

CMX Design Study Phase

05 April 2012

Contents

1	Introduction	2
1.1	Goals	2
1.2	Definitions	2
1.2.1	Base-CMX functionality	2
1.2.2	TP-CMX functionality	2
1.3	Observations on locally and globally derived quantities	3
1.4	Starting point: Guidance from Stockholm Review	3
2	Implementing the Base-CMX functionality	4
2.1	General design criteria	4
2.1.1	Main signals of the extended CMM functionality	4
2.1.2	Optical outputs to TP	5
2.1.3	Multiple optical outputs	6
2.1.4	FPGA package considerations	6
2.2	Evaluating the FF1924 package	6
2.3	Evaluating the FF1923 package	7
2.4	Evaluating the FF1759 package	7
3	Strategies for providing some TP capability	8
3.1	Additional requirements	8
3.2	Single FPGA solution	8
3.3	Dual FPGA solution	9
3.4	Modified Dual FPGA solution	11
4	Appendix A: FPGA Cost estimate	12
5	Appendix B: Cost/Benefit of 10Gb output from CMX to a TP	12
5.1	Motivation	12
5.2	Common comparison aspects for the Single and Dual FPGA cases	13
5.2.1	Pros	13
5.2.2	Cons	13
5.3	Comparison aspects specific to the Single FPGA solution	13
5.4	Comparison aspects specific to the Dual FPGA solution	13
5.5	Cost comparison	14
5.6	Conclusion	14

Rev: 4-Apr-2012

1 Introduction

1.1 Goals

This note summarizes the arguments and conclusions of the feasibility study prompted by the CMX design review of June 2011 in Stockholm.

This study addressed two goals:

1. Determine the best FPGA design choices for implementing the Basic CMX functionality
2. Explore strategies, costs and risks of adding a Topological Processing capability onto the CMX platform

1.2 Definitions

1.2.1 Base-CMX functionality

The Base-CMX functionality includes all aspects required to become a direct replacement of the CMM card with the added ability to become a source of information for transmitting input information to a Topological Processor (abbreviated as TP). The Topological Processor referenced at this point is a generic concept, with no assumption made as to its architecture or location. All functionality requirements listed here would apply equally to a separate standalone TP or a TP based on the CMX platform.

The requirements for the Base-CMX functionality thus can be summarized as:

1. Provide all of the CMM functionality including both the Crate CMM and System CMM aspects * All CMXs must receive and process 400 backplane input signals from the JEM or CPM processor modules in the crate. * A Crate CMX must send its local crate summary data to a System CMX via the backplane over LVDS cables * A System CMX must gather and send its overall triggering information to the CTP over LVDS cables * All CMXs must send ROI and DAQ information over G-links
2. Send JEM or CPM input information out optically to a Topological Processor * This information should be sent over 12-fiber ribbons * 6.4Gbps is nominally sufficient to send all raw data from all JEM or CPM sources over one 12-fiber bundle (using 8b/10b encoding) * The CMX should be able to send multiple copies of this information

1.2.2 TP-CMX functionality

TP-CMX functionality here means adding some Topological Processing capability onto the CMX platform.

This is in addition to, and not to be confused with, the Base-CMX requirement of sending information out to a generic Topological Processor. ATLAS is planning on building a Standalone Topological Processor Sub-System (referred as Standalone TP in this document) in which case the CMX would not necessarily need to provide any TP functionality. Providing some TP functionality on the CMX platform may still be desired as a transitional or permanent capability of the overall ilcalo system. It could be viewed as a backup option in case a Standalone TP does not become available or if its construction is delayed. TP functionality on the CMX platform could also be viewed as additional flexibility or insurance to prepare for future ideas or future requirements that may be asked of the overall system.

The requirements for the TP-CMX functionality can be summarized as: * Receive optical input from each of 12x CMX (all CMXs, itself plus 11 others) * Run multiple Topological Algorithms * Send Topological Triggering information to the CTP * Send the relevant TP information as ROI and DAQ over G-Links

1.3 Observations on locally and globally derived quantities

There is a difference between the type of triggering information achievable using the Crate CMX to System CMX connections and the type of triggering information achievable using the Topological Processor functionality.

Some triggering information is first locally derived by processing the energy deposited in neighboring calorimeter cells, and this information is received from the JEM or CPM cards by a CMX in a given crate. This locally derived trigger information then needs to be summed or gathered across the full calorimeter geographic coverage. This gathering can be achieved over the Crate CMX to System CMX data path and the resulting trigger information can then be sent to the CTP by the System CMX of a given information type (Electron, Tau, Jet, Energy). For example implementing some additional thresholds or additional types of thresholds could be achieved by this method alone, without requiring any TP capability on the CMX, within the limits of the bandwidth available between Crate and System CMX.

