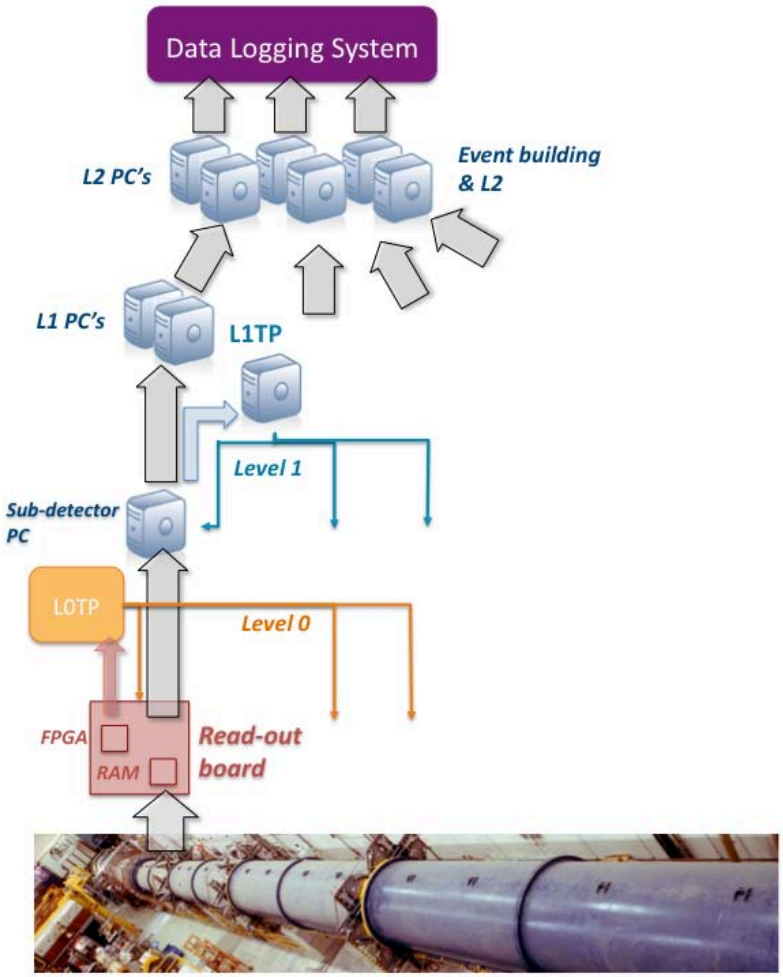

- *Figure 1 Logical scheme of the data-flow, from the detector to the permanent storage.*

1.1 Online Computing System

Data processing model

The data processing model is strongly connected with the Trigger & DAQ architecture, described in the previous sections. The general scheme of the data-flow, depicted in Figure 1 is summarized in the following.

The flow of data, from the sub-detectors to the permanent data storage, can be schematically summarized in the following steps

- The **readout boards** – for most sub-detectors the TEL62 – get data from the front-end electronics of the detectors, store the information in buffer memories (RAM), and (if the sub-detector is contributing to the L0 trigger) generate Level 0 Trigger primitives, as shown in Figure 2 the GTK is indeed not partecipating to the L0 trigger and has its own readout system. The LKr has a different data-flow, due to the large amount of data produced by the 40 MHz flash ADCs. The calorimeter readout modules (CREAM) are capable of holding events without (or with a very loose) zero-suppression in large RAM memories, that are downloaded to the online farm only on a positive Level-1 trigger decision. In order to partecipate to the Level-0

trigger, reduced granularity information, coming from a separate readout is used, based on sums of super-cells continuously digitized and pipelined in order to perform cluster-finding and photon-counting algorithms and send information to the Level 0 Trigger Processor (L0TP).

- The **Level 0 Trigger Processor** (**L0TP**) elaborates the L0 primitives from different sub-detectors, and produces the L0 trigger decision.
- The L0 positive decision is distributed to all the readout boards, and when this is received, the data are transferred to dedicated PCs (**sub-detector PCs**).

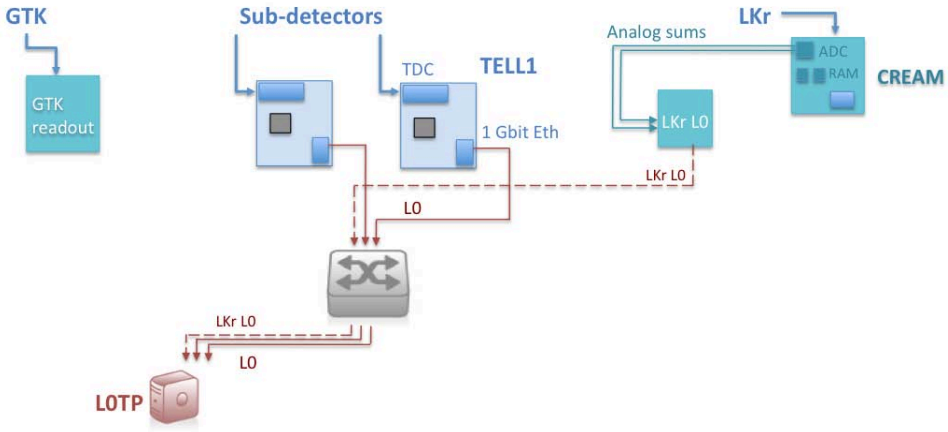

Figure 2 Logical scheme of the data-flow from the detectors to the readout boards and of L0 primitives to the L0 trigger processor.

