HF Jet Trigger Upgrade R&:D Project

Drew Baden, Tullio Grassi, Jeremy Mans *University of Maryland*
Chris Tully, Dan Marlow *Princeton University*

Introduction

We investigate the possibility of adding functionality to the existing CMS Level 1 trigger for the LHC upgrade. The new functionality is primarily concerned with the trigger rate for the "W-Boson Fusion" (WBF) diagrams resulting in high energy back-to-back jets in the forward region (HF), and a Higgs decaying more centrally. Monte Carlo simulations indicate that an improvement in the Level 1 trigger efficiency can be obtained with only additions (no changes) to the existing hardware. We present the simulation results, and a path towards the new hardware. It is important to keep in mind that the current Level 1 acceptance rate of 100kHz means a rejection factor of only 400 in this trigger, and that this will be maintained in the upgrade.

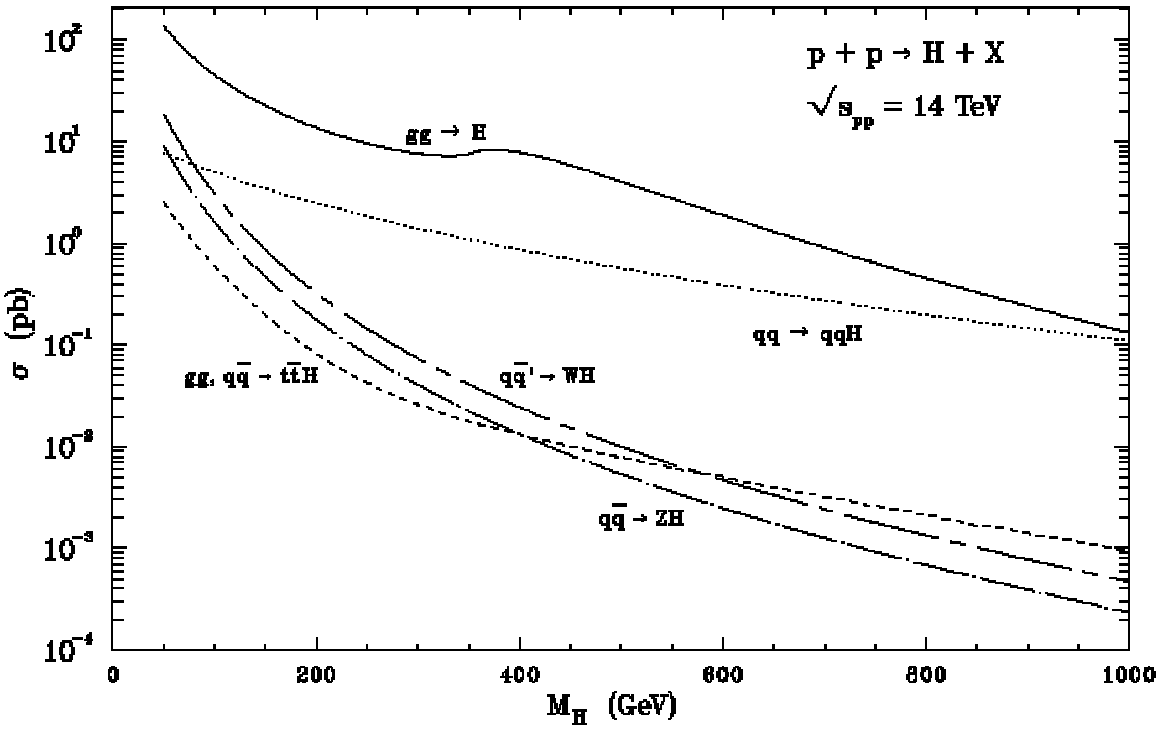

Figure 1 Higgs production cross-section at the LHC

As shown in Figure 1, the dominant mode of Higgs production is via the gluon fusion diagram via a top quark triangle. The next highest, down by a factor of ~10 at low Higgs masses, is via W boson fusion (WBF) where the radiating quark lines turn into jets in the forward/backward going direction. The additional jets can be tagged to enhance the Level 1 trigger efficiency of these events; unlike the gluon fusion process where there are no additional tags (other than ISR, similar for these two processes). The goal of this project is to try to enhance the Level 1 trigger efficiency, and thus lowering of the trigger threshold on the Higgs decay daughters, increasing signal/background.

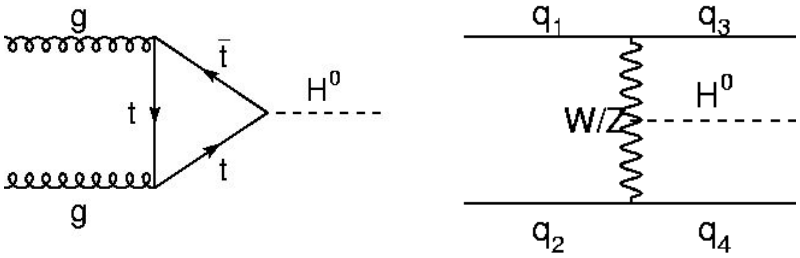

Figure 2: "gluon fusion" and "W-boson fusion" diagrams for Higgs production

HCAL Triggering, and the HF Detector

The HF detector covers the pseudo-rapidity region from $3<|\eta|<5$. The detector is made with steel absorber and radiation hard quartz fibers. Shower energy is measured from the Cerenkov light collected in 8-stage phototubes. From the output of the phototubes, the same electronics are used uniformly in HCAL.

