optical link protocol for back-end data transfer. The initial aim of the protocol was its adoption by the Serenity family boards, or, more specifically, by the users of the EMP Framework. However, it was soon modified for usage by the Ocean board, and later for the BMTL1 demonstrator board. This work reflected the contribution of the University of Ioannina to the Serenity collaboration. The resulting protocol, called Hermes, was essentially proposed as a solution to the back-end communication between Level-1 Trigger processors for Phase-2. It was being used during the research and development stages of the boards mentioned above until it was replaced by its successor, the CSP protocol, described in detail in the Chapter 8.

The development of Hermes [36] was a joint project between the Imperial College of London and the University of Ioannina. It was based on the asynchronous architecture of the Phase-1 Trigger protocol but targeted the new generations of Xilinx FPGAs, Ultrascale and Ultrascale Plus. The line rates supported by the protocol are 16G and 25G, following the specification of the Samtec Firefly modules. Its key objectives are to keep the total latency as short as possible while maintaining the utilization footprint considerably low. Furthermore, it was developed to be able to be used by both the GTH and GTY MGT types.

### 7.4.1 Encoding Layer

Hermes protocol is designed to transmit two types of 64-bit words, Data and Control words. The encoding scheme used by the first version of the protocol was the 64b/66b that is used by the Aurora 64b/66b protocol [37]. This scheme attaches 2 additional bits, called Header, to every 64-bit word that is transmitted. The Header defines the nature of every word, whether it is Data or Control word. The final version of Hermes, however, adopted the 64b/67b scheme, where the Header is made of 3 bits. Without changing the scope of the Header, the additional bit is used as a parity check bit, as described in 7.4.7. The overhead of the chosen encoding scheme is 4.68% over the line rate of the link.

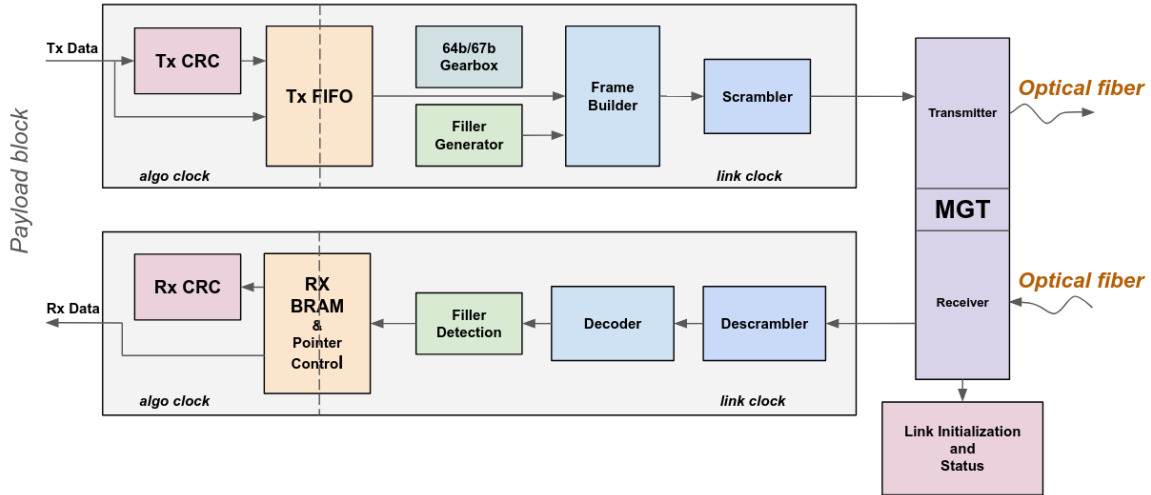
The required DC balance is achieved at Hermes by scrambling the 64-bit words. DC balance refers to the transmission of equal, or almost equal, numbers of Zeros and Ones over the medium, while in parallel avoiding long sequences of the same binary value. If DC balance is not implemented, the receiver CDR circuit will not be able to determine the phase of the bitstream and fail to recover the clock of the data. The scrambler used by Hermes for every 64-bit word is the one defined by IEEE 802.3 and introduced by 10 Gb/s Ethernet. Regarding the Header, the supported values of Hermes are either "010" or "101", resulting to an imposed imbalance of 1 in a sequence of 67 bits. Extended testing under normal conditions showed that the additional Header bit did not cause any degradation in the operation of the link. However, methods to equalize the polarity of the transmitted Header were studied and implemented.

### 7.4.2 Asynchronous Architecture

The back-end protocols used by the L1 Trigger processors operate asynchronously with respect to the LHC clock, as described in 7.2.4. Usually, the trigger algorithms run with a multiple of the LHC 40.078 MHz clock. Data words are handled by the transceiver using the parallel link clock. The frequency at this domain is defined by the line rate of the transceiver and the width of the parallel data bus. The reason of implementing an asynchronous architecture is to allow the transmitting line rates and the reference clock frequency be chosen freely, following the standards set by industry, while in parallel maintaining synchronization of physics data with the LHC clock.

The implementation of the asynchronous architecture was introduced in CMS already since 2006 and is currently used by the Phase-1 Trigger link protocol. The architecture instantiates a FIFO and a dual port Block RAM (BRAM) in the trans-

mitter and receiver data paths respectively. As illustrated in Figure 7.6, both of them are placed between the algorithm (payload) block and the MGT transceiver. This architecture can operate successfully only when the algorithm clock rate is lower than that of the link clock. Since the latter's rate is fixed by the line rate of the link, the algorithm can run with maximum frequency which is multiple of the LHC clock and yet remains lower than the link clock. This technique guarantees that no data loss would ever occur due to FIFO overflow.



**Figure 7.6:** Block diagram of Hermes protocol. Arrows illustrate data propagation through the transmitter and receiver data paths, flowing between the payload block and the MGT.

This difference in clock domain rates also results to a bandwidth difference. The additional bandwidth is called filler bandwidth and it is compensated by transmitting padding words, generated when the Tx FIFO is empty. Its size is the subtraction of the algorithm bandwidth from the total link bandwidth. For example, in 25.78125 Gb/s links the maximum payload clock can be 360 MHz resulting in 23.04 Gb/s of algorithm bandwidth. The remaining 2.74125 Gb/s is the filler bandwidth. Since the latter is generated in the link domain of the transmitter, it is imperative that it will be detected and removed in the link domain of the receiver. Otherwise, data coming out of the Rx BRAM will go out of sync. To avoid such occurrences, the usage of two FEC schemes has been introduced, as described in section 7.4.7.

### 7.4.3 Framing Layer

The kinds of 64-bit words, or frames, specified by Hermes are shown in Table 7.2. They consist of Data and Control words, with the latter consisting of two categories: Idle and Filler control words. The Idle control word is a null frame that is transmitted when the payload block has no valid data to send. Thus, it belongs to the algorithm domain, same way as Data frames. Filler control words are generated in the Filler Generator block and belong to the Filler bandwidth. The nature of every frame (Data or Control) is defined by the value of the Header, which is set by a user indicator, called *Valid bit*. When the Valid bit is zero, the Frame Builder block generates and

transmits Idle words. When one, the Frame Builder transmits the data coming from the Payload block. The kind of Control words is defined in the Control Word Type (CWT) field. The transmission of Filler words is invisible to the user. As illustrated in the block diagram of Figure 7.6, a Filler detection block exists in the receiving datapath and removes any Filler before they are written inside the Rx BRAM. The only frames written inside it are Data and Idle control words.

Header	Byte 7	Byte 6 <>	Byte 0
Data	Data		
Control	CWT	Idle Payload	
Control	CWT	Filler Payload	

**Table 7.2:** *Hermes frames specification. The Header specifies the transmission of either Data or Control words, the type of which is defined in the Control Word Type field.*

In order to devote the whole algorithm bandwidth solely to physics data, the filler bandwidth is utilized to transmit information such as link metadata, Cyclic Redundancy Check (CRC) checksums and alignment markers. This is achieved by artificially generating two kinds of filler words that propagate the above information, the CRC and the Align Marker Filler words. The format of Hermes’s filler words can be seen in Table 7.3. The least significant byte is used to carry link id information, while the next 3 bytes are user defined and could contain such information as board id, crate id, subsystem kind, etc.

Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0
CWT	Reserved	CRC	CRC	User Info	User Info	User Info	Link Id

**Table 7.3:** *Format of the Padding, CRC and Align Marker Filler words*

#### 7.4.4 Data transmission modes

The main objective of Hermes protocol is to support all physics data packet types that meet the requirements of CMS Level-1 Trigger, while at the same time maintaining ease of use. These requirements can be addressed by implementing two transmission modes, called Packet and Streaming mode.

Packet mode is used when data transmission by the algorithm block includes Data and Idles. The packets can be of any length and have at least one Idle word between them. The user defines a packet by asserting the Valid bit to declare the start and by de-asserting it at the end. A *Data Start* marker is generated on the receiver side at the rising edge of Valid bit and is used as the alignment marker. The falling edge triggers the transmission of the CRC checksum.

Streaming mode is introduced to facilitate the transmission of back to back packets, i.e. packets that are not separated by Idle words. In this scenario, the algorithm

is constantly streaming valid data frames through the link, hence Valid bit is always asserted. The boundaries of a packet are defined by the “End of Packet” bit. When asserted, it declares the last word to be used in the CRC calculation and also triggers its transmission. An alignment marker can be transmitted by asserting the “Align Marker” bit

### 7.4.5 Cyclic redundancy check

Cyclic Redundancy Check (CRC) is the method used by Hermes to determine the validity of received data. This method feeds every transmitted data word of every packet to a polynomial on the transmitting side. At the end of every packet, the result of the CRC calculation is a specific word, called checksum. The checksum used by Hermes is 16-bits long and it is transmitted at the end of every packet. On the receiving side, the exact same calculation takes place for every frame of a packet that is received. At the end, the CRC that is calculated on the receiver side is compared to the one calculated on the transmitter side. If the value of the two matches, data transmission is performed without errors. In case they are different, at least one bit flit has occurred. The Tx and Rx CRC blocks can be seen in Figure 7.6. The Tx block is located in the algorithm clock domain exactly before data are written to the Tx FIFO. Similarly, the Rx block is placed right after the Rx BRAM at the algorithm clock domain.

The CRC calculation and transmission was one of the most challenging parts in the design of Hermes protocol. One option was to execute the calculation and check in the link clock domain. This way, the CRC checksum could directly be sent as a deterministic Filler word and the CRC error indicator (1 bit) would cross through the available space of the Rx BRAM. This design choice, however, does not include in the calculation the path that data follow to cross clock domains, both in the transmitting and in the receiving side. For this reason, it was decided that CRC calculation and check will take place in the LHC clock domain.

At 25.78125 Gb/s links the maximum frequency of the algorithm clock is 360 MHz, resulting to the transmission of 9 frames per LHC bunch crossing (BX). In the case where Packet mode is used to transmit a packet of 8 frames, the CRC checksum could be transmitted as part of the first Idle word at the end of every packet, as shown at the left of Figure 7.7. However, this option cannot work in Streaming mode. In this case, the user would have to acquire part of the last data word (i.e. 16 out of the 64 bits) of each packet in order to transmit the CRC, as shown at the right of Figure 7.7.

Both methods of transmitting the CRC share one common feature, they use part of the Payload bandwidth. While this is not prohibitive, it interferes with the physics data transmission and for this reason none was chosen. Instead, Hermes transmits the CRC by using the available Filler bandwidth. This is done by crossing 16-bit Checksum to the link clock domain through 4 out of the 8 available bits of the Tx FIFO and RX BRAM, and then transmitting it as a Filler word.

Since the default width of BRAM Tiles (used to create the Tx FIFO and Rx BRAM) is 72, the 8 Most Significant Bits can be used to cross metadata information in parallel with every data frame. For example, the Valid bit value of every word crosses domains as the 65th FIFO bit. The method of crossing the CRC through

0 to 63		64	65	66	67	68 to 71	0 to 63		64	65	66	67	68 to 71
DATA 0		1	-	-	-	-	DATA 0		1	-	-	-	-
DATA 1		1	-	-	-	-	DATA 1		1	-	-	-	-
DATA 2		1	-	-	-	-	DATA 2		1	-	-	-	-
DATA 3		1	-	-	-	-	DATA 3		1	-	-	-	-
DATA 4		1	-	-	-	-	DATA 4		1	-	-	-	-
DATA 5		1	-	-	-	-	DATA 5		1	-	-	-	-
DATA 6		1	-	-	-	-	DATA 6		1	-	-	-	-
DATA 7		1	-	-	-	-	DATA 7		1	-	-	-	-
IDLE	CRC	0	-	-	-	-	DATA 8	CRC	1	-	-	-	-

**Figure 7.7:** *Left: Packet mode example of transmitting the CRC with the first Idle word at the end of every packet. Right: Streaming mode example of transmitting the CRC with the last data word at the end of every packet.*

FIFO splits the Checksum value in 4 chunks of 4bits each, in a way that the whole CRC value crosses clock domain after 4 clock cycles. The process is illustrated in Figure 7.8. A deterministic Filler word is then generated to carry the CRC to the receiver side. This Filler word is called *CRC Value* filler. The same amount of clock cycles is needed to cross domains in the Receiver using the same technique. Checksum Value is then checked with the value calculated on the Rx CRC block.

0 to 63	64	65	66	67	68 to 71
DATA 0	1	-	-	-	-
DATA 1	1	-	-	-	-
DATA 2	1	-	-	-	-
DATA 3	1	-	-	-	-
DATA 4	1	-	-	-	-
DATA 5	1	-	-	-	-
DATA 6	1	-	-	-	-
DATA 7	1	-	-	-	-
DATA 8	1	-	-	-	CRC3
DATA 0	1	-	-	-	CRC2
DATA 1	1	-	-	-	CRC1
DATA 2	1	-	-	-	CRC0

**Figure 7.8:** *Example of transmitting the CRC value in 4 chunks of 4 bits utilizing space 68 to 71 of the Tx FIFO and Rx BRAM.*

### 7.4.6 Link Alignment

Link alignment, or channel bonding, in Hermes is implemented by using the alignment marker bits of all receiving channels and by controlling the read pointer of each Rx BRAM separately. Both Packet and Streaming modes transmit this marker through all links on the same clock cycle. Due to differences in propagation delay, these markers may not arrive simultaneously but in different clock cycles, placed between the first and the last received markers. The method of aligning, or syncing, all channels together counts for every channel the number of clock cycles required after the reception of its own marker and until the reception of the last. It then subtracts this number from the read pointer of the BRAM of the individual channel, so that all channels are aligned to the link with the largest latency.



### 7.4.7 Protection Mechanism

Assuming a channel is aligned to a specific bunch crossing and frame, the receiving algorithm expects a strict sequence of incoming physics data. If the alignment of an individual link is lost, it results in misinterpretation of all data frames until the link is re-aligned. The above scenario takes place in cases where fillers are not properly recognized and get mixed with words of the algorithm bandwidth, or more specifically, when either the Header or the Control Word Type are received incorrectly. In order to protect the link from such occurrences, Hermes encodes both of them into FEC codes, each of which is capable of correcting up to one bit flip of the original code. The 3rd Header bit is used as a secondary check to determine the original value and the 4-bit CWT is encoded in Hamming (7, 4) codes, able to correct up to 1 bit flip. If two or more consecutive bit errors occur, the synchronicity is lost and the link needs to get re-aligned.

# Chapter 8

## The CSP Trigger Link protocol

The CMS Standard trigger link Protocol (CSP) is the optical link protocol that will be used for the interface between Phase-2 Level-1 Trigger processors. The protocol reflects the efforts of two groups in the development of a robust trigger protocol that satisfies the needs and requirements of the system. Before CSP was introduced, the two groups were independently exploring ideas and solutions by implementing their own protocol version. One group was based at University of Wisconsin, Madison and their protocol, called Iridis, was targeting the APx ATCA boards family. The second protocol was the Hermes protocol, developed by the Imperial College of London and University of Ioannina and described in section 7.4.

The two protocol versions shared many common features and were trying to resolve the same problems, but in some cases using different approaches. However, the final Trigger system will consist of a mix of all four ATCA processors, transferring physics data between each other. For this reason, it was decided that a common protocol should be implemented that would share the common characteristics and implement the most efficient solutions of each of the two.

The effort of implementing CSP started by the summer of 2021 and was finalized by the summer of 2023. The protocol was specified by both groups in a shared 20 pages document describing the syntax of CSP. The firmware implementation, however, is developed by the two groups independently, following their own coding approach on the different board frameworks (APx firmware shell and EMP Framework).

Implementing CSP for EMP was held exclusively under the context of this thesis. The work involved exploring ideas and solutions on aspects of the protocol, writing of the VHDL code inside the framework of EMP, maintaining the corresponding online repository and testing its performance. In addition, the CSP code was modified to be used in the Ocean and X2O boards, as well.

This chapter begins with a description of the common CSP syntax and continues with sections about the transmitter and receiver datapaths, as implemented by the Hermes firmware.

### 8.1 Protocol syntax

The size of a CSP word is 67 bits. The 3 Most Significant Bits (MSB) form a Header and the remaining 64 bits are the payload, or frame. The Header bits are used to

characterize the kind of the frame. Value "101" defines a Data word and value "010" defines a Control word.

The definition of the bit fields of a CSP Control word is shown in Table 8.1. There are five Control word types supported by CSP. Four of them belong to the Filler bandwidth and one, the Idle word, belongs to the Algorithm bandwidth. The four filler words are: The *BC0/OT* that is generated deterministically to transmit the BC0 information through the link. The *LID0* and *LID1* (LinkID) are non-deterministic padding words that are sent when the Tx FIFO is empty. They are also used to carry meta-data information of the transmitter to the receiver end. The fourth Filler word is called *CRCV*. It is a deterministic filler word that sends the CRC value at the end of every packet. It also contains a special status field. Differentiation between the five Control words happens through the Type Field value, specific for each type. Moreover, the Control Word Payload is defined only for the LID0, LID1 and CRCV words. BC0/OT and Idle should hard code this field to zero. The payload of each word is described later in this section.

Bits 63-56	Bits 55-52	Bits 51-32	Bits 31-0
Type Field	Index Number	Zeros (rsrved)	Control Word Payload

**Table 8.1:** *Definition of the CSP Control word bit fields.*

### 8.1.1 Index Number

The Index Number field is a functionality implemented in the Filler bandwidth of the link. It provides a reference waveform to the receiver end that can be used to detect missing or miss-identified Filler words. On the transmitter, the value of this field must increment by 1 on every transmitted Filler word, that is the BC0/OT, LID0, LID1 and CRCV. This 4-bit counter should return to 0 when it reaches the value of 15. By using this information, the receiver can lock to the counter value and be aware of errors in the reception of Filler words. In addition, the counter concept can be used to implement a logic that informs the Rx BRAM about the mistake and corrects the position of the read address pointer. The implementation of the Index Correction Mechanism is described in section 8.3.5.

### 8.1.2 Control Word Payload

The Control Word Payload field is used to transmit important information in the Filler bandwidth of the link. This way, the Algorithm bandwidth remains devoted solely to the transmission of physics data. This field is not defined in the payload of BC0/OT and Idle words. Link meta-data information, such as the origin of the individual channel (region and channel number), Crate and Slot ID, information about the packet in transmission, and others, is propagated through the LID0 (Link ID 0) and LID1 (Link ID 1) words. These are transmitted when the Tx FIFO is empty. The number of meta-data bits could not fit in one such word and hence, the two are transmitted alternately. The Payload field definition of LID0 and LID1 is described in Table 8.2 and Table 8.3, respectively.

Link ID Word 0				
Bits 31-20	Bits 19-12	Bits 11-8	Bits 7-2	Bits 1-0
Rsrved	Crate/System ID	Slot ID	Channel Region	Channel Number

**Table 8.2:** *Payload fields of the LID0 Control word.*

Link ID word 1					
Bits 31-29	Bits 28-26	Bit 25	Bit 24	Bit 23-12	Bits 11-0
"001"	Rsrved	Idle-as-Filler	Idle Mthod	Pckt Size	Pckt Intrvl

**Table 8.3:** *Payload fields of the LID1 Control word.*

The payload of the CRCV word is described in Table 8.4. Apart from transmitting the 16-bit CRC checksum, it also carries a 4-bit field called *Special Status Bits*. At the time of writing, only one of these bits is defined, even though not actually used by the Hermes implementation of the protocol. Bit d3, position 27, is defined as Upstream Link Error Flag. This flag is set to 1 if the packet being transmitted at that time is associated with a CRC error during its transmission.

CRC Value Word		
Bits 31-24	Bits 23-16	Bits 15-0
Special Status	Rsrved	CRC16

**Table 8.4:** *Payload fields of the CRCV Control word.*

### 8.1.3 Data Integrity

The CSP protocol defines a combination of methods to either monitor or sustain the integrity of transmitted data. In the case of Data words, the latency constraints imposed by the L1 Trigger system do not permit the implementation of FEC techniques that could correct transmission errors. However, CRC checksums are calculated for every CSP packet to provide a monitoring on the quality of physics data. Experience of Phase-1 links show that bit flips during the optical transmission between L1 processors in the CMS counting room is negligible. Thus, not implementing FEC codes is considered as a good trade off for the obtained link latency.

On the other hand, as described in the Hermes protocol, FEC codes are implemented in specific fields of the CSP words. While an error in the payload of Data words might have a minor effect in the operation of the Trigger subsystems, a single error either on the Header or the Type Field of Control words can cause a significant dysfunction. For this reason, the Header value is protected by using an extra parity bit and the Type Field by implementing Hamming (8,4) codes, as described below.

## Header parity

The CSP Header is used to define two kinds of words. Thus, only two values are specified by the protocol, one to tag Data words and one to tag Control words. While 1 bit would be enough for this operation, the 64b/66b scheme uses 2 bits in order to maintain DC balance by transmitting either the value "10" or "01".

The reason of using a 3-bit Header in CSP is to create a simple FEC technique, using the extra bit as a parity check bit. This protection is mandatory in order to avoid miss-identification of a Filler word due to single bit errors in the Header. To implement the 3-bit Header, the 64b/67b encoding is used instead of 64b/66b. This extra bit, however, adds imbalance in the serial bit stream. To correct this occurrence, the extra bit is scrambled together with the 64-bit words, as explained in section 8.2.8. The implementation of the parity check is shown in Table 8.5.

p	h1	h0	Decoded meaning
x	0	1	Data word
x	1	0	Control word
1	0	0	Data word
1	1	1	Data word
0	0	0	Control word
0	1	1	Control word

**Table 8.5:** Table for decoding the Header 2-bits (h1,h0) and the 3rd parity (p).

## Hamming codes

Hamming codes are used to encode the Type Field of CSP Control words. The actual size of the usable Type Field is 4-bits. Before transmission, it is encoded into Hamming (8,4) codes of 8-bits. When decoded on the receiver side, they are able to correct up to 1-bit error and determine but not correct 2-bit errors.

If we assume the bits of Type Field as "d3,d2,d1,d0" and parity bits as "p3,p2,p1,p0", the parity bits are calculated on the transmitter as follows:

$$\begin{array}{l} p3: \quad d0 \text{ xor } d1 \text{ xor } d2 \\ \hline p2: \quad d1 \text{ xor } d2 \text{ xor } d3 \\ \hline p1: \quad d0 \text{ xor } d2 \text{ xor } d3 \\ \hline p0: \quad d0 \text{ xor } d1 \text{ xor } d3 \end{array}$$

**Table 8.6:** Parity bits calculation of Hamming (8,4) codes.

The encoded Type Field contains the parity bits at the 4 MSB and the type bits at the 4 LSB. Decoding of the Hamming code on the receiver side is performed using the so called syndrome bits. The value of these 4 bits declares whether no error has occurred, the position of the error in case of single-bit errors and the occurrence of

double errors. Table 8.7 shows the logical functions that are used to calculate the 4 syndrome bits.

s3:	p0 xor p1 xor p2 xor p3 xor d0 xor d1 xor d2 xor d3
s2:	p2 xor d1 xor d2 xor d3
s1:	p1 xor d0 xor d2 xor d3
s0:	p0 xor d0 xor d1 xor d3

**Table 8.7:** *Syndrome bits calculation to decode the Hamming (8,4) codes on the receiver side.*

## CRC

The CSP protocol uses the CRC (Cyclic Redundancy Check) method to determine the integrity of every transmitted Packet. A checksum is calculated two times. On the transmitter side Data words of every packet are fed into the CRC block. At the end of a Packet, a reset is issued in the CRC block and the checksum is transmitted through the CRCV Filler word. The receiving end calculates the same checksum for every Packet as well. At the end of every packet, the Rx block is being reset and the Rx CRC value is compared with the Tx value. If the two match, the transmission of both the Packet and the CRC checksum is done with no errors. Otherwise, at least one bit error has occurred.

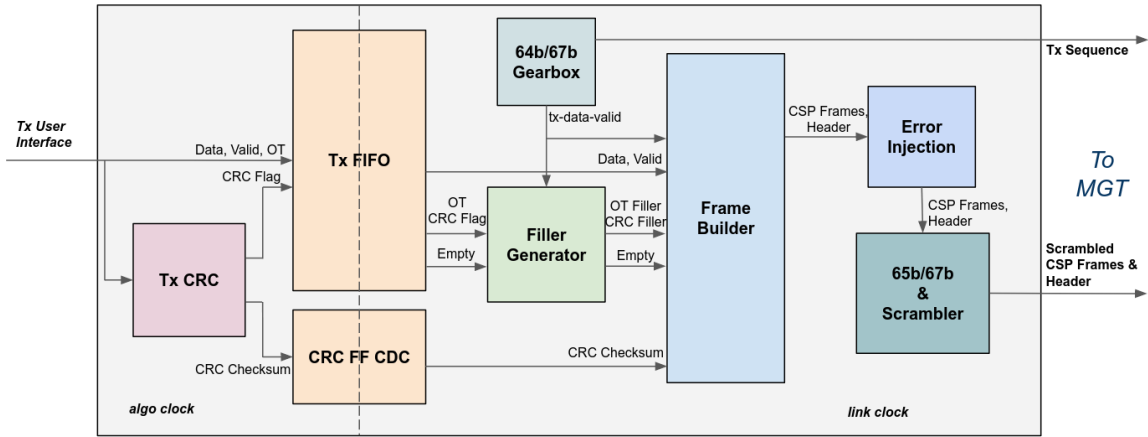
In CSP, the CRC-16-ANSI standard is used. The CRC block is generated by an online generator, following the polynomial:  $1 + x^2 + x^{15} + x^{16}$ . Furthermore, the CRC code has been modified in order to support on-the-fly resetting and feeding of the new calculation with the data stream of the next Packet. This modification is required for the transmission of back-to-back packets.

### 8.1.4 Scrambling logic

DC balance of the CSP serial stream is achieved by using a standard scrambling method. The polynomial used by the scrambler block is  $G(x) = 1 + x^{39} + x^{58}$ , as defined by IEEE 802.3av. The same polynomial is used on the descrambling block at the receiving end to decode data back to their original form. The CSP specification defines that the 3rd Header bit is also scrambled together with the 64-bit words. The implementation of the resulting scheme, called 65b/67b, is described in section 8.2.8.

## 8.2 Transmitter Datapath

This section describes the Hermes implementation of the CSP transmitter datapath. The block diagram of the firmware blocks is shown in Figure 8.1. A description of each module and its main connections is given in the sub-sections below.