This is shown in Figure 3 where the flow of data packets is shown, from different readout boards of a given sub-detector until they reach a PC in the L1 farm, that will assemble all the fragments of a given event for the sub-detector in a “sub-event”.

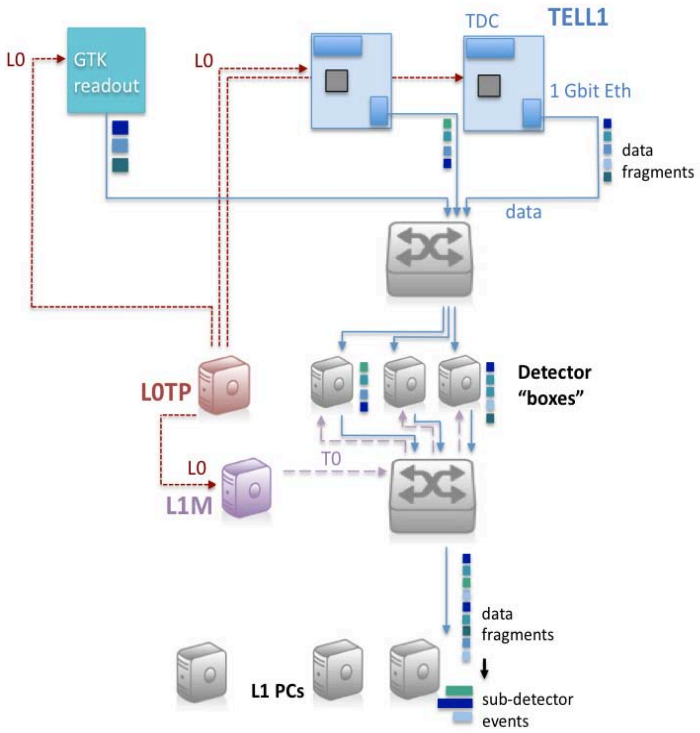

Figure 3 Logical scheme of the data-flow from the readout boards to the L1 farm, upon arrival of positive L0 decision.

The "detector boxes" shown in such figure can be either ordinary PCs or intelligent network boards dispatching the data fragments of a given event all to a single L1 PC, according to an algorithm (that could even be as simple as a round-robin table) produced by a Level-1 manager (L1M) machine. This task can even be performed by the readout boards (*e.g.* the TEL62 credit-card-PC) or by the main switch (placed in the electronics barrack and connecting all the readout boards of all sub-detectors) using IP packet inspection and redefinition of the final address of the data fragments.

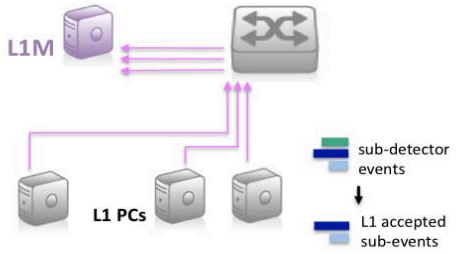

Figure 4 Logical scheme of assembly of sub-events in the L1 farm and building of L1 decision in the L1 Manager (L1 Trigger Processor).

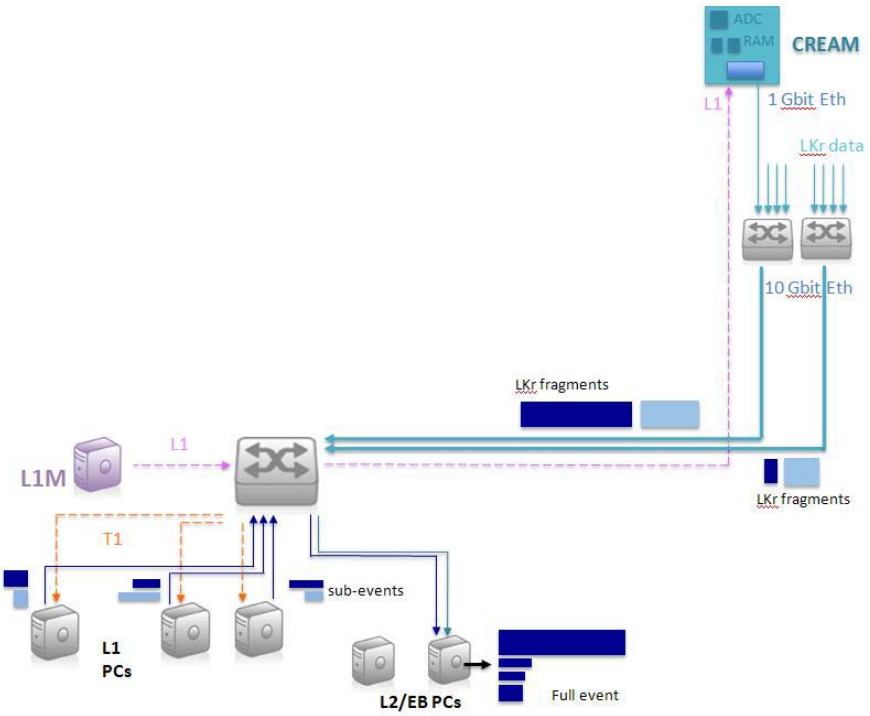

Figure 5 Logical scheme of the L1 distribution, data-flow from the L1 to the L2 machines and full readout of LKr data from the CREAM modules memories to the L2 farm.