The figure shows the HF segmentation. The detector has 11 towers in η with a ϕ modularity of 36 ($\Delta\phi$ = 10°, or 0.175 radians), and 2 innermost (highest η) towers with a ϕ modularity and 18. All towers have a segmentation of $\Delta\eta$ = 0.166. In each tower, a long and short fiber runs parallel to the z-axis. The short fiber in the rear (further from the interaction point) can be thought of as containing only hadronic energy, whereas the longer fiber, running along the entire detector, contains both hadronic and electromagnetic energy. The total number of HF channels per side is thus

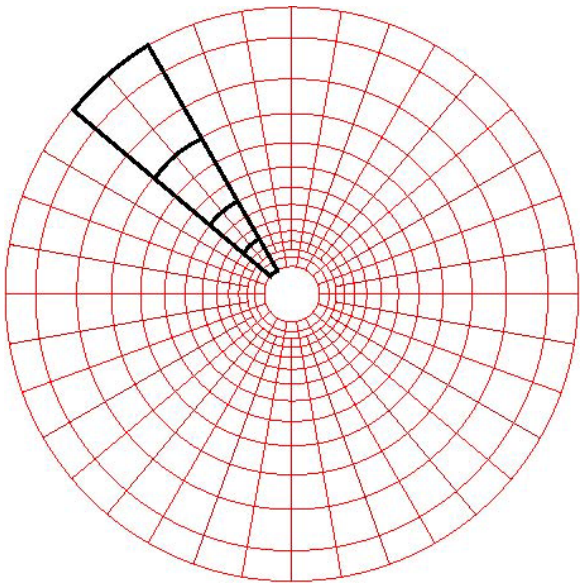

2(fibers)x[11(η)x36(ϕ)+2(η)x18(ϕ)]=2x432=864 channels. For most counting purposes, and for the rest of this document, we can think of the innermost 2 η rings with the 20° segmentation as a single η ring with a 10° segmentation, giving us virtually 12 η rings in with a ϕ segmentation of 36 (10°).

For the the existing Level 1 trigger, "Trigger Primitives" (TPGs) made from physical towers are sent to the Regional Calorimeter Trigger (RCT) receiver cards. For the HB/HE region, a TPG is the same as a physical tower in ηxϕ space ($\Delta\phi$ = 0.087 and $\Delta\eta$ = 0.087 in the HB, and $\Delta\eta$ = 0.09 and increasing to 0.15 through η band 27 in the HE). In the HF, the TPGs are made up of only the data from the long fibers. These TPGs are the result of a sum over 3ηx2ϕ physical towers, resulting in a width of $\Delta\phi$ = 20° (0.35) and

$\Delta\eta = 0.5$ (notice the heavy outline in the figure below). Table 1 summarizes the physical towers and the TPGs for the different regions, and the picture following shows the HCAL segmentation.

Tower #	η_{max}	$\Delta\eta$	$\Delta\phi$ .087=5°	TPG ($\Delta\eta$x $\Delta\phi$)	Cal Regions	Description
1	0.087	0.087	0.087	0.087 x 0.087	Region 1 0.000<η<0.348 $\Delta\eta$=0.348, $\Delta\phi$=0.348	HB
2	0.174					
3	0.261					
4	0.348					
5	0.435				Region 2 0.348<η<0.695 $\Delta\eta$=0.348, $\Delta\phi$=0.348	
6	0.522					
7	0.609					
8	0.695					
9	0.783				Region 3 0.695<η<1.044 $\Delta\eta$=0.348, $\Delta\phi$=0.348	
10	0.870					
11	0.957					
12	1.044					
13	1.131				Region 4 1.044<η<1.392 $\Delta\eta$=0.348, $\Delta\phi$=0.348	
14	1.218					
15	1.305					
16	1.392					HB/HE overlap
17	1.470	0.087	0.087	0.087 x 0.087	Region 5 1.392<η<1.740 $\Delta\eta$=0.348, $\Delta\phi$=0.348	HE
18	1.566					
19	1.653					
20	1.740					
21	1.830	0.09	0.174	0.090 x 0.087	Region 6 1.740<η<2.172 $\Delta\eta$=0.432, $\Delta\phi$=0.348	HE, split into 2 equal half energies in ϕ
22	1.930	0.1		0.100 x 0.087		
23	2.043	0.113		0.113 x 0.087		
24	2.172	0.129		0.129 x 0.087		
25	2.322	0.15		0.150 x 0.087	Region 7 2.172<η<3.0 $\Delta\eta$=0.828, $\Delta\phi$=0.348	
26	2.500	0.178		0.178 x 0.087		
27	2.650	0.15		0.150 x 0.087		
28	3.000	0.35		0.350 x 0.087		
29	3.167	0.167	0.174	0.348 x 0.5	Region 8 3.0<η<3.5 $\Delta\eta$=0.5, $\Delta\phi$=0.348	HF
30	3.333					
31	3.5					
32	3.667			0.348 x 0.5	Region 9 3.5<η<4.0 $\Delta\eta$=0.5, $\Delta\phi$=0.348	
33	3.833					
34	4.0					
35	4.167			0.348 x 0.5	Region 10 4.0<η<4.5 $\Delta\eta$=0.5, $\Delta\phi$=0.348	
36	4.333					
37	4.5					
38	4.667			0.348 x 0.5	Region 11 4.5<η<5.0 $\Delta\eta$=0.5, $\Delta\phi$=0.348	
39	4.833					
40	5.0					

