

# High Channel Density Fault Recording with IEC61850 Sampled Values - Challenges and Benefits

Abhishek Madhyastha  
AMETEK Power Instruments  
[abhishek.madhyastha@ametek.com](mailto:abhishek.madhyastha@ametek.com)  
India

## Abstract

The advent of IEC61850 sampled values has created several opportunities in the digital substation space for efficient management of equipment and data. A key benefit from the adoption of this technology is the reduction in physical space required for IEDs through high density installations. Digital fault recorders typically comprise of a large number of analog and digital input channels to monitor several signals throughout a substation. The installation space required for a fault recording system is heavily dependent on the analog & digital channel counts in both a single unit as well as the overall system. IEC61850 sampled values help circumvent this greatly by eliminating the space required for analog & digital channel connections within the DFR. This allows higher channel counts to be accommodated in a single unit & thereby achieve high channel density. However, there're a variety of challenges and limitations imposed by Ethernet communication itself that impede effective implementation. This paper discusses the advantages of using IEC61850 sampled values in DFR applications as well as the challenges in achieving high channel density.


## 1 Channel Density - What is it?

A substation can have several analog and digital measurement points of interest for monitoring and analysis. Unlike devices such as protection relays, bay controllers, meters, etc. which are typically installed on a per feeder basis or a small group of feeders, digital fault recorders can be expected to monitor multiple feeders and potentially multiple types of CTs (protection/metering) based on the range of applications. As such, there's a need for supporting large number of analog and digital input channels on such devices. Naturally, this implies sufficient space required for installing potentially several units depending on the number of analog and digital channels that can be accommodated in each unit.

Channel density is a measure of the number of channels that can be supported in each unit with reference to the dimensions of the unit. This plays a key role in determining in the physical space required within a substation and contributes to the overall "footprint" of the substation itself.

## 2 Impact of Channel Density in Substations

In a conventional installation involving hardwired analog and digital inputs, channel density is often heavily dependent on the physical space for voltage and current measurement inputs within a DFR. Shown below is a snapshot of a 6 analog & 12 digital inputs acquisition module –




**Fig.1: 6A12D Input Acquisition Module.**

A module such as this can typically be used to monitor inputs from 1 feeder. Several such modules may be installed within a single unit to meet the channel count requirements for monitoring all points. Although the unit itself may have sufficient resources (CPU, RAM, storage, etc.) to be able to support more channels than this, the lack of physical space for input connections can be an impediment towards achieving higher channel density.

The need for high channel density poses multiple challenges in a conventional installation –


## **2.1 Complex Cabling and Routing**

The following block diagram shows the typical cabling methodology from the yard to the control house. Such an architecture requires significant amount of cabling both within the substation yard as well as the control house. Additionally, large cable trenches are required, especially between the yard and the control house.



**Fig.2:** Cable routing stages within a substation

Within the control house, the same cables are often connected to many instruments. For instance, three phase voltages and currents from protection PTs/CTs can be fed to relays as well as DFRs. Shown below is a snapshot of a DFR installation monitoring many feeders –



**Fig.3:** TR3000 DFR installation with large number of analog & digital inputs

Depending on the total number of feeders, digital contacts, analog points, etc. & the number of devices being fed with inputs, huge number of cables may have to be routed to multiple panels within the control house.

The scope of work for such an installation covers trenching, wiring labour, preparation & maintenance of design drawings, wiring verification, etc.

## 2.2 Panel Space

In addition to the cabling work, the sheer number of panels required to host these devices is determined by the number of units that need to be installed. Discrete “boxes” for various functions would increase the number of panels and subsequently the area required to host them. Integrated devices help reduce the impact but conventional analog and digital wiring sets limitations both within the device and outside.

The overall number of panels play a role in determining the “footprint” of the control house.

### 2.3 Commissioning Time & Documentation Effort