- The **Level 1** software trigger is based on primitives produced by the sub-detector PCs, each handling the complete event data from a single sub-detector; the **L1M** trigger manager (or L1 Trigger Processor) will then merge those primitives in order to assert a Level 1 decision, as shown in Figure 4.
- The L1 trigger decision is distributed to the sub-detector PCs and (assuming a positive decision) sub-event data (L1 accepted sub-events) is transferred to the **Event Builder PCs.** These processes, running on part of the online-farm machines (labelled **L2/EB PCs** in Figure 5, assemble full events by merging together all the information coming from the different sub-detectors. Data from the CREAM (calorimeter readout modules) are also downloaded on a positive L1 decision, from the buffers onto the appropriate L2 machine. Once events are fully built, a **Level 2** trigger can be evaluated and issued.

- At this point, L2-accepted events can be transferred to the Data Logging System (DLS), possibly running remotely (e.g. in the CERN main data-center or even off-site), in order to be registered on a permanent storage (e.g. tapes), as schematically shown in Figure 6.

Putting together all the logical steps of the trigger and event-building, the general picture of the data-flow is the one shown in Figure 7.

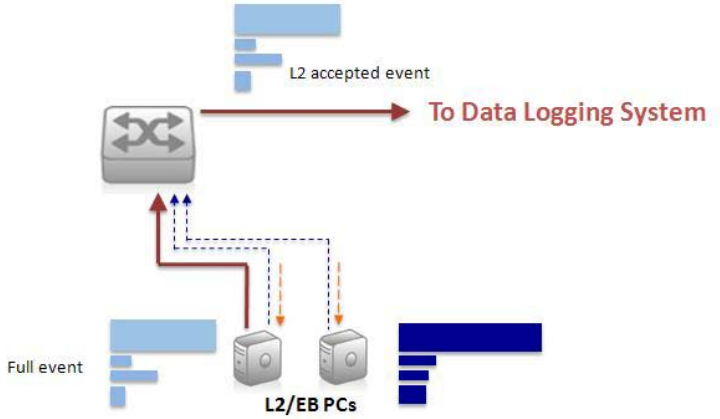

Figure 6 Logical scheme of the data-flow of fully reconstructed and L2-selected events from the L2 farm to the Data Logging System.

Event data model

The Event Data Model defines the data formats for the experiment. We can identify at least two basic data formats:

- **RAW**: the binary data readout by the different sub-detectors and assembled by the event-building processes in the on-line farm. It contains all the information produced by the sub-detectors electronics and the trigger modules, possibly pre-processed by dedicated processors (*e.g.* the Liquid Krypton calorimeter system farm) and filtered by the Level 2 trigger. The size of the event is estimated on the basis of the average number of active channels in each sub-detector and on the number of words needed to store the data of each channel, that we generically call **element**: an element is a single hit in the spectrometer chambers, a fired cell in the Liquid Krypton calorimeter, a fired pixel in one station of the GTK, etc. In general, the minimal information for one element is given by a time measurement word.
- **RECO**: the output of the reconstruction of the elements: tracks from the tracker hits, clusters from the calorimeters cells, etc.

Other data formats will be needed by the experiment, in particular we can identify at least the following additional formats:

- **THIN**: summary data extracted from the fully reconstructed events, easier to handle by analysis programs, complemented with meta-data coming from the conditions data-base and calibration tasks; and possibly even more skimmed events, that we can call **Super-THIN**, that can be used for specific analysis purposes.

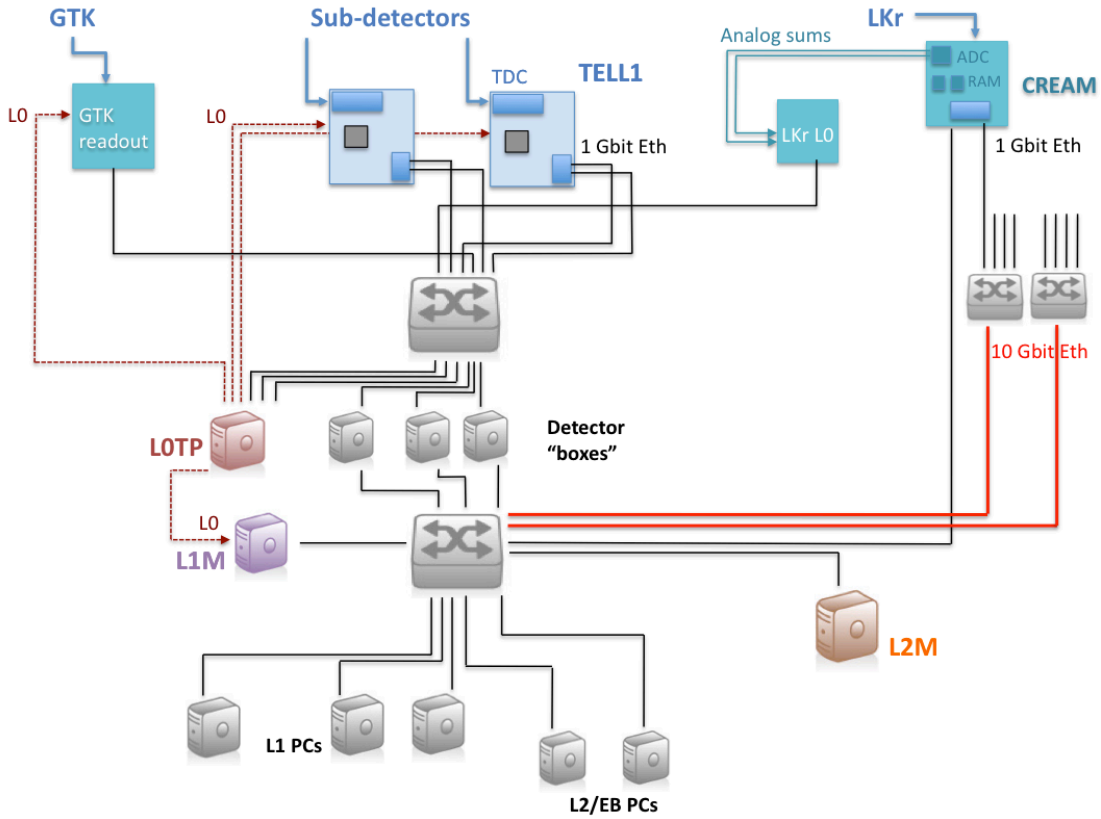

Figure 7 Logical scheme of the data-flow, trigger distribution and online farm.