Table 1 HCAL ηx ϕ details for physical towers, TPGs, and Calorimeter Regions

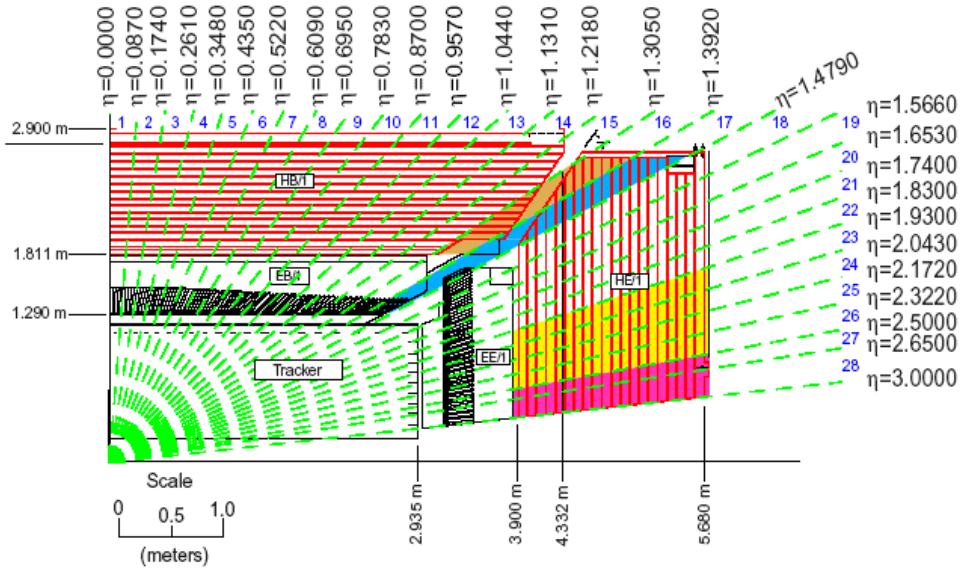

In Level 1, jets are formed from "Calorimeter Regions" (CR), defined as a 4x4 grouping of TPGs (each TPG is a yellow square in the picture) in the HB/HE but only a single TPG in the HF. These CRs (4x4 set of yellow squares) are sent to the Global Calorimeter Trigger (GCT) for jet determination using a "sliding window" algorithm. This algorithm considers each set of 3x3 CRs, and declares a jet to be present if the energy in the central CR is larger than in the 8 surrounding regions. The jet E_T is then the sum of the E_T in the 9 CRs, and the jet coordinates in η and ϕ are the coordinates of the central tower, chosen to be the CR with the highest E_T sum of the 9 CRs in the 3x3 window. Note that a 3x3 CR covers a square of $\Delta\eta = 1.05$ to 1.50 and $\Delta\phi = 1.05$ depending on location in CMS. Jets with diameters of order $\Delta R = 0.5$ should be very well contained in the CR, unless the jet appears right at the corner of an intersection of 4 CRs (a corner of one of the 4x4 yellow squares), where it would be spread out and have a lower efficiency of being captured. As part of the jet-finding algorithm, each CR has a bit which is set if more than 2 of the 4x4 TPGs in the CR are above a threshold. This bit is called the "τ-veto bit". Jets are labeled a "τ-jet" if none of the 9 CRs has a τ-veto bit set, and is only applied to the non-HF jet candidates. In the HF, the TPG "feature bit" will be set (in the HTR cards) if some fraction of the total E_T of the 6 contributing towers is contained in 1 tower. This will be used by the GCT in an attempt to distinguish narrow forward jets

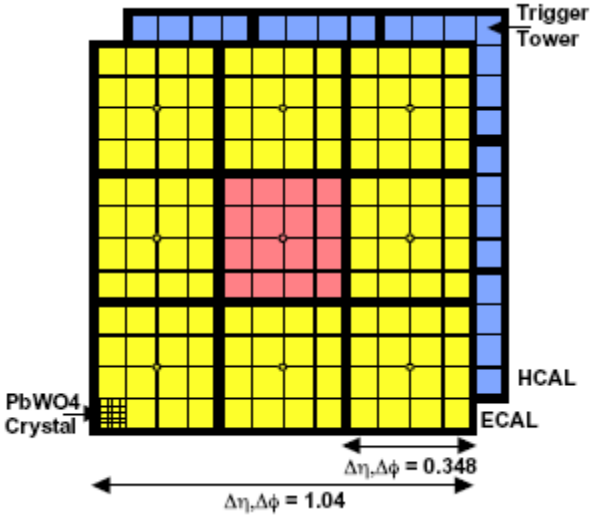

from general pileup of energy due to the large average number of interactions per crossing.

Simulation Studies

We have studied the performance of the feature bit and alternative definitions for the feature bit defined within the 3ηx2ϕ HF towers for luminosities up to $10^{35}cm^{-2}s^{-1}$. In all of the following plots, we assumed a 25ns RF structure at the upgraded luminosities, mostly for lack of any real guidance, but also for a worst-case scenario.

The discrimination power of the feature bit can be seen in the following figures where the ratio of the highest energy cell over the sum of the highest 2x2 sum is plotted for signal and background at a luminosity of $10^{34}cm^{-2}s^{-1}$ and $10^{35}cm^{-2}s^{-1}$. The shapes are not all that different as far as signal/background is concerned. However, at the high luminosities, one can clearly see the effect of the multiple interactions adding energy uniformly, and in order to preserve signal/background for the cut one would have raise the threshold up past 0.6 to get decent signal/background. We calculate that this would drop the efficiency of the signal by a factor of 4. The bottom line is that at high luminosities, the feature bits become almost useless. This is no surprise – at high luminosity, there is so much energy spread uniformly over the detector that single tower thresholds become much more difficult to use for discrimination purposes.