**Figure 8.1:** Block diagram of the CSP transmitter datapath.

### 8.2.1 Transmission modes - User Interface

The transmission modes that CSP supports are similar to those introduced by the Hermes protocol. The user can define packets in two ways. One is by using the *Valid bit*. This mode assumes that at least one Idle word exists between two packets. Their length is defined solely by the value of Valid bit: a rising edge indicates the start of a packet and the next falling edge the end of this packet. In cases where Idle words are not transmitted, i.e. Valid bit constantly asserted, back-to-back packets are being sent. The packet separation in back-to-back packets is performed by asserting the End of Packet (EoP) bit for one clock cycle. Assertion of EoP indicates the last frame of the packet and the next frame constitutes the first frame of the next packet. Algorithm data are interfaced to CSP using the *payload-clock*. In the case of 25 Gbps, its maximum value is 360 MHz, corresponding to 9 frames per BX.

### 8.2.2 CRC generation

The CRC calculation is always enabled by the Valid bit. In the case of packets defined by Idle words, the calculation stops when the value of Valid bit becomes zero. A reset is issued on the CRC block and calculation begins again when the Valid bit is re-asserted. When back-to-back packets are used, the EoP signal defines the last frame to be included in the calculation. On the next clock cycle, the CRC block is being reset and also fed with the first data word of the next packet. In both cases, the checksum along with a CRC Flag signal are sent to the Clock Domain Crossing block to be crossed to the link clock domain.

### 8.2.3 Data Domain Crossing

Following the asynchronous architecture, the CSP implementation uses a FIFO block to cross clock domains on the transmitter side. For every Tx channel, a 36K BRAM Tile is used from the FPGA resources, configured as  $72 \times 512$ . In addition, Flip-Flop (FF) registers are used in parallel to cross the CRC checksum value. Out of the 72 bits of the FIFO width, the 64 are used to cross the data payload, 1 bit is used for

the Valid bit, 1 bit for the CRC Flag, and one for the BC0, or Orbit Tag, flag.

The choice of crossing the CRC through FF registers is made for two reasons. On the one hand, the CRC value crosses domains without interfering with the crossing of data words. Furthermore, crossing using FFs adds no additional delay than that already added by the FIFO block to cross the data words. The amount of this delay is enough to resolve meta-stability phenomena of the link domain CRC registers. The value of these is read when the CRC Value bit is crossed to the link domain.

## 8.2.4 Synchronous Gearbox

The Tx Synchronous Gearbox is a feature of Xilinx MGTs, provided to support the implementation of 64b/66b and 64b/67b encoding schemes. It runs on the link clock domain. The CSP protocol is based on the 64b/67b scheme, meaning that for every 64-bit payload word an additional 3-bit header is transmitted. The resulting 67-bit word cannot be handled by the internal MGT circuits, since their interfaces support data widths multiple of 16 bits. For this reason, the Synchronous Gearbox is used to accommodate the implementation of the schemes mentioned above.

At its CSP configuration, the inputs to the Tx Gearbox are a 64-bit port for the payload data and a 3-bit port for the Header. Furthermore, a 7-bit *sequence* port is included. The sequence is a counter logic that is implemented by the user. The GTY transceivers support 64-bit internal interface. In this case, the sequence counts from 0 to 66, which are the total number of cycles needed to transmit the payload plus header information. The operation of the Gearbox is to concatenate the 67 bits in a way that its output is always 64 bits of combined payload and header bits. This results in 3 additional 64 bit words that transmit the header information, or, in other words, 3 extra clock cycles. During these three extra clock cycles the MGT ignores data that are written to its tx-data and tx-header ports. This information is used everywhere in the CSP firmware implementation. It is propagated by a single-bit indicator, called *tx-data-valid*.

## 8.2.5 Deterministic Filler Generation

The CSP protocol uses the Filler bandwidth to transmit information without interrupting the transmission of physics data over the Algorithm bandwidth. This information includes: link meta-data, the CRC checksum and the Orbit Tag. As described in 8.1.2, link meta-data are propagated using two kinds of padding words, LID0 and LID1. These are labelled as non-deterministic filler words, since their transmission is triggered by the Empty flag of the Tx FIFO block. On the other hand, the CRC Value and BC0/OT Fillers are deterministic words, generated on demand when the user logic needs to transmit either the CRC, or the BC0 at the start of every orbit. This procedure is performed at the *Filler Generator* component.

The Filler Generator supports the generation of two types of Filler words, the CRCV and the BC0/OT. The generation is issued either by the CRC Flag and the OT bits, as soon as they arrive at the read port of the Tx FIFO (link clock domain). In case both of them arrive at the same clock cycle, the CRCV Filler is generated first. The logic takes into account the Empty FIFO flag and the tx-data-valid indicator to judge when the deterministic Filler can be generated. Upon generation, it asserts for



one clock cycle either a CRC Filler flag of an OT Filler flag, which connects directly to the Frame Builder block. In addition, it issues a pause to the FIFO Read port to provide the necessary space for the transmission of a deterministic filler word. In cases a padding Filler (non-deterministic) and a deterministic word are to be sent at the same time, only the non-deterministic is transmitted.

## 8.2.6 Frame Builder

The Frame Builder block builds the CSP 64-bit frames by implementing the protocol's specification. It combines all Tx pieces to decide what kind of frame is going to be written to the MGT, along with the corresponding 3-bit Header. The supported CSP frames are: Data, Idle, LID0, LID1, CRCV and BC0/OT. The Header value is "101" for data frames and "010" for Control words. The process is executed by a state machine that is updated by the Valid bit, the CRC Filler flag, the OT flag and the Empty flag. In the case of Data frames, the output of the FIFO is directly transmitted with the corresponding Header. In case of Control words, the payload of each word is being built and sent. Furthermore, the Hamming codes are calculated on the fly and transmitted at the Type Field of every Control word.

## 8.2.7 Error Injection

An Error Injection module is placed right after the Frame Builder. Its function is the injection of errors, when it is demanded by the user. The list of supported errors is shown in Table 8.8. Their main purpose is to provide a mean of validating the operation of many functionalities of the protocol. Those can be: error detection on the Data Packets, single-error detection and correction on the Header and the Type Field, double error detection on the Header and the Type Field and the operation of the Index Correction Mechanism.

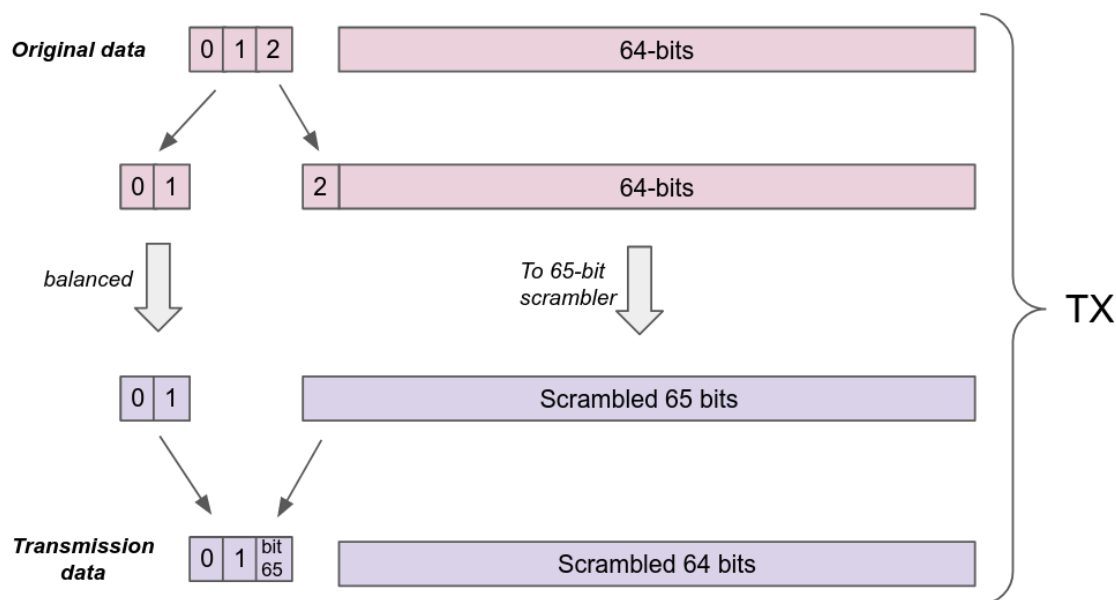
<b>Error Types</b>
Header-1bit
Header-2bit
Control Type Field 1-bit
Control Type Field 2-bit
CRC Error
Index Error 1-bit

**Table 8.8:** *List of the Error Injections supported by CSP.*

## 8.2.8 Scrambler - 65b/67b implementation

The scrambling method used at 64b/66b or 64b/67b encoding schemes feeds the Tx polynomial with 64-bit words, prior to their transmission. In the CSP case, the 3rd Header bit disrupts the running disparity of the link. To accommodate this issue, this

bit gets scrambled alongside the 64 word bits. The process is illustrated in Figure 8.2. The Header bit-2 is glued together with the 64-bit words at bit position 65, and the new 65-bit word is fed to the scrambler block. The value of the remaining Header bits 0 and 1 are always balanced (either "01" or "10"). On the output of the scrambler block, the 65th bit is attached again as the 3rd Header bit in order to be transmitted by the 64b/67b Gearbox logic. On the receiver side, the same but opposite implementation is used to decode the data. This implementation scheme used by the CSP has taken the name *65b/67b encoding*.



**Figure 8.2:** CSP scrambling method. Figure illustrates the transmitter side logic. Similar technique is implemented on the receiver side.

### 8.3 Receiver Datapath

This section describes the Hermes implementation of the CSP receiver. The receiver blocks inside the MGT are in charge of sampling to the correct phase of the received serial stream and of extracting the recovered link clock, as described in section 4.3. Once the CDR is locked, clock and data are directed to the SIPO (Serial-In-Parallel-Out) circuit and the parallel bus flows to the PCS blocks of the receiver. A reset has to be issued on the receiver datapath before it can be reliably used. Upon completion of the internal reset procedure of the MGT, the *rx-reset-done* signal is asserted, indicating that the user can access the data that appear to the MGT output ports. These ports are: the 64-bit rx-data bus, the 3-bit rx-Header and the rx-data-valid.

Initially, the MGT PCS blocks output 64-bit words with random boundaries. For the parallel data to be sampled to the correct word boundaries, they have to be aligned to the form the transmitter has sent them. This action is performed by the user logic during the link initialization and alignment procedure. This procedure uses known and recognizable sequences the transmitter sends. In the CSP case, these patterns

have to first be decoded before they can be used for the alignment of the link. Figure 8.3 shows the blocks of the CSP receiver datapath and their basic connections. A description of their operation is given in the sub-sections that follow.

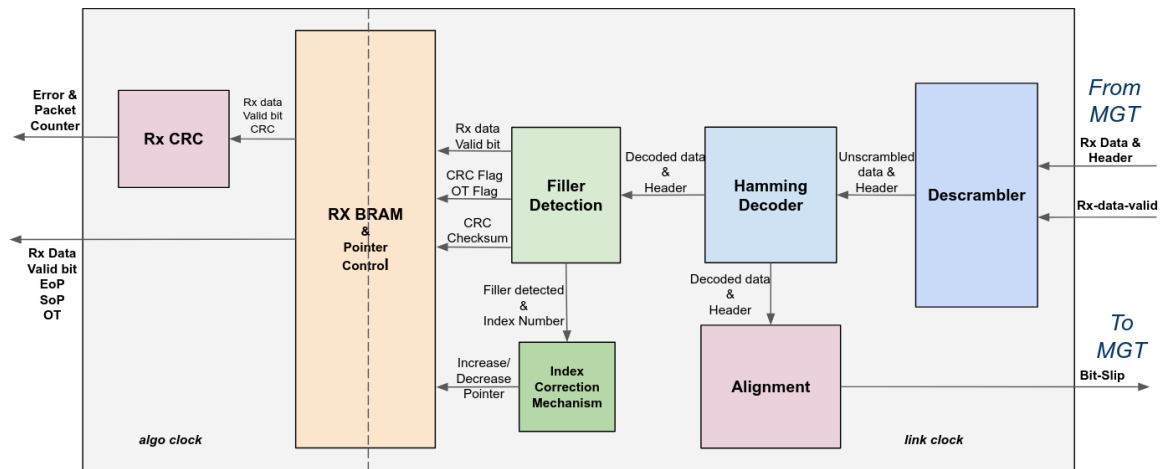


Figure 8.3: Block diagram of the CSP receiver datapath.

### 8.3.1 Descrambler

The received stream is a scrambled version of the original data words. The CSP transmitter implements the 65b/67b scrambling method, as discussed in 8.2.8. A similar but reverse implementation is implemented on the receiver side, as illustrated in Figure 8.4. Initially, the 3rd Header bit, which is the 65th bit, is attached to the 64-bit data to re-create the 65-bit scrambled word. This word is then passed through the Descrambler block, which uses the same polynomial equation, to bring the 65-bit words to their original form. Then, the 65th bit is re-attached to the Header as the 3rd Header bit, and both data words and the Header are routed to the next firmware block.

### 8.3.2 Hamming Decoder

The Hamming Decoder block implements the decoding of the Control word Type Field and the correction in cases of single-bit errors. Initially, the syndrome bits are calculated for every received word, as described in the Hamming codes part of section 8.1.3. Then, by storing the parity check matrix in Look-Up Tables, the firmware logic takes one of the following actions. If all syndrome bits are zero then the decoded Type has no errors. The case where the s3 binary value is 1 means that 2-bit errors have occurred. In this case, no action can take place and the corresponding double-bit indicator is asserted. The remaining cases correspond to single-bit errors that are being corrected and the Hamming Decoder block outputs the updated Type Field.

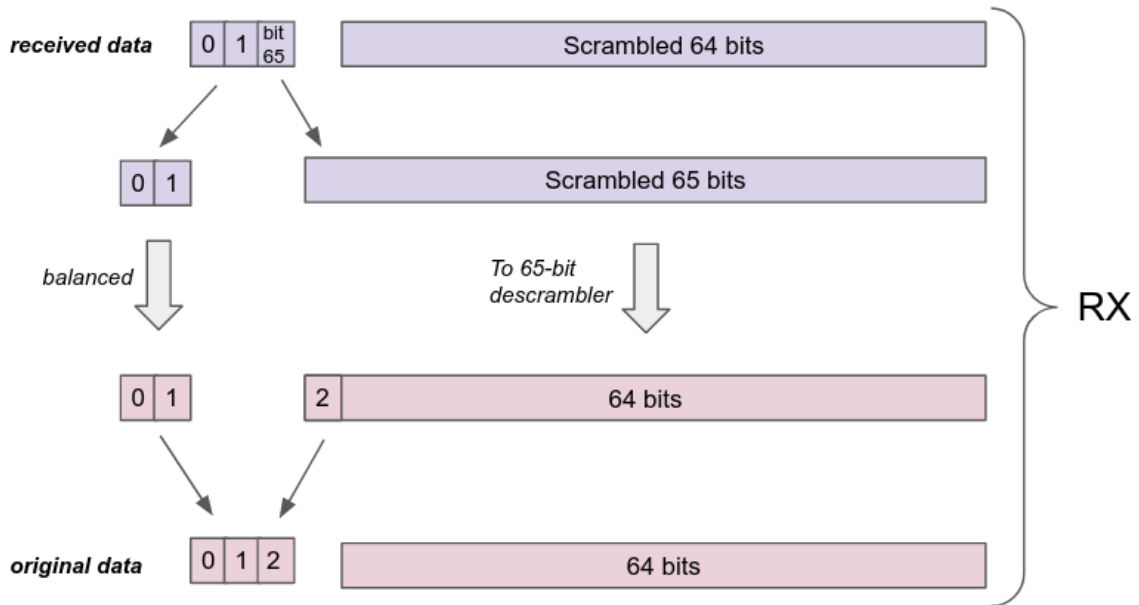


Figure 8.4: CSP descrambling method. Figure illustrates the receiver side logic.

### 8.3.3 Link Initialization and Alignment

At the point where data are decoded by both the Descrambler and Hamming Decoder blocks, they can be used by the Alignment block to initialize and align the link. There are two processes that operate in parallel for this purpose. The first process generates the *check-data* and *data-good* signals. These two are defined by the combination of the received Header and Type Fields. When Data words arrive, no check is performed (*check-data*=0). Otherwise (*check-data*=1), check is performed for every Control word, since the Type Field is the only known sequence of the 64-bit data. If the Type value matches that of Idle, LID0, LID1, BC0/OT or CRCV, the *data-good* signal is asserted. If not, *data-good* is zero.

The second process is mainly driven by the *check-data* and *data-good* indicators. It controls the *bit-slip* port of the MGT transceiver and declares the status of the link. When the slip signal is asserted, the boundaries of the parallel data bus slide by one bit. The next bit slip can occur after at least 32 clock cycles. The flow chart of this process is depicted in Figure 8.5. When a slip can occur, the first indicator that is taken into account is the *check-data*. The process continues only when it is asserted. Next, there are two cases for the link, either the sequence header is locked, which means that the status of the link is *Up*, or the header is unlocked, stating that the link is *Down*. When the link is *Down*, the *data-good* indicator defines whether a bit-slip should be requested. A request is not needed when the Type Field and Header values match one of the expected. If this case occurs more than 64 consecutive times, the Link is declared as *Locked*. If not, even one wrong combination resets *header-valid* counter to zero. In the case where the link is *Locked*, the *data-good* signal is constantly monitored. When a wrong sequence arrives, the *header-valid* counter is subtracted by 8. When a correct sequence arrives, it increases by 1. If its value fall below 8, the sequence header is considered unlocked and the status of the link goes *Down*.

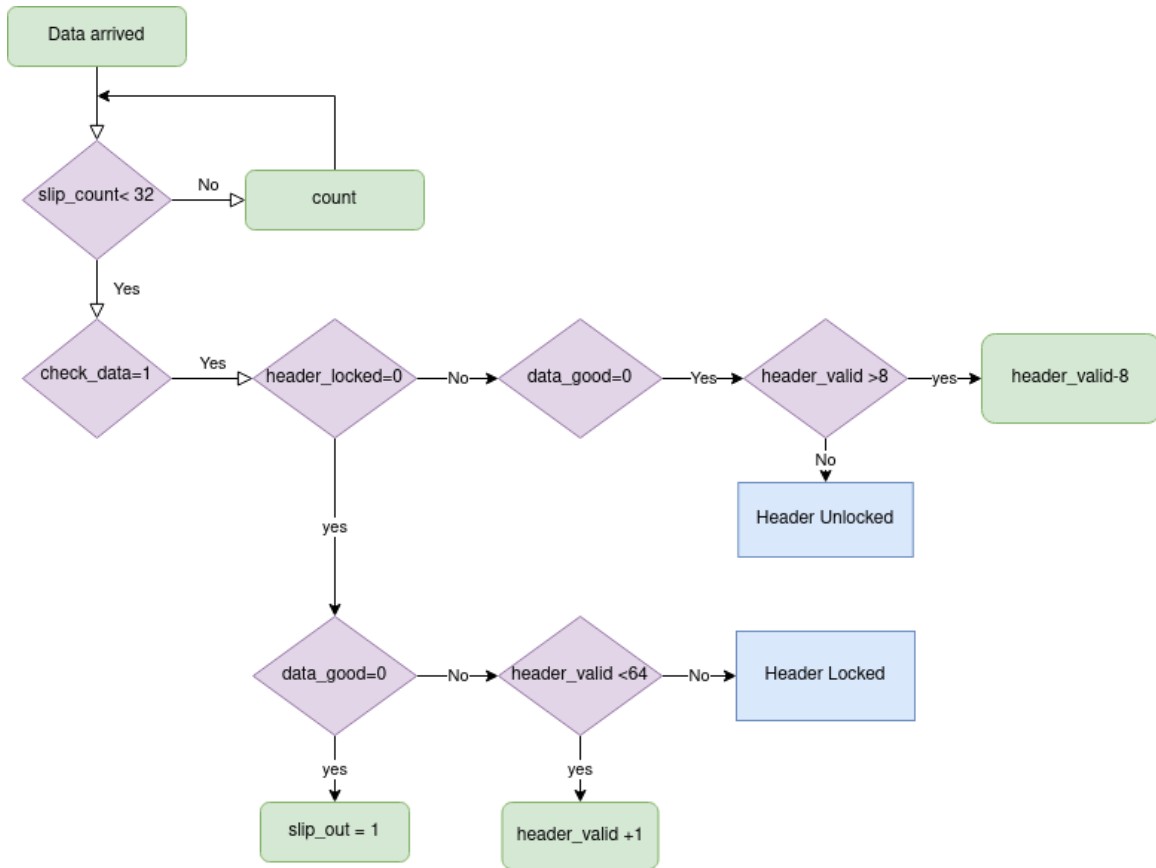


Figure 8.5: Flow chart of the CSP alignment procedure.

### 8.3.4 Filler detection and BRAM control

The decoded CSP data remain in the link clock domain until they are written to the Rx BRAM. At this stage, the Filler words have to be detected and not be written to the BRAM. The only words that cross to the algorithm domain are Data, Idles and the CRC.

Filler detection is performed using the Type Field of every Control word. For every type, a corresponding indicator is asserted when such a word is received. For Filler words that carry information in their payload, their value is read and written to registers. These are the CRC checksum and link meta-data information from the two link ID words. The BC0/OT word asserts an **orbit-tag** signal. In parallel, the separate Filler indicators pass through a logical OR gate to generate the *filler-detected* signal. It is used to drive the *write-enable* port of the BRAM. Thus Filler words never cross to the algorithm clock domain.

The *write* and *read* pointers of the Rx BRAM are controlled at this stage as well. The write pointer is a free-running counter that counts only when the write enable is High. The read pointer is controlled by the enable signal. It also consists of a free-running counter but there are a few additional signals that can manipulate its value. Two of them can either pause the increment for one clock cycle or increase by one additional value and one of them resets the counter to a specific value. These three signals are controlled by the EMP Software in order to direct the Data coming out of the BRAM to a specific clock cycle. Controlling the BRAM read pointer of

a certain number of links can align their output to the same clock. This process is one of the basic functionalities of EMP and is controlled by the Align Module of the Region block, as described in section 7.2.4.

Another piece of information extracted by the Filler words is the Index Number. The value of this field together with the filler-detector signal are sent to the Index Correction Mechanism component. Its functionality is described in the next subsection.

### 8.3.5 Index Correction Mechanism

The Index Correction Mechanism (ICM) has been introduced to correct potential flaws in the detection of Filler words. Miss-identification can occur in cases of more than 1 bit errors on either the Header or the Type Field. The ICM functions by tracking the Index Number value of every received Filler word. The CSP transmitter sends a 4-bit counter that increases with the transmission of every Filler. Similarly, the receiver locks to the first Index Number value and creates a separate counter that increases with the reception of a Filler. If the values of the two counters match, it means that the operation of the link is valid. If not, the receiver has either missed a Filler word, or confused a Data or Idle word with a Filler. In both cases, the ICM block takes action and issues either an increase or decrease in the value of the BRAM read pointer. This way, the loss of alignment that has occurred is corrected automatically only a few clock cycles after the incident.

### 8.3.6 Domain Crossing

Domain crossing on the receiver datapath is performed using a dual port BRAM. The size of the block is identical to that of the Tx FIFO, which is 36K configured as 72×512. The reason of implementing a BRAM instead of a FIFO is to have absolute control over the write and read pointers for applications such as the link alignment. The only kinds of words written to the BRAM are Data and Idles and they occupy 64 out of the 72 available width bits. The remaining are used for information that also needs to cross to the algorithm domain. Similar to the Tx FIFO, this information includes the Valid bit, the Orbit Tag bit and a CRC Flag bit. In addition, contrary to the Tx case, the CRC value crosses in the 4 available MSBs. As described in 7.4.5, the 16-bit checksum breaks into 4 chunks of 4 bits each and crosses this way to the algorithm clock domain. This process, however, requires 4 clock cycles to be completed. For this reason, the CRC flag is crossed immediately in order to issue a reset on the Rx CRC block.

### 8.3.7 CRC Check

The CRC calculation on the Rx side is also performed in the algorithm clock domain. Data words from the BRAM are directed to the Rx CRC block, together with the CRC checksum, the CRC flag and the EoP bit. The calculation is performed by the same exactly CRC block and for the same Packet boundaries as the Tx side. The Valid bit enables the calculation for Packets separated by Idle words and the EoP declares the end for back-to-back packets.

As discussed in the previous sub-section, the checksum arrives to the CRC block 4 clock cycles after the actual end of calculation. This information is propagated to the block by the CRC flag bit. Its reception issues a reset on the CRC block. The checksum value is latched and the Tx CRC is compared to the Rx CRC when the former arrives. If the two are different, an indicator is asserted and the corresponding error counter is increased. The CRC flag is also used by the packet counter, since the reception of the CRC is always equal to the end of a packet.

### 8.3.8 User Interface

The user interface on the Rx datapath is similar to that on the Tx side. The output of the CSP component, or the input to the payload block, is a record that includes: the 64-bit data words, the Valid bit, the start of packet, the end of packet and the start of orbit. When packets are defined by the Valid bit, the start of packet is asserted with the first frame of the packet and the end of packet with the last frame of the packet. When back-to-back packets are transmitted, the Valid bit is constantly high and the end of packet is always followed by the start of packet bit. In both cases, the start of orbit arrives alongside the data frame that the transmitter has asserted the corresponding bit. All above signals are clocked to the payload-clock, which in the 25 Gbps links is 360 MHz.

## 8.4 Control and Status registers

The Hermes implementation of CSP contains a wide set of control and status registers, separated in common (Quad-oriented) and in channel (channel-oriented). They contain registers that originate either from the CSP protocol implementation, or directly from the MGT IP core. All of them are connected to IPbus read&write slaves, so that they can be accessible by software. The complete list of the CSP control and status registers containing a description of their operation can be seen in Appendix A.

The control signals include the set of resets that are used to configure the CSP firmware, ports of the MGT that provide direct access to functionalities of the transceiver, and other control options of the protocol. There is a number of different reset registers that apply to: the Tx Datapath and QPLL of the MGT, the Rx Datapath of the MGT, the Tx and Rx firmware blocks of every channel separately, the Latched status signals, as well as the error counters. Some of the MGT IP registers provide access to settings such as: power on and off the QPLL, Tx and Rx of the MGT, Loopback selection, PRBS data selection, Eye-diagram ports, Tx and Rx polarity, and others. In addition, the injection of CSP errors, as listed in Table 8.8, is performed by the relevant control registers.

The status registers of CSP include indicators that inform about the status of the Tx and Rx firmware blocks, as well as that of the MGT. Some examples are: the MGT Tx and Rx Reset Done, the Link Status and Link Down Latch indicators. Moreover, there is a wide set of error counters, such as the CRC counter and the Packet counter. Counters are also implemented for: Header single-bit errors, Header double-bit errors, Control Word Type single-bit errors, Control Word Type double-bit errors, errors in

the received Index Number. Furthermore, the number of times the Index Correction Mechanism has been engaged is also registered to the corresponding counter.

## 8.5 Software Interface and Configuration

The control and status registers of CSP are accessible to IPbus transactions through their corresponding address tables. Control of the links is performed using the emp-butler software. It contains a set of commands that specifically target the configuration and monitoring of CSP links. The basic steps that are followed to configure a set of 20 links in Loopback mode is described below. Prior to running these commands, the Tx Channel Buffer is filled with counter data patterns.