In contrast, some triggering information can only be acquired by correlating energy deposits geographically separated across multiple crates, or correlating different types of triggering information (e.g. Electron vs. Jet). All these sources of information need to be collected in one place before any triggering information can be derived. Such types of trigger capability would need a TP infrastructure, either from a Standalone TP or by building some TP capability onto the CMX platform.

There is also some overlap, as the TP capability can be viewed as a superset of the capability of the Crate to System communication path. The TP input capability is more generic, as it can accept inputs from more sources, and can offer higher bandwidth. More complex or bandwidth intensive algorithms could be implemented on the CMX platform by using the optical path instead of the LVDS cables to communicate the Crate level information and by using the TP capability to implement the System CMX algorithm.

1.4 Starting point: Guidance from Stockholm Review

The CMX design review panel made a number of recommendations that directly affect the questions addressed here, and form the starting point for this study.

The panel recommended the use of the Virtex 6 family of FPGAs, as the Virtex 7 family would not be available in the desired timescale to guarantee the production of the CMX card within the time constraints. This conservative requirement continues to be a valid argument.

The review suggested a comparison of the relative advantages of the FF1924 & FF1923 packages for the Virtex 6 Family of FPGAs.

The Virtex 6 Family of FPGAs offers Multi-Gigabit Transceivers (MGT) of two types: GTX and GTH. GTX transceivers are specified at up to 6.6Gbps and GTH transceivers are specified at up to ~11Gbps (but over limited frequency ranges). The review panel recommended that the CMX be designed to support 6.4Gbps of output bandwidth which is sufficient to transfer the full raw backplane data over a 12-fiber bundle using the 8b/10b encoding protocol.

The review panel also recommended to implement the ROI and DAQ output G-link encoding in the FGPA.

2 Implementing the Base-CMX functionality

2.1 General design criteria

2.1.1 Main signals of the extended CMM functionality

The constraints that will drive the choice of FPGA for the Base-CMX functionality will come from the principal data signals involved. The main signal data path related to the emulation and extension of the CMM functionality starts with the JEM or CPM Processor inputs, includes the communication between Crate and System CMX cards, and ends with the outputs to the CTP.

Here are the main characteristics of each group:

1. Processor input signals from the 16 JEM or CPM modules in the crate
 - they are received through the backplane connectors
 - they will be sent at 160Mbps (instead of 40Mbps for the CMM) in order to transfer four times as much information from the processor modules to the CMX than was achieved with the CMM.
 - one signal line is used to carry a merged parity and clock signal (this was a simple parity line on the CMM)
 - these signals will be routed directly to the FPGA. There will be no additional buffer chip on the CMX.
 - the trace line impedance needs to be 60 Ohms
 - the signals are series terminated at the source and are not terminated at the CMX end.
 - the clock/parity signal may need to be treated differently and terminated on the CMX card (the results of the BLT studies need to inform this decision)
 - these signals use 2.5 Volt CMOS logic level
 - we cannot afford 1x IO block per source processor but these signals will need regrouping in neighbor IO blocks that can share FPGA regional clocking resources
 - the signals from a given processor source module are spread out on the backplane connectors and will need regrouping

2. Crate CMX to System CMX IO cables
 - these cables are connected to Rear Transition Modules (RTM) adapters and accessed through the backplane connectors
 - there are 3 cables with only 27 differential LVDS signals from each cable routed through the backplane
 - some Crate CMXs send one cable worth of information to the System CMX in their group, others send two
 - some System CMXs receive two cables worth of information from the Crate CMXs in their group, others receive three
 - for maximum flexibility each cable should be able to independently act as a driver or receiver of LVDS information
 - only 24 signals plus a parity signal are used on the CMM, but all 27 should be routed to the CMX FPGA

- these signals were exchanged at 40 Mbps on the CMM, but the CMX should plan for sending at a higher rate, to be determined. This would increase the amount of information sent from the Crate CMXs to a System CMX for possible future usage (e.g. for additional thresholds as mentioned earlier)
- the CMX should provide LVDS transceivers near the backplane to protect the FPGA pins from static discharges
- the LVDS differential line impedance is 100 Ohm
- the LVDS differential lines will be parallel terminated at the CMX and the LVDS drivers used must thus be able to drive doubly terminated lines.

3. CTP output

- two connectors should be accessible on the front-panel for output to the CTP
- each connector will carry 33 differential LVDS signals
- the CMX needs to send these signals at 40 Mbps, like the CMM
- the LVDS differential line impedance is 100 Ohm
- the LVDS differential lines will not be terminated at the CMX source end

2.1.2 Optical outputs to TP

The other major category of Base-CMX signals is for the high speed serial signals using the Multi Gigabit Transceiver (or MGT) Resources of the Virtex 6 FPGA. These MGT signals are used for the optical output to a TP. The serial MGT differential line pairs will need a careful layout to operate at 6.4Gb. Xilinx design notes give guidelines and recommend using field simulation tools.