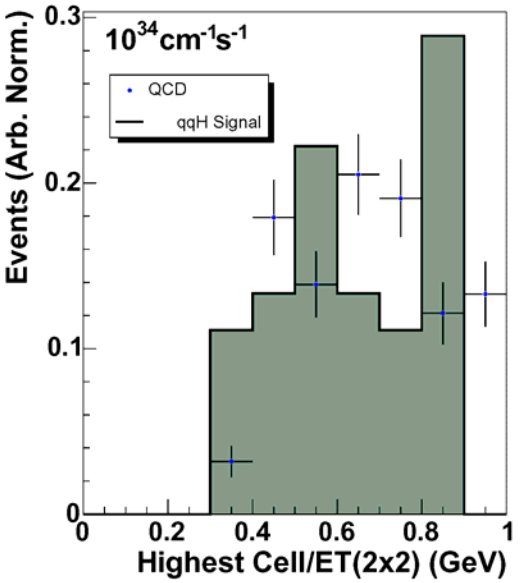

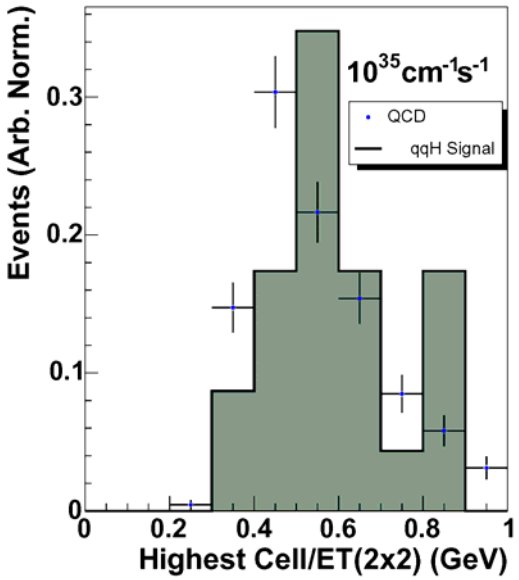

The additions to the trigger that we are proposing would enable us to trigger on narrow jets in the forward/backwards region with high efficiency and as low a background rate as possible, using the full granularity of the HF as opposed to the current trigger, which uses a granularity of 6 towers (3x2 in ηxϕ) for both TPGs and CRs. We propose to do this without any changes to the existing HCAL TriDAS hardware, however there will be

some extra boards to build, and some changes to the HTR firmware (firmware changes will have no latency impact). These proposed additions are explained below.

HF Jet Trigger

The goal for this trigger enhancement is to implement a real sliding window jet algorithm with an isolation measurement, similar to what CMS is now doing for the isolated e/γ component of Level 1. We propose to make use of the full granularity of the HF, specifically to:

- Construct jet candidates from sums of 4x4 towers in ηx ϕ resulting in a width of 0.67 x 0.70. This is compared to the current scheme, which uses 3x2 sums for CRs and 3x3 CRs for jet candidates, or 9x6 physical HF towers in ηx ϕ for jet candidates with a total width of 1.5 x 1.0 in ηx ϕ. All possible 4x4 sums will be constructed, and since the HF has a ϕ segmentation of 36, there would be 36 such sums in ϕ. Also since the HF has 12 η rings, there would be 9 such sums in η (12-4 + 1) which include only towers in HF (we have not yet considering adding towers from HE, but this will have to be done eventually).
- Construct a feature bit that computes the narrowness of the jet based on all of the cells contributing to the jet candidate. The most effective quantity found in this study was the minimum number of cells that are needed to contain some fraction (we used 60% for our studies) of the energy of the 4x4 sum, denoted as n60.
- Require that the jet-candidate be above an E_T threshold, and require that n60<7 for these studies. Any jet-candidate passing this cut is considered an isolated forward jet.

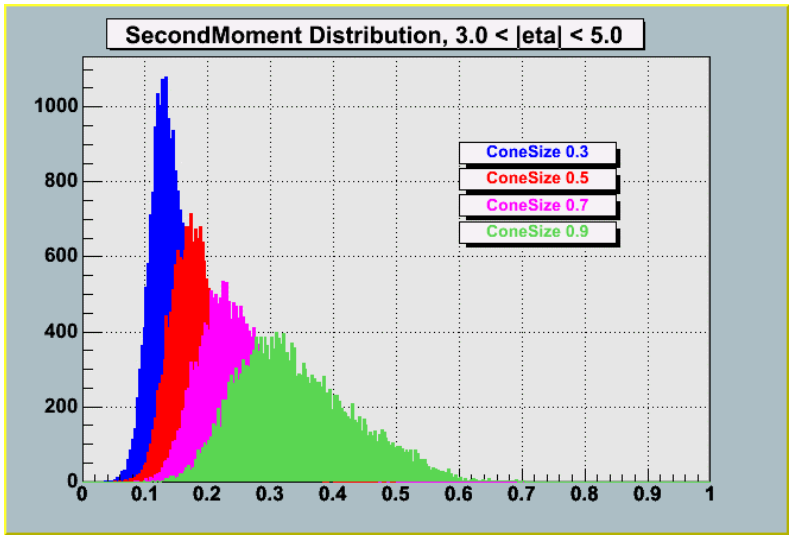

With this algorithm, we are able to capture any isolated jet with a width that fits inside our 4x4 sum. It is interesting to consider the size of this 4x4 area, which is 0.67 x 0.70 in ηx ϕ. At the LHC where jets have enormous energy compared to the jet mass, one would expect the jet 2nd moment $\delta R^2 = \delta\eta^2 + \delta\phi^2$ to be rather small. In the figure, we see results from a simulation of forward jets from QCD events with min bias overlap. The figure shows the 2nd moment (δR as defined above) for jets from different cone clustering algorithms. The cone size of 0.5 (a common choice) shows that a vast majority of all jets have a second moment less than 0.3, which is about the radius of our 4x4 tower sum, showing that this sum describes the core of the jet. The possibility of extending the sum

and the n60 computation to the size of the current GCT jets (9x6) is included in the design.

The discrimination power of the feature bit in the new design can be seen in the following figures where the narrowness of the jet is estimated by the minimum number of cells that are needed to contain 60% of the energy of the 4x4 sum (n60). A cut of n60<7 has a high efficiency for forward jets at luminosities of $10^{34}cm^{-2}s^{-1}$ and $10^{35}cm^{-2}s^{-1}$. At a luminosity of $10^{35}cm^{-2}s^{-1}$, the efficiency is above 80% with a background rejection of a factor of 3.