Initially, the transmitter side has to be configured. The emp-butler command used here is called *mgts configure tx*. It powers on the QPLL and Tx MGT of the specified channels, configures the Loopback mode (optional) and runs the appropriate reset sequence. In the end, it executes a status check and informs the user about the result. The output of running this command in *Near-end PMA* Loopback mode can be seen in Figure 8.6.

```
[root@bmtl1-2 fpga-ctrl]# empbutler -c connections.xml do bmtl1_fpga mgts configure tx --loopback nearPMA
12-12-23 14:43:15.140589 [281473107263488] WARNING - Address overlaps observed - report file written at "/tmp/connections.xml.txt"
12-12-23 14:43:15.146361 [281473107263488] NOTICE - mmap client with URI "ipbusmmap-2.0:///dev/mem?offset=0x00000000"
Firmware revision - design: 0x0a infra: 0x000801 algo: 0x00000200
Channels - [0-127]
Powering up QPLL and Tx links in regions [4-8]
Configuring channels [16-35]
Resetting links in regions [4-8]
Resetting latched + counters

Running post-configure status checks ...
... all channels report good status!
[root@bmtl1-2 fpga-ctrl]#
```

**Figure 8.6:** Example of running the *configure tx* command of emp-butler.

The receiver side is configured in a similar way. The emp-butler command is now called *mgts configure rx*. It powers on the Rx MGT and configures the corresponding channels. The appropriate reset sequence is performed and the value of all status registers is checked. The result of running this command in the same set of links is shown in Figure 8.7. As can be seen, the status of all channels is reported as "good".

```
[root@bmtl1-2 fpga-ctrl]# empbutler -c connections.xml do bmtl1_fpga mgts configure rx
12-12-23 14:48:42.120136 [281473373597696] WARNING - Address overlaps observed - report file written at "/tmp/connections.xml.txt"
12-12-23 14:48:42.125905 [281473373597696] NOTICE - mmap client with URI "ipbusmmap-2.0:///dev/mem?offset=0x00000000"
Firmware revision - design: 0x0a infra: 0x000801 algo: 0x00000200
Channels - [0-127]
Powering up QPLL and Rx links in regions [4-8]
Enabling eyescan circuitry for channels [16-35]
Configuring channels [16-35]
Resetting links in regions [4-8]
Resetting link status registers and alignment monitor for channels [16-35]

Running post-configure status checks ...
... all channels report good status!
[root@bmtl1-2 fpga-ctrl]#
```

**Figure 8.7:** Example of running the *configure rx* command of emp-butler.

Detailed information about the status of all CSP registers can be printed by running the *mgts status rx* command. The output of this command for the same set



of 20 links can be seen in Figure 8.8. It contains a list where every row corresponds to one MGT receiver link. The columns contain the value of the corresponding registers, such as the Status, the Packets in, CRC errors, etc. Furthermore, the position of the received Orbit Tag is printed.

```
[root@bmt11-2 fpga-ctrl]# empbutler -c connections.xml do bmt11 fpga mgts status rx
12-12-23 14:53:11.170364 [281472972771328] WARNING - Address overlaps observed - report file written at "/tmp/root/uhal/OverlapReport-root-fpga-ctrl-decoders.xml.txt"
12-12-23 14:53:11.176404 [281472972771328] NOTICE - mmap client with URI "ipbusmmap-2.0:///dev/mem?offset=0xa1010000" : Address offset set to 0xA1010000
CSP status:
Channel  Config  QPLL_CDR  Init  Status  Packets  CRC  Header  CtrlType  errors  Index  ICM  FIFO  Orbit tag  OT  OT counters
          locks  (latch)  in      errors  errors  Single,Double  error  depth  position  status  Good:Bad:Absent
-----
16      LPM      1         1     1 (1)   3015763  0    0        0,0      0    0    5    32    11111111  3015964:0:0
17      LPM      1         1     1 (1)   3015764  0    0        0,0      0    0    5    32    11111111  3015962:0:0
18      LPM      1         1     1 (1)   3015766  0    0        0,0      0    0    8    33    11111111  3015961:0:0
19      LPM      1         1     1 (1)   3015769  0    0        0,0      0    0    6    33    11111111  3015960:0:0
20      LPM      1         1     1 (1)   3015779  0    0        0,0      0    0    5    31    11111111  3015959:0:0
21      LPM      1         1     1 (1)   3015781  0    0        0,0      0    0    5    32    11111111  3015958:0:0
22      LPM      1         1     1 (1)   3015783  0    0        0,0      0    0    7    33    11111111  3015957:0:0
23      LPM      1         1     1 (1)   3015785  0    0        0,0      0    0    5    32    11111111  3015956:0:0
24      LPM      1         1     1 (1)   3015797  0    0        0,0      0    0    5    33    11111111  3015955:0:0
25      LPM      1         1     1 (1)   3015798  0    0        0,0      0    0    6    32    11111111  3015955:0:0
26      LPM      1         1     1 (1)   3015801  0    0        0,0      0    0    6    33    11111111  3015953:0:0
27      LPM      1         1     1 (1)   3015802  0    0        0,0      0    0    5    32    11111111  3015953:0:0
28      LPM      1         1     1 (1)   3015812  0    0        0,0      0    0    6    32    11111111  3015951:0:0
29      LPM      1         1     1 (1)   3015814  0    0        0,0      0    0    6    33    11111111  3015951:0:0
30      LPM      1         1     1 (1)   3015816  0    0        0,0      0    0    5    32    11111111  3015949:0:0
31      LPM      1         1     1 (1)   3015818  0    0        0,0      0    0    6    32    11111111  3015949:0:0
32      LPM      1         1     1 (1)   3015831  0    0        0,0      0    0    5    32    11111111  3015947:0:0
33      LPM      1         1     1 (1)   3015833  0    0        0,0      0    0    6    33    11111111  3015946:0:0
34      LPM      1         1     1 (1)   3015836  0    0        0,0      0    0    4    32    11111111  3015945:0:0
35      LPM      1         1     1 (1)   3015837  0    0        0,0      0    0    4    32    11111111  3015944:0:0
[root@bmt11-2 fpga-ctrl]#
```

Figure 8.8: Example of running the *mgts status rx* command of *emp-butler*.

As can be seen, prior to aligning all links to the same bunch crossing number the Orbit Tag of every link is received randomly. The software command that aligns all links together is called *mgts align*. This command also supports an optional setting that defines the BX number and the frame where the Orbit Tag, or BC0, will be placed. The result of the status command after performing alignment on the links is shown in Figure 8.9. As can be seen, the Orbit Tag position is now the same for every channel, at place 33.

```
[root@bmt11-2 fpga-ctrl]# empbutler -c connections.xml do bmt11 fpga mgts status rx
12-12-23 15:02:50.780196 [281473297661952] WARNING - Address overlaps observed - report file written at "/tmp/root/uhal/OverlapReport-root-fpga-ctrl-decoders.xml.txt"
12-12-23 15:02:50.785964 [281473297661952] NOTICE - mmap client with URI "ipbusmmap-2.0:///dev/mem?offset=0xa1010000" : Address offset set to 0xA1010000
CSP status:
Channel  Config  QPLL_CDR  Init  Status  Packets  CRC  Header  CtrlType  errors  Index  ICM  FIFO  Orbit tag  OT  OT counters
          locks  (latch)  in      errors  errors  Single,Double  error  depth  position  status  Good:Bad:Absent
-----
16      LPM      1         1     1 (1)   234311   0    0        0,0      0    0    6    33    11111111  234399:0:0
17      LPM      1         1     1 (1)   234314   0    0        0,0      0    0    8    33    11111111  234400:0:0
18      LPM      1         1     1 (1)   234316   0    0        0,0      0    0    6    33    11111111  234401:0:0
19      LPM      1         1     1 (1)   234318   0    0        0,0      0    0    7    33    11111111  234401:0:0
20      LPM      1         1     1 (1)   234328   0    0        0,0      0    0    6    33    11111111  234403:0:0
21      LPM      1         1     1 (1)   234331   0    0        0,0      0    0    4    33    11111111  234403:0:0
22      LPM      1         1     1 (1)   234332   0    0        0,0      0    0    6    33    11111111  234404:0:0
23      LPM      1         1     1 (1)   234334   0    0        0,0      0    0    7    33    11111111  234405:0:0
24      LPM      1         1     1 (1)   234344   0    0        0,0      0    0    7    33    11111111  234406:0:0
25      LPM      1         1     1 (1)   234346   0    0        0,0      0    0    7    33    11111111  234406:0:0
26      LPM      1         1     1 (1)   234348   0    0        0,0      0    0    6    33    11111111  234407:0:0
27      LPM      1         1     1 (1)   234349   0    0        0,0      0    0    6    33    11111111  234408:0:0
28      LPM      1         1     1 (1)   234360   0    0        0,0      0    0    7    33    11111111  234409:0:0
29      LPM      1         1     1 (1)   234361   0    0        0,0      0    0    7    33    11111111  234409:0:0
30      LPM      1         1     1 (1)   234363   0    0        0,0      0    0    7    33    11111111  234410:0:0
31      LPM      1         1     1 (1)   234364   0    0        0,0      0    0    6    33    11111111  234411:0:0
32      LPM      1         1     1 (1)   234375   0    0        0,0      0    0    5    33    11111111  234411:0:0
33      LPM      1         1     1 (1)   234377   0    0        0,0      0    0    6    33    11111111  234412:0:0
34      LPM      1         1     1 (1)   234378   0    0        0,0      0    0    6    33    11111111  234413:0:0
35      LPM      1         1     1 (1)   234380   0    0        0,0      0    0    6    33    11111111  234414:0:0
```

Figure 8.9: Example of running the *mgts status rx* command of *emp-butler*, after running the *align* command.

## 8.6 Firmware Performance

There are two main kinds of tests that validate the operation and robustness of the links. The first kind refers to multiple re-configurations of the links, followed by

checks of its basic status reports. This test targets the validation of the firmware and its ability to bring Up the links after every configuration. The number of repeats that set this test successful should be in the order of thousands. The second kind of test validates the links operation for prolonged periods of time, by configuring the firmware and have it running for days. During the development stages of the CSP firmware, both of the above tests took place, using the emp-butler software. It assisted the execution of these tests, by providing immediate status reports after the configuration of the links.

Further validation involved tests between two ATCA boards. These kinds of tests have been conducted multiple times, by using different Phase-2 platforms. The basic test platform has been Serenity. However, extended testing has been performed by the Ocean and X2O boards as well. Nevertheless, the Serenity, X2O and BMTL1 use the same firmware implementation of CSP.

### 8.6.1 Final CSP Testing

The final testing of the implementation of the CSP syntax involved communication between hardware that uses the two different CSP implementations. These were the APd1 board (described in 5.1.7) and the Serenity board (described in 5.1.7). The tests were held in the standard integration facility of the Phase-2 Level-1 Trigger, where the two boards were installed in the same crate and connected with 4 optical bi-directional links operating at 25.78125 Gbps.

The CSP validation procedure consisted of three sets of tests, as depicted in the list of Figure 8.10. The first set targeted normal operation of the links, where five Packet structures were transmitted through the links. The second set, called *forced errors* test, involved the deliberate injection of errors to the links in both directions. The error injection mechanism of the firmware was used for this test, as well as the corresponding error counters to measure the result. The final, third test is a normal operation long duration test, targeting the endurance validation of the links

The APd1-Serenity tests were held two times in 2023. During the first series, communication between the two boards was established, and the tests mentioned above were performed normally. The outcome of most of them was successful. However, a few issues appeared that had to be resolved. The second series, held in the summer of 2023, repeated the same test sequence. This time the modified firmware was used, where all previous issues had been resolved. During this round of tests, both firmware implementations performed as expected. The most notable test took place after the execution of the three sets mentioned above. This last test would attenuate one of the 4 channels in order to force errors in the link. The result of the Serenity Rx status after running for 23 hours is depicted in Figure 8.11. As can be seen, all error counters have accumulated a high count of errors. In addition, the Index Correction Mechanism (ICM) had to engage 930 times in order to prevent loss of the links alignment.

1.1	Normal operation: TM2 back-to-back packets
1.2	Normal operation: TM6 back-to-back packets
1.3	Normal operation: TM6 with gap between packets
1.4	Normal operation: TM18 back-to-back packets
1.5	Normal operation: TM18 with gap between packets
2.1	Forced error: Single bitflip in header
2.2	Forced error: Double bitflip in header
2.3	Forced error: Single bitflip in control word type
2.4	Forced error: Double bitflip in control word type
2.5	Forced error: Deliberate CRC mismatch
2.6	Forced error: Bitflip in index number
2.7	Forced error: Incorrect polarity
2.8	Blind misconfigurations & error injections
3.1	Normal operation: Long duration "endurance run"

**Figure 8.10:** List of the three sets of tests performed between an APd1 and a Serenity board to validate the CSP implementation.

CSP status:														
Channel	Config	QPLL,CDR locks	Init	Status (latch)	Packets in	CRC errors	Header errors	CtrlType errors Single,Double	Index error	ICM	FIFO depth	Orbit tag position	OT status	OT counters Good:Bad:Absent
56	LPM	1	1	1 (1)	4.295e+09	0	0	0,0	0	0	7	195	11111111	951074282:0:0
57	LPM	1	1	1 (1)	4.295e+09	2.410e+08	65535	65535,4047	65535	930	8	196	11111111	951074277:654:0
58	LPM	1	1	1 (1)	4.295e+09	0	0	0,0	0	0	5	195	11111111	951074273:0:0
59	LPM	1	1	1 (1)	4.295e+09	0	0	0,0	0	0	5	195	11111111	951074268:0:0
60	LPM	1	1	1 (1)	4.295e+09	0	0	0,0	0	0	6	48	11111111	951074263:0:0
61	LPM	1	1	1 (1)	4.295e+09	0	0	0,0	0	0	5	47	11111111	951074259:0:0

**Figure 8.11:** Result of the APd1-Serenity tests using attenuated channel for 23 hours.

# Chapter 9

## The BMTL1 ATCA Trigger Processor

### 9.1 Introduction

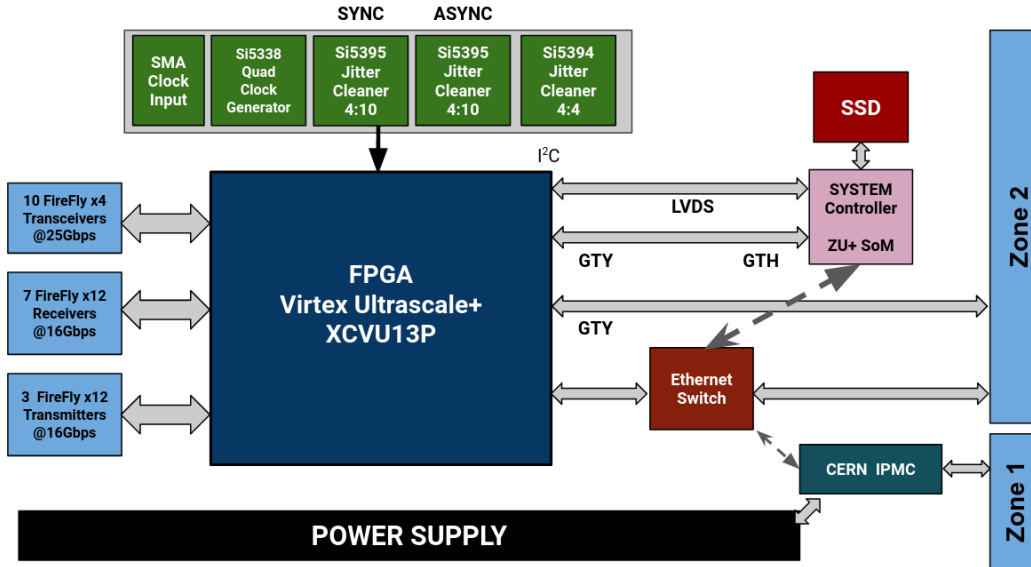
The BMTL1 ATCA is the trigger processor that will instrument the Barrel Muon Trigger Layer-1 subsystem, as described in Chapter 6. The BMTL1 subsystem receives data from the DT and RPC detectors through 10 Gbps optical links, originating from the OBDT and MLB boards, respectively. The information arrives in the form of TDC hits, produced by muons that cross the detector chambers. Processing of the DT TDC hits is performed by the Analytical Method algorithm. It runs in the firmware framework of the BMTL1 board to generate Trigger Primitives in the form of track segments, also called stubs. The output is transmitted to the next processing layer, the Global Muon Trigger, via 25 Gbps optical links. The Kalman Filter algorithm runs at the GMT boards to match track segments and reconstruct muon particles that have crossed the barrel region of CMS.

This chapter describes work that has been conducted for the BMTL1 board, as part of this theses. It includes a description of high level design concepts of the board, such as the optical interfaces and the clocking network, and tests that have been performed upon the reception of the first prototype. Furthermore, the integration of the firmware that runs on the FPGA is also described.

### 9.2 Hardware Design

A high level block diagram of the BMTL1 board is depicted in Figure 9.1[38]. The design is based on a powerful VU13P FPGA (*XCVU13P-1FLGA2577E*) that constitutes the central processing unit. This device was chosen as it combines a large number of logical resources, high count of serial transceivers and, most importantly, reasonable market price and purchase availability. The resources of the VU13P include: 1,728,000 LUTs, 3,456,000 FF, 12,228 DSPs, 128 GTY transceivers, 94.5 Mb Block RAM and 360 Mb UltraRAM [39]. The power supply of the board is capable of delivering up to 420 Ampere to the core voltage of the VU13P FPGA through seven phases. Its optical connectivity is facilitated by 20 Samtec Firefly modules (see 4.3.2) running at both 16 Gbps and 25 Gbps. In addition, there is a complicated clocking network that provides both synchronous and asynchronous (with respect to

LHC) reference clocks to the Multi Gigabit Transceiver (MGT) Quads, as well as free-running clocks.



**Figure 9.1:** Block diagram of the BMTL1 ATCA board.

Control and monitoring of the board is performed through a Xilinx ZYNQ Ultra-scale Plus System-on-Module, the ZU5EG. The Processing System (PS) of the ZYNQ is supplemented by an SSD (Solid State Disk) and an SD (Secure Digital) card. It runs a CentOS 7 operating system (OS) that can access the Internet through an Ethernet Switch, also included on the board. The switch provides several connectivity options, that include a standard RJ-45 Ethernet connector at the front panel and connection to the Ethernet signals coming from the Zone 2 backplane connector. The OS of the ZYNQ is the essential system controller of the board and the user can connect to it over the network. Additional connectivity options include a USB 3.0, a Display Port and a UART port. The FPGA, ZYNQ and IPMC can be programmed using three JTAG headers attached on the board. Interface to all programmable peripherals, such as clock synthesizers, Firefly modules and the Power modules, is performed by I2C lanes that connect to the ZYNQ. Moreover, there is extended connectivity between the ZYNQ Programmable Logic (PL) and the FPGA through four MGT channels operating at up to 10 Gbps, and also by 21 differential pair signals, one of which is dedicated for clock.

Following the ATCA standard, the board includes an IPMC module to facilitate the corresponding transactions with the crate. The CERN-IPMC device [21] is instrumented for this purpose. It connects to the shelf manager and communicates information of the power supply status and temperature. It also connects to the Ethernet Switch and the user can perform actions from distance, such as power on and off, when it is installed inside a crate. Furthermore, the CMS-defined Zone 2 signals are connected to the FPGA, either directly or through the chips of the clocking network. These include the LHC 40.078 MHz clock, a High Precision (HP-LHC) multiple of it and the TCDS2 serial stream.

### 9.2.1 High Speed Serial Links

During the time of designing the BMTL1 board neither the final number of input and output links was exactly defined, nor the type of the links. Furthermore, it was not clear whether the BMTL1 board will process data from one or two DT chambers. As a result, the optical interfaces of the board are designed in a way to be able to serve all potential use case scenarios existed at that time, including the most demanding ones. That being said, the number of 25 Gbps links at the board is 40. The known number needed at that time was 18 outputs going to GMT and 2 going to the DTH for readout. For contingency reasons their amount was doubled. On the receiving side, the maximum number of OBDT links was calculated to be 43. In addition, 5 links were required for the RPC links. Hence, the scenario of running 2 Sectors in 1 BMTL1 board required 96 lpGBT links. However, the remaining free MGT channels of the VU13P FPGA is 80, after we exclude the 25 Gbps links, the ZYNQ links and the Zone 2 links. Thus, these 80 links were instrumented with modules that operate at line rate up to 16 Gbps. In case more than 80 Rx links are needed, the usage of the receiving sides of the 25 Gbps links is possible. Furthermore, it was decided to instrument 36 transmitter channels for front-end communication, in case they are needed.

The optical modules used at BMTL1 are Samtec Fireflies. More specifically, links that are foreseen to operate front-end protocols are routed to receiving x12 (Rx12) and transmitting x12 (Tx12) connectors. The modules attached on these connectors can support line rates of up to 10, 14 or 16 Gbps. The links that will operate the CSP back-end protocol are routed to x4 bi-directional connectors. The Firefly parts foreseen here can operate to up to 25 Gbps. Routing of the VU13P MGT Banks to Firefly connectors is depicted in Figure 9.2. The 25 Gbps channels are placed closer to the FPGA to maintain short PCB traces and reduce signal losses. There are 10 such parts, labeled a  $TxRx_*$  (light blue color). The 7 receiving parts are labeled as  $Rx_*$  (green color) and the 3 transmitting parts as  $Tx_*$  (light pink color). The Figure also illustrates the ZYNQ (128) and Zone 2 (220) Banks.

### 9.2.2 Clocking Network

The BMTL1 ATCA board, as described in this thesis, is the first revision of an ATCA processor for the BMTL1 subsystem. As such, multiple use case scenarios have been considered during the design of its clocking network. First of all, it was foreseen that it would operate in different environments, such as the bench of a laboratory or inside an ATCA crate. Furthermore, the CMS guidelines point to very specific implementations regarding the delivery of the LHC clock to the FPGA logic, as discussed in 7.2.1. It also foresees two different kinds of reference clocks routed to the MGT QPLLs, one synchronous and one asynchronous with respect to the LHC clock. To accommodate all of the above requirements, the BMTL1 ATCA adopted the clocking network illustrated in Figure 9.3.

The network is based on one Si5344 and two Si5345 jitter cleaner (JC) clock multiplier chips from Skyworks [40]. The functionality of the two chips is essentially the same, with the only difference being the number of output clocks. The Si5344 chip features 4 input and 4 output clocks, while the Si5345 features 4 input and 10 outputs. As a result, both of them also act as clock multiplexers. The chips

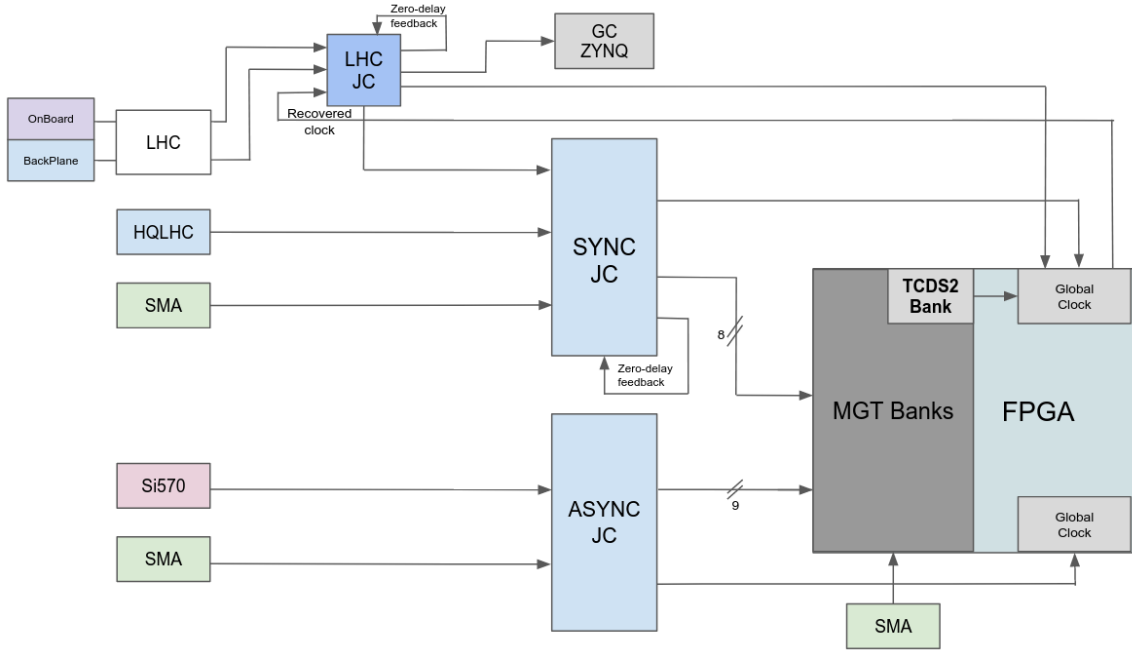
	Rx4	BANK 135	SLR 3	BANK 235	Rx7	
		BANK 134		BANK 234		
	Rx3	BANK 133		BANK 233		
		BANK 132		BANK 232		TxRx_9
		BANK 131	SLR 2	BANK 231		TxRx_8
TxRx_10		BANK 130		BANK 230		TxRx_7
TxRx_3		BANK 129		BANK 229		TxRx_6
ZYNQ		BANK 128		BANK 228		TxRx_5
		BANK 127	SLR 1	BANK 227		TxRx_4
TxRx_2		BANK 126		BANK 226	Rx6	
TxRx_1		BANK 125		BANK 225		
	Rx2	BANK 124		BANK 224		
Tx2		BANK 123	SLR 0	BANK 223		Rx5
		BANK 122		BANK 222		
	Rx1	BANK 121		BANK 221		
Tx3		BANK 120		BANK 220	Zone 2	

**Figure 9.2:** Connections of the BMTL1 ATCA MGT Banks to Samtec Firefly modules.

support free-running, synchronous and holdover operation modes and can be easily programmed through a standard I2C interface. In synchronous mode, all or some of the output clocks can be synchronized with one of the four inputs. In the clocking network of BMTL1 ATCA, the Si5344 is referred to as *LHC-JC*, one of the two Si5345 is referred to as *Sync-JC* and the other as *Async-JC*.

The main consumers of low jitter output clocks of the Sync and Async JCs are the FPGA's MGT Banks. The VU13P device (described in section 4.2.3) contains 128 MGT channels grouped in 32 Banks. Furthermore, the chip logic is divided into 4 SLRs (Super Logic Region), with each SLR containing 4 MGT banks at its left side and 4 MGT Banks at its right side, as shown in Figure 9.4. On each side, a reference clock can be shared between Banks of the same SLR. For line rates below 16.375 Gbps, the reference clock can be sourced from up to two Quads above or two Quads below. For line rates between 16.375 Gbps and 28.21 Gbps, sharing is possible between up to one Quad above and one below [41]. As a consequence, the front-end channels (maximum 10.24 Gbps) can be served by one synchronous reference clock connected to each side of every SLR, making the number of sync reference clocks be 8, labeled as *CLK\_SYNC\_\**. The asynchronous clocks target 25 Gbps links. In most cases the number of such Quads per SLR side is three or less, except from the right side of SLR 2 that contains 4 25 Gbps Quads. As a result, nine outputs of the Async-JC are connected in total, labeled as *CLK\_ASYNC\_\**. The Sync and Async clock connections to the relevant MGT Banks and SLRs are also depicted in Figure 9.4.