We can expect that these MGT traces will compete for access to, and real estate near, the FPGA.

We will need to define a protocol to transfer the data payload for each beam crossing over the 12x 6.4Gbps fiber ribbon. This transfer can be made synchronously to the beam crossings by providing an MGT reference clock synchronous to the beam crossing clock. This transfer could also be made asynchronously by choosing a line rate slightly higher than the desired 6.4 Gbps worth of data payload, and using additional control characters plus idle characters to identify and separate the data from each beam crossing. Xilinx design notes contain a simple example for such asynchronous transfer. The preferred choice is to transfer the data synchronously. We might use some of the unpopulated bunches in the beam structure to transmit control characters suited for link synchronization and allow each link to automatically synchronize at startup or in case of problem, without needing to switch to a special data pattern. There are multiple aspects to link synchronization. At the most basic level, the source serializer and sink deserializer must be synchronized and aligned so that the user data can be recovered from the 8b/10b encoded serial data. There also needs to be some mechanism to synchronize the user data stream to identify the beginning of the data from each beam crossing, and know which beam crossing the data corresponds to. The CMX output link protocol needs to be studied and defined.

The FPGA MGT Output pins will be connected by signal trace pairs to optical transmitters. One standard of optical transceiver is the SNAP12 pluggable form factor identified in the CMX project specification document reviewed in June 2011. The original SNAP12 specification is limited to a line rate of 2.72 Gbps. Avago together with Emcore has defined a modified version of the SNAP12 specification for devices with line rate up to 10Gbps. This modified specification uses the same physical MEG Array board connector as SNAP12 for its pluggable electrical interface, but with a different pinout. Avago sells a AFBR-810BxxxZ Twelve-Channel Transmitter for line rate up to 10.0 Gbps which would be appropriate for the CMX 6.4 Gbps requirement. Newer form factors have also become available for 12-fiber transceiver components since June 2011. Avago's product line includes a pluggable form factor

called MiniPOD that is quite interesting here. MiniPOD uses the same family of mechanical interface to the circuit board as SNAP12, a MEG Array with a 9x9 contact matrix (instead of 10x10 for SNAP12). The footprint and clearance required for MiniPOD devices is smaller than for SNAP12. Avago sells a MiniPOD AFBR-81uVxyZ Twelve-Channel Transmitter for line rate up to ~10.3 Gbps which would be appropriate for the CMX 6.4 Gbps requirement. The optical power output of this MiniPOD device is also notably higher than the corresponding SNAP12-like device above. The MiniPOD devices however cannot be mounted at the edge of the card for direct optical cable connection. They will thus require “pigtail” sections of optical 12-fiber ribbon cables to connect between the MiniPODs and MPT adapters mounted on the card front-panel. MiniPOD devices are called mid-board components as they can be placed closer to the FPGA than SNAP12 devices thus needing relatively shorter traces and potentially providing better electrical signal integrity and lower trace attenuation.

2.1.3 Multiple optical outputs

Requiring the Base-CMX functionality to send multiple copies of its output to a TP will bring flexibility for using the CMX in the overall l1calo system.

A Topological Processor may end up having several modules operating in parallel which would thus benefit from receiving two identical (or maybe different, tailored) copies of the backplane information. Having multiple outputs would also make it possible to operate a Standalone TP and a CMX-TP in parallel at the same time, should this become necessary.

2.1.4 FPGA package considerations

The first consideration faced is deciding on the FPGA orientation. Priority should be given to most sensitive signals listed above, i.e. the processor input signals and MGT signals. The orientation of the FPGA should try to position the FPGA edge with the most Select IO banks facing the backplane. The opposite FPGA edge of the FF1924 and FF1923 packages also present Select IO banks and will thus be facing the front panel. The other two edges of these FPGA packages carry the Multi Gigabit Transceiver IOs and will be facing the top and bottom sides of the board. This arrangement seems natural and the best fit to match the card input and output constraints.

2.2 Evaluating the FF1924 package

This package is interesting because it offers the most MGT IOs (48 GTX and 24 GTH) within the Virtex 6 family, while the trade-off is that it has fewer Select IOs resources than other packages (640 in 16 banks).

One of FPGA edge with select IOs can be facing the backplane side of the CMX card but only provides 320 IO pads. This is not enough to receive all 400 processor inputs. The rest of the processor inputs (80 of 400 inputs = 20%) will thus need to be routed around the FGPA to reach the opposite side of the FPGA facing the front-panel. This means these signals will use comparatively longer traces, and it also means that they will have to pass across the MGT traces connected to the other two edges of the FF1924 FPGA package or pass near the optical transceivers.