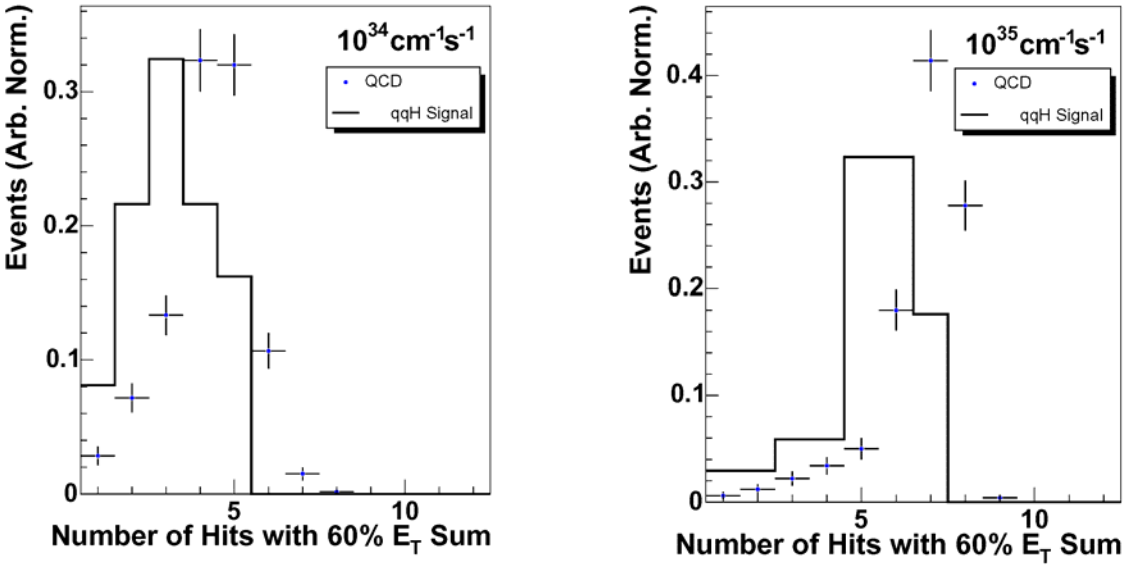

The simulation studies show that the addition of the forward jet trigger term to a central dilepton trigger term is stable with luminosity when the sliding window 4x4 sum is applied with the n60<7 cut. The following table shows the percentage the events accepted by central τ-τ, τ-e, and e-e triggers and the logical OR of these triggers for qqH events, m_H=130 GeV/c^2 and H→WW→ℓνℓν. The percentage of QCD events passing the trigger is also listed. The QCD events are taken from a uniform sampling from 30 GeV to 300 GeV. At $10^{35}cm^{-2}s^{-1}$ the τ fake-rate is substantial, and the E_T threshold is increased to reduce the rate. The E_T thresholds are on reconstructed quantities.

Lumens (/cm^2/s)		τ-τ	τ-e	e-e	OR
10^{34}	Signal	12% τ E_T>30	33% τ E_T>30, e E_T>20	11% e E_T>20	35%
	QCD	11%	10%	4%	20%
10^{35}	Signal	85% τ E_T>50	34% τ E_T>50, e E_T>20	4% e E_T>20	90%
	QCD	100%	17%	4%	100%

Percentage of Triggered qqH and QCD Events by the Central Dilepton Trigger

In addition, we list the percentage of events when we require an additional L1 forward jet from the current trigger system and the corresponding number from the forward jet term from the new design. The signal purity is defined in this case to be the percentage of triggered signal events where the trigger forward jet object corresponds to the parton-level forward jet tag.

Lumens ($/cm^2/s$)		OR-FJT (E_T>20 GeV)	OR-NewFJT (E_T>10 GeV)
10^{34}	Signal Purity	59%	65%
	Signal Eff.	13%	12%
	QCD	49%	30%
10^{35}	Signal Purity	41%	68%
	Signal Eff.	37%	22%
	QCD	100%	30%

Percentage of Triggered QCD Background, Signal Efficiency and Purity. The left column is for the current trigger system, and the right column for the new design.

At a luminosity of $10^{34}cm^{-2}s^{-1}$, the signal efficiencies and purities are comparable between the old and new design, while the new design has a higher QCD background rejection factor. At $10^{35}cm^{-2}s^{-1}$, there is a substantial difference in performance. The new design has a stable QCD rejection factor and stable and high signal purities and efficiencies, as shown in the right column. For fixed forward-jet E_T thresholds, the signal efficiency increases due to pile-up events increasing the reconstructed forward jet E_T values. In contrast, the current trigger fails to have any rejection for QCD and drops by 50% in signal purity. This effect is due primarily from an ineffective feature bit. Many feature bits were defined to try to make the current trigger system more effective, but the hardware limitation of the tower size (3x2 in ηx ϕ) is too restrictive to adequately characterize a narrow forward jet. The new design makes use of all the cells in a region large enough to contain the core of the forward jet, and the region slides on a cell-by-cell granularity to narrowly define the location of the jet shower maximum. This provides a more powerful feature bit and a sharper turn-on curve. Based on these studies, CMS will not have forward jet triggering capabilities for the second highest Higgs boson production cross section at luminosities of $10^{35}cm^{-2}s^{-1}$ without an upgrade of the trigger system as described here.

Design Concept

The following figure shows a schematic for the HTR boards. Note the 6 daughterboard sites labeled “SLB”. SLB stands for Serial Link Boards, and are responsible for the sending of TPGs to the RCT such that all TPGs from the same interaction arrive at the RCT at the same time (same 25ns clock tick). These SLB boards are produced by the LIP group.