The LHC-JC implements the clocking circuitry that is needed to deliver the LHC clock both to the FPGA and to one input of the Sync-JC. Two of the LHC-JC inputs are connected to LHC clock sources. One is the 40.078 MHz clock coming from the



**Figure 9.3:** Block diagram of the clocking network of the BMTL1 ATCA board.

DTH through Zone 2, and the other is a free-running source produced on the board. The third input is connected to a Global Clock (GC) differential pin pair of the FPGA. This connection is provisioned to serve the standard TCDS2 implementation, which requires the recovered clock of the TCDS2 stream to be routed outside the FPGA and inside a jitter cleaner chip. It is then cleaned and get re-routed back to the FPGA logic, while in parallel used as reference clock of the Sync-JC. This implementation also requires one I/O pair of the LHC-JC to be configured in zero-delay feedback mode, as can be seen in Figure 9.3. The fourth output is connected to a GC pair of the ZYNQ in case the LHC clock is needed for any application.

The BMTL1 clocking network also includes SMA (SubMiniature version A) connectors to assist the operation of the board either on the bench or in other work case scenarios. Five such pairs are connected in total to provide external clock source to different consumers on the board. One is routed to a GC pin pair of the FPGA logic, two are connected to QPLL pins of Bank 130 and Bank 222, one is routed as input to the Sync-JC and one as input to the Async-JC. Furthermore, one output of the Sync-JC and one of the Async-JC are connected to GC pin pairs to be used by the FPGA logic in cases it is needed.

### 9.3 Hardware Testing

The first BMTL1 prototype was received on April 2022. Upon its arrival, every hardware component and circuitry had to be verified and stress tested. The first action was the configuration of the power supply. There is a number of different voltages created on the board that deliver power to different chips. Once their operation was validated, the VU13P FPGA and ZYNQ devices were powered on and ready to be used.



	BANK 135	SLR 3	BANK 235	
CLK_ASYNC_4	BANK 134		BANK 234	CLK_ASYNC_9
CLK_SYNC_4	BANK 133		BANK 233	CLK_SYNC_8
	BANK 132		BANK 232	
	BANK 131	SLR 2	BANK 231	CLK_ASYNC_8
CLK_SYNC_3	BANK 130		BANK 230	CLK_SYNC_7
CLK_ASYNC_3	BANK 129		BANK 229	CLK_ASYNC_7
	BANK 128		BANK 228	
	BANK 127	SLR 1	BANK 227	
CLK_ASYNC_2	BANK 126		BANK 226	CLK_ASYNC_6
CLK_SYNC_2	BANK 125		BANK 225	CLK_SYNC_6
	BANK 124		BANK 224	
	BANK 123	SLR 0	BANK 223	
CLK_SYNC_1	BANK 122		BANK 222	CLK_ASYNC_5
CLK_ASYNC_1	BANK 121		BANK 221	CLK_SYNC_5
	BANK 120		BANK 220	

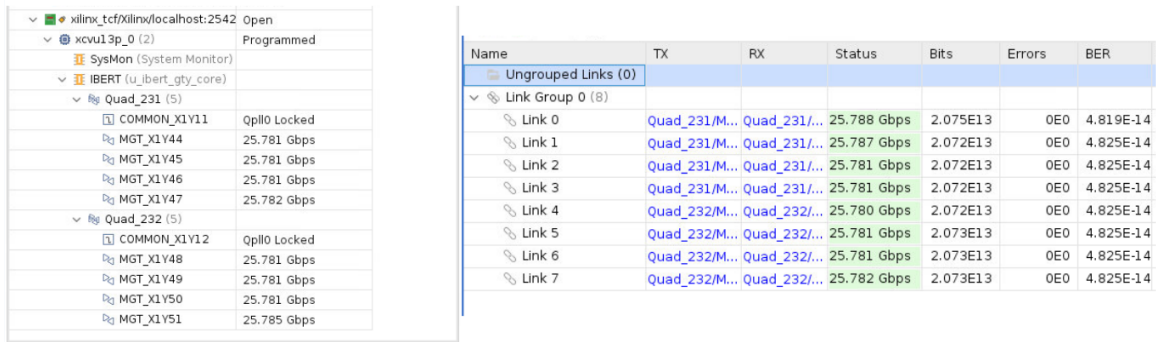
**Figure 9.4:** Connections of Sync and Async reference clocks to the BMTL1 ATCA MGT Banks.

One important requirement at this testing stage was the installation of CentOS 7 at the ZYNQ's processor. An OS (Operating System) is needed here in order to interface with all the on-board peripherals through the I2C protocol. The CentOS was chosen because it is required by the software that runs on the ZYNQ and controls the FPGA. Once installed, parts like the clocking network Jitter Cleaners could be configured and provide clocks to the FPGA and ZYNQ logic.

### 9.3.1 Optical links


The optical interfaces of the BMTL1 were verified using the IBERT IP, described in 4.3.1. It provides a tool to validate the operation of reference clock sources, stability of the transceiver channels and quality of the links. The clocking sources, as described in section 9.2.2, are divided in Synchronous and Asynchronous, originating from the corresponding JC chips. To validate every output, an IBERT project was created that instantiates all 32 MGT Banks of the VU13P FPGA using all of the 17 in total reference clocks (8 Sync and 9 Async). Reports of the IBERT project are shown in Figure 9.5. The left image illustrates the status of QPLLs that are used to drive each Quad. The status in both is reported as *QPLL Locked*, indicating that stable reference clock is provided in the correct frequency. The same status was true for the remaining, not depicted quads. The status of active links is shown both on the left, and on the right of the same Figure (green color and line rate measurement), together with the number of transmitted bits, errors detected in the reception and the BER estimation.

The links are created using Firefly modules and by connecting the transmitter side of one Bank to the receiver side of another via optical fiber cables. A connection setup can be seen in Figure 9.6, where the BMTL1 board is instrumented with 8 Firefly modules. Next testing step included scans to produce eye diagrams for every link.




**Figure 9.5:** *Left: IBERT interface indicating that QPLLs are Locked. Right: IBERT interface showing the status of the links, bits transmitted and errors detected.*

The IBERT settings of these tests were: BER  $1^{-8}$ , PRBS-31 sequences, horizontal and vertical step increment 1. The procedure was repeated for every MGT Bank by changing each time the position of the Firefly modules. The front-end links were tested to their maximum supported line rate of 16 Gbps. The resulting eye diagrams can be seen in Figure 9.7. Figure 9.8 illustrates eye scans taken for the 40 links operating at 25.78125 Gbps.




**Figure 9.6:** *The BMTL1 board instrumented with 8 Firefly modules.*

The resulting eye-scans of both 16G and 25G links indicate very good signal integrity for every link of the board. All diagrams have an extended open area in which errors are not produced (blue color). This concludes to an expected error-free performance for extended periods of time. To assist this statement, endurance tests were held where links were left running for hours. Figure 9.9 illustrates 24 channels operating at 16G and connected with optical fibers. As shown in the corresponding bar, the BER of this test reached more than  $10^{-15}$  without any error being observed in any of the links.



**Figure 9.7:** *Eye diagrams of all 16G links running in external loopback using fiber cables.*



**Figure 9.8:** *Eye diagrams of all 25G links running in external loopback using fiber cables.*

Similar endurance tests were held for 25G links. Figure 9.10 illustrates an IBERT screen-shot of 8 links running at 25.781 Gbps. The test used two bi-directional x4 Firefly parts connected with optical fiber. The accumulation of bits led to a BER of  $10^{-16}$  with no errors observed in any link.

Name	TX	RX	Status	Bits	Errors	BER	BERT Reset	TX Pattern	RX Pattern	Loopback Mode
Ungrouped Links (0)										
Found Links (24)										
Auto detected link 0	Quad_125/MGT_X0Y23/TX (xcvu13p_0)	Quad_120/MGT_X0Y0/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 1	Quad_124/MGT_X0Y18/TX (xcvu13p_0)	Quad_120/MGT_X0Y1/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 10	Quad_124/MGT_X0Y17/TX (xcvu13p_0)	Quad_122/MGT_X0Y10/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 11	Quad_123/MGT_X0Y12/TX (xcvu13p_0)	Quad_122/MGT_X0Y11/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 12	Quad_121/MGT_X0Y4/TX (xcvu13p_0)	Quad_224/MGT_X1Y16/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 13	Quad_120/MGT_X0Y0/TX (xcvu13p_0)	Quad_224/MGT_X1Y17/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 14	Quad_120/MGT_X0Y3/TX (xcvu13p_0)	Quad_224/MGT_X1Y18/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 15	Quad_121/MGT_X0Y5/TX (xcvu13p_0)	Quad_224/MGT_X1Y19/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 16	Quad_120/MGT_X0Y1/TX (xcvu13p_0)	Quad_225/MGT_X1Y20/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 17	Quad_121/MGT_X0Y6/TX (xcvu13p_0)	Quad_225/MGT_X1Y21/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 18	Quad_122/MGT_X0Y0/TX (xcvu13p_0)	Quad_225/MGT_X1Y22/RX (xcvu13p_0)	15,998 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 19	Quad_121/MGT_X0Y7/TX (xcvu13p_0)	Quad_225/MGT_X1Y23/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 2	Quad_125/MGT_X0Y22/TX (xcvu13p_0)	Quad_120/MGT_X0Y2/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 20	Quad_122/MGT_X0Y8/TX (xcvu13p_0)	Quad_226/MGT_X1Y24/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 21	Quad_122/MGT_X0Y9/TX (xcvu13p_0)	Quad_226/MGT_X1Y25/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 22	Quad_122/MGT_X0Y11/TX (xcvu13p_0)	Quad_226/MGT_X1Y26/RX (xcvu13p_0)	15,994 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 23	Quad_120/MGT_X0Y2/TX (xcvu13p_0)	Quad_226/MGT_X1Y27/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 3	Quad_124/MGT_X0Y19/TX (xcvu13p_0)	Quad_120/MGT_X0Y3/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 4	Quad_125/MGT_X0Y20/TX (xcvu13p_0)	Quad_121/MGT_X0Y4/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 5	Quad_125/MGT_X0Y21/TX (xcvu13p_0)	Quad_121/MGT_X0Y5/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 6	Quad_123/MGT_X0Y15/TX (xcvu13p_0)	Quad_121/MGT_X0Y6/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 7	Quad_124/MGT_X0Y16/TX (xcvu13p_0)	Quad_121/MGT_X0Y7/RX (xcvu13p_0)	15,998 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 8	Quad_123/MGT_X0Y14/TX (xcvu13p_0)	Quad_122/MGT_X0Y8/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None
Auto detected link 9	Quad_123/MGT_X0Y17/TX (xcvu13p_0)	Quad_122/MGT_X0Y9/RX (xcvu13p_0)	16,000 Gbps	3.073E14	0E0	3.254E-15	Reset	PRBS 31-bit	PRBS 31-bit	None

**Figure 9.9:** Long runs of 24 16G links using Samtec Fireflies. No errors counted for a BER of  $10^{-15}$ .

Name	TX	RX	Status	Bits	Errors	BER	BERT Reset	TX Pattern	Loopback Mode	TX Pre-Cursor	TX Post-Cursor	TX Diff Swing	DFE Enabled	Inject Error	TX Reset	RX Reset	RX PLL Status	TX PLL Status
Ungrouped Links (0)																		
Found Links (8)																		
Found 0	MGT_X1465/MGT_X1465R	25.781 Gbps	1.114E15	0E0	8.977E-16	Reset	PRBS 31-bit	None	None	0.00 dB (000000)	0.00 dB (000000)	950 mV (11.0000)		Inject	Reset	Reset	Locked	Locked
Found 1	MGT_X1479/MGT_X1479R	25.781 Gbps	1.114E15	0E0	8.977E-16	Reset	PRBS 31-bit	None	None	0.00 dB (000000)	0.00 dB (000000)	950 mV (11.0000)		Inject	Reset	Reset	Locked	Locked
Found 2	MGT_X1483/MGT_X1483R	25.781 Gbps	1.114E15	0E0	8.977E-16	Reset	PRBS 31-bit	None	None	0.00 dB (000000)	0.00 dB (000000)	950 mV (11.0000)		Inject	Reset	Reset	Locked	Locked
Found 3	MGT_X1497/MGT_X1497R	25.782 Gbps	1.114E15	0E0	8.977E-16	Reset	PRBS 31-bit	None	None	0.00 dB (000000)	0.00 dB (000000)	950 mV (11.0000)		Inject	Reset	Reset	Locked	Locked
Found 4	MGT_X1426/MGT_X1446R	25.781 Gbps	1.114E15	0E0	8.977E-16	Reset	PRBS 31-bit	None	None	0.00 dB (000000)	0.00 dB (000000)	950 mV (11.0000)		Inject	Reset	Reset	Locked	Locked
Found 5	MGT_X1430/MGT_X1450R	25.781 Gbps	1.114E15	0E0	8.977E-16	Reset	PRBS 31-bit	None	None	0.00 dB (000000)	0.00 dB (000000)	950 mV (11.0000)		Inject	Reset	Reset	Locked	Locked
Found 6	MGT_X1462/MGT_X1462R	25.781 Gbps	1.114E15	0E0	8.977E-16	Reset	PRBS 31-bit	None	None	0.00 dB (000000)	0.00 dB (000000)	950 mV (11.0000)		Inject	Reset	Reset	Locked	Locked
Found 7	MGT_X1411/MGT_X1479R	25.781 Gbps	1.114E15	0E0	8.977E-16	Reset	PRBS 31-bit	None	None	0.00 dB (000000)	0.00 dB (000000)	950 mV (11.0000)		Inject	Reset	Reset	Locked	Locked

**Figure 9.10:** Long runs of 8 25G links using Samtec Fireflies. No errors counted for a BER of  $10^{-16}$ .

### 9.3.2 ZYNQ to FPGA interface

The interface between the FPGA and the ZYNQ PL includes 1 MGT Bank (4 channels), 20 general purpose differential pairs and 1 differential pair dedicated for clock.


The MGT channels are tested using two independent IBERT projects, one on each side. On the FPGA side Bank 128 was included in the project using the Async clock of that SLR side. A similar project was created for the ZU5EG ZYNQ device. The line rate of the link was configured at 12 Gbps. Eye-diagrams of the 4 links as captured in both sides can be seen in Figure 9.11.

The differential pairs were tested by transmitting a counter pattern from the ZYNQ to the FPGA. The pattern was clocked by a free-running clock of the ZYNQ, which was also propagated to the FPGA using the clock pins. On the FPGA side, a logic was written that locked to the incoming counter and then started a local counter on the FPGA side. The value of the received counter was then compared with the value of the local counter at the FPGA side. The test showed that the two values were indifferent for many hours of running the test. This way, the 21 differential traces between the ZYNQ and the FPGA were also validated.


### 9.3.3 Zone 2 connections

The BMTL1 board interfaces the ATCA backplane through the Zone 2 connector. The interface consists of one differential pair connected to one of the inputs of the LHC-JC and one differential pair connected to one MGT channel. The DTH board delivers the LHC 40.078 MHz clock to the relevant Zone 2 pins of every slot inside the crate.

FPGA side



ZYNQ side

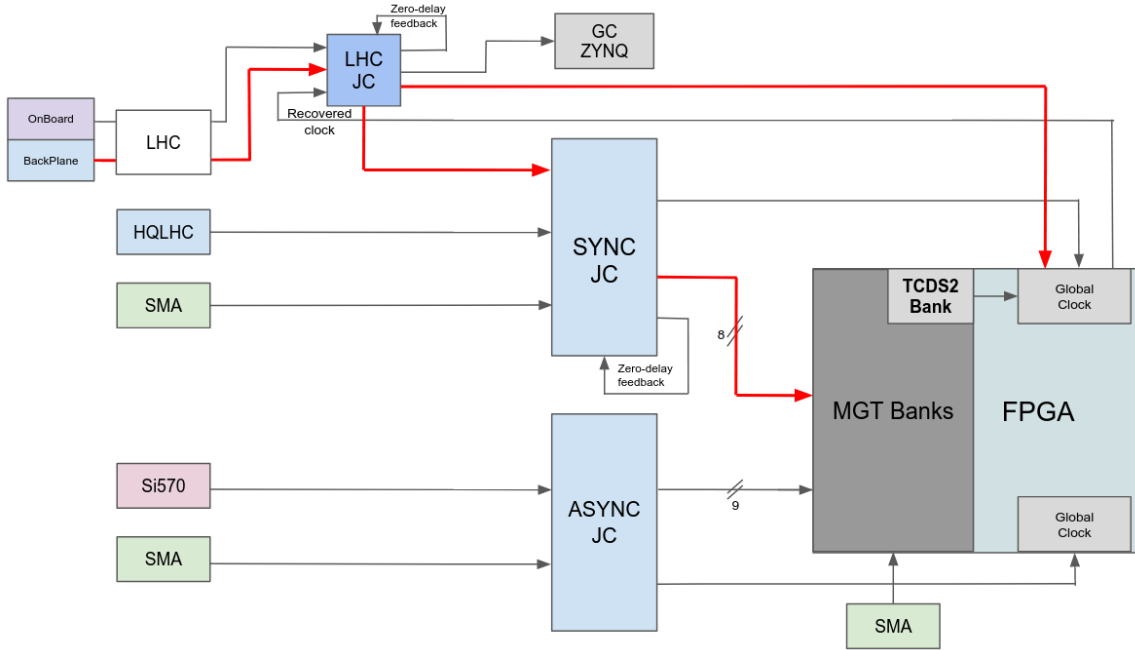


**Figure 9.11:** *Eye diagrams in both directions of the channels connecting the FPGA with the ZYNQ.*

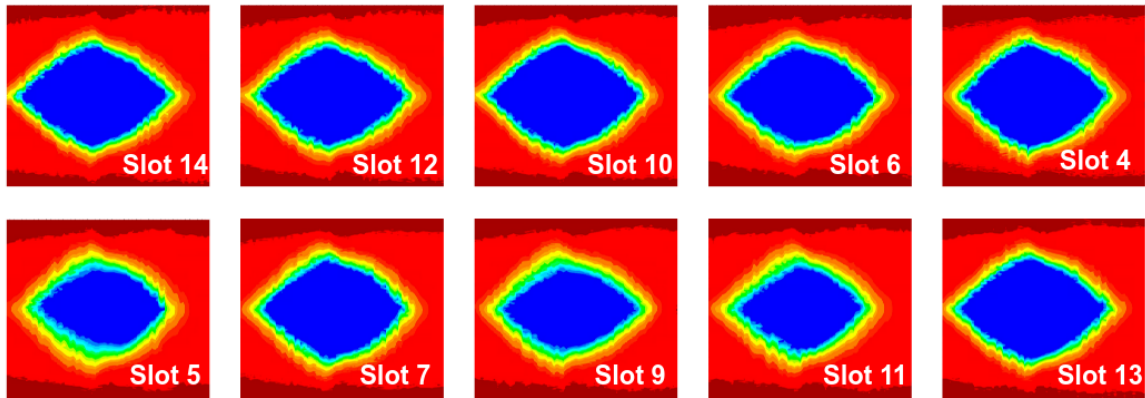
In addition, it delivers the TCDS2 stream through a high speed serial link to every backplane connector, as well.

The first method of testing both interfaces was by using an IBERT project. The custom TCDS2 protocol operates synchronously with respect to the LHC clock that is delivered to the backplane. Thus, the IBERT link can only be established if the FPGA TCDS2 channel uses a reference clock in phase with DTH. To achieve this, the LHC and Sync jitter cleaners had to be configured properly. The routing of the clocking network for this scenario is illustrated in Figure 9.12, where the clocking path is highlighted in red color. The LHC-JC locks to the backplane input and delivers a 40.078 MHz clock both to the FPGA Global Clock (GC) and to the corresponding input of the Sync-JC. The Sync-JC is configured to lock to that input and output reference clock with value  $8 \times 40.078$  MHz, which results to 320.642 MHz. This way, both the FPGA logic and the Sync reference clocks are in phase with the DTH.

Once the clocking network was configured, eye diagrams of the TCDS2 link (Bank 220, channel 1) were captured in the IBERT project. The result is illustrated in Figure 9.13. The open are of these eyes is smaller, since these signals have to travel through copper-copper backplane connections, both from the DTH to the backplane and from the backplane to BMTL1. Using this method, eye scans were taken for every slot of the ATCA crate. This method was used to verify the Zone 2 connections at the BMTL1 board, but also to verify the delivery of the signals from the DTH board to every crate slot. The validation of the actual TCDS2 firmware was concluded with the firmware framework of the board, as described in section 9.4.4.



**Figure 9.12:** Configuration of the clocking network (highlighted in red) to sync to the LHC clock coming from the backplane.



**Figure 9.13:** Eye diagrams of the Zone 2 TCDS2 link taken for every slot of the ATCA crate.

## 9.4 EMP Framework at BMTL1 board

After the operation of the BMTL1 hardware was verified, the next stage involved the firmware that will carry out the operation of this board, as well as that of the subsystem overall. The development of the infrastructure firmware was based on the EMP Framework, in order to profit from the functionalities it provides, as described in section 7.2. Namely, those are firmware and software infrastructure for controlling and monitoring the FPGA using IPbus, all TTC and TCDS2 interfaces to the experiment and a flexibility in using any optical protocol required by L1 Trigger. All these constitute the perfect environment for the placement of the algorithms that operate at BMTL1, with the largest one being the Analytical Method algorithm. Since EMP

is developed around the Serenity boards, it had to be tuned and modified to fit the characteristics of the BMTL1. This process is described in the next sub-sections.

### 9.4.1 BMTL1 In and Out ports

Some of the most important modifications that are required to use the EMP Framework at BMTL1 involve the Input and Output ports. In the BMTL1 case, they contain inputs of the clocking network, the SMA connectors and the chip-to-chip (c2c) differential pairs. They are all listed in the table of Figure 9.14 and are declared in the entity section of the top BMTL1 firmware component. Their declaration required the creation of the corresponding constraint files, that bind every port to the correct pins of the VU13P chip.

Port (diff pairs)	Size	Direction	Source/Target
CLK_100	1	IN	FREERUN_SYNTH
TTC_CLK[0]	1	IN	LHC_JC
TTC_CLK[1]	1	IN	SYNC_JC
TTC_CLK[2]	1	IN	ASYNC_JC
TCDS_REC_CLK	1	OUT	LHC_JC
REFCLK[0-8]	9	IN	ASYNC_JC
REFCLK[9-16]	8	IN	SYNC_JC
SMA_REFCLK	2	IN	SMA CONNECTOR
C2C_CLK	1	INOUT	ZYNQ
C2C	20	INOUT	ZYNQ
SMA_GEN	1	INOUT	SMA CONNECTOR

**Figure 9.14:** List of the In/Out ports of BMTL1 firmware. The MGT channels are not included in the list.

To constraint the MGT transceivers that are used each time by the firmware, EMP uses a dedicated constraint file called *mgt-constraints* and described in section 7.2.7. Transceivers that are not located by this file, such as the ZYNQ Bank and the TCDS2 channel, are locked to their corresponding Banks through a separate constraints file.

A description of every BMTL1 port group is given below.

#### Free-running clocks

BMTL1 uses the *clk\_100* as its free-running clock source. It is produced by a corresponding clock synthesizer that exists on the board and its frequency is 100 MHz. It is routed to the FPGA through to a global clock pin and connects directly to a MMCM chip inside the FPGA logic. This free-running MMCM outputs 4 additional clocks. One at 40 MHz to act as a free-running LHC clock, one at 31.25 MHz and one at 50 MHz used by the IPbus logic, and one at 125 MHz that is used by the TCDS2 and lpGBT firmware.

## TTC clocks

There are three different clock sources that can provide the LHC 40.078 MHz clock to the BMTL1 framework. The first, *ttc\_clk[0]*, is the output of the LHC-JC that can originate from one of the three inputs of this chip: the ATCA backplane, a free-running on-board clock and the recovered clock of the TCDS2 serial stream. The second TTC clock, *ttc\_clk[1]*, is one of the outputs of the Sync-JC. This can either be free-running or can be locked to an external input source connected to the corresponding SMA connector. The third TTC clock, *ttc\_clk[2]*, is one of the outputs of the Asyn-JC. Similar to *ttc\_clk[1]*, it can either be a free-running clock or can be locked to an external source connected to the corresponding SMA. One of these three can be used at a time to generate the LHC multiple clocks of 160, 240, 360 MHz.

## TCDS2 Recovered clock

The *TCDS\_REC\_CLK* port connects the recovered clock of the TCDS2 link to one of the inputs of the LHC-JC.

## Reference clocks

The *refclk* port is an array of clocks that contains the reference clocks used by every MGT channel. For protocols that operate asynchronously with respect to LHC there are 9 sources (*refclk[0-8]*) originating from the Async-JC. Additionally, the remaining clocks (*refclk[9-16]*) deliver the 8 synchronous reference sources that are generated by the Sync-JC.

## SMA reference clocks

The BMTL1 board includes 2 pairs of SMA connectors that are connected directly to MGT Banks. They can be used for any application, if needed.

## Chip-to-chip clock

The ZYNQ to FPGA interface includes one differential pair connected between global clock pins of the two devices. It can be used to deliver clock either from the ZYNQ to the FPGA or vice versa.

## Chip-to-chip pins

The ZYNQ to FPGA interface consists of 20 pairs of differential pins that can be used to transfer data between the two chips for any application.

## SMA general purpose

The board contains two SMA connectors that can be used for any application. They are connected to global clock pins and can either input external clock sources to the firmware logic or output clocks to a different system.

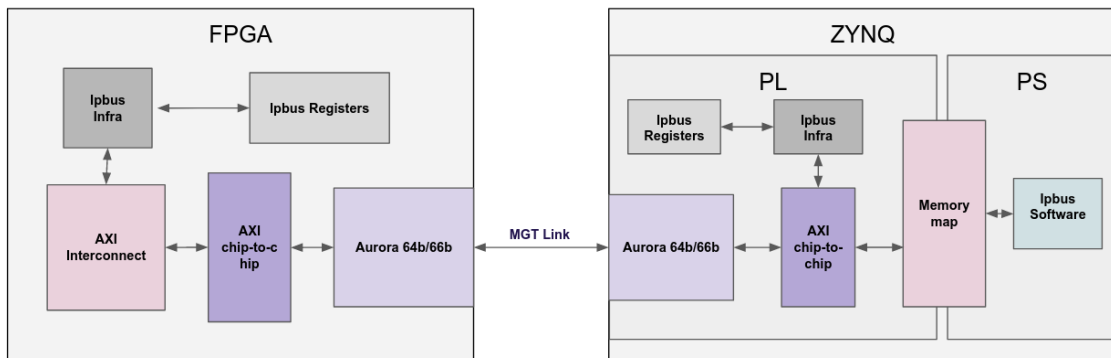


## 9.4.2 FPGA to ZYNQ IPbus interface

The BMTL1 firmware is controlled by software that runs on the processing system of the ZYNQ. As described in 7.2.2, the EMP Framework utilizes the IPbus protocol to facilitate data transfer between a software application and the firmware logic. Its high level implementation is inherent in the framework but the physical layer communication is specific and depends on the interface between the FPGA and the software host.