The total number of select IO resources available on the FF1924 is also not sufficient to support all expected IO needs. We have already listed 400 processor inputs, 81 cable IOs, and 66 CTP outputs leaves 93 IO signal. The VME bus interface will need about 42 signals. Each 12-fiber transceiver might need about 7 control/status signals. ROI and DAQ activity control may require 5 more. The G-Link transmitters will require some 4 to 6 control/status signal each. We can also estimate the need for another twenty FPGA IO signals for clocks, control and status information from the VME and TCM interface as well as some status for front-panel LEDs. This rough estimate is already 7 over the maximum number of available Select IOs.

The HX565T in a FF1924 package is thus not an acceptable solution for implementing the Base-CMX functionality.

2.3 Evaluating the FF1923 package

This FF1923 package has more Select IOs (720 in 18 banks) than the FF1924 but fewer MGT IOs (40 GTX and 24 GTH).

Another important difference is that one of the two FPGA edges with select IOs has exactly 400 IO resources, i.e. exactly enough for all processor inputs, assuming no termination is needed (i.e. if no reference pin is needed for such termination). This simple count matching avoids the crossing of processor traces near or under MGT traces or transceivers that was described for the FF1924.

The Select IO signal count for the FF1924 package above came close the total 640 count available on the FF1924 package. With 80 more Select IO resources, the FF1923 package should thus provide enough IO resources.

The HX565T in a FF1923 package is thus an acceptable solution for implementing the Base-CMX functionality.

2.4 Evaluating the FF1759 package

Since the Base-CMX functionality doesn't need nor use all the available MGT resources of the FF1923 package, and since we have so far assumed that the CMX does not need to operate as a TP, there is another interesting Virtex 6 package to consider: the FF1759.

This package has yet fewer MGT IOs(36 GTX and 0 GTH) but even more select IOs (840 in 21 banks) than the FF1923.

The Base-CMX functionality only requires two (one plus a second for the desired duplication) 12-fiber outputs to the TP. Two individual GTX drivers are also required for the ROI and DAQ G-link implementations. No MGT inputs are used for normal operation, but one 12-fiber input would still be desired for testing and commissioning of the CMX outputs. 36 GTX outputs is thus sufficient for all known Base-CMX optical inputs and outputs requirements.

Having an excess of Select IO resources available has many advantages. This margin allows for more flexibility in grouping processor inputs. This package also inherently has less rows/depth of IO pins counted from the edge toward the center of the part (16 for FF1923/1924 vs ~13-14 for FF1959), which will translate into less signal layers necessary to route the board. We can push the idea of using less trace layers further. Having some 30% more Select IO resources available than the estimated need would allow to use even less pins in each access row by being "carefully wasteful" of the extra IO pins. This is expected to translate into even less trace layers. The FF1923 has 16 rows of IO pins which would require a minimum of 14 trace layers. We can estimate that using the FF1759 package will reduce this number by 4 to 6 trace layers which would be a significant simplification.

Both the LX550T and SX475T are available in a FF1959 package. The SX475T has a large number (~2,000) of DSP slices and somewhat fewer logic resources (~15%) than the HX565T. The SX475T is also twice as expensive as the LX550T. The LX550T has almost the same amount of logic resources as the HX565T. The LX550T is about one third cheaper than the HX565T FF1923 or FF1924 above.

The LX550T Virtex 6 in a FF1759 package is thus the preferred choice for implementing the Base-CMX functionality.

3 Strategies for providing some TP capability

3.1 Additional requirements

Some additional requirements are needed to implement the TP-CMX functionality on the CMX platform.

The main additional requirements are on MGT resources. The TP-CMX functionality will need as many sets of 12-fiber optical inputs available as possible. It will also need another two individual GTX outputs to implement two more ROI and DAQ G-links to send out TP-related information.

The TP-CMX functionality needs to receive information from the Base-CMX functionality of each of the 12 CMX cards in l1calo. We also know that the maximum number of MGT resources available on the FF1923 or FF1924 package limits the maximum number of 12-fiber inputs to five (for the FF1923) or six (for the FF1924).

One CMX cards is able to send up to 12 fibers to a Standalone TP, assuming the TP has the capacity to receive that many fibers from that many sources. One CMX card would however only be allowed to send a maximum of 5 or 6 individual fibers worth of data to a Topological Processor based on a TP-CMX because it could only field a total of 5 or 6 12-fiber ribbons. This means that a TP-CMX will never be able to receive the full raw data from all 12 CMX cards. The input JEM and CPM data will thus need to be zero-suppressed to reduce the number of output fibers by a factor 2 or more. This limitation is inherent to using a CMX to operate as a TP, but was expected as an acceptable restriction.

This also means that l1calo will need to build a “CMX Fiber Re-Bundler Box”, i.e. a patch panel that can take twelve 12-fiber optical cables as inputs and split and re-arrange them into five or six 12-fiber optical outputs to send to one CMX card with TP-CMX capability.