- **CALB**: data from single sub-detectors or group of sub-detectors, possibly produced by earlier stages of the trigger (*e.g.* event-fragments from the output of the Level 1 trigger), needed for monitoring and calibration purposes. A typical example are events flagged as containing muon candidates by the Level 1 trigger of the photon or muon veto detectors, and partially reconstructed in order to perform monitoring with straight muons.

In addition to data coming from the readout of the sub-detectors, meta-data should also be considered. The three main types of meta-data produced by the experiment will be:

- **COND**: conditions of the detector, stored in the appropriate data-base, connected to the Detector Control System (DCS), such as high- and low-voltage settings of the sub-detector elements, temperature and pressure readouts, status of the readout boards and crates, etc.
- **CONF**: configurations of the run, including active detectors, trigger configurations, beam conditions, etc.
- **COST**: calibration parameters computed by calibration tasks running on (partially or fully) reconstructed data.

Event reconstruction model

The event data model is strongly connected with the model for event reconstruction: they are schematically shown in Figure 8.

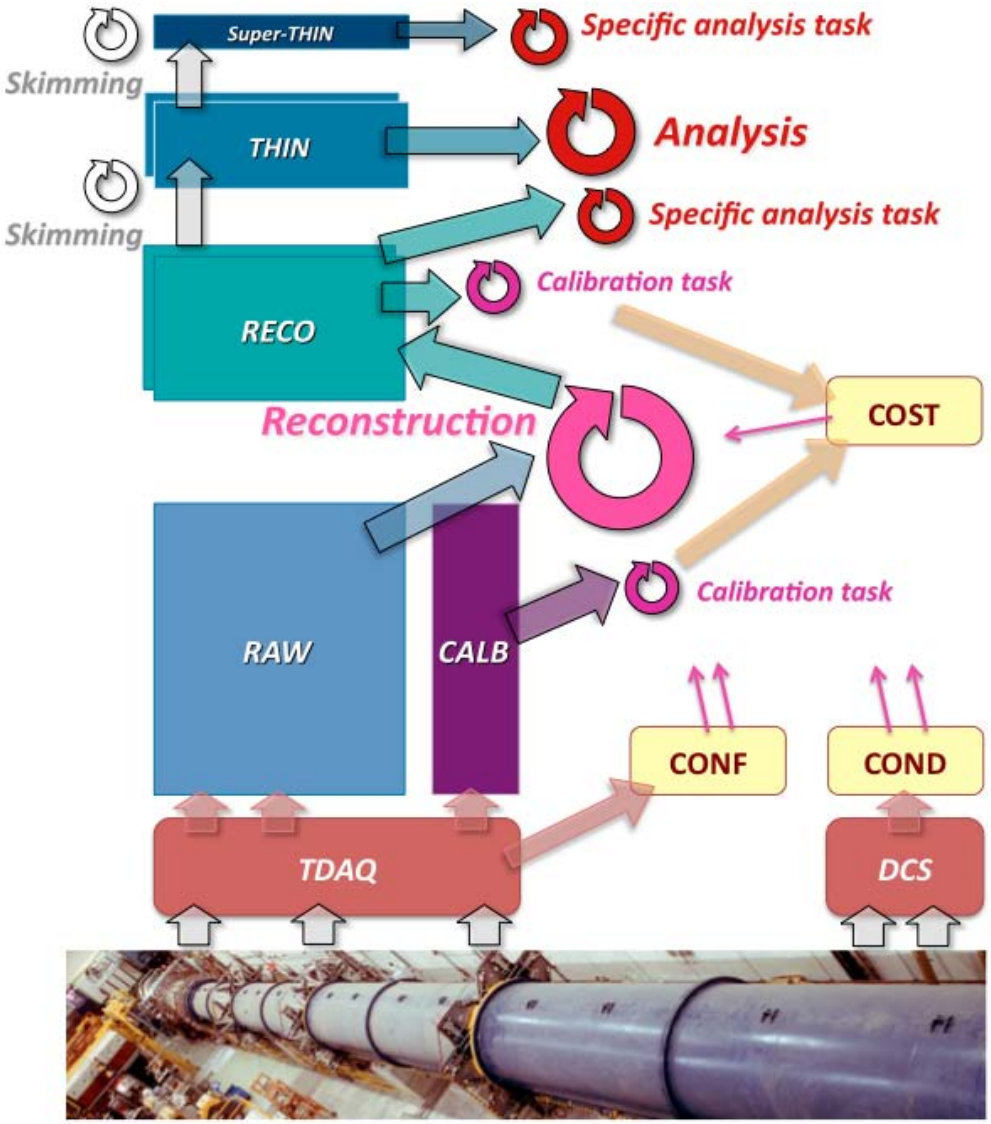

Figure 8 Event data model and event reconstruction model.