In the BMTL1 board, the interface is performed through a high speed link. One out of the four available links that connect the FPGA with the ZYNQ is used for this purpose. An IPbus example implementation of a similar circuitry is provided at the official IPbus repository [42]. This example was used as guideline to implement the firmware logic, both in the FPGA side and the ZYNQ side.

The IPbus low level communication at BMTL1 is performed using the AXI interface. The AXI specifies a point-to-point protocol for interfaces that involve a master and a slave. In our case, the ZYNQ acts as the master that issues read or write commands to the FPGA, the slave. Thus, in both ends the IPbus words are converted into AXI bus. The interface between the two devices is operated by the AXI chip-to-chip IP that acts as a bridge for AXI transactions [43]. The data transfer through the serial link is handled by a different IP core, the Aurora 64b/66b [44]. This IP implements the serial Aurora 64b/66b protocol to establish a link using one MGT transceiver in each side. The parallel bus interface of the Aurora IP are configured to AXI and hence, it directly connects to the chip-to-chip IP to handle the AXI physical communication. A simple schematic of the IPbus circuitry in the BMTL1 board is shown in Figure 9.15. The IPbus Infra block implements the translation of IPbus packets to AXI, and vice versa. Finally, both read and write operations are directed to the corresponding IPbus register, defined and placed in different places of the firmware logic by the user.



**Figure 9.15:** Schematic of the VU13P to ZYNQ IPbus implementation through AXI.

The IPbus over AXI implementation is similar on both the VU13P FPGA and on the Programmable Logic (PL) of the ZYNQ. The interface between the processor and the AXI chip-to-chip is performed through memory mapped addresses. The IPbus software that runs on the processor system initiates transactions to the memory addresses that are specified by the interconnect between the PS and the PL inside the

ZYNQ. Each of these addresses connects to IPbus slaves that the user can interact with.

### 9.4.3 Declaration Files

The *declaration files* of EMP had to be adapted to the BMTL1 board as well. One of them is called *device-declaration*. It declares generic constants of the firmware project, defines the kind of MGTs that are supported by every Bank and facilitates the reference clock distribution to optical protocols of these Regions. The device-declaration file of BMTL1 is shown in Figure 9.16. Line 19 of the VHDL file defines the number of regions that are supported by the FPGA. This number is usually equal to the number of MGT Banks, hence, in the VU13P case it is set to 32. The line below declares the number of reference clocks supported by the device, which in this case is 17. The *IO\_REGION\_SPEC* array includes 32 rows, each corresponding to one Region block. In every row, the setting of the first column can be one of the *no\_mgt*, *io\_gth* or *io\_gty*, declaring whether an optical protocol can be implemented in that region and if yes, the supported MGT kind. For example, Region 0 cannot implement an MGT protocol since it corresponds to the Bank of the Zone 2 connections. The second and third columns connect the asynchronous and synchronous reference clocks, respectively. The numbers are referring to the corresponding vector of the *refclk* port, as discussed in 9.4.1. Clocks 0 to 8 are generated by the Async-JC and clocks 9 to 16 by the Sync-JC. As can be seen, all Regions-Banks are provided with both types of reference clocks that are able to be used by both kinds of optical protocols.

```

18
19 constant N_REGION      : integer := 32;
20 constant N_REFCLK     : integer := 17;
21 constant CROSS_REGION : integer := 15;
22
23 constant IO_REGION_SPEC : io_region_spec_array_t(0 to N_REGION - 1) := (
24   0 => (io_nogt, -1, -1),      -- Bank 220 -- Right column -- Zone 2
25   1 => (io_gty, 9, 9),       -- Bank 221
26   2 => (io_gty, 9, 9),       -- Bank 222
27   3 => (io_gty, 9, 9),       -- Bank 223
28   4 => (io_gty, 1, 10),      -- Bank 224
29   5 => (io_gty, 1, 10),      -- Bank 225
30   6 => (io_gty, 1, 10),      -- Bank 226
31   7 => (io_gty, 1, 10),      -- Bank 227
32   8 => (io_gty, 2, 11),      -- Bank 228
33   9 => (io_gty, 2, 11),      -- Bank 229
34  10 => (io_gty, 3, 11),      -- Bank 230
35  11 => (io_gty, 3, 11),      -- Bank 231
36  12 => (io_gty, 4, 4),       -- Bank 232
37  13 => (io_gty, 4, 4),       -- Bank 233
38  14 => (io_gty, 4, 4),       -- Bank 234
39  15 => (io_gty, 4, 4),       -- Bank 235
40  -- cross chip
41  16 => (io_gty, 5, 13),      -- Bank 135 -- Left column
42  17 => (io_gty, 5, 13),      -- Bank 134
43  18 => (io_gty, 5, 13),      -- Bank 133
44  19 => (io_gty, 5, 13),      -- Bank 132
45  20 => (io_gty, 6, 14),      -- Bank 131
46  21 => (io_gty, 6, 14),      -- Bank 130
47  22 => (io_gty, 6, 14),      -- Bank 129
48  23 => (io_nogt, -1, -1),    -- Bank 128 -- ZYNQ
49  24 => (io_gty, 7, 15),      -- Bank 127
50  25 => (io_gty, 7, 15),      -- Bank 126
51  26 => (io_gty, 7, 15),      -- Bank 125
52  27 => (io_gty, 7, 15),      -- Bank 124
53  28 => (io_gty, 8, 16),      -- Bank 123
54  29 => (io_gty, 8, 16),      -- Bank 122
55  30 => (io_gty, 8, 16),      -- Bank 121
56  31 => (io_gty, 8, 16),      -- Bank 120
57  others => kIONoGtRegion
58 );

```

**Figure 9.16:** The device declaration file of the BMTL1 framework.

The second declaration file is called *project-declaration* and is also described in

7.2.6. It defines the protocol kind that will be implemented in every region of the project. The user can select whether or not an MGT is going to be implemented (*no\_mgt*), or select one of the supported protocols, named as *gty16*, *gty25*, *gbt* and *lpGBT*. All actions that are executed by the framework to implement the selected configuration are invisible. The user, however, can interface with data of the enabled Regions inside the Payload block.

An example of a declaration-file targeting the BMTL1 board is shown in Figure 9.17. The project of this file will implement GTY CSP links running at 25 Gbps in Regions 7 to 10, receiving lpGBT links at Regions 18-20 and GBT links at regions 26 to 29.

```

38  constant REGION_CONF : region_conf_array_t := {
39  0    => kDummyRegion, -- Zone 2 -- Bank 220
40  1    => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 221
41  2    => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 222
42  3    => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 223
43  4    => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 224
44  5    => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 225
45  6    => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 226
46  7    => {gty25, buf, no_fmt, buf, gty25}, -- Bank 227
47  8    => {gty25, buf, no_fmt, buf, gty25}, -- Bank 228
48  9    => {gty25, buf, no_fmt, buf, gty25}, -- Bank 229
49  10   => {gty25, buf, no_fmt, buf, gty25}, -- Bank 230
50  11   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 231
51  12   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 232
52  13   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 233
53  14   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 234
54  15   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 235
55  -- Cross-chip
56  16   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 135
57  17   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 134
58  18   => {lpGBT, buf, no_fmt, buf, no_mgt}, -- Bank 133
59  19   => {lpGBT, buf, no_fmt, buf, no_mgt}, -- Bank 132
60  20   => {lpGBT, buf, no_fmt, buf, no_mgt}, -- Bank 131
61  21   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 130
62  22   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 129
63  23   => kDummyRegion, -- ZYNQ -- Bank 128
64  24   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 127
65  25   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 126
66  26   => {gbt, buf, no_fmt, buf, gbt}, -- Bank 125
67  27   => {gbt, buf, no_fmt, buf, gbt}, -- Bank 124
68  28   => {gbt, buf, no_fmt, buf, gbt}, -- Bank 123
69  29   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 122
70  30   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 121
71  31   => {no_mgt, buf, no_fmt, buf, no_mgt}, -- Bank 120
72  others => kDummyRegion
73  };
74  ;

```

Figure 9.17: The project declaration file of a BMTL1 example project.

## 9.4.4 TCDS2 interface

The firmware logic of the TCDS2 interface is included by default in the EMP Framework. The circuitry implements the custom TCDS2 protocol, built around an MGT channel. The TTC and TTS information is delivered by the DTH board to every ATCA blade inside a crate through the payload of this link. At the current development stage, the main signal that is delivered to the framework is the BC0.

The TCDS2 firmware in the BMTL1 board had to be located at MGT Bank 220 and channel 1. Furthermore, the LHC clock that arrives from the LHC-JC had to be used as the reference clock for this link. The valid operation of the firmware was demonstrated using the EMP-butler, and the corresponding reset command. The *reset TCDS2* option would configure the framework to use the backplane 40.078 MHz clock and use the BC0 that arrives from the TCDS2 stream to lock the corresponding bunch counters. The result of running this command at the BMTL1 board is shown

in Figure 9.18. The output prints the status of the steps that are executed until the completion of the command. Two statements that inform about its successful completion are the *Clock 40 locked after 0 ms* and the *TTC BC0 locked*. Furthermore, status registers are printed, providing information such as the frequency of the clock and measurement of received BC0s, as well as potential missing or invalid receptions of the BC0 signal.

```
[root@bmtl1-2 fpga-ctrl]#
[root@bmtl1-2 fpga-ctrl]# empbutler -c connections.xml do bmtl1_fpga reset tcds2
05-12-23 09:51:30.811488 [281473781272576] WARNING - Address overlaps observed - report file writ
05-12-23 09:51:30.817323 [281473781272576] NOTICE - mmap client with URI "ipbusmmap-2.0:///dev/me
Resetting device 'bmtl1_fpga'
Changing clock and TTC source to: TCDS2
Clock 40 locked after 0 ms
TTC BC0 locked
Global BC0 locked

STATUS
  Clock source: TCDS2
  TTC source: TCDS2
  Internal BC0: False

  Clock locked: True
  Frequency: 40.079 MHz

  BC0 locked: True
  Dist locked: True

  Bunch count: 0x00089c
  Orbit count: 0x001606
  Event count: 0x000000
  BC0s recvd: 5639 valid, 0 missing, 0 invalid


External interface (rx, tx):
  Mode: Normal
  Power good: True
  PLL lock: True, True
  Ref clock: 320.628 MHz
  Reset done: True, True
  Ready: True, True
  Frame lock: True
  Unlock count: 0
[root@bmtl1-2 fpga-ctrl]#
```

**Figure 9.18:** *Output of the emp-butler reset tcds2 command.*

### 9.4.5 QPLL based lpGBT

The firmware blocks of the FPGA version of the lpGBT protocol are provided by the corresponding group. The group also provides an example project that was used as a reference by the people who ported the lpGBT links in the EMP Framework. This example implements the protocol logic based on a single MGT channel that utilizes the Channel PLL (CPLL) of the MGT Quad. When multiple channels are used, identical copies of it are generated. While this firmware is used by other EMP (Serenity) users without issues, this CPLL based implementation cannot be used by FPGAs that are of speed grade 1. This case is also true for the BMTL1 device. When the CPLL is used to drive the MGT reference clock, the maximum supported line rate is lower than the QPLL one. Specifically, in speed grade 1 Virtex Ultrascale Plus devices, the maximum line rate a CPLL can support is 8.5 Gbps [45]. The corresponding line rate when QPLL is used is 12.5 Gbps. Thus, by using this lpGBT version the BMTL1 cannot operate the lpGBT protocol at line rate of 10.24 Gbps. For this reason, the lpGBT firmware at EMP was modified to use QPLL instead of CPLL.

A simple schematic of the CPLL lpGBT implementation is shown on the left of Figure 9.19. The MGT IP core is configured to use 1 channel and its corresponding CPLL. The firmware code, consisting of the lpGBT Uplink and Downlink blocks, is structured around it. The three of them form one lpGBT channel and identical versions of this code are generated 4 times to create an lpGBT Quad.



**Figure 9.19:** Left schematic illustrates the the block diagram of the single channel CPLL based implementation of the lpGBT firmware. The four channel QPLL based implementation is shown on the right.

By using a QPLL, the MGT IP is configured as a Quad. This results to the conversion of some of its internal signals and external ports from channel based to Quad based. Some examples are the resets of the Tx and Rx datapaths and the Tx and Rx *reset-done* signals. For this reason, the single-channel structure could no longer be used and it was modified, as shown on the right of Figure 9.19. In the new, QPLL-based architecture, the MGT IP core is configured to use a QPLL and instantiate all four channels of the Quad. Thus, its code is implemented once and it had to be moved outside the *Channel* block. The new channel component that wraps the lpGBT Uplink and Downlink blocks is modified to match the new structure. The internal lpGBT files have remained untouched. Moreover, the VHDL files of the QPLL based lpGBT wrapper are modified so that they connect the corresponding ports of the channel wrapper to the QPLL MGT block.

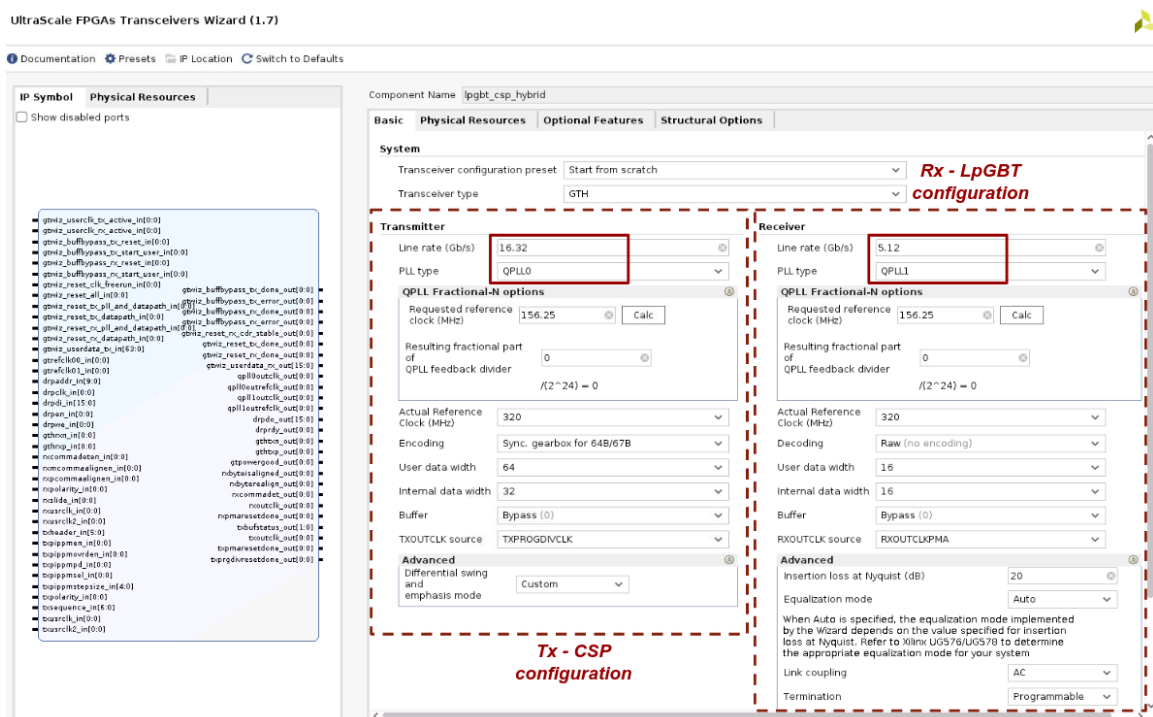
The new lpGBT structure was tested and its operation validated. The test was performed using the BMTL1 board and one OBDT-theta, connecting four lpGBT channels. Since then, the QPLL lpGBT is the default firmware used by the BMTL1 board at its integration tests and will also replace the CPLL version at the original EMP Framework.

#### 9.4.6 lpGBT In CSP Out Hybrid links

The architecture of the BMT Layer-1 is such that different protocols are used to interface the detector and to transmit data to GMT. The result of this asymmetry is that the transmitter of channels that receive lpGBT data remains unused and, similarly, the receiver of channels that transmit CSP data also remains unused. This

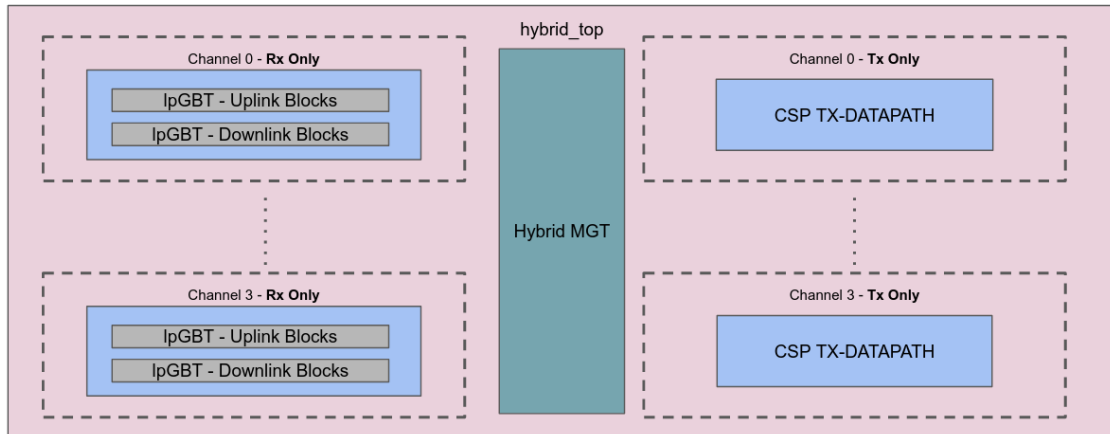
scenario is possible to affect the final BMTL1 architecture, but it can also be met in other subsystems that receive large number of detector links. For this reason, it was decided that the development of a hybrid link version in EMP will benefit many of the framework users.

The link version described here is called *hybrid*, due to the fact that it receives lpGBT data and transmits CSP data at the same MGT channel. This architecture is able to be implemented by profiting from the ability of the MGT Quad to use both of its QPLLs at the same time. The configuration of a Hybrid IP core is depicted in Figure 9.20. The left part of the screen contains the transmitter settings and the right part the receiver settings. The transmitter is configured to use QPLL0 with the corresponding CSP settings and the receiver to use QPLL1 with the lpGBT settings. At this case, the Tx line rate is set to 16.32 Gbps and the Rx line rate at 5.12 Gbps. The line rates can be of any Tx and Rx combination, but different instances of the IP core are required for each of them.



**Figure 9.20:** Configuration of the Hybrid MGT IP core. In this example, the transmitter uses QPLL0 and is configured with the CSP settings at 16.32 Gbps. The receiver uses QPLL1 with the lpGBT settings at 5.12 Gbps.

The Hybrid IP instance is the central block of the Hybrid firmware. Around it are placed the lpGBT and CSP firmware blocks. A simple block diagram of this structure is shown in Figure 9.21. On the left are depicted the lpGBT blocks, that follow the QPLL based architecture. Only the Rx Uplink datapath blocks are used and are connected to the corresponding ports of the IP core. On the right is the transmitter datapath blocks of the CSP firmware. They are as well connected to the Tx ports of the Hybrid IP. The two QPLLs are sourced by two separate clock sources. One of the asynchronous reference clocks of the board is connected to QPLL0 and one of the synchronous to QPLL1.



**Figure 9.21:** Block diagram of the Hybrid link, including the receiver *lpGBT* blocks, the transmitter CSP blocks and a hybrid MGT instantiation.

The Hybrid links firmware was demonstrated in SX5 using the BMTL1 ATCA board. For the input links, an OBDT-theta board was used to connect *lpGBT* input links to BMTL1. Outputs from the same Quad were connected to an Ocean board that run the CSP protocol. The test showed that both directions of the firmware were operating successfully. The BMTL1 links would lock to *lpGBT* links and receive the expected data. In addition, the Ocean CSP input links were locked and received data without errors.

#### 9.4.7 Constraint files

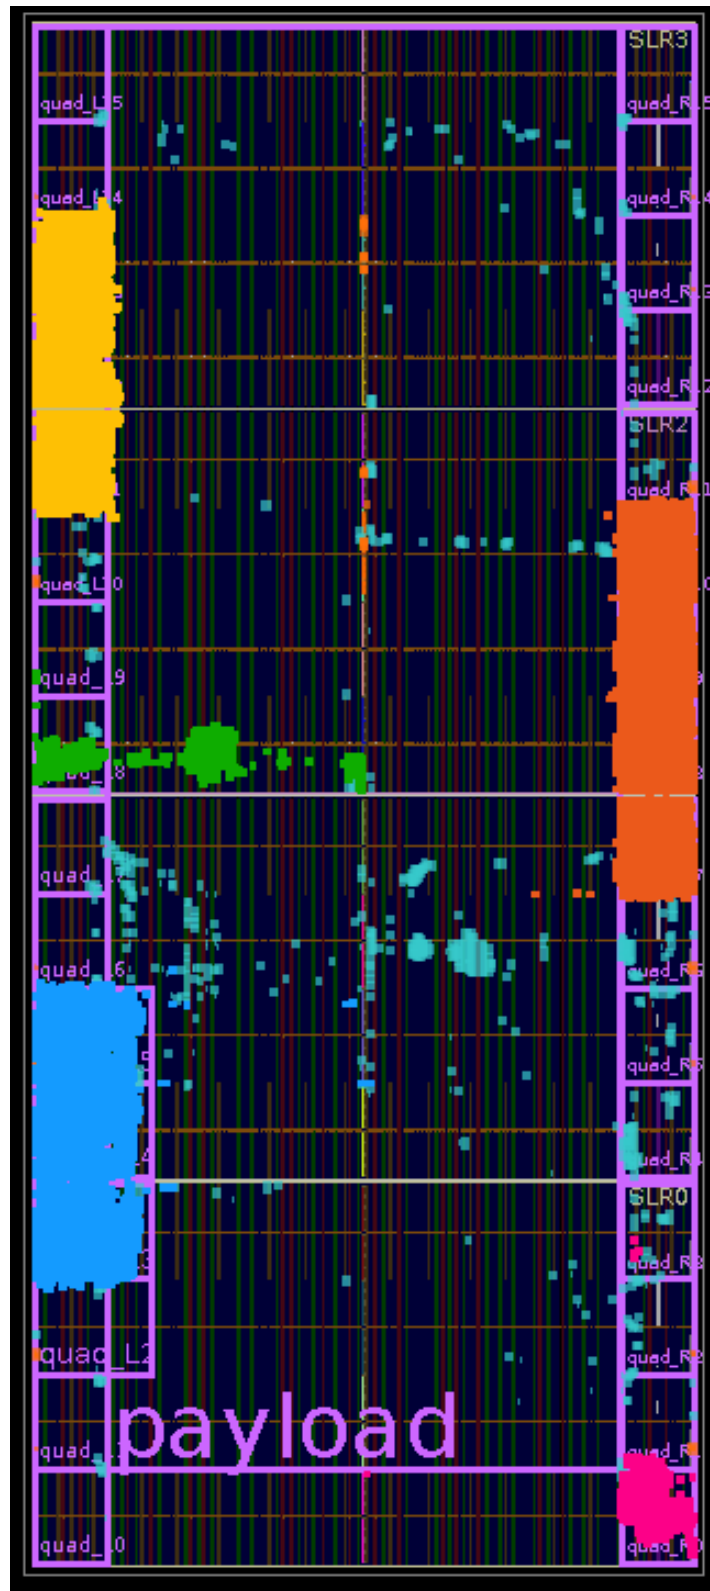
The constraint files of the EMP Framework are described in section 7.2.7. Some of them are board or project specific and others are completely agnostic. Those that correspond to the former category are replaced in the BMTL1 by new, matching the characteristic of this board. For example, files that connect the top ports of the project are replaced by the ones that correspond to the I/O ports listed in Figure 9.14. On the other hand, the file that constraints the MGT channels has remained untouched. Other files that were modified include those used to declare the clocks that are created and used by the framework, as well as those that manipulate the floorplaning of the chip to create the *pblocks*.

#### 9.4.8 Building the project

The BMTL1 firmware files are stored in an online repository. This is the standard method of storing files for every project in the experiment, as it assists code sharing between people and development in common projects. The BMTL1 repository structure follows that of the EMP repository. Main reason is the usage of the same software tool to include all files and build the Vivado project automatically. This tool, called *IPBB* (IPBus Builder), is a command-line firmware management and build tool [46]. It expects a specific organization of the files in the repository and defines the structure and names of the basic folders in it. The files of a project are

defined inside declaration files that are managed by the firmware maintainer. When a project is built using IPBB, the software creates a Vivado project, configures it to the user settings and adds the declared files to it. The user can then either develop on top of that version of the firmware, or use it to generate a bit file. The implemented design of a BMTL1 project using the project-declaration file of Figure 9.17 is illustrated in Figure 9.22.





**Figure 9.22:** Floorplanning of a BMTL1 project using the VU13P FPGA. This project instantiates 3 lpGBT regions (yellow), 3 GBT regions (blue) and 4 CSP regions (orange). The TCDS2 firmware is highlighted in purple (bottom right) and the IPbus infrastructure in green.

## 9.5 Analytical Method algorithm Integration

The Analytical Method (AM) algorithm is the main processing unit of the BMTL1 subsystem. It receives TDC hits from DT chambers of the CMS barrel and generates trigger primitives in the form of muon track segments, as described in section 6.3.1. The algorithmic logic is developed by the DT group. The BMTL1 group develops the hardware and firmware framework of the subsystem. In addition, it is responsible for the algorithm integration to the corresponding infrastructure.

### 9.5.1 Single-Chamber Integration

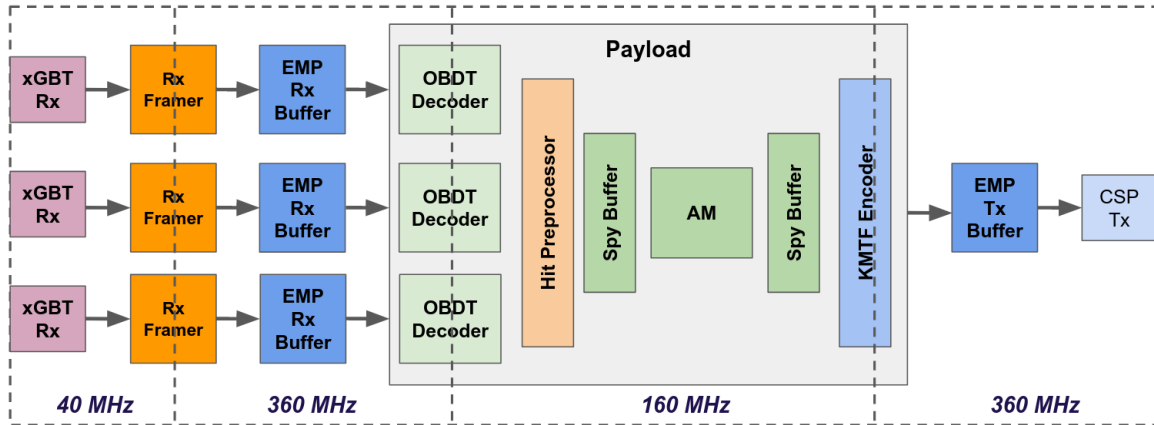
The AM code contains a set of VHDL files that are stored online and maintained by the DT group. The BMTL1 framework, stored in a different repository, calls the AM files and includes them during the building of the project. All AM related logic is placed inside the Payload block. The interface between the framework and the algorithm is facilitated by a corresponding wrapper module, written specifically for this task.