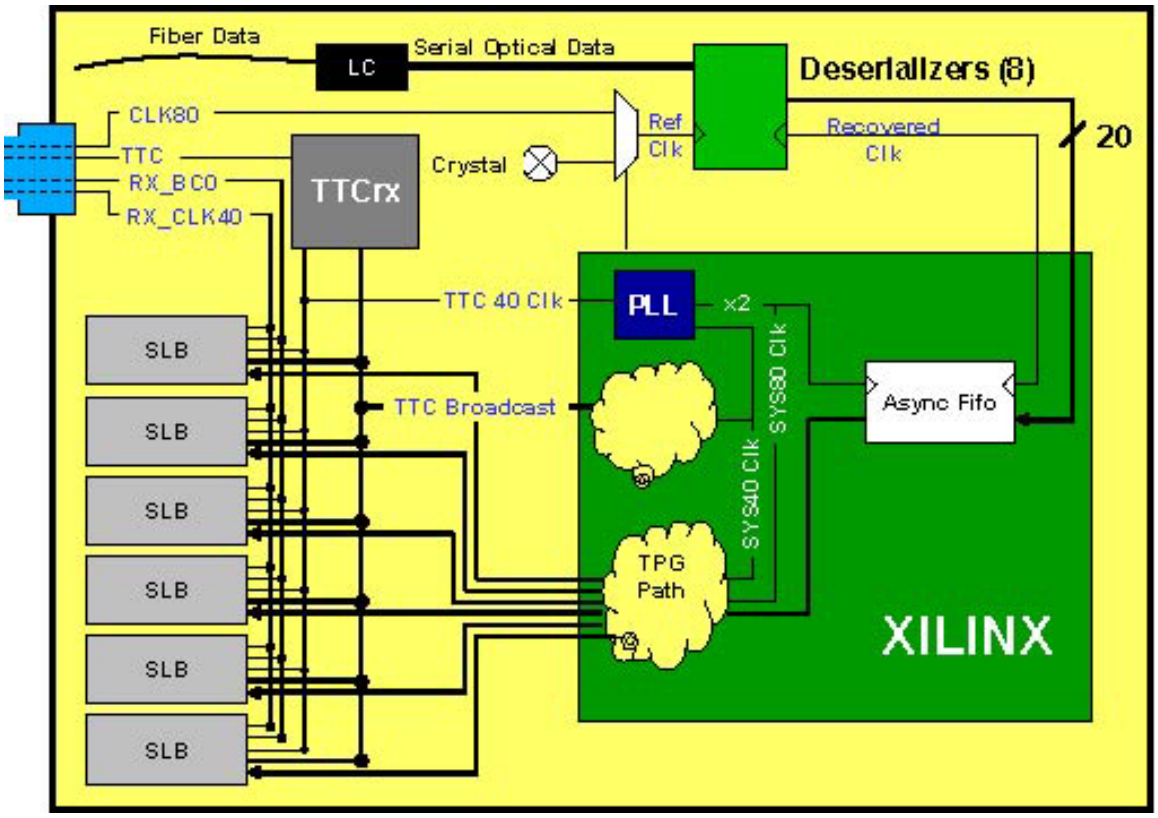

Each HTR card has 2 fiber ribbons worth of data coming in, with 8 fibers per ribbon. Each fiber has 3 QIEs worth of data, for a total of 2x8x3=48 QIE channels per HTR input. In the HB and HE, each QIE (each HB tower) is turned into a TPG and sent out the SLB to the RCT. There are 6 SLBs per HTR, which means that each SLB delivers 8 TPGs to Level 1 for HB and HE. In the HF, however, the TPG is made from a sum of 6 HF long fiber towers (3 in η, 2 in ϕ), so the number of TPGs coming out of the HF HTR boards is only 8, therefore these boards will only use 1 SLB. We propose to take advantage of the 5 empty SLB sites to house a new daughterboard that will send data to an HF jet trigger, described above.

Since each HF HTR will have 5 empty SLB slots, we propose to build 2 additional daughterboards (for both the HF jet trigger, and for the HF luminosity determination, described elsewhere):

1. a 4-SLB-slot board to collect data from physical HF towers to send to a processing board to implement the HF jet tagging trigger
2. a 1-SLB-slot board to send information for instantaneous luminosity determination

To be sure that we have enough pins on the existing HTR to implement the HF jet tagging, we have to consider how the processing would occur, and what data is needed. Each HTR participating in the jet trigger will be responsible for receiving 48 HF long

fibers, and will be cabled so that the data belong to 4 consecutive ϕ bins and all 12 η bins. The figure below shows that each side of the detector (HF+ and HF–) will be serviced by 9 HTR boards.

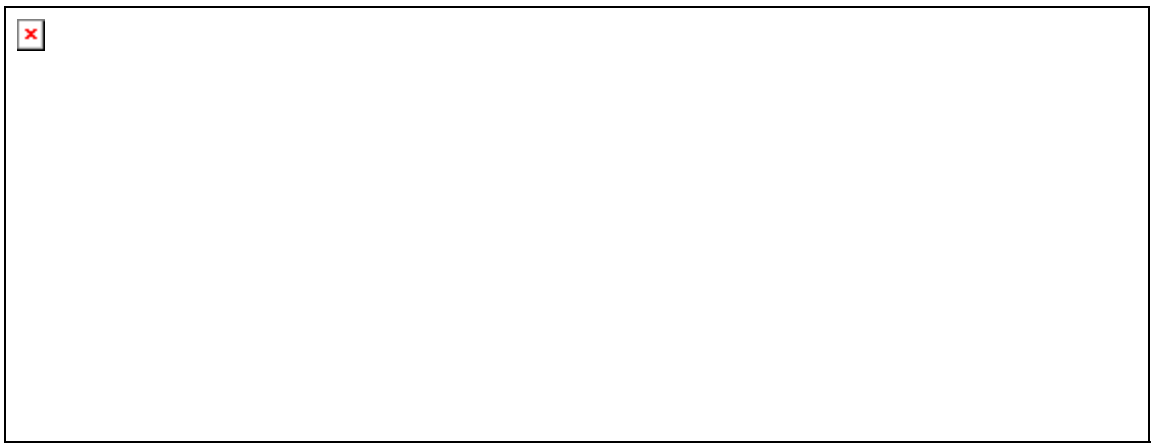

The goal will be to study implementing a 4x4 sum for the jet, with an extended 6x6 isolation window. Other possible configurations can also be considered, and studied. The isolation window will give the flexibility to extend the range of the sum and/or the feature bit calculation to the full range of the current GCT jet.