The Event Reconstruction Model indeed defines the architecture, both from the logical and implementation points of view, with which data coming from the sub-detectors are first assembled (event-building, EB), reconstructed and possibly filtered (for L2 trigger purposes or for the generation of data streams), and finally permanently stored by the data-logging system (DLS). The model is strictly interconnected with the trigger scheme of the experiment, since some of elaboration and filtering of the partial or full events can be done at the earlier stages of the trigger and EB.

The data coming from the readout of the sub-detector are in general stored in buffer memories until a positive Level 0 trigger decision is taken by the Level 0 trigger processor or the L0 trigger latency has elapsed; in the former case the readout boards send data packets to the sub-detector/L1 PCs, where a first, rough reconstruction of the sub-detector is performed. The maximum rate at the output of the L0 is defined to be **1 MHz**.

The implementation of the Level 1 will be fully in software, and the machines receiving the data and running the L1 reconstruction and decision will constitute the Level 1 farm. The L1 farm will be connected to the readout boards through Gigabit Ethernet links. In order to assemble the data fragments coming from different readout boards of a sub-detector in a given PC, the data packets from the readout boards belonging to a specific event should be routed to the same machine. This can be done assigning the destination IP address of the packets on the fly on a event-by-event basis (*e.g.* using a round-robin mechanism and a look-up-table for the sub-detector PCs to be used by the readout boards). An alternative scheme can be that of concentrating the links from the readout and to the L1 processors into a dedicated switch, handling the packet traffic. At the output of the L1 stage, we consider a maximum rate of **100 kHz**.

The implementation of the L1 farm will be done using rack-mounted PCs, placed in the experiment computing room, upstairs in the Building 918. Network switches concentrating Gigabit Ethernet links from the sub-detector readout systems in the experimental hall will be connected through 10 Gigabit Ethernet fibres to a main concentrator switch upstairs in the computing room. An additional, similar system of network concentrator switches connected by high-speed fibres (from downstairs in ECN3 to upstairs in the computing room) will be dedicated to the readout of the Liquid Krypton calorimeter, which is handled differently from most other sub-detectors because of its large raw data volume. The network layout is described in the following.

On the sub-events selected by the L1 trigger further operations should be performed:

- The sub-events should be correctly re-assembled in complete events by an **event-building** process (EB); this will include the assembling of the (possibly non zero suppressed) LKr data;
- A **fast reconstruction** (FR) software, not necessarily using the most detailed information from the detector (*e.g.*, possibly using heavily zero-suppressed data from the Liquid Krypton calorimeter), nor calibrations, and based on simplified reconstruction algorithms, should run on the assembled events;
- On a positive L1 decision, the complete Liquid Krypton calorimeter data should be downloaded from the readout buffers and appended to the complete event data structure;
- A **L2 trigger** decision could be taken for further rejection of background events;
- At this stage event selection algorithms can be used on the reconstructed events, in order to have one or more **data streams** for storage. A minimal set of streams should include:
    - Golden events, fulfilling a first, loose selection on the main signal;
    - Muon events, selecting decay or halo muons with a (variable) pre-scaling factor.

The EB, FR, L2 and streaming processes, should run on a sub-set of the online-farm.

Additional machines will be hosted in the experimental hall, in order to handle the DCS of the different sub-detectors, and will also be connected by Ethernet links to a network switch downstairs, linked with 10 Gigabit/second (Gbps) links to the main network switch in the computing room.

All the machines in the L1 trigger farm will be on an experiment private network, while dedicated gateway machines will be connected to the CERN General Purpose Network (GPN).

Table 1: Estimated event sizes and data rates at the input of the event building/L2 farm. Additional data from L0 trigger systems and Trigger Processors is not included.

Detector	Event size (B)	L1 trigger rate	Through-put at L2 input (MB/s)
CEDAR	216	100 kHz	21
GTK	2250		215
CHANTI	192[1]		18
LAV	160[2]		15
STRAW	768[3]		73
RICH	160[4]		15
CHOD			
MUV	768[5]		73
IRC & SAC	576[6]		55
Total	≈5k		≈500
LKr	222 k[7]		

Event size and rates

The maximum L0 trigger rate is fixed to be 1 MHz, and we assume 100 kHz at the output of the Level 1 trigger. On the other side of the data processing chain, the Data Logging System, the maximum speed of event logging to tape can be fixed at **100 MB/s** (per tape drive, we assume one fully dedicated tape drive for NA62). The maximum event rate of the DLS, in an operational model in which raw data is assembled by the event-building processes and then directly transferred to tape (with suitably dimensioned disk buffers in order to profit from the low duty cycle of the SPS), is fixed then by the event size:

DLS logging rate (events/s) = tape logging bandwidth (B/s) / event size (B).

The maximum Level 2 output rate is then fixed by the duty-cycle of the slow extraction of the SPS:

6 stations * 2 views * 2 strips/view * 2 words * 32bit
20 fired crystals * 2 words * 32bit
4 stations * 4 views * 4 straws * 1.5 average multiplicity * 2 words * 32bit
20 fired PMTs * 2 words * 32bit
48 scintillator planes * 2 tiles/plane * 2 words * 32bit
32 IRC sectors + 4 SAC PMTs * 8 samples *16bit
[7] 13248 cells*8 samples*16bit

L2 rate (events/s) = DLS logging rate (events/s) / SPS duty-cycle

Assuming a SPS cycle made of a 9.6 flat-top spill each 40 seconds, a duty-cycle of **0.25** can be considered.