The structure of the algorithm is chamber-oriented, since muon stubs at this processing layer are generated for every DT station independently. Its inputs are DT TDC hits, received directly from the detector through the OBDT boards. As discussed in section 6.5, in most cases one OBDT is dedicated to process DT cells from one Super-Layer (SL). Thus, three OBDT boards cover the two  $\phi$ -SLs and the one  $\theta$ -SL at MB1 and MB2. At MB3, due to its larger size, 4 OBDTs are required in total. The same applies for MB4. Even though it does not contain a  $\theta$ -SL, two OBDTs are needed for each of the two  $\phi$ -SLs. At this stage, the current algorithm implementation expects one optical link per SL. In the current algorithm implementation, processing of the  $\theta$ -view is not yet fully supported, even though its connections are taken into account in the firmware logic.

Reception of TDC data is facilitated using front-end protocols. Even though the final OBDT version is instrumented with the lpGBT ASIC, its first prototype uses the Phase-1 GBTX chip. Hence, the BMTL1 firmware infrastructure supports both optical protocols for data reception. Following the GBT frame structure, on every bunch crossing (BX) the OBDT transmits a payload of 84-bits, consisting of up to 3 muon hits. The size of every hit (described in 6.5), is 25-bits and contains the TDC hit value, cell number and BX number.

There are three main clock domains that data follow inside the BMTL1 framework. They are the LHC 40 MHz clock (40.078), the framework clock at 360 MHz and the algorithm clock at 160 MHz. All three are generated from the same source and are integer multiples of it. Thus, data crossing between domains can be performed simply using flip-flop registers. A block diagram of the processing path inside BMTL1 is illustrated in Figure 9.23. The *xGBT Rx* receives GBT or lpGBT frames at the LHC clock. In the current implementations, frames that arrive from both OBDT versions have a size of 84 bits (in the lpGBT case the remaining MSBs are not used). These frames have to cross to the framework domain, that supports 64-bit frames clocked at 360 MHz, in order to reach the Payload block. The transition is implemented inside the *Rx Framer*. The 64 LSB of an xGBT frame are delivered to the first 64-bit frame of the 360 MHz clock, the 20 MSB are delivered to the second 64-bit frame and the

remaining 7 frames of the 360 MHz clock are filled with zeros. The same framing method is applied for both GBT and lpGBT data.




**Figure 9.23:** Block diagram of the processing path inside BMTL1. The clock domains of every block are also depicted.

The EMP channel buffers operate using the framework clock domain. The Rx buffer is located right after the receiving xGBT block and the Tx Buffer right before the transmitting CSP block. Their usage is optional and both of them can direct data to/from either the payload module or the links. When used to test the Payload block, the user can load the Rx Buffer with test patterns, play them through the algorithm and read the result in the Tx Buffer. This way, they provide a realistic mean of testing the algorithm operation, as well as the connections to and from it. EMP buffers are bypassed when not configured.

The basic algorithmic logic exists inside the *AM* component. The firmware blocks around it are preparing the data to be consumed by the algorithm, and also prepare the result for its transmission. The *OBDT Decoder* receives TDC hits on the 360 MHz domain, crosses them to the algorithm domain and performs a minor decoding of the OBDT data. The current version of the AM algorithm (v1) operates at 160 MHz clock. The new version (v2), currently being finalized, will operate to frequency up to 400 MHz. After the OBDT decoder there is another preparation stage, performed by the *Hit Preprocessor*. The output of this block is then routed to the AM block, where TDC hits from SL1 and SL3 are used for the generation of DT primitives. The result at the output of the AM block is delivered to the *KMTE Encoder*. Inside it, TP data are encoded to their final format and cross to the 360 MHz clock in order to be transmitted. Right before the AM block there is an additional buffer logic, called *Spy Buffers*. Their operation is similar to that of the EMP Buffers. However, they operate at the 160 MHz clock domain.

The floorplanning of an implemented BMTL1 firmware is illustrated in Figure 9.24. The project contains one instance of the AM algorithm, contained in one SLR (SLR 1) of the VU13P FPGA. A GBT Quad is implemented to receive data and a CSP Quad is implemented to transmit data.



**Figure 9.24:** Floorplanning of a BMTL1 implemented design. The project includes 1 instance of the AM algorithm (1 Chamber), one GBT quad and one CSP quad.

### Configuration and validation

The Payload block contains a number of IPbus registers that control and monitor the operation of the algorithm. Moreover, they control the operation of the Spy Buffers. All AM related operations are facilitated by the software of the algorithm, called *DTU*. It executes different commands that configure the registers of the algorithm and also write or read and print the content of the Spy Buffers. This can be either the TDC hits, in their decoded form, or the generated TPs. Thus *DTU* is extensively used, both for the configuration of the algorithm but also for debugging and validation purposes.

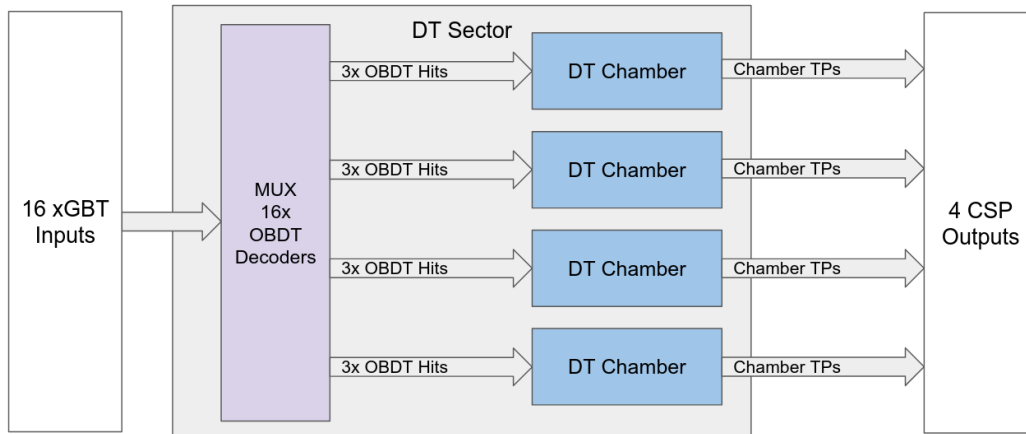
The first step of validating the AM operation inside the BMTL1 framework was done using Spy Buffers. Simulation data were used, played in the algorithm logic and the TPs were read out by the buffers at the output. The result was compared to the expected and no miss-matches were observed in the generated stubs.

Next step was similar, but this time using the EMP buffers. This way, larger part of the data path inside the firmware was tested. In addition, all clock domain crossings were included. A generated muon pattern was loaded to the Rx EMP buffer, in the form that the Rx Framer would output its data. The pattern was played to the Payload block and the result was read at the Tx Buffer. Again, the generated TPs were exactly equal to the expected ones.

The tests described above are single-board tests. They were followed by simple multi-board tests. An OBDT board was used for this purpose. The same muon pattern was loaded to corresponding Tx buffers at the OBDT firmware and transmitted through the GBT link. The pattern was received at the BMTL1 board and the result of its processing by the AM algorithm was read at the Tx Buffer. The TPs had again the same value as the expected ones. By performing such tests, the operation of the AM algorithm inside the BMTL1 board was considered valid.

## 9.5.2 Sector Implementation

The next step after integrating and validating the single-Chamber firmware was the implementation of a Sector-based firmware. Towards this version of the BMTL1 Payload, all components that relate to one chamber, apart from the OBDT-decoder, were placed inside a new block, called *dt-chamber*. This new block is now generated four times in order to create a Sector, with each chamber instance being identical to the other. This logic is now included into another new block, called *dt-sector*. A block diagram of the sector firmware is depicted in Figure 9.25.





**Figure 9.25:** Block diagram of a *dt-sector* block.

The BMTL1 architecture at this stage expects one xGBT link to be connected from every OBDT board. In addition, the maximum number of OBDTs in one Sector is going to be 16, as discussed in 6.5. Thus, the maximum number of xGBT links for one DT Sector in the BMTL1 firmware is also 16, facilitated by 4 MGT Quads. These links are directly connected to the OBDT decoders. However, in order to allow flexibility in the mapping of links to the final *dt-chamber*, the inputs to these 16 OBDT decoders are multiplexed. The structure can be seen in Figure 9.25. All 16 inputs are delivered to the multiplexer. Then, the DTU software maps the outputs, that can be up to 3 OBDT Decoder outputs, to every *dt-chamber*.

The output of every *dt-chamber* is 64-bit TP frames. These are going to be transmitted to the Barrel Filter layer. However, the final architecture of this interface is not finalized yet. For this reason, at this development stage we assume no time-multiplexed architecture and direct the outputs of every chamber to one CSP transmitting link. The maximum number of TPs generated per bunch crossing by AM is four. The available number of CSP frames at 25 Gbps per one link per bunch crossing is 9. Thus, one link per chamber is more than capable of transmitting the generated data.



The implemented design of a Sector BMTL1 project is illustrated on the right of Figure 9.26. There are four instances of the AM algorithm (orange, yellow, green and purple) that have been placed in three SLRs. One instance in SLR 0, two in SLR 1 and 1 in SLR 2. In addition, there are implemented 4 GBT Quads for the inputs and 1 CSP Quad for the outputs. The resource utilization of this project is depicted on the bottom left of Figure 9.26. It consumes 20 % of the available LookUp Tables, 9 % of the Flip Flops and 34 % of the chip's BRAM.



**Figure 9.26:** *Left: Resource utilization of a Sector BMTL1 project. Right: Floorplanning of an implemented BMTL1 design consisting of one DT sector.*

### Implementation of two Sectors

In the final system, every BMTL1 board will process data from two DT Sectors. Thus, a firmware that implements the AM algorithm for two Sectors, or 8 chambers, has been demonstrated. The floorplanning of the project can be seen on the right of Figure 9.27. The logic of the first Sector is placed in SLRs 0 and 1, while the logic for the second Sector in SLRs 2 and 3. Furthermore, the inputs of the first Sector are 16 GBT links, while the inputs of the second Sector are 16 lpGBT links. There are also 2 CSP quads to handle the output of each Sector. The resource utilization of this project is depicted in the bottom left of Figure 9.27. The logical resources used by this project are: 40 % of Look-Up Tables, 18 % Flip-Flops, 61 % BRAM.



**Figure 9.27:** Floorplanning of an implemented BMTL1 design consisting of two DT sectors.

# Chapter 10

## Barrel Muon Trigger Slice Tests

The integration and validation of the Analytical Method firmware inside the BMTL1 framework was followed by multi-board tests. The conduction of these tests started at the surface area (SX5) of the CMS experiment at the interaction point 5 (P5). The Drift Tube groups maintained a lab at SX5 (which has moved to SXA5) and it is there where the integration tests took place. To assist testing of Phase-2 hardware, an operational DT chamber was placed in the area. It was installed parallel with the ground in order to detect hits from cosmic muons that crossed its surface. The signals from the hits were then connected to OBDT boards and were used to validate their operation. In addition, the setup used to include Phase-1 hardware, such as a uTCA crate and corresponding blades, as well as a Phase-2 ATCA crate. The ATCA was used for the first tests of the BMTL1 board inside a real crate. It also hosted an Ocean board, provided by the GMT group, and a DTH board.

### 10.1 SX5 Single Chamber Slice Test

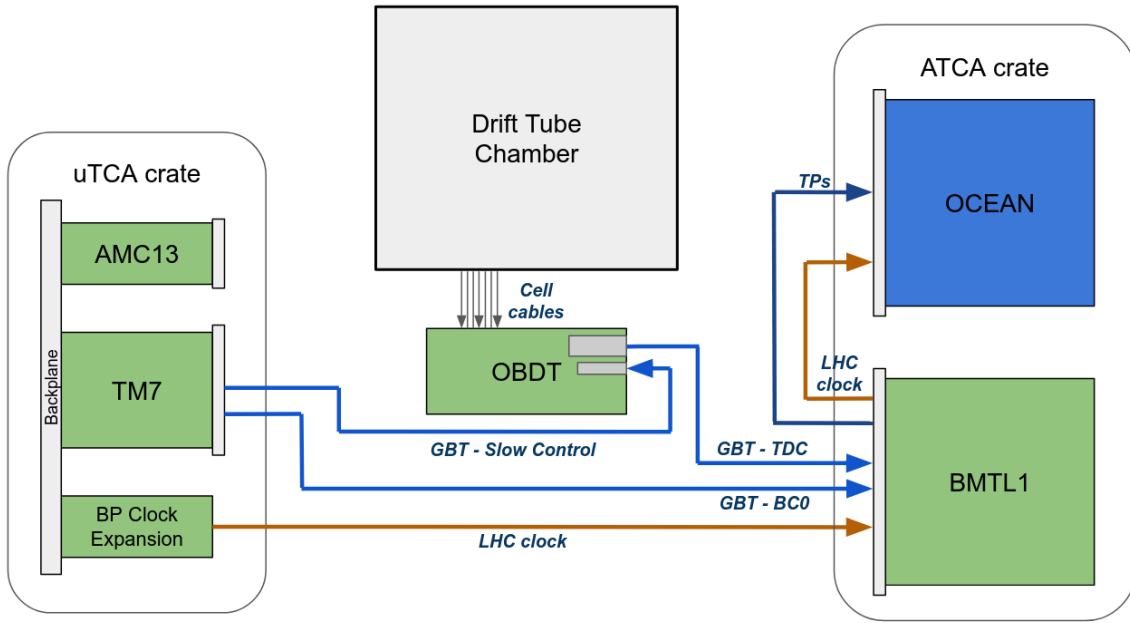
The setup at the DT integration area was used to carry out the first tests of the BMT system. The tests included only Phase-2 hardware, that is an OBDTv1 board, the BMTL1 ATCA and the Ocean card. In addition, the DT chamber was used to generate real cosmic muon data.

#### 10.1.1 Setup description

A block diagram of the setup can be seen in Figure 10.1. Cables from DT cells of the chamber were connected to the OBDT, transmitting pulses of muons that cross the gas volume. The OBDT converted them to TDC hits and transmitted them to the BMTL1 board using the GBT protocol. Data were transferred via one optical fiber, connected to the QSFP module of the OBDT and to an Rx12 Firefly of the BMTL1 board using a patch panel. Moreover, optical fibers would connect one TxRx Firefly of the BMTL1 to another TxRx Firefly module at the Ocean card, to transmit data using the CSP protocol at 16 Gbps.

Valid data transfer between all boards require that they are all synchronized with each other. Furthermore, a BC0 tag should be present in order to synchronize the reception of frames. These two were facilitated using Phase-1 hardware. Firstly, an





**Figure 10.1:** Block diagram of the SX5 setup.

AMC13 card (distributes TCDS during Phase-1) which was installed inside the uTCA crate generated a local LHC clock and distributed it to the backplane of the crate. This clock was extracted using a custom PCB that exposes the relevant backplane pins to SMA connectors. From there, the clock was directed both to the OBDT board and to the BMTL1 board. To Ocean received the LHC clock from the BMTL1, which outputs the 40.078 MHz clock to one of the general purpose SMA connectors. The uTCA crate also contained a Phase-1 blade, called TwinMux (TM7). This board is used for the slow control of the OBDT board. This way, it delivers both the LHC clock and the BC0 information through the Phase-1 TTC stream. This information, however, had to also be propagated to BMTL1 and Ocean. This was done using an additional GBT link from the TM7 board to the BMTL1. This link was only used to transmit the BC0 signal, once every 3564 clock cycles of the LHC clock (one orbit). From BMTL1 the BC0 information was propagated to Ocean using the *OT/BC0* functionality of the CSP protocol. This way, all three boards were synchronous with each other.


Hit reception in the BMTL1 board can be observed using the DTU software. The software runs on the ZYNQ device and reads the input spy buffers of the AM firmware. The output at the terminal is a list of hits that have been received, their TDC time and BX number. This software is also responsible for configuring the algorithm and can be used to print a list of Trigger Primitives that have been generated by it. Thus, it provides a first check that hit reception and TP generation is operating.

The generated data are transmitted to the Ocean board using one CSP link at 16 Gbps. The link also transmits the BC0 information. Since only one chamber was used, reconstruction of muon tracks could not be performed at Ocean. However, the board here is used in order to validate TP reception at the GMT system. In addition, and most importantly, it performs readout of large amounts of data for long periods of time. This is achieved by a local readout system that has been implemented in the ZYNQ device that is hosted on the board. The system utilizes an AXI interface

configured for Direct Memory Access between the FPGA and the processor system. Data are received, transferred to a DDR4 memory and then written to the SSD of the board. From there one can access their binary form and process them using high-level software tools.

### 10.1.2 Validation of optical interfaces

The setup of the slice test at SX5 is illustrated at Figure 10.2. The left image shows the front view of the DT chamber and the cell cables. As can be seen, the cables are connected to the OBDT board, shown at the bottom left of that image. The image on the right shows the ATCA crate, installed right next to the chamber. The leftmost board inside the crate is Ocean, in the center at slot 1 is a DTH board (even though not used by this test) and at its right is the BMTL1 card. An optical fiber (pink color) connects the Ocean with the BMTL1. Furthermore, an SMA cable (light gold color) is connected to BMTL1 to deliver the LHC clock.

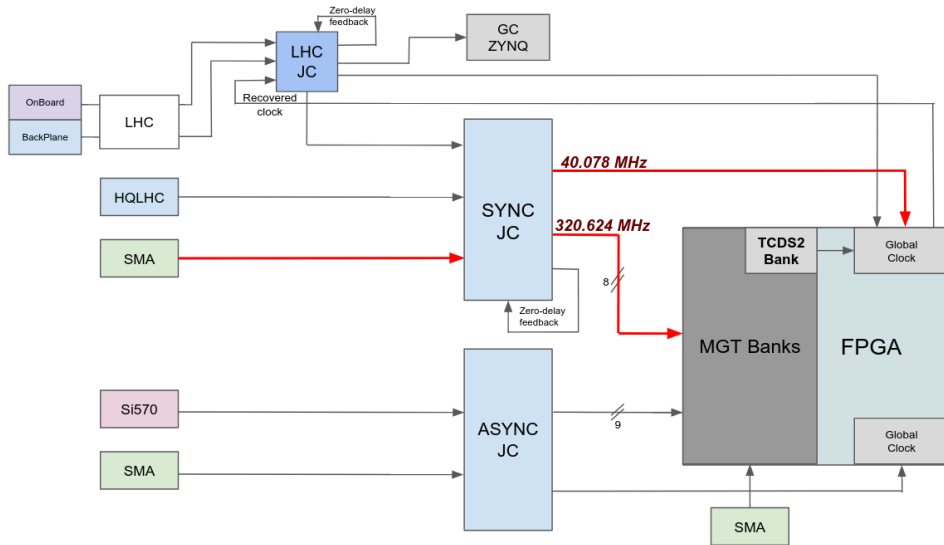


**Figure 10.2:** *The SX5 slice test setup.*

The first action after the setup was connected included the validation of the optical interfaces. Initially, the GBT link that connects OBDT to BMTL1 was tested. The validation was performed by transmitting counter data at the payload, incrementing on every LHC clock. At the BMTL1 side, a similar counter would lock to the received one and count independently. The two values were checked after every received word. After running the test for two days, no miss-match of the two values was detected, indicating error-less data reception. Next, the operation of the CSP links between BMTL1 and Ocean was tested. The CSP protocol contains error detection mechanisms that constantly report the status of the link. After running the 4 connected links for 2 days, no error was detected at the receiving side of the Ocean board.

### 10.1.3 Configuring the slice test

First action towards setting the test was the configuration of the DT chamber and the OBDT board. Once ready, the configuration of the BMTL1 could follow. This process includes configuring the clocking network and the firmware of the FPGA. In this setup, the DTH board could not be used, as it did not support sourcing the LHC clock externally from SMA connector. Instead, the cable from the backplane clock expansion board was delivered to the Sync-JC through the corresponding SMA. The jitter cleaner used a configuration in which it would lock to the SMA input and output a 40.078 MHz clock to the FPGA global clock and 320.624 MHz to the MGT Banks. The clocking network configuration is illustrated with red in Figure 10.3.



**Figure 10.3:** Block diagram of the BMTL1 clocking network illustrating the configuration of the Sync-JC for the SX5 slice test.

The next step was to use the EMP-butler *reset* command to configure the framework. The option used here was the *reset legacy*, which uses the external clock source as the LHC clock. This command also expects the BC0 to arrive from an external source. However, the method used here, that is receiving the BC0 through a GBT link, was not supported by the EMP Framework. For this reason, the TTC block was modified in order to be able to lock from the GBT BC0 signal when the reset legacy command is used.


Once the BMTL1 framework was locked to the LHC clock and to BC0 coming from the TM7 board, the AM algorithm could be configured. The DTU software and the corresponding command were used for this purpose, as well as to validate the TP generation. Next, data were transmitted to Ocean through one 16 Gbps optical link. Data reception at Ocean was validated following instructions that were provided by the GMT group.

### 10.1.4 Results

The purpose of the SX5 setup was to demonstrate a full slice of the BMT datapath, hence called *slice test*. The path begins from real muon data produced by a DT

chamber, sent to the OBDT to be converted to TDC hits, transmitted to BMTL1 and be processed by the AM algorithm to produce TPs, and finally be received at GMT. The readout system at the Ocean board was a key development that assisted the analysis of generated primitives. The duration of capturing data at Ocean was about 30 minutes. Here are presented the results of one such run. The analysis included plotting of the basic parameters of DT primitive payload fields.

Figure 10.4 depicts a plot of the bunch crossing number, as produced by AM during the cosmic muons test. The horizontal axis contains the BX number and the vertical the number of TPs. It can be seen that the distribution is flat for the whole range of possible BX values, that is from 0 to 3563. Here, TPs are generated by cosmic muons that cross the chamber randomly in time. Hence, the flat distribution is totally expected and can be considered as an indicator of reasonable operation for the whole slice test.



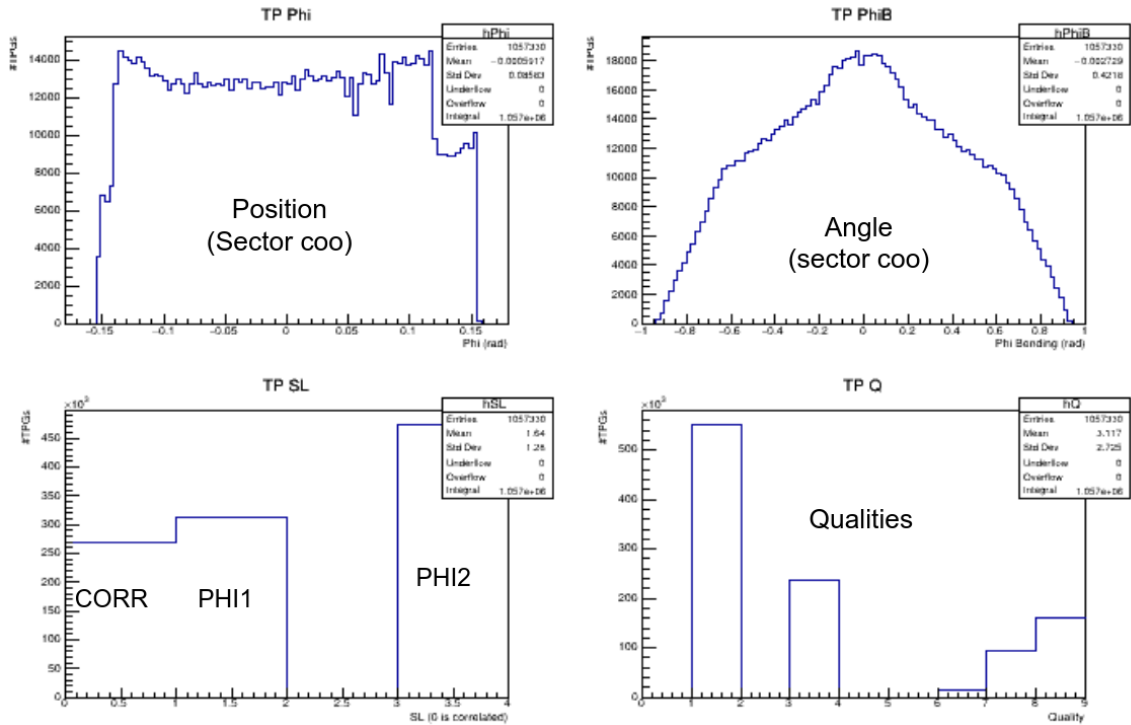
**Figure 10.4:** Results of the SX5 slice tests. Figure plots the bunch crossing number of the DT primitives.

Plots of four additional DT primitive fields are illustrated in Figure 10.5. The top left plot depicts the position of the muon in Sector coordinates. The distribution of this plot is also flat, indicating that muons have crossed the chamber with no spacial preference. On the top right plot the bending angle of the muons can be seen. This distribution is symmetric around the Sector coordinate 0, again indicating muons crossing the surface of the chamber in random angles. Both results are the expected ones for data produced by cosmic muons.

The bottom left plots give information about the super-layer. Value 0 reports TPs where the two phi SLs have been correlated. Values 1 and 3 report TPs that are produced by SL1 and SL3, respectively. The plot at bottom left contains the quality field.

## 10.2 USC Multi Chamber Slice Test

The SX5 slice test, described above, took place the first days of November 2022. At that time the LHC machine was operating, producing proton-proton collisions at the center of CMS. During long shutdown 2 (2019-2022), the DT group had installed 13 OBDT boards at all four chambers of a detector Sector. This Sector is S12 of wheel



**Figure 10.5:** Results of the SX5 slice tests. Figure illustrates some of the fields of the DT primitives.


+2. The S12 OBDTs were operating in parallel with the Phase-1 system, processing hit data from muons that crossed the CMS detector. Up to that time, they were used by the DT slice test in which the OBDTs were transmitting TDC hits to TM7 boards. The purpose of this setup was the validation of both the OBDT operation, but also that of the AM algorithm, which was running in the Phase-1 TM7 boards.

Once the BMT chain was validated at the surface of CMS, the next step was to move the setup downstairs at the CMS counting room, USC. The plan was to re-run the slice test, but this time with the S12 OBDTs and running the test using real collisions data.

### 10.2.1 The USC setup

The setup for the USC slice test was similar to that of SX5. The main differences were the LHC clock source and the origin of the OBDT links. A basic requirement for the execution of this test was the installation of an ATCA at USC. However, this action was completed already by July of the same year, in preparation for this test. Hence, the two ATCA cards could move downstairs. An image of the rack inside which the ATCA is installed is shown in Figure 10.6. On the bottom of the same rack the BMFT subsystem is installed, currently reconstructing muons for the Phase-1 L1 Trigger system. The ATCA crate can be seen above, with the BMTL1 and the Ocean boards installed in it.

The block diagram of this test is illustrated in Figure 10.7. This time, the setup is synchronized with the real LHC clock that originates from a similar expansion board.