Sliding Window Jet Finding

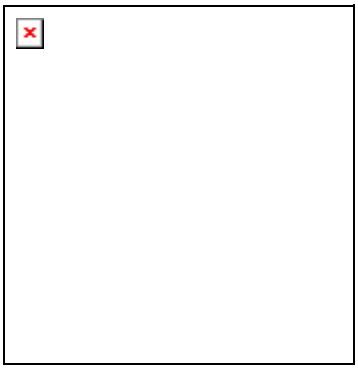

In order to implement a full 4x4 sliding window algorithm, we need to consider the region on the "edges" of each HTR. In the figure we see 3 HTR cards worth of data (12 bins in η, 4 bins in ϕ per HTR card). The heavy dashed lines show the "edges" of the HTR cards. Considering the HTR in the center, we will need to form 4x4 candidates from the data sent by this card. The solid lines show the 5 possibilities for forming combinations of 4 towers in consecutive ϕ using a "half-HTR" worth of data for the "left" and "right" edges. To illustrate better, consider the next figure where we see two 4x4 sum surrounded by gray ϕ-slices on either side for the isolation window. The upper 4x4 sum is made from towers from a single HTR, whereas the bottom 4x4 sum have half towers from 2 separate HTRs. The HTR that contributes to the sum from the left side of the figure also sends its data to the board that forms sums from these same towers, analogous to the 4x4 sum at the top of the diagram. We therefore conclude that each HTR will have to send its data to 2 different sources, doubling the number of cables needed to send data to a single source. Note, however, that these two cables will be sending in parallel. Therefore the IO bandwidth needed has to accommodate 48 towers, but the number of cables

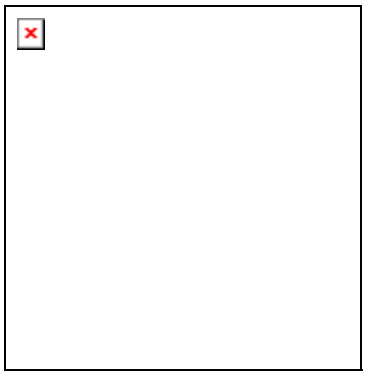

is double that.

Bandwidth and Cabling Requirements

The daughterboard that will transmit this data will occupy, as said above, 4 SLB slots. We call this daughterboard the HFT (for HF Tower) card for this discussion. We are assuming that each HF tower will need to send a single byte of data to the clustering circuits. To have enough lines from the existing HTR FPGA, we will have to send data to the HFT card at twice the current rate of 40MHz. We believe that the HTR board will be able to do this – a lot of care has been put into the layout of the existing board in order to be able to send data at 40MHz with a large margin for jitter. This assumption will be tested with the current Rev4 boards.

It would be convenient if we could send the data over standard cat6 or cat7 quad twisted pair using the same Ethernet 8B/10B protocols that we use now. The first consideration is the length of the cables. We are proposing to the CMS rack/installation group that we add another VME crate to the rack on the HCAL aisle in the middle of the row. This would be rack 5, and would contain the electronics for the HF jet processing cards and the luminosity measurement electronics (discussed elsewhere). This crate will be in a rack that has no more than 2 racks between it and the farthest existing HF rack. The figure below shows the proposal as submitted in July of 2004 and was later approved. The crate marked "luminosity" will house the receiver electronics. The distance from the HF crate in row 2 (or 8) to the crate in row 5 will be no more than 4 crates wide. Each crate is about 0.5 m, so the total horizontal distance will be about 2m. If we put the luminosity crate at the same level as the HF crates in rows 2 and 8, then we will minimize the vertical climb. If the vertical climb is no more than 2m up and 2m down, then altogether, we believe we can route whatever cables we need in 6m lengths. This sets the scale for how fast we can transmit using current technological experience (high speed differential links over cat6/7 twisted pair).

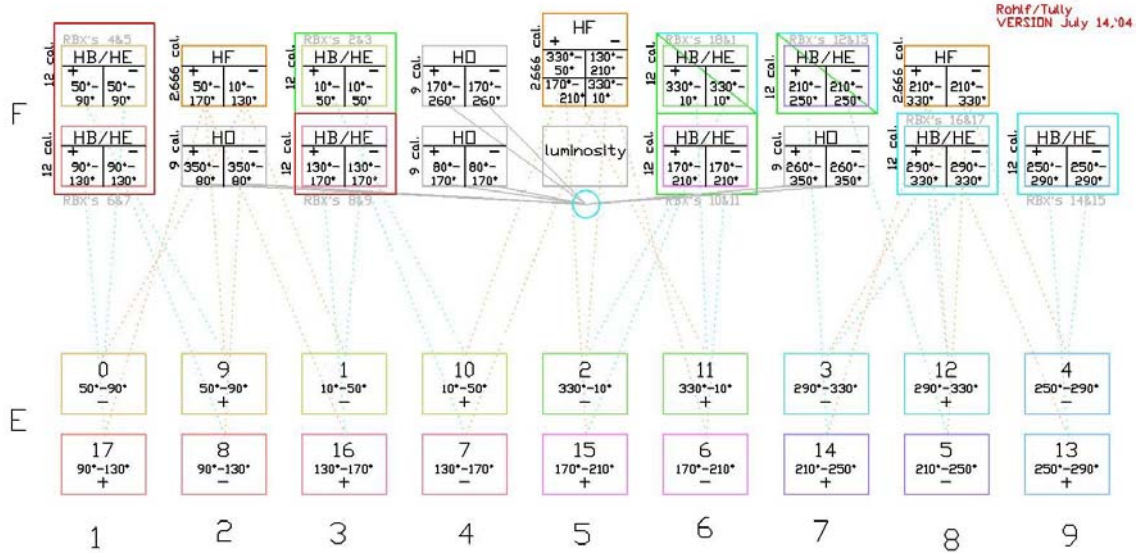

The next thing to consider is the link speed, and the number of cables per HTR going out to the receiver boards that will form the sums. In the existing barrel system (HB) where all 48 towers are transmitted to the trigger as TPGs, the SLBs use a Vitesse transceiver which transmits a single byte at 120MHz per link, or 1.2Gbps using the 8B/10B protocol (10 bit frames at 120 MHz frame rate). Each SLB has a single quad twisted pair cable going to Level 1, and there are 6 SLBs per HTR for a total of 6 cables running at 1.2Gbps over a distance of 10m. In this new system, we will have the same number of channels, but the cables will be limited to a card that will span 4-SLB sites as opposed to 6. On the HTRs, 4 SLBs take up approximately 15.5cm of space. The Princeton fanout card has 2 8-way RJ45 connectors, each taking up 11.5cm of space, which will fit nicely into the 4-SLB site HFT. If we could send the data from all 48 channels on 4 of the quad twisted pair cables, then we can easily fit the required 8 connectors onto the HFT card. The bandwidth calculation is then straightforward:

1. 4 cables means 16 twisted pair
2. 48 channels per 16 twisted pair mean 3 channels/pair at 40MHz
3. 1 byte per channel with no hamming codes means 3 bytes @ 40MHz or 120MByte/sec or 960Mbps (8 bit bytes)
4. If we add hamming codes, that will increase the number of bits from 24 to 29, or 32 if we round off. The required data bandwidth will be 4 bytes @ 40MHz or 1280 Mbps. The number of bits per second is given by: 4 byte/frame x 10 bits/byte x 40MHz = 1.6Gbps.