In order to estimate the size for the different formats of data, we start from the **RAW** and make educated guesses for **RECO**, **THIN** and the other data formats described above.

The data throughput at the input of the event building/L2 farm can be computed from the sub-event data format and the expected rates (100 kHz L1 trigger rate) for each of the sub-systems. This is summarized in Table 1.

Networking

NA62 plans to use switched Ethernet technology to transfer data from the sub-detector electronics to the online processing farm. Several network links will be needed in the detector area: links for connecting readout boards to sub-detector/L1 PCs, to transmit data, and to the L0 Trigger Processor, to transmit L0 Trigger primitives; links for connecting L1 PCs to the L1 Trigger Processor PC, among themselves and to the L2 farm fabric; and links for slow control of the readout boards themselves. To have a uniform system we intend to use Gigabit Ethernet (GbE) everywhere, even for tasks that require less bandwidth (DCS for example). One possible exception to this rule will be the embedded PCs on TEL62 boards for which 100Mbps links are used.

An overall sketch of the TDAQ system and of the network interconnections can be seen in **Error! Reference source not found.**. To facilitate the maintenance of the computers, we plan to keep as many as possible of them upstairs, *i.e.* in a dedicated computing room in building 918 next to the NA62 main control room (approximately half of the surface formerly dedicated to the NA60 control room, as shown in Figure 9 rather than downstairs, *i.e.* in the ECN3 experimental area.

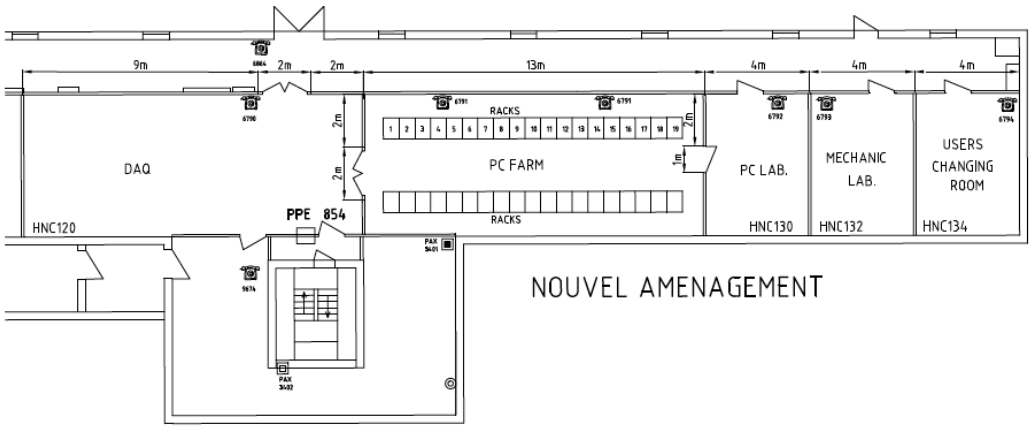

Figure 9 Proposed layout of building 918. The NA62 upstairs computing room will be hosted in a part (approximately one half) of the area formerly used as NA60 control room.

This requires a large amount of fibres connecting the experimental area to the computer room in building 918. Moreover, due to the position of the sub-detectors all along the decay region, the readout boards are not concentrated in the same area, but rather spread along the more than 100 m long underground experimental zone. For this reason we plan to split the network in two branches: an

experimental hall (ECN3) networking structure and a ground level (building 918) one, with optical fibres connecting the two. In the experimental hall we plan to have a central switch in the electronics barrack sitting on top of the Liquid Krypton calorimeter. Starting from there we will have connections to all sub-detectors along the NA62 detector setup.

The following table summarizes the expected number of network links required up to the level of L1 PCs.

We are planning to have several fibres per station/sub-detector running 1 Gbps Ethernet instead of just one (or more) fiber running 10 Gbps Ethernet. The 10 Gbps approach would require to place a concentrator switch at each station/sub-detector (TEL62 network cards are 1 Gbps) and we feel the combined cost would be greater compared to the 1 Gbps solution.

The central switch will act as a sort of aggregator/transceiver, allowing us to reduce the number of fibres connecting the experimental hall to the control room using 10 Gbps ones. The estimated number of 10 Gbps fibers needed is 60.

Table 2: Expected number of network links up to the L1 PC farm. Links internal to sub-system hardware not included.

Detector	**GbE links** to L1 PCs (**data**)	**GbE links** to other detector boards (**trigger primitives**)	**10 GbE links** to farm **(data)**	**GbE links** to L0TP **(trigger primitives)**	**DCS FE links** (100 Mbps)
CEDAR	8	0	0	0	1
GTK	60	0	0	0	
CHANTI	1	0	0	0	1
LAV	48	11	0	1	12
STRAW	64	0	0	0	2
RICH	32	3	0	1	4
CHOD	4	0	0	1	1
MUV	2	0	0	1	1
IRC/SAC		0	0	0	1
LKr			28		
LKr L0 trigger	1	0		1	35
L0TP	1	0		0	1

The ground level (building 918) network infrastructure will consist of a main switch placed inside the computer room. This will be connected to the 10 Gbps fibers going to the experimental hall and to the roughly 300 computers forming the online system (L1 PCs, L2 and event reconstruction farm). The switch will also be connected to a 10 Gbps optical fibre going to the main IT facility, used to transfer the data to be logged onto tape.