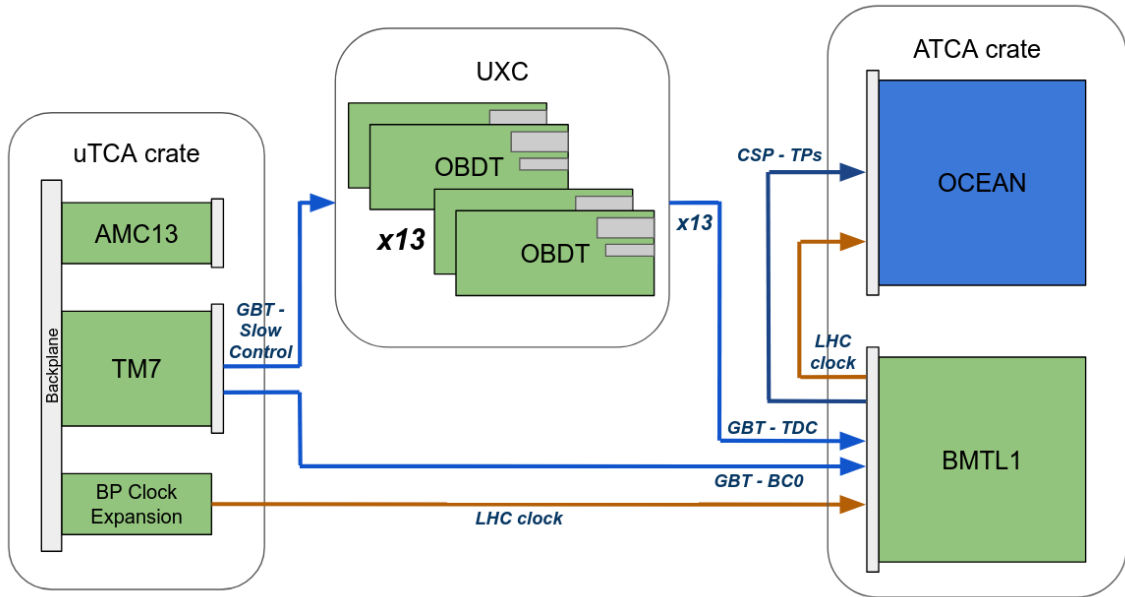
**Figure 10.6:** *The USC slice test setup, showing the Phase-2 ATCA with BMTL1 and Ocean cards. The Phase-1 BMTF subsystem is installed in the same rack.*

This board was installed inside an uTCA crate, used for testing purposes by the DT group. The crate also includes a TM7 board that would transmit the real LHC BC0 through a GBT link to the BMTL1, similar to the SX5 setup. Then, both the LHC clock and BC0 are sent from BMTL1 to Ocean, using an SMA cable for the former and through the CSP protocol the latter. Four optical links were also connected between MGTs of the two boards.

The 13 OBDTs of S12 would deliver 13 optical fibers (1 from each) from the experimental cavern, UXC, directly to USC. In order to keep the original DT slice test intact, these fibers are split at USC and connected both to the TM7 boards and to BMTL1. The reception at BMTL1 was performed using one x12 fiber cable connected to one Rx12 Firefly and by using a patch panel for the remaining two (13th UXC fiber and BC0-fiber).

### 10.2.2 Configuring the slice test

Validation of the 14 in total OBDTs was not needed. The GBT *Link Status* indicator was used instead. The origin of the 13 detector links are the  $\phi$  and  $\theta$  SLs of MB1, MB2 and MB3 chambers, with one OBDT in each SL, and two OBDTs per  $\phi$  SL at the MB4 chamber. All links were mapped during their connection to BMTL1. Since the GBT protocol does not transmit information about the transmitter's origin, the links were connected one by one in order to determine which physical input corresponds to



**Figure 10.7:** *Block diagram of the USC slice test.*

which framework channel. Then, at any time one could be aware of which link comes from each SL/chamber.

The USC test was operating when the LHC beams were present and when real proton-proton collisions were taking place at CMS. On every LHC fill, the setup had to be re-configured and re-synchronized, due to stability loss of the LHC clock during ramp up of the beams. The configuration process was similar to that of the SX5 test. The OBBDT boards were automatically configured, since they were included in the central control software of CMS. The BMTL1 had to re-configure the clocks after every fill and the FPGA get reprogrammed. The clocking network for this test was configured as shown in Figure 10.3.

The firmware used this time did not include one AM chamber, but four. It is similar to the implementation described in section 9.5.2. Even though only one link was sending data from one chamber to Ocean, four instances of the algorithm were implemented in BMTL1. This way, the firmware of the USC slice test was closer to the real Sector-wise firmware of the final system. In addition, the firmware implementation of all four AM chambers could be tested, by delivering hit data to all four AM instances using the OBBDT Decoder multiplexer.

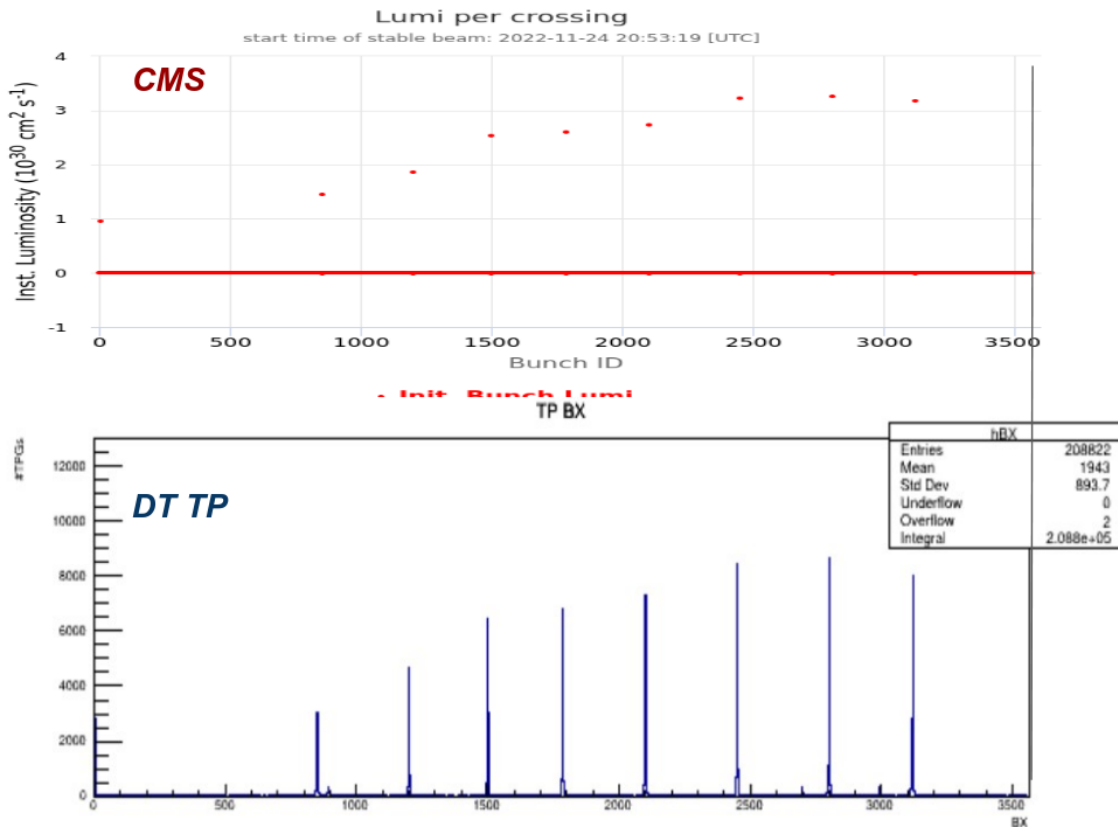
### 10.2.3 Results

The 2022 run of LHC (the first year of Run 3) ended on the 28th of November and the Year End Technical Stop (YETS) started. The BMT slice test was moved downstairs about 10 days prior to this. Hence, the available time window for performing this test was short. The next chance of collisions data would be March of 2023. During the available time, the setup was installed, validated and data were captured for a number of LHC runs. However, there was not enough time for extended analysis of data originating from all chambers.

Presented here are DT primitives, as generated by the AM algorithm running at

BMTL1 board. The primitives were sent to Ocean and captured using the readout system. The capture duration was about 15 minutes, and the chamber under study is MB3. At this chamber, the calculated rate was about 4 KHz of TP generation.

The first results shown at Figure 10.8 illustrate the BX number, as produced by the DT primitives and extracted from the corresponding payload field. The orbit at which this data were captured corresponds to a calibration fill, in which only 9 bunches were circulating the LHC rings. The above plot (in red) is taken from the CMS control software. Below (in blue) is the plot of the DT primitives. It can clearly be seen that the BX number distribution of the two are very similar.

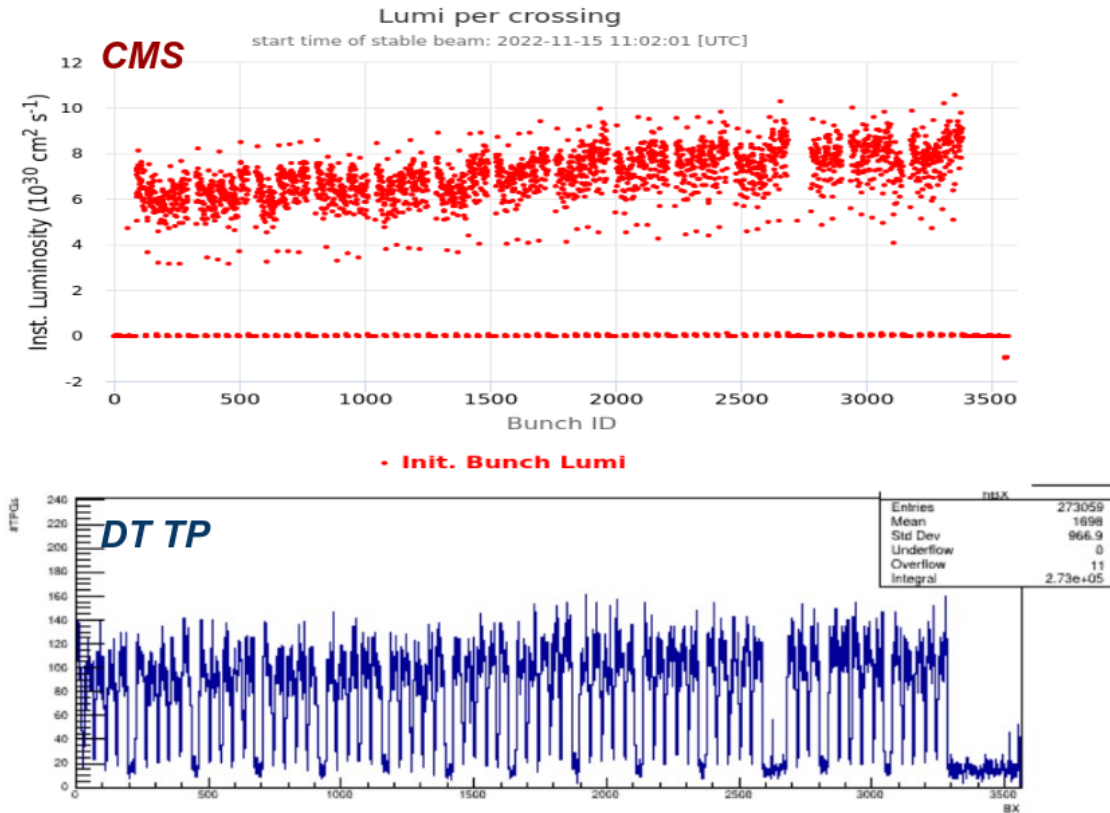


**Figure 10.8:** *BX number of a calibration LHC fill. Above: Plot of the central CMS software. Below: BX number of the DT primitives.*

The plots of Figure 10.9 plot the same BX number value, but this time data are taken from a regular LHC orbit. In both of them, the LHC bunch structure can be seen. The trains of bunches are spaced by empty regions, and the end of the orbit is empty for an extended number of bunch crossings.

The plots of Figure 10.10 illustrate the remaining important fields of the DT primitive. They were captured at Ocean during a standard LHC fill. The top left plot depicts the location of muons in the  $\phi$  coordinate. The muons generated by proton-proton collisions at CMS are expected to follow a random distribution in space. Hence, the flat distribution in this plot is expected. The top right figure plots the bending angle in the  $\phi$  direction. The plot is symmetric around 0 (in Sector coordinates), as expected. The two horns at the top correspond to the two charges of muons that bend in opposite directions due to the magnetic field of the CMS solenoid.

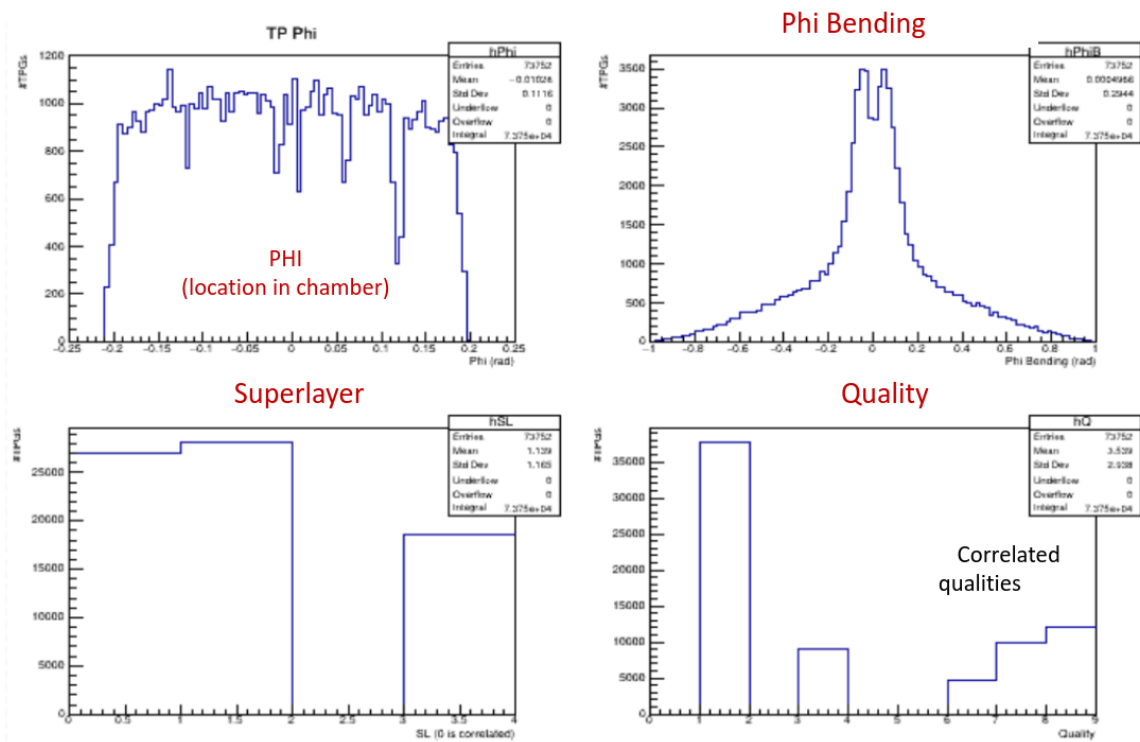




**Figure 10.9:** *BX* number of a standard LHC fill. Above: Plot of the central CMS software. Below: *BX* number of the DT primitives.

The two plots below depict the Superlayer (left) and Quality (right) values, both of which are not far from their expected values.

All of the above plots demonstrate a good performance of the AM algorithm operation. Furthermore, they validate the operation of the hardware chain, starting from the detector with the DT chamber and the OBDT boards, up to the counting room with the BMTL1 and Ocean boards. The BMT slice test using Phase-2 hardware and real muon data from proton-proton collisions at 2022 can be considered successful.



**Figure 10.10:** *Plots of the DT primitive fields, as produced by muons generated in LHC collisions.*

# Chapter 11

## Conclusions

The work conducted for this thesis during the last years refers to the upgrade of the CMS Level-1 Trigger. Specifically, it contributed to the development of the hardware and firmware infrastructure of the BMTL1 sub-system, as well as on common solutions for system-related developments, such as the CSP link. The final version of the CSP link required almost four years to be completed, combining ideas and inputs from many researchers of L1T, and presentations and discussions in many meetings of the trigger project. Its final version, however, manages to accommodate all requirements imposed by the experiment and the technology that is used. It has been extensively tested and its performance is demonstrated using multiple ATCA boards.

The BMTL1 subsystem has evolved significantly during the last years and the basic elements of its final version are in place. A hardware platform adequate to instrument this system has been designed, produced and extensively tested. The firmware infrastructure for this board is in place and the Analytical Method algorithm has been integrated inside of it. The operation of the firmware has been validated in multiple cases, with most notable being the BMT slice test at the underground USC room of CMS. This test was conducted using real data originating from proton-proton collisions at the center of the detector. The result validated the operation of the BMTL1 hardware and firmware, that of the AM algorithm, as well as the operation of the whole BMT system.

During the time of writing this thesis, developments were ongoing towards future tests and the implementation of missing parts of the final BMT system. A new set of OBDv2 boards have been installed in Sector S1 of the CMS barrel, aiming for a slice test using the production version of this board. A first set of tests has been conducted, where data using lpGBT links were received at BMTL1 in the USC setup. The second version of the AM algorithm is in place and almost finalized inside the BMTL1 framework. Also, the KMTF firmware is ready for tracking and our team is preparing the local readout system for further tests. All these developments lead to the next set of USC tests, when the LHC will start its operations once again on March 2024.

In addition, the BMTL1 ATCA board is being revised. The current revision is the first of this board. Even though it is able to operate in realistic conditions, it remains a prototype for the subsystem. Thus, the second revision has been defined and its design is underway, targeting production at the first half of 2024.

# Appendix A

## CSP Control and Status Registers

### A.1 Channel Control Registers

Channel Control (Read-Write) Registers		
Name	Bitfield (Rn[MSB:LSB]) - # of bits	Description
loopback	R0[2:0] 3	Controls the loopback mode of the channel. x0: Normal Operation x2: Near-End PMA Loopback (More at ug578 pg. 88)
reset_counters	R0[3:3] 1	Resets all the status counters of the channel.
tx_polarity	R0[4:4] 1	Inverts the polarity of outgoing data. 0: Not inverted. 1: Inverted. (More at ug578 pg. 156)
rx_polarity	R0[5:5] 1	Inverts the polarity of incoming data. 0: Not inverted. 1: Inverted. (More at ug578 pg. 238)
tx_usrst	R0[6:6] 1	Resets the transmitter firmware blocks of the channel. No reset is issued on the MGT.
rx_usrst	R0[7:7] 1	Resets the receiver firmware blocks of the channel. No reset is issued on the MGT.
rx_align_marker_sel	R0[8:8] 1	Selects the output of the align_marker_out signal of the channel. 0: Rising edge of Valid Bit. 1: Marker generated by the align_tag signal.
rx_align_marker_dis	R0[9:9] 1	Disables the output of the align_marker_out signal.

Figure A.1: Channel Control Registers - 1.

<b>reset_latched_signals</b>	<b>R0[10:10]</b> 1	Reset the latched indicators of the channel.
<b>rx_eyescan_reset</b>	<b>R0[11:11]</b> 1	Used by software to start the EYESCAN reset process. <i>(More at ug578 pg. 73,224)</i>
<b>tx_prbs_sel</b>	<b>R0[15:12]</b> 4	Transmitter PRBS generator test pattern control. <i>(More at ug578 pg. 155)</i>
<b>rx_prbs_sel</b>	<b>R0[19:16]</b> 4	Receiver PRBS checker test pattern control. <i>(More at ug578 pg. 240)</i>
<b>rx_prbs_error_reset</b>	<b>R0[20:20]</b> 1	Resets the PRBS error counter. <i>(More at ug578 pg. 239)</i>
<b>tx_prbs_force_error</b>	<b>R0[21:21]</b> 1	When this port is driven High, a single error is forced in the PRBS transmitter for every <u>TXUSRCLK2</u> clock cycle that the port is asserted. <i>(More at ug578 pg. 155)</i>
<b>rxpmareset</b>	<b>R0[22:22]</b> 1	This port is driven High and then deasserted to start RX PMA reset process. <i>(More at ug578 pg. 73)</i>
<b>rxlpm_en</b>	<b>R0[23:23]</b> 1	0: DFE 1: LPM <i>(More at ug578 pg. 197)</i>
<b>tx_idle_method</b>	<b>R0[24:24]</b> 1	Defines the Idle Method of the Transmitter. 0: Idle Method 1 - Transmits Idles as Control words 1: Idle Method 2 - Transmits Idles as Data words
<b>rx_disable_icm</b>	<b>R0[25:25]</b> 1	Disables Index Correction Mechanism (ICM)
<b>diffctrl</b>	<b>R1[4:0]</b> 5	Driver Swing Control. <i>(More at ug578 pg. 169)</i>
<b>postcursor</b>	<b>R1[9:5]</b> 5	Transmitter post-cursor TX pre-emphasis control. <i>(More at ug578 pg. 172)</i>

Figure A.2: Channel Control Registers - 2.

<b>precursor</b>	<b>R1[14:10]</b> 5	Transmitter pre-cursor TX pre-emphasis control. <i>(More at ug578 pg. 174)</i>
<b>tm_interval</b>	<b>R2[4:0]</b> 5	User defined 5-bits allocated for TM Interval information.
<b>packet_size</b>	<b>R2[12:5]</b> 8	User defined 8-bits allocated for packet payload size information.
<b>header_1bit</b>	<b>R3[0:0]</b> 1	Injects forced error to transmitting data. 1 bit error to the 3bit Header on every assertion of this register.
<b>header_2bit</b>	<b>R3[1:1]</b> 1	Injects forced error to transmitting data. 2 bits error to the 3bit Header on every assertion of this register.
<b>CC_1bit</b>	<b>R3[2:2]</b> 1	Injects forced error to transmitting data. 1 bit error to the Control Word Code on every assertion of this register.
<b>CC_2bit</b>	<b>R3[3:3]</b> 1	Injects forced error to transmitting data. 2 bits error to the Control Word Code on every assertion of this register.
<b>CRC_1bit</b>	<b>R3[4:4]</b> 1	Injects forced error to transmitting data. 1 bit error to CRC checksum on every assertion of this register.
<b>index_1bit</b>	<b>R3[5:5]</b> 1	Injects forced error to transmitting data. 1 bit error to the index number of filler words on every assertion of this register.

**Figure A.3:** *Channel Control Registers - 3.*

## A.2 Common Control Registers

Common Control (Read-Write) Registers		
Name	Bitfield (Rn[MSB:LSB]) - # of bits	Description
reset_tx_pll_and_datapath	R0[0:0] 1	Resets the transmitter side of the MGT core and the corresponding QPLL. A reset is also issued to the firmware transmitter blocks. <i>* After <u>resetting QPLL</u> the rx_datapath must issue a reset in order to operate.</i>
reset_rx_datapath	R0[1:1] 1	Resets the receiver side of the MGT core. A reset is also issued to the firmware receiver blocks.
enable_power_pll	R0[2:2] 1	Enables power on the Quad PLL. <i>(Powered down on start-up)</i> 0: Power-down mode. 1: Normal mode. <i>(More at ug578 pg. 83)</i>
enable_power_rx	R0[3:3] 1	Enables Power on the receiver lane. <i>(Powered down during start-up)</i> 0: Power-down mode. Rx is idle 1: Normal mode. Rx is active. <i>(More at ug578 pg. 83)</i>
enable_power_tx	R0[4:4] 1	Enables Power on the transmitter lane. <i>(Powered down during start-up)</i> 0: Power-down mode. Tx is idle. 1: Normal mode. Tx active. <i>(More at ug578 pg. 83)</i>
crate_id	R0[12:5] 8	User defined 8-bits allocated for crate/subsystem ID information.
slot_id	R0[16:13] 4	User defined 4-bits allocated for slot ID information.

Figure A.4: Common Control Registers.

## A.3 Channel Status Registers

Channel Status (Read) Registers		
Name	Number of bits	Description
<b>channel</b>	<b>R0[1:0]</b> 2	Framework channel number (0-3).
<b>region</b>	<b>R0[7:2]</b> 6	Framework region of the channel.
<b>slod</b>	<b>R0[11:8]</b> 4	Slot ID as written in the corresponding Tx register of the received channel/link.
<b>crate</b>	<b>R0[19:12]</b> 8	Crate/subsystem ID as written in the corresponding Tx register of the received channel/link.
<b>clock_multiplier</b>	<b>R1[3:0]</b> 4	Multiplier of the 40 MHz clock. (Algo clock)
<b>packet_size</b>	<b>R1[11:4]</b> 8	Packet size as written in the corresponding Tx register of the channel.
<b>tm_interval</b>	<b>R1[16:12]</b> 5	Time Multiplex period of the received data.
<b>idle_method</b>	<b>R1[17:17]</b> 1	Idle method used by the transmitter of the receiving channel. 0: Idle as Control Word. 1: Idle as Data Word.
<b>rx_control_word_index_lock_lost</b>	<b>R2[0:0]</b> 1	Latches if the rx index lock is lost for at least one clock cycle.
<b>rx_init_done</b>	<b>R2[1:1]</b> 1	When asserted indicates that the GTX transceiver RX has finished reset and is ready for use. <i>(More at ug578 pg. 74)</i>
<b>tx_init_done</b>	<b>R2[2:2]</b> 1	When asserted indicates that the GTX transceiver TX has finished reset and is ready for use. <i>(More at ug578 pg. 65)</i>
<b>link_down_latched</b>	<b>R2[3:3]</b> 1	Latched if the link is down for at least one clock cycle.

Figure A.5: Channel Status Registers - 1.



<b>rx_status</b>	<b>R2[4:4]</b> 1	0: Link is Down. 1: Serial receiving stream is aligned and link is Up.
<b>tx_fifo_full</b>	<b>R2[5:5]</b> 1	Latched indicator of the Tx FIFO full port.
<b>tx_fifo_almost_full</b>	<b>R2[6:6]</b> 1	Latched indicator of the Tx FIFO almost full port.
<b>rx_rxcdrlock</b>	<b>R2[7:7]</b> 1	Indicates that the Rx CDR circuit is locked to the receiving serial stream. UG 576&578 refer to this signal as Reserved. Experience has shown that in GTH transceivers this register toggles when the channel receives optical signal. In GTY transceivers it seems to operate as expected.
<b>rx_prbs_locked</b>	<b>R2[8:8]</b> 1	Output to indicate that the RX PRBS checker has been error free for RXPBRS_LINKACQ_CNT XCLK cycles after reset. <i>(More at ug578 pg. 240)</i>
<b>rx_prbs_error</b>	<b>R2[9:9]</b> 1	Non-sticky status output indicates that PRBS errors have occurred. <i>(More at ug578 pg. 240)</i>
<b>tx_invalid_signal_count_start_or_last_when_valid_low</b>	<b>R2[15:12]</b> 4	Counts invalid combinations of data to be transmitted - Start or last bits asserted when Valid Bit is low.
<b>tx_invalid_signal_count_start_or_last_continuous</b>	<b>R2[19:16]</b> 4	Counts invalid combinations of data to be transmitted - Start or last bits are asserted for more than 1 clock cycle.
<b>tx_invalid_signal_count_start_and_last_simultaneous</b>	<b>R2[23:20]</b> 4	Counts invalid combinations of data to be transmitted - Start and last bits are asserted simultaneously
<b>tx_invalid_signal_count_valid_before_start_without_last</b>	<b>R2[27:24]</b> 4	Counts invalid combinations of data to be transmitted - No last bit asserted prior to assertion of start bit, while Valid Bit is high
<b>tx_invalid_signal_count_valid_after_last_without_start</b>	<b>R2[31:28]</b> 4	Counts invalid combinations of data to be transmitted - Valid bit is high a last is asserted without the assertion of start bit

Figure A.6: Channel Status Registers - 2.