Such fiber re-bundler box or patch panel could as well provide two sets of five (or six) 12-fiber outputs, for example by splitting the upper half of each 12-fiber input into one set and the lower half into another set. The two sets of outputs from this fiber re-bundler could be sent to two separate CMX cards acting as CMX-TP. This could be useful for a number of reasons. Two CMX-TP would be able to implement twice as many Topological Algorithms and both could send their outputs to the CTP. Alternatively one CMX-TP could run the production version of the TP-CMX firmware while the other could be used to debug and compare the next version of the TP-CMX firmware. Another usage for the second output from this fiber re-bundler is to operate both the CMX-TP and a Standalone TP in parallel on zero-suppress data, should this become desirable. Such operational flexibility may become an important feature, as we may not yet know how the system will be used.

There is an important limitation of the MGT resources, indirect but intrinsic, that needs to be taken into account. The Virtex 6 GTH inputs (i.e. two of the five or six 12-fiber inputs) used for the TP-CMX inputs will not be capable of 6.4 Gbps operation. The GTH circuitry relies on an internal PLL with a narrow frequency range of operation which only support specific ranges of line rates. 5.5 Gbps is the highest frequency that would be compatible between GTX outputs and GTH inputs (the specification specifies 5.591 Gbps, rounded here to 5.5 Gbps as a mildly conservative rounded number below that limit).

Another downside of using GTH transceivers is that they require two more power supplies to operate (1.1V and 1.8V). GTH transceivers also have a higher power consumption, with for example GTH receivers using about twice as much power at 5.5 Gbps (3.6 W per 12-lane) as GTX transceivers (1.9 W per 12-lane).

3.2 Single FPGA solution

If a maximum number of MGT inputs needs to be added to the FPGA already providing the Base-CMX functionality, the FF1759 would be drastically limiting the TP-CMX functionality to only 3x 12-fiber ribbons. The FF1924 package would allow 6x 12-fiber inputs but would not offer enough Select IO pins

for the Base-CMX operation. Choosing the FF1923 package would be the best choice and allow 5x 12-fiber inputs.

A single FPGA solution would add an incremental cost, going from the LX550T-FF1759 to the VX565T-FF1923 for approximately 30% increase in FPGA cost.

There are several disadvantages to such single FPGA solution. The first disappointment is the need to forego the FF1759 packages with all the features that made it the preferred choice for implementing the Base-CMX functionality (the FF1759 being ruled out here because it does not have enough MGT resources).

The second class of difficulty comes with the 50 additional MGT resources used over the Base-CMX-only requirements (i.e. 4x12 additional MGT inputs and 2x additional MGT outputs for the 2x additional G-links). There now is a total of 88 MGT signal pairs and a total of 7 optical transceiver devices which all compete for real estate near the FPGA and for access to its pads. This means that the layout of the TP-CMX functionality will compete with the layout of the Base-CMX functionality. The area around the FPGA will be quite crowded. This will make the layout of the Base-CMX functionality more complicated. There is also an increased risk of electrical interference between the 400x 160 MBps processor single ended input lines and the 88x pairs of 6.4 Gbps traces.

In addition to the routing challenges above, we will also have to deal with interference between Base-CMX and TP-CMX functionality in firmware management. The single FPGA solution ties the firmware of the Base-CMX to the firmware of the TP-CMX firmware. The first effect is that the total available FPGA logic resources will be split between the two separate functionalities, which was of course expected. A single FPGA solution would also make commissioning of both aspects of the firmware more complicated. One might expect that the Base-CMX firmware will become stable early during commissioning while the TP-CMX firmware might evolve for a longer period of time, maybe even during beam physics while the TP functionality is being integrated in the trigger list. Any change of TP-CMX firmware would require compiling a new bitstream which could easily affect the Base-CMX logic resource allocation. It would be impossible to be guarantee that a change in TP-CMX firmware have no effect on the Base-CMX firmware operation which had already been tested and commissioned. It seems unwise to take such an operational risk if such risk could be avoided.

These two functions form two very different aspects of l1calo triggering, and the ultimate usage of the TP-CMX functionality is not fully defined. It is not clear when or exactly how the TP-CMX feature will be used. The Base-CMX functionality however is very well defined, and very much mission-critical. Having both functions in the same FPGA means that they would be competing for the same FPGA resources and would be sharing the same bitstream. Any change in the firmware for one functionality could clearly affect the operation of the other if the overall firmware is recompiled and a new bitstream is generated. One can envision a time when the Base-CMX is stable and finalized while the TP-CMX firmware is still evolving and improving, maybe even during beam physics. Running the risk of affecting the Base-CMX operation, which is critical to beam physics, every time a change to the TP-CMX firmware is being tested is not a comforting situation.

The single FPGA solution is achievable, in principle, but with a number of drawbacks.