Level 1 farm

At this level, the data are merged for each sub-detector. In the general case sub-detectors will be readout following a L0 trigger by separate, independent devices, so that a first building of the event at the sub-detector level is needed. The L1 farm is responsible for receiving data from the sub-detectors front-end electronics, for generating trigger primitives for the Level 1 Trigger Processor (L1TP), and for storing the data pending L1TP response. The data selected by the L1TP will then be passed on to the Level 2 farm, where the event building for the entire detector will be performed and a more complex selection will take place. The only exception to this scheme is given by the LKr sub-detector where, due to the very large event data size involved, a dedicated farm is used.

The L1 farm needs to provide enough aggregate network bandwidth and CPU resources to handle the foreseen data rate, and of course we plan to take full advantage of modern multi-core CPUs and fast 10Gb network links. Still, depending on sub-detector, in can be difficult to accommodate the data from one entire burst in just one computer, generating the trigger primitives at the same time. For this reason we plan to use several computers for each sub-detector, subdividing the load between them with a plain round robin algorithm. To reach this target we count on the ability of the front-end electronics to send different events to different computers. This should be easily accomplished by the TEL62 readout boards used by the majority of the sub-detectors.

Table 3: Estimation for the number of sub-detector/L1 PC and links.

Sub-detector	Number of readout boards	Number of readout links	Number of sub-detector PCs
CEDAR	2	2	1
GTK	tbd	30(60)	15(30)
CHANTI	1	tbd	tbd
LAV	12	24	12
STRAW	2	8	tbd
RICH	4	8	200 cores
CHOD	1	4	1
LKr	432	28 (10GbE)	30 (12-core)
LKr L0 trigger	35	7	tbd
MUV	1	2	1
IRC & SAC	1	tbd	tbd

In case this solution is not viable, we plan to investigate the possible use of network processors to be placed in between the front-end electronics and the L1 farm to modify the ethernet packet at wire speed.

Both solutions guarantee an easily scalable system, where more computers may be added depending on the needs, the main limitation being the central switch.

The number of sub-detector/L1 PCs and links is roughly estimated Table 3.

1.2 Online Control System

The online control system will be a loosely connected set of hardware and software with the main task of controlling the initialization and monitoring the progress of a data-taking run ("run control"), as a thin layer interfacing the sub-system specific electronics on one side and the common infrastructure - such as the SPS timing signals and the PC farms - on the other. The online system should normally not interfere with the sub-detector electronics during a data-taking run, but it may be allowed some action during the inter-spill period. The online system should also interface to the Detector Control System (*a.k.a.* "slow control"), and to the data monitoring system, and present a common user interface for all these items. It is expected that the DIM software (1) will be heavily used for the development of most parts of the online control system.

The configuration, initialization, monitoring and control of sub-detector specific hardware is of course responsibility of each sub-detector group. However, when a sub-detector is included in the active list of sub-detectors for running, each sub-system (thus including the Trigger Processors, etc.) should be able to respond to a common set of commands issued by the NA62 TDAQ system via network. Each sub-detector will identify a *single* Linux machine which will handle the communication tasks with the common NA62 TDAQ system. The remote commands to which each sub-detector system should respond must include the following:

- **Enable/disable global NA62 control**: when enabled the sub-system can only be controlled from the NA62 TDAQ system, for normal running; when disabled the sub-system is under local control (for tests and debugging) and is completely ignored for running and data collection purposes.
- **Perform cold start initialization**: this task forces all basic (possibly time consuming) sub-detector configuration actions to be started, and must be used after a power-up or a reset. No parameter is passed from the TDAQ to the sub-system and therefore no conditional configurations can be performed at this time.
- **Start a new run**: three parameters are passed, the first one being the run number, the second being a common (equal for all sub-system) TDAQ configuration string, and the third being a sub-system specific configuration string. After successfully completing this task, the sub-detector is in running mode from the point of view of the TDAQ system, and it becomes ready to receive triggers as soon as a "start of burst" trigger is received through the TTC system.
- **End a run**: after completing this task, the sub-detector system replies with a set of summary data which the TDAQ system will take care to have stored together with the data.
- **Query status**: this command can be dispatched at any time, and the system should reply with the information on its current state, which should include at least the number of the last run started and the number of the last burst started.
- **Reset**: this command can be dispatched at any time, and forces the system to go back to an uninitialized state (its main use being to get back to a known situation in case some serious error occurred). Dispatching this command during a run can have unpredictable effects on the data being collected.

Most likely, during the development phase or for testing purposes, sub-detectors will also want to be able to run also with different setups (*e.g.* in the lab, or at the experiment using a separate "private" clock/trigger distribution system), and therefore they might be interested in running parts of the NA62

online software on different machines. The online group will provide support for the common NA62 online software to run on “NA62 standard configuration” hardware.

Bibliography

1. **C. Gaspar et al.** *DIM: Distributed Information Management System.* CERN. Geneva : s.n. http://dim.web.cern.ch/dim/.