The current Level 1 system uses Vitesse VSC7216 transceivers which can run up to 1088 Mbps data links (1.36 Gbps overall). This is slightly below what we would need if we were to add hamming code data, however this is an older chip and there are new ones on the market. In fact, we would probably use the Xilinx FPGAs which have built-in gigabit transmitters. The input clock will be the high precision 40MHz clock that is cleaned up by the QPLL, and the Xilinx has internal DLLs for the serializers. For the receiver

section, we could use the current HTR deserializers by TI, which we know how to work at 1.6Gbps. The R&D needed here is to test a serial link at this speed running over a ~6m copper link.

The QPLL has an intrinsic jitter measured to be around 20ps RMS. An alternate data transmission scheme would consist of a fixed (nonLHC) crystal oscillator frequency that is greater than the LHC frequency, but stable to better than 20ps (easy to do) on both the transmitter and receiver. This would necessitate asynchronous FIFOs on both the sending and receiver side, with appropriate logic, to ensure that data is synchronized to the LHC frequency. The only drawback to this scheme would be that it would incur some latency. However, keep in mind that the latency for the jet clustering is equal to the total latency of the RCT + GCT combined, since it is in the GCT that jets are assembled and sent to the global trigger for processing. This latency is around 40 clock ticks total, a very large amount of time compared to the ~12 clock ticks taken up by the entire HTR firmware as it exists now.

In summary, the HFT will have 2 sets of 4 RJ45 connectors, with each RJ45 connected to a quad twisted pair cable. The data on each of the quad links will be running at 1.6 Gbps over a distance of about 6m, using the QPLL clock or a crystal oscillator with asynchronous FIFOs for the frame clock.

Receiver Board Design

Our preliminary design for the receiver board (HF jet board, or HFJ) is straightforward – all of the links come in on the same RJ45 connectors, is deserialized, and fanned out to the FPGAs. Each twisted pair would need its own deserializer, so there would have to be 4 per quad cable. There are 9 HTRs servicing HF+ and 9 for HF−. If each HFJ board serviced 1/3 of the HF per side, then we would need 3 HFJs for HF+ and 3 for HF− for a total of 6 boards. If each HTR has 8 output cables, and there are 9 HTRs per side, then the total number of cables is 72, or 24 per HFJ. This would probably necessitate a double width VME board, which would take up 6 slots per HF side, or 12 VME slots (out of 21 total). We anticipate using a single VME crate, so slots will be a critical factor.

We want to minimize the number of FPGAs per HFJ board. Each of the 24 cables coming into each HFJ will have 4 signals, or 96 total serial data streams, which means that we need 96 deserializers per HFJ. The current HTR boards use 16, which means a factor of 6 more per HTR, but these chips are small (roughly 1cm square) and 96 is probably doable. Each of the 96 deserializers would have roughly 20 pins of output, or 1920 pins total. If we use 3 FPGAs/HFJ, then each FPGA would need to service 32 inputs, and have 640 pins just for the output of the deserializers. There are many choices for FPGAs that would meet this requirement, e.g. the XC2V4000 in the flip-chip package has 912 user I/O pins and the XC2V6000 has 1104 user I/O pins. These chips are from the Virtex-2 family, a relatively old technology, and would probably cost less than $1000/chip. If there are 3 per HFJ, and 6 HFJ for the entire project, then that amounts to only $18k worth of FPGA chips.

It might also be possible to use the newer Virtex-2 Pro chips with the build-in deserializers. Xilinx already has announced a chip with 24 deserializers, so perhaps one with 32 is on the horizon, or we could try to use 6 FPGA/board instead of 3 and eliminate the discrete deserializers. Note that the latency for deserializers is not small – current estimates are that it is no smaller than about 5 frame clocks. This option will have to be seriously studied, as it would make the board layout of the HFJ quite a bit simpler if the deserializers were incorporated into the FPGA. However, the point here is that a conservative board using the technology that we have been working with already would be buildable.

Output to Global Trigger

We envision a single VME board that would collect the jet candidates from the 6 HFJ board, and format and transmit to the Global Trigger. Transmission into this board can be over high speed LVDS on very short cables since this board will live in the same VME crate, and we believe that this board can be limited to a single-width VME board with a single FPGA. We have not considered any further details in the design of this board at this time.

Cost and Schedule

The project would consist of building the following boards:

1. 18 HFT boards (these are the 4-SLB slot daughterboards on each of the 18 HF HTR boards for the long fibers)
2. 6 HFJ boards (double-width VME)
3. 1 HF jet trigger board to collect the jet candidates and transmit to the global trigger

This is a modest amount of hardware, and will surely not cost more than $5k/board. We believe the total hardware costs of the project will be in the $100k range.

It will be important to have good engineering at the beginning of this project. Tullio Grassi is the obvious candidate. Since he has other HCAL responsibilities, we believe that we would need about half of his yearly salary, which comes to around $80k/year (including overhead and benefits). We would need other engineering resources for this project, and this is difficult to estimate at this time but probably in the same range. Total engineering costs would be in the $150k/year range, covering all engineering, in the peak years.

The goal would be to have something running as close to the 2008 startup date as possible. We believe that if we start designing and testing in 2005, after the HTR production, we might would have a good chance at success.