3.3 Dual FPGA solution

The first observation is that the Base-CMX functionality and the TP-CMX functionality are totally separate functions. They are logically separated and don't share inputs or outputs, with the particularity that the local Base-CMX is an input to the TP-CMX like the other eleven non-local Base-CMXs in the system. Both TP-CMX and Base-CMX may need to send data to the CTP, but never at the same time and furthermore never on the same CMX card (this aspect is further explained below). These functions also operate in separate latency segments, with the TP-CMX receiving and, de-serializing and processing beam crossing data which had previously been processed, serialized, and transmitted

by the Base-CMX functionality from all CMX cards.

Merging Base-CMX and TP-CMX firmware in one design and one bitstream should be avoided, and one early idea to address this concern was to use an additional and separate CMX card (i.e. in addition to the 12 CMX cards of the existing CMX) that would be used to provide just the TP-CMX functionality. Such proposal had the inconvenient of requiring an additional crate to host just that one additional CMX card, and this wasn't easy to implement.

There is a simpler solution which will achieve essentially the same result as having a separate CMX card to act as a TP-CMX, and it is to have two FPGAs on the CMX card. One FPGA would implement the full Base-CMX functionality and one FPGA would implement just the TP-CMX functionality. As noted above these two functions share very few resources, and the only aspect to consider in more detail is the output to the CTP. The one (or more than one) CMX card that will use its TP-CMX functionality will need to send the triggering information it generated to the CTP but this does not mean the CMX card needs to have a second set of output cables to the CTP. It is not the case that the Base-CMX functionality of every CMX card needs to send data to the CTP. It is only the 4 System CMX cards in 11calo (i.e. the electron, tau, jet, and energy System CMX cards) that gather and form the overall trigger information which need to use the output to CTP. We need simply to insure that no System CMX card is chosen to make use of its TP-CMX functionality. We can simply choose one (or more) of the 8 Crate CMX cards to make use of its TP capability. The CMX card thus only needs to provide one set of two output connectors to the CTP. On a given CMX card the CTP output could be used by the Base-CMX FPGA or by the TP-CMX FPGA or not be used at all. A multiplexer between the two FPGAs and the CTP output drivers will need to be included on the CMX card.

Having two separate FPGAs also allows to independently choose the optimal package type and device for each functionality. This means we can use the FF1759 package for the Base-CMX FPGA for all the reasons described above. We had rejected the FF1924 package earlier because it didn't provide enough pins to support the Base-CMX functionality, but if we need to choose a package for just the TP-CMX functionality, we can now select the FF1924 package that will support the largest number of optical inputs.

The board resources that these two functions actively share are the VME bus interface, the TTC interface, the System ACE and JTAG resources. There is no particular difficulty expected in sharing these resources.

In summary, the advantages of a dual FPGA design are:

1. We create a physical separation of two logically separate tasks
2. This greatly reduces the chance for interference between these two functions during all phases of design and operation
3. The firmware management effort will be easier during commissioning, and maintenance
4. In particular we will be able to upgrade the TP-CMX firmware without fear of affecting the Base-CMX operation
5. We can leverage the advantages of the FF1959 package that the one-FPGA solution had precluded, i.e. more IO pins, more flexibility, less trace layers
6. It is now possible to support one additional 12-fiber ribbon of input to TP-CMX (6x 12-fiber ribbon inputs for the FF1924 instead of 5x for the FF1923)
7. We can separate the types routing challenges. The high density of backplane inputs are connected to one FPGA.

8. Most of the high frequency layout of MGT resources are in one area, and this should translate to more flexibility and shorter MGT traces.
9. There are lower chances for electrical interferences between MGT resources and Select IO resources
10. We also separate the MGT transmitters from the MGT receivers as the Base-CMX FPGA will only use its MGT transmitters and the TP-CMX will only use its MGT receivers (except for 2x G-link), thus reducing the density of MGT traces near each FPGA
11. If the outlook for a Standalone TP changes we could abandon the TP-CMX part of the layout during the design phase with little impact on the Base-CMX layout. We would then either rip out the TP-CMX traces or leave the site unpopulated. This means we could adapt to a final decision on TP-CMX desirability up to the time the prototype cards are being built.
12. We could even stop development of the TP-CMX functionality after the prototype phase and before production assembly phase by either not populating the TP-CMX FPGA location for production boards or by modifying the layout at that phase. With this possibility in mind, we should make an explicit effort to be prepared for such eventuality by planning a layout compatible with such option. This further means we could still adapt to a final decision on TP-CMX desirability up to the time the production cards are being built.