<b>crc_errors</b>	<b>R3[31:0]</b> <b>32</b>	Counts the number of CRC errors.
<b>packets_in</b>	<b>R4[31:0]</b> <b>32</b>	Counts the number of received packets / number or checked CRCs.
<b>typeField_double_errors</b>	<b>R5[31:16]</b> <b>16</b>	Counts double errors that occur in the Type Field of the Control Words. Double errors cannot be corrected by the Hamming 8,4 code.
<b>typeField_single_errors</b>	<b>R5[15:0]</b> <b>16</b>	Counts single errors that occur in the Type Field of the Control Words. Type Field is encoded using Hamming 8,4 codes hence single errors are corrected.
<b>control_word_index_lock_lost</b>	<b>R6t[31:16]</b> <b>16</b>	Counts the number of times the rx index lock has been lost.
<b>hard_errors</b>	<b>R6[15:0]</b> <b>16</b>	Hard error counter. Hard errors are the errors occurring in the 3-bit Header (value not equal to "101" or "010").
<b>rx_pointer_diff</b>	<b>R7[24:16]</b> <b>9</b>	Difference between the read pointer and the write pointer of the Rx Buffer.
<b>wrong_index</b>	<b>R2[15:0]</b> <b>16</b>	Counts errors that occurred on the Index number of the filler words. These errors do not cause link misalignment.

Figure A.7: Channel Status Registers - 3.

## A.4 Common Status Registers

Common Status (Read) Registers		
Name	Bitfield (Rn[MSB:LSB]) - # of bits	Description
qplllock	R0[0:0] 1	Frequency lock signal indicates that the QPLL frequency is within the predetermined tolerance.

Figure A.8: *Common Status Registers.*

# Bibliography

- [1] URL: <https://home.cern/news/press-release/cern/cern-experiments-observe-particle-consistent-long-sought-higgs-boson>.
- [2] Oliver Sim Brüning et al. *LHC Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: [10.5170/CERN-2004-003-V-1](https://doi.org/10.5170/CERN-2004-003-V-1). URL: <https://cds.cern.ch/record/782076>.
- [3] Maurizio Vretenar et al. *Linac4 design report*. Vol. 6. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. DOI: [10.23731/CYRM-2020-006](https://doi.org/10.23731/CYRM-2020-006). URL: <https://cds.cern.ch/record/2736208>.
- [4] “The Proton Synchrotron Booster”. In: (2012). URL: <https://cds.cern.ch/record/1997372>.
- [5] K. Aamodt et al. “The ALICE experiment at the CERN LHC”. In: *JINST* 3 (2008), S08002. DOI: [10.1088/1748-0221/3/08/S08002](https://doi.org/10.1088/1748-0221/3/08/S08002).
- [6] ATLAS Collaboration. “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *JINST* 3 (2008). Also published by CERN Geneva in 2010, S08003. DOI: [10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003). URL: <https://cds.cern.ch/record/1129811>.
- [7] CMS Collaboration. “The CMS experiment at the CERN LHC. The Compact Muon Solenoid experiment”. In: *JINST* 3 (2008). Also published by CERN Geneva in 2010, S08004. DOI: [10.1088/1748-0221/3/08/S08004](https://doi.org/10.1088/1748-0221/3/08/S08004). URL: <https://cds.cern.ch/record/1129810>.
- [8] FASER Collaboration. *Technical Proposal for FASER: ForwArd Search Experiment at the LHC*. Tech. rep. Geneva: CERN, 2018. arXiv: [1812.09139](https://arxiv.org/abs/1812.09139). URL: <https://cds.cern.ch/record/2651328>.
- [9] LHCb Collaboration. “The LHCb Detector at the LHC”. In: *JINST* 3 (2008). Also published by CERN Geneva in 2010, S08005. DOI: [10.1088/1748-0221/3/08/S08005](https://doi.org/10.1088/1748-0221/3/08/S08005). URL: <https://cds.cern.ch/record/1129809>.

- [10] Alessio Tiberio. “The LHCf experiment at the Large Hadron Collider: status and prospects”. In: *PoS ICRC2023* (2023), p. 444. DOI: [10.22323/1.444.0444](https://doi.org/10.22323/1.444.0444). URL: <https://cds.cern.ch/record/2868230>.
- [11] MoEDAL Collaboration. *Technical Design Report of the MoEDAL Experiment*. Tech. rep. 2009. URL: <https://cds.cern.ch/record/1181486>.
- [12] G. Anelli et al. “The TOTEM experiment at the CERN Large Hadron Collider”. In: *JINST* 3 (2008), S08007. DOI: [10.1088/1748-0221/3/08/S08007](https://doi.org/10.1088/1748-0221/3/08/S08007).
- [13] HL-LHC Collaboration. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. DOI: [10.23731/CYRM-2020-0010](https://doi.org/10.23731/CYRM-2020-0010). URL: <https://cds.cern.ch/record/2749422>.
- [14] *The Phase-2 Upgrade of the CMS Tracker*. Tech. rep. Geneva: CERN, 2017. DOI: [10.17181/CERN.QZ28.FLHW](https://doi.org/10.17181/CERN.QZ28.FLHW). URL: <https://cds.cern.ch/record/2272264>.
- [15] Paulo Moreira et al. *lpGBT documentation: release*. 2022. URL: <https://cds.cern.ch/record/2809058>.
- [16] S Meroli et al. “Development and prototyping of the Versatile Link + optical fibre cabling plants for the HL-LHC upgrades of the ATLAS and CMS experiments at CERN”. In: *JINST* 17.12 (2022), p. C12012. DOI: [10.1088/1748-0221/17/12/C12012](https://doi.org/10.1088/1748-0221/17/12/C12012). URL: <https://cds.cern.ch/record/2861851>.
- [17] *The Phase-2 Upgrade of the CMS Barrel Calorimeters*. Tech. rep. This is the final version, approved by the LHCC. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2283187>.
- [18] *The Phase-2 Upgrade of the CMS Endcap Calorimeter*. Tech. rep. Geneva: CERN, 2017. DOI: [10.17181/CERN.IV8M.1JY2](https://doi.org/10.17181/CERN.IV8M.1JY2). URL: <https://cds.cern.ch/record/2293646>.
- [19] *The Phase-2 Upgrade of the CMS Muon Detectors*. Tech. rep. This is the final version, approved by the LHCC. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2283189>.
- [20] Collaboration CMS. *A MIP Timing Detector for the CMS Phase-2 Upgrade*. Tech. rep. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2667167>.
- [21] Julian Maxime Mendez et al. “CERN-IPMC solution for AdvancedTCA blades”. In: *PoS TWEPP-17* (2018), p. 053. DOI: [10.22323/1.313.0053](https://doi.org/10.22323/1.313.0053).
- [22] Xilinx. “Vivado Design Suite User Guide: Getting Started (UG910)”. In: (). URL: <https://docs.xilinx.com/r/en-US/ug910-vivado-getting-started>.

- [23] Xilinx. “UltraScale FPGA Product Tables and Product Selection Guide (XMP102)”. In: (). URL: <https://docs.xilinx.com/v/u/en-US/ultrascale-fpga-product-selection-guide>.
- [24] Xilinx. “UltraScale+ FPGAs Product Selection Guide (XMP103)”. In: (). URL: <https://docs.xilinx.com/v/u/en-US/ultrascale-plus-fpga-product-selection-guide>.
- [25] *The Phase-2 Upgrade of the CMS Level-1 Trigger*. Tech. rep. Final version. Geneva: CERN, 2020. URL: <https://cds.cern.ch/record/2714892>.
- [26] CMS. “Particle-flow reconstruction and global event description with the CMS detector”. In: *Journal of Instrumentation* 12.10 (2017), P10003–P10003. DOI: [10.1088/1748-0221/12/10/p10003](https://doi.org/10.1088/1748-0221/12/10/p10003). URL: <https://doi.org/10.1088/1748-0221/12/10/p10003>.
- [27] Daniele Bertolini et al. “Pileup Per Particle Identification”. In: *JHEP* 10 (2014), p. 059. DOI: [10.1007/JHEP10\(2014\)059](https://doi.org/10.1007/JHEP10(2014)059). arXiv: [1407.6013 \[hep-ph\]](https://arxiv.org/abs/1407.6013).
- [28] CMS. *40 MHz Level-1 Trigger Scouting for CMS*. Tech. rep. Geneva: CERN, 2020. DOI: [10.1051/epjconf/202024501032](https://doi.org/10.1051/epjconf/202024501032). URL: <https://cds.cern.ch/record/2798134>.
- [29] Samtec. “Micro Flyover On-Board Optical Engine, FireFly”. In: (). URL: <https://www.samtec.com/optics/optical-cable/mid-board/firefly>.
- [30] CMS Collaboration. *The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger*. Tech. rep. This is the final version of the document, approved by the LHCC. Geneva: CERN, 2021. URL: <https://cds.cern.ch/record/2759072>.
- [31] Javier Sastre Alvaro. *The OBDT board: A prototype for the Phase 2 Drift Tubes on-detector electronics*. Tech. rep. Geneva: CERN, 2020. URL: <https://cds.cern.ch/record/2797780>.
- [32] CMS DT Group. “The Analytical Method algorithm for trigger primitives generation at the LHC Drift Tubes detector”. In: (2023). ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2023.168103>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900223000931>.
- [33] R. Frühwirth. “Application of Kalman filtering to track and vertex fitting”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 262.2 (1987), pp. 444–450. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/0168-9002\(87\)90887-4](https://doi.org/10.1016/0168-9002(87)90887-4). URL: <https://www.sciencedirect.com/science/article/pii/0168900287908874>.

- [34] M. Bachtis et al. “Upgrade of the CMS Barrel Muon Track Finder for HL-LHC featuring a Kalman Filter algorithm and an ATCA Host Processor with Ultrascale+ FPGAs”. In: *PoS TWEPP2018* (2019), p. 139. DOI: [10.22323/1.343.0139](https://doi.org/10.22323/1.343.0139).
- [35] C. Ghabrous Larrea et al. “IPbus: a flexible Ethernet-based control system for xTCA hardware”. In: *JINST* 10.02 (2015), p. C02019. DOI: [10.1088/1748-0221/10/02/C02019](https://doi.org/10.1088/1748-0221/10/02/C02019).
- [36] K. Adamidis et al. “Hermes — a robust, low latency, optical link protocol for synchronous data transfer at commercial asynchronous line rates”. In: *Journal of Instrumentation* 17.06 (2022), p. C06002. DOI: [10.1088/1748-0221/17/06/C06002](https://doi.org/10.1088/1748-0221/17/06/C06002). URL: <https://dx.doi.org/10.1088/1748-0221/17/06/C06002>.
- [37] Xilinx. “Aurora 64B/66B Protocol Specification (SP011)”. In: (). URL: [https://docs.xilinx.com/v/u/en-US/aurora\\_64b66b\\_protocol\\_spec\\_sp011](https://docs.xilinx.com/v/u/en-US/aurora_64b66b_protocol_spec_sp011).
- [38] I. Bestintzanos et al. “An ATCA processor for Level-1 trigger primitive generation and readout of the CMS barrel muon detectors”. In: *Journal of Instrumentation* 18.02 (2023), p. C02039. DOI: [10.1088/1748-0221/18/02/C02039](https://doi.org/10.1088/1748-0221/18/02/C02039). URL: <https://dx.doi.org/10.1088/1748-0221/18/02/C02039>.
- [39] Xilinx. “UltraScale+ FPGAs Product Selection Guide (XMP103)”. In: (). URL: <https://docs.xilinx.com/v/u/en-US/ultrascale-plus-fpga-product-selection-guide>.
- [40] Skyworks. “Si5345/44/42 Rev D Data Sheet”. In: (). URL: <https://www.skyworksinc.com/-/media/SkyWorks/SL/documents/public/data-sheets/si5345-44-42-d-datasheet.pdf>.
- [41] Xilinx. “UltraScale Architecture GTY Transceivers, UG578 (v1.3.1) September 14, 2021”. In: (). URL: <https://docs.xilinx.com/v/u/en-US/ug578-ultrascale-gty-transceivers>.
- [42] IPbus. “IPbus page”. In: (). URL: <https://ipbus.web.cern.ch/>.
- [43] Xilinx. “AXI Chip2Chip LogiCORE IP Product Guide (PG067)”. In: (). URL: <https://docs.xilinx.com/r/en-US/pg067-axi-chip2chip>.
- [44] Xilinx. “Aurora 64B/66B LogiCORE IP Product Guide (PG074)”. In: (). URL: <https://docs.xilinx.com/r/en-US/pg074-aurora-64b66b>.
- [45] Xilinx. “Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics (DS923)”. In: (). URL: <https://docs.xilinx.com/v/u/en-US/ds923-virtex-ultrascale-plus>.

- [46] IPBB. “IPBB primer”. In: (). URL: <https://ipbus.web.cern.ch/doc/user/html/firmware/ipbb-primer.html>.



# List of Figures

1	Σχηματική απεικόνιση του ανιχνευτή CMS. . . . .	iv
2	Σχηματικό διάγραμμα της αρχιτεκτονικής του Συστήματος Σκανδαλισμού Επιπέδου-1. . . . .	vi
3	Σχηματικό διάγραμμα του πρωτοκόλλου Hermes. . . . .	viii
4	Σχηματικό διάγραμμα του κυκλώματος αποστολέα του CSP. . . . .	ix
5	Η κάρτα BMTL1 ATCA με οπτικές διασυνδέσεις. . . . .	ix
6	Διαγράμματα ματιού των οπτικών διασυνδέσεων που λειτουργούν στην ταχύτητα των 25 Gbps. . . . .	x
7	Διαγράμματα από ίχνη μιονίων που παρήχθησαν από αληθινές συγκρο- ύσεις μιονίων. . . . .	xi
1.1	The Standard Model of elementary particle physics. . . . .	7
2.1	The total integrated luminosity delivered by the LHC and recorded by the CMS experiment during Run 1, Run 2 and the beginning of Run 3. . . . .	10
2.2	Schematic view of the accelerator complex at CERN. The protons used by LHC originate from LICAC4 and are gradually accelerated by the Booster (PSB), the PS and SPS. They are injected to LHC at 450 GeV energy. Beams from the accelerator complex are also used by other experiments operating at CERN. . . . .	11
2.3	The bunch structure of a nominal LHC fill. The total number bunches fitting in an orbit is 3564. A typical orbit consists of about 2808 proton bunches, arranged as shown in the image. . . . .	12
2.4	The LHC and HL-LHC timeline plan, as of 2023. . . . .	16
3.1	Schematic view of the CMS detector. . . . .	18
3.2	Left: The CMS detector coordinate system. Right: The relation be- tween the pseudorapidity, $\eta$ , with the polar angle $\theta$ . . . . .	19
3.3	Left: Artistic view of the CMS solenoid magnet. Right: The magnet during the construction of CMS at UXC. The red iron layers outside the coil are the return yokes. . . . .	20

3.4	The CMS Phase-1 Tracker. . . . .	21
3.5	View of one quarter of the CMS Phase-2 Inner Tracker detector. . . .	22
3.6	The r-z view of one quarter of the CMS Phase-2 Outer Tracker. . . .	23
3.7	CMS Phase-2 Outer Tracker hybrids. Left: The 2S module. Right: The PS module, having half the size of 2S. . . . .	24
3.8	Schematic of the front-end electronics architecture for Phase-2. Left: The connections between ECAL crystals, the APDs, the new VFE card and the new FE card. Right: The architecture of data path between the new FE card and the Level-1 Trigger and DAQ systems. . . . .	25
3.9	Left: The Hadronic Calorimeter at the surface assembly area of CMS. Right: The Electromagnetic Calorimeter installed inside the HCAL.	26
3.10	Layout of one half of the HGCAL detector. . . . .	27
3.11	Left: Drawing of Layer-9 of CE-E, made of HGCAL hexagonal sensors. Right: Layout of layer 24 of CE-H, made of both silicon and scintillator sensors. . . . .	28
3.12	A quadrant of the muon system. . . . .	29
3.13	Layout of the CMS barrel muon DT chambers and sectors in one of the 5 wheels. . . . .	30
3.14	Left: Structure of a DT Chamber consisting of three Super Layers. Each SL is made of four layers of DT cells, two oriented parallel to z-axis and one vertically to it. Right: View of a Drift Tube cell. . . .	31
3.15	The ME2 station of CSCs. . . . .	32
4.1	Drawing of the side view of an ATCA blade attached with an RTM. The figure highlights the backplane connectors of the ATCA standard.	35
4.2	Data interface between the hub blades and the processor boards inside an ATCA crate, as defined by CMS. . . . .	35
4.3	Left: Block design of fundamental building blocks and programmable interconnects of an FPGA. Right: Photo of a Spartan FPGA produced by Xilinx. . . . .	37
4.4	Product table of FPGAs produced by Xilinx and are used by the CMS Level-1 Trigger for Phase-2. . . . .	38
4.5	Floorplanning of SLR 0 of a VU13P FPGA, extracted from the Vivado software. . . . .	39
4.6	Block design of the MGT Quad structure. . . . .	40
4.7	MGT Channel Block design showing the PCS and PMA of the TX and RX functional blocks. . . . .	41
4.8	Left: Data Sampler and Offset Sampler. Right: Eye-diagram 2D plot.	42

4.9	Optical Transceiver Modules. Left: A Quad Short Form-factor Plug-gable (QSFP). Right: Samtec Fireflies. . . . .	43
5.1	Functional block diagram of the CMS Level-1 Trigger system for Phase-2. . . . .	46
5.2	Block diagram of the Phase-2 Calorimeter trigger. . . . .	47
5.3	Block diagram of the Phase-2 Muon Trigger system. . . . .	48
5.4	Channel correlation at the two sides of a $P_T$ module. Green channels show hits that pass the 2 GeV threshold. . . . .	50
5.5	Track Trigger architecture . . . . .	51
5.6	Block design of the Phase-2 Level-1 Trigger including the 40 MHz scouting system. . . . .	54
5.7	The APd1 ATCA board featuring a VU9P FPGA. . . . .	55
5.8	The Serenity ATCA board hosting two daughter-cards. . . . .	56
5.9	The Ocean ATCA board featuring a ZU19EG ZYNQ. . . . .	57
5.10	The Barrel Layer-1 demonstrator board featuring a KU040 FPGA. . . . .	57
5.11	Block design of the structure of the Data Acquisition system. . . . .	58
5.12	Structure of the TCDS2 system. . . . .	59
5.13	Left: Block design of the DTH-400 board. Right: The prototype DTH-400 ATCA board. . . . .	60
6.1	Barrel Muon Trigger Architecture. . . . .	63
6.2	The first version of the OBDT board. . . . .	64
6.3	Data format of a single TDC hit. . . . .	64
6.4	Block diagram of the RPC Phase-2 system. The Barrel part is high-lighted in the blue lines. . . . .	65
6.5	Left: Groups of 10 cells where combinations of hits are searched for. Right: All cell layouts compatible with a muon straight line inside a SL. . . . .	66
6.6	Graphical illustration of a slice of the CMS in the transverse view. The figure highlights the steps followed by the KMTF algorithm to reconstruct a muon track. . . . .	69
6.7	Left: The BMTL1 ATCA board. Right: The X2O board used by GMT. . . . .	70
6.8	A schematic illustration of the interface architecture between BF and GMT. . . . .	72
7.1	Block diagram with the basic firmware components of the EMP Frame-work. . . . .	75

7.2	Left: Block diagram of the basic components inside the Datapath block. The Region block is implemented N-1 times, where N is the total number of Regions supported by the FPGA type. Right: Block diagram of the Region block. . . . .	78
7.3	Example of a declaration file for an EMP project. . . . .	80
7.4	Output of the reset internal command using the emp-butler software. . . . .	81
7.5	Block diagram of the lpGBT-FPGA firmware. . . . .	83
7.6	Block diagram of Hermes protocol. Arrows illustrate data propagation through the transmitter and receiver data paths, flowing between the payload block and the MGT. . . . .	85
7.7	Left: Packet mode example of transmitting the CRC with the first Idle word at the end of every packet. Right: Streaming mode example of transmitting the CRC with the last data word at the end of every packet. . . . .	88
7.8	Example of transmitting the CRC value in 4 chunks of 4 bits utilizing space 68 to 71 of the Tx FIFO and Rx BRAM. . . . .	88
8.1	Block diagram of the CSP transmitter datapath. . . . .	95
8.2	CSP scrambling method. Figure illustrates the transmitter side logic. Similar technique is implemented on the receiver side. . . . .	98
8.3	Block diagram of the CSP receiver datapath. . . . .	99
8.4	CSP descrambling method. Figure illustrates the receiver side logic. . . . .	100
8.5	Flow chart of the CSP alignment procedure. . . . .	101
8.6	Example of running the configure tx command of emp-butler. . . . .	104
8.7	Example of running the configure rx command of emp-butler. . . . .	104
8.8	Example of running the mgts status rx command of emp-butler. . . . .	105
8.9	Example of running the mgts status rx command of emp-butler, after running the align command. . . . .	105
8.10	List of the three sets of tests performed between an APd1 and a Serenity board to validate the CSP implementation. . . . .	107
8.11	Result of the APd1-Serenity tests using attenuated channel for 23 hours. . . . .	107
9.1	Block diagram of the BMTL1 ATCA board. . . . .	109
9.2	Connections of the BMTL1 ATCA MGT Banks to Samtec Firefly modules. . . . .	111
9.3	Block diagram of the clocking network of the BMTL1 ATCA board. . . . .	112
9.4	Connections of Sync and Async reference clocks to the BMTL1 ATCA MGT Banks. . . . .	113

9.5	Left: IBERT interface indicating that QPLLs are Locked. Right: IBERT interface showing the status of the links, bits transmitted and errors detected. . . . .	114
9.6	The BMTL1 board instrumented with 8 Firefly modules. . . . .	114
9.7	Eye diagrams of all 16G links running in external loopback using fiber cables. . . . .	115
9.8	Eye diagrams of all 25G links running in external loopback using fiber cables. . . . .	115
9.9	Long runs of 24 16G links using Samtec Fireflies. No errors counted for a BER of $10^{-15}$ . . . . .	116
9.10	Long runs of 8 25G links using Samtec Fireflies. No errors counted for a BER of $10^{-16}$ . . . . .	116
9.11	Eye diagrams in both directions of the channels connecting the FPGA with the ZYNQ. . . . .	117
9.12	Configuration of the clocking network (highlighted in red) to sync to the LHC clock coming from the backplane. . . . .	118
9.13	Eye diagrams of the Zone 2 TCDS2 link taken for every slot of the ATCA crate. . . . .	118
9.14	List of the In/Out ports of BMTL1 firmware. The MGT channels are not included in the list. . . . .	119
9.15	Schematic of the VU13P to ZYNQ IPbus implementation through AXI. . . . .	121
9.16	The device declaration file of the BMTL1 framework. . . . .	122
9.17	The project declaration file of a BMTL1 example project. . . . .	123
9.18	Output of the emp-butler <i>reset tcds2</i> command. . . . .	124
9.19	Left schematic illustrates the the block diagram of the single channel CPLL based implementation of the lpGBT firmware. The four channel QPLL based implementation is shown on the right. . . . .	125
9.20	Configuration of the Hybrid MGT IP core. In this example, the transmitter uses QPLL0 and is configured with the CSP settings at 16.32 Gbps. The receiver uses QPLL1 with the lpGBT settings at 5.12 Gbps. . . . .	126
9.21	Block diagram of the Hybrid link, including the receiver lpGBT blocks, the transmitter CSP blocks and a hybrid MGT instantiation. . . . .	127
9.22	Floorplanning of a BMTL1 project using the VU13P FPGA. This project instantiates 3 lpGBT regions (yellow), 3 GBT regions (blue) and 4 CSP regions (orange). The TCDS2 firmware is highlighted in purple (bottom right) and the IPbus infrastructure in green. . . . .	129
9.23	Block diagram of the processing path inside BMTL1. The clock domains of every block are also depicted. . . . .	131

9.24	Floorplanning of a BMTL1 implemented design. The project includes 1 instance of the AM algorithm (1 Chamber), one GBT quad and one CSP quad. . . . .	132
9.25	Block diagram of a dt-sector block. . . . .	133
9.26	Left: Resource utilization of a Sector BMTL1 project. Right: Floorplanning of an implemented BMTL1 design consisting of one DT sector.	134
9.27	Floorplanning of an implemented BMTL1 design consisting of two DT sectors. . . . .	135
10.1	Block diagram of the SX5 setup. . . . .	137
10.2	The SX5 slice test setup. . . . .	138
10.3	Block diagram of the BMTL1 clocking network illustrating the configuration of the Sync-JC for the SX5 slice test. . . . .	139
10.4	Results of the SX5 slice tests. Figure plots the bunch crossing number of the DT primitives. . . . .	140
10.5	Results of the SX5 slice tests. Figure illustrates some of the fields of the DT primitives. . . . .	141
10.6	The USC slice test setup, showing the Phase-2 ATCA with BMTL1 and Ocean cards. The Phase-1 BMTF subsystem is installed in the same rack. . . . .	142
10.7	Block diagram of the USC slice test. . . . .	143
10.8	BX number of a calibration LHC fill. Above: Plot of the central CMS software. Below: BX number of the DT primitives. . . . .	144
10.9	BX number of a standard LHC fill. Above: Plot of the central CMS software. Below: BX number of the DT primitives. . . . .	145
10.10	Plots of the DT primitive fields, as produced by muons generated in LHC collisions. . . . .	146
A.1	Channel Control Registers - 1. . . . .	148
A.2	Channel Control Registers - 2. . . . .	149
A.3	Channel Control Registers - 3. . . . .	150
A.4	Common Control Registers. . . . .	151
A.5	Channel Status Registers - 1. . . . .	152
A.6	Channel Status Registers - 2. . . . .	153
A.7	Channel Status Registers - 3. . . . .	154
A.8	Common Status Registers. . . . .	155

# List of Tables

6.1	Qualities of the DT trigger primitive. . . . .	67
6.2	Data format of a DT Trigger Primitive. . . . .	67
6.3	The number of OBDTs for every Sector of a wheel of the DT system. The numbers do not change in any of the 5 wheels. . . . .	71
7.1	Size of the different lpGBT Uplink frame bit fields. . . . .	83
7.2	Hermes frames specification. The Header specifies the transmission of either Data or Control words, the type of which is defined in the Control Word Type field. . . . .	86
7.3	Format of the Padding, CRC and Align Marker Filler words . . . . .	86
8.1	Definition of the CSP Control word bit fields. . . . .	91
8.2	Payload fields of the LID0 Control word. . . . .	92
8.3	Payload fields of the LID1 Control word. . . . .	92
8.4	Payload fields of the CRCV Control word. . . . .	92
8.5	Table for decoding the Header 2-bits (h1,h0) and the 3rd parity (p). .	93
8.6	Parity bits calculation of Hamming (8,4) codes. . . . .	93
8.7	Syndrome bits calculation to decode the Hamming (8,4) codes on the receiver side. . . . .	94
8.8	List of the Error Injections supported by CSP. . . . .	97