NA62 Acronyms and Abbreviations

ADC	Analog to Digital Converter
AM	Absorber Module
APD	Avalanche PhotoDiode
ASIC	Application Specifc Integrated Circuit
BCRST	Bunch Counter ReSeT
BEATCH	Progran to provide coordinates of all beam elements as input for alignment
BEND	BENDing magnet or dipole
BIF	Barrier Improvement Factor
CALB	Sub-detector data format for calibration
CCPC	Credit-Card PC: commercial processor on TELL1/TEL62 boards
CEDAR	Cerenkov Differential counter with Achromatic Ring Focus: differential Cerenkov detector developed at CERN
CHANTI	Charged ANTI
CHOD	Charged HODoscope
CKM	Cabibbo–Kobayashi–Maskawa matrix
CM	Circulation Module
COLL	COLLimator
COND	Condition data of the detector (extracted from DCS)
CONF	configurations of the run, including active detectors, trigger configurations, beam conditions, etc.
COST	calibration parameters computed by calibration tasks running on (partially or fully) reconstructed data.
COTS	Commercial-Of-The-Shelf
CPD	Calorimeter Pipeline Digitizer module
CREAM	Calorimeter Readout Module
CTL	Chamber Trigger Logic for straws L0 trigger system
DAQ	Data Acquisition System
DCS	Detector Control System
DDR2	Double Data Rate SDRAM (memory chips)
DIM	Distributed Information Management system
DLL	Delay Locked Loop
DLS	Data Logging System
DM	Distribution Module
DPRAM	Dual Ported RAM
DRAM	Dynamic RAM
EB	Event-Building
ECN3	Experimental Cavern housing the NA62 experiment
ECRST	Event Counter ReSeT
EDX	Energy-Dispersive X-ray spectroscopy
EE	End of Ejection
EOB	End Of Burst
EOC	End Of Column option for GTK chip architecture

EoC	End-of-Column part in the P-TDC chip architecture of the GTK
EOF	End Of Frame
FADC	Flash Analog to Digital Converter
FE	Front-End
FEE	Front-End Electronics
FEM	Finite Element Model
FIFO	First In First Out buffer
FISC	Filament Scanner, a beam profile detector inside the beam vacuum system
FNAL-NICADD	Fermi Nat Lab - Photo injector Lab (18Mev electron linac)
FPGA	Field Programmable Gate Array
FR	Fast Reconstruction
GbE	Gigabit Ethernet
GIM	Glashow–Iliopoulos–Maiani mechanism which suppresses flavour-changing neutral currents)
GOL	Gigabit Optical Link transmitter
GPN	General Purpose Network
GPU	Graphic Processing Unit
GTK	GigaTracker
HALO	A beam simulation program to calculate muon HALO rates
HPTDC	High Performance Time to Digital Converter
HV	High Voltage
IRC	Intermediate Ring Calorimeter
JTAG	Joint Test Action Group protocol
L0	Level 0 Trigger
L0TP	Level 0 Trigger Processor
L1	Level 1 trigger
L1TP	Level 1 Trigger Processor
L2	Level 2 trigger
LAV	Large Angle Veto
LED	Light-Emitting Diode
LG	Lead Glass
LGTS	Lead Glass Test Station
LKr	Liquid Kripton calorimeter
LTU	Local Trigger Unit
LV	Low-Voltage
LVDS	Low-Voltage Differential Signalling
MBPL-TP	Dipole Bending Magnet with Tapered Pole
MEPs	Multi-Event Packets
MIP	Minimum ionising particle
MNP33	NA62 Experimental Magnet
MUV	Muon Veto System
NAHIF	North Area High Intensity Facility
NIM	Nuclear Instrumentation Module
NINO	Fast front-end preamplifier-discriminator chip developed by ALICE
NNLO	Next-to-next -to-leading order
PCB	Printed Circuit Board

PDE	Photon Detection Eff.
Pe	Photo-electron
PECL	Positive Emitter-Coupled Logic
PEI	PolyEtherImide
PET	PolyEthylene Terephthalate
PMT or PM	Photomultiplier Tube
POPOP	1,4-bis(5-phenyloxazol-2-yl) benzene organic scintillator
PP-FPGA	Pre-Processing FPGA in the TELL1/TEL62 boards
PPO	2,5-Diphenyloxazole organic scintillator
PTP	Para-TerPhenyl
PVSS	Object-oriented process visualization and control system by ETM (a commercial SCADA system)
QCD	Quanten Chromodynamik
QDR-II	Quad Data Rate II memory
QPLL	Quartz-crystal based Phase-Lock Loop
QUAD	QUADrupole
RECO	Data format for fully reconstructed events
RICH	Ring Imaging Cherenkov
SAC	Small Angle Calorimeter
SAV	Small Angle Veto
SDRAM	Synchronous Dynamic Random Access Memory
SEM	Scanning Electron Microscope
SiPM	Silicon PhotoMultiplier
SOB	Start of Burst
SL-FPGA	Sync-Link FPGA in the TELL1/TEL62 boards
SLM	Smart Link Modules
SM	Standard Model of particle physics
SPI	Serial Port Interface
SPR	Single Photoelectron Response
SRAM	Static Random Access Memory
SRB	Straw Readout Board
TCC8	Target Chamber Cavern upstream of ECN3
TDAQ	Trigger and Data Acquisition system
TDC	Time to Digital Converter
TDCB	Time to Digital Converter Board
TDCC-FPGA	TDC Controller FPGA in the TDC boards
TEL62	Trigger and Data Acquisition board developed for NA62, based on TELL1 design
TELL1	Trigger ELectronics for L1 trigger: readout board developed by LHCb
THIN	Data format for summary data from fully reconstructed events
TRIM	Steering Magnets for the Beam
TTC	Timing, Trigger and Control
TTCex	TTC encoder VME board developed by CERN
TTCrq	TTC receiver mezzanine card developed by CERN
TTCrx	TTC receiver ASIC developed by CERN
TURTLE	Trace Unlimited Rays Through Lumped Elements, a beam tracking and simulation program

VME	Electronic bus and rack standard
VTL	View Trigger Logic for straws L0 trigger system
WE	Warning of Ejection
WLS	Wave-Length Shifting
WWE	Warning of Warning of Ejection