In contrast, there is only one perceived disadvantage:

1. Two FPGAs makes the card more expensive by a factor of $\sim 2x$ in FPGA cost when comparing a dual FPGA design with a FF1923 Base-CMX FPGA plus a FF1924 TP-CMX FPGA to a single FPGA design with one FF1923. The factor is only ~ 1.5 when a FF1759 is used for the Base-CMX FPGA instead of a FF1923. cf. appendix for the cost table

3.4 Modified Dual FPGA solution

The Dual FPGA solution is quite attractive, but it comes with an increased cost. It should also be noted that the TP-CMX FPGA would only be used on only one or possibly very few CMX boards of the running system, while all other CMX boards would have their TP-CMX FPGAs remain unused. The initial goal of the CMX project was to build only one CMX card type which can support all necessary functionality. The dual FPGA solution implemented with a single card type translates into an increased cost in components and a waste of resources for the majority of the cards where the TP functionality is left unused.

We could instead build two types of CMX cards, with a few (the exact number needs to be defined) cards with the TP-capability and the majority of the CMX cards with only the Base-CMX functionality. Two cards would not mean two separate design efforts, since the dual FPGA solution proposed has already maximally separated these two functions.

These two card types could use the exact same circuit board design with some cards leaving the TP-CMX FPGA location unpopulated and possibly also leaving unpopulated other components related to the TP functionality. If this approach causes more difficulties than it solves, two close derivatives of the same card layout could be produced. However the intent is to have only one CMX circuit board type and assemble two CMX card types from it, and aim toward that goal from that start.

In addition to keeping all the advantages listed for the Dual FPGA solution, the Two CMX Type solution will mitigate the cost escalation of needing two FPGA devices on the CMX card, and the wastefulness of leaving the TP-CMX unused on the bulk of the CMX cards in the system.

The disadvantage of this Two Card Type solution is that it breaks with l1calo "tradition" of using all identical boards in l1calo. The systems responsible for programming and monitoring the CMX cards will need to be aware of the location of the one (or more) CMX card including a TP-CMX FPGA.

4 Appendix A: FPGA Cost estimate

This is a preliminary table for the FPGA cost of the FPGAs needed for the options explored above. This list is based on the quantity-one price from the DigiKey website and for the “-2” speed commercial grade, and rounded up to nearest \$0.5k. We should expect a quantity discount when we order the FPGA for the prototype and production boards. Digikey and AVnet are the two distributors for Xilinx.

Option	Base-CMX FPGA Device-Package	Base&TP FPGA Device-Package	TP-CMX FPGA Device-Package	Per Card	Total Q=15	Total Q=4+11
1 Base-CMX only 6.4Gb out	LX550T-FF1759			\$5.5k	\$83k	
2 Single FPGA solution Base: 6.4Gb out TP: 5.5Gb in		LX565T-FF1923		\$8.5k	\$128k	
3 Dual FPGA solution Base: 6.4Gb out TP: 5.5Gb in	LX550T-FF1759		HX565T-FF1924	\$14k	\$210k	
4 Two CMX Types e.g. 2-4 with TP and 11-13 without TP Base: 6.4Gb out TP: 5.5Gb in	LX550T-FF1759		HX565T-FF1924 but not on all cards	3 or 1 above		\$117k

Note: The HX565T device could be replaced with the HX380T, which offers less logic resources but is available at a higher speed grade (~+10% performance in delays and setup times for an additional cost of ~\$3k per part).

5 Appendix B: Cost/Benefit of 10Gb output from CMX to a TP

5.1 Motivation

A total of 12 fibers transmitting at 6.4 Gbps is sufficient to send out the full raw data received by the CMX on its backplane, and there are no known additional sources of information that the CMX would need to send. It is still worthwhile to consider the possibility of designing the CMX to be capable of sending data at a higher rate, namely at 10 Gbps, as supported by the Virtex 6 GTH resources.

From the point of view of a standalone TP, receiving CMX data at 10Gbps could have several advantages. Instead of needing to receive 12 fibers from each CMX card, 8 fibers might then be sufficient to send all raw backplane data. This would reduce the total number of fibers that the standalone TP needs to receive from all CMX cards plus other sources of TP input information. The CMX cards could be sending zero-suppressed information on a reduced number of fibers (less than 12), and increasing the line rate would also further reduce the total number of fibers in the system. From a different perspective and for a fixed number of fibers, increasing the line rate by ~30% would decrease the latency of the data transfer by a similar factor of ~30%.

From the point of view of the CMX project, it is of course desirable to plan ahead and provide as much functionality as possible to provide maximum flexibility in the future. From that point of view, and as long as the FPGA on the CMX card has unused GTH resources, it would make sense to use of such resources and provide 10Gbps output. We should also expect that a card that is capable of operating at 10Gbps could also be made to operate at half that speed when needed.

5.2 Common comparison aspects for the Single and Dual FPGA cases

5.2.1 Pros

As described above, the optical transceivers (SNPA12 or MiniPOD) that we need to use to support 6.4 Gbps are already capable of running at 10Gpbs at no additional cost.

5.2.2 Cons

Raising the line rate of the optical outputs to 10Gbps means that we likewise increase the operating frequency of the signal driven on the 12 differential trace pairs connecting the Virtex 6 GTH output to the optical transmitter input pins. Increasing this operating frequency will increase the attenuation per unit length along the trace, while the FR-4 circuit board material is already near the top of its usage range. Increasing the line rate will also increase the degree of precision required in layout. The relative length of both traces in a differential pair will need to be matched to an even higher degree. The scale at which imperfections can contribute to signal degradation is scaled down correspondingly and will affect vias, trace direction changes (45 degree bends) or even the solder pads for decoupling capacitors will be having an increased impact on signal integrity. The layout would have to be modeled and fine tuned to a significantly higher degree, which will take time and resources away from the other aspects of board layout. This might translate into a delay in the production of the boards.

Requiring the CMX card to support 10Gbps output in addition to 6.4 Gbps output will probably also require an additional low jitter oscillator, as well as additional power supply voltages.

GTH transmitters also draw more current than GTX transceivers, namely ~4.25 W per 12 lanes of GTH at 10Gb (cf. ds152-p16 or 4.3W from XPE spreadsheet) versus ~1.5W per 12 lanes of GTX at 3Gb (cf. ds152-p9 or 1.9W at 5.5Gb 2.1W at 6.6Gb).

As was already described, 10Gb GTH outputs are not compatible with TP-CMX 6.6Gb inputs (but 5.5 Gb ok) and this might matter in some future use case if the propose 10Gb output is replacing one of the 2x12 GTX outputs.

5.3 Comparison aspects specific to the Single FPGA solution

Adding some 10Gbps Base-CMX outputs to the Single FPGA solution does not require a change for the choice of FPGA described above as the FF1923 package was already desired to maximize the number of inputs to the TP-CMX function.

Since there is no change in FPGA used, there is thus no change in FPGA part cost, and no additional pros or cons from what was described above.

5.4 Comparison aspects specific to the Dual FPGA solution

Adding 10Gbps capability to the Base-CMX functionality precludes using the FF1759 package which has only GTX. The Base-CMX FPGA thus would have to be based on a FF1923 package. The TP-CMX FPGA choice would not be affected anyway since the 10 Gbps output capability is relating only to the Base-CMX functionality.

Switching from the FF1759 to the FF1923 package also implies using more trace layers as was described earlier, as well as FPGAs using this package being more expensive than FPGAs using the FF1759 package (~\$8.5k instead of ~\$5.5k on Digikey website).

GTH resources require 2 additional power supply voltages (1.1 and 1.8V) for the Base-CMX FPGA but those voltages were already needed to support the GTH resources on the TP-CMX FPGA.

5.5 Cost comparison

	Option	Base-CMX FPGA	Base & TP FPGA	TP-CMX FPGA	Per Card	Total 15	Total 4+11	Delta 10Gb
1a	Base-CMX only 6.4Gb out	LX550T-FF1759			\$5.5k	\$83k		
1b	Base-CMX only with 10Gb out	LX565T-FF1923			\$8.5k	\$128k		+\$45k
2a	Single FPGA Base: 6.4Gb out TP: 5.5Gb in		LX565T-FF1923		\$8.5k	\$128k		
2b	Single FPGA Base: 10Gb out TP: 5.5Gb in		LX565T-FF1923		\$8.5k	\$128k		+\$0k
3a	Dual FPGA Base: 6.4Gb out TP: 5.5Gb in	LX550T-FF1759		HX565T-FF1924	\$14k	\$210k		
3b	Dual FPGA Base: 10Gb out TP: 5.5Gb in	LX565T-FF1923		HX565T-FF1924	\$17k	\$255k		+\$45k
4a	Two CMX Types Base: 6.4Gb out TP: 5.5Gb in	LX550T-FF1759		HX565T-FF1924	1a or 3a		\$117k	
4b	Two CMX Types Base: 10Gb out TP: 5.5Gb in	LX565T-FF1923		HX565T-FF1924	1b or 3b		\$161k	+\$45k

note: “10Gb out” corresponds to the maximum line rate achievable with an Avago AFBR-810B SNAP12-like transmitter (specified at up to 10.0 Gbps) or an Avago AFBR-81uVxyZ MiniPOD transmitter (specified at up to ~10.3 Gbps).

5.6 Conclusion

Even though adding 10Gbps output capability to the CMX card is attractive as it would bring additional options for future usages of the CMX cards, it would however automatically prevent using the simpler FF1759 Virtex 6 FPGA package, which would add more trace layers to the board, and would require a more complicated layout process to support these additional 10Gb traces.

Adding some 10Gb output capability to the CMX project would have a significant impact on the limited engineering resources available and would thus detract from producing and delivering the minimum required functionality on schedule.

The perceived benefits of requiring 10Gb output from the CMX thus seem to be outweighed by the associated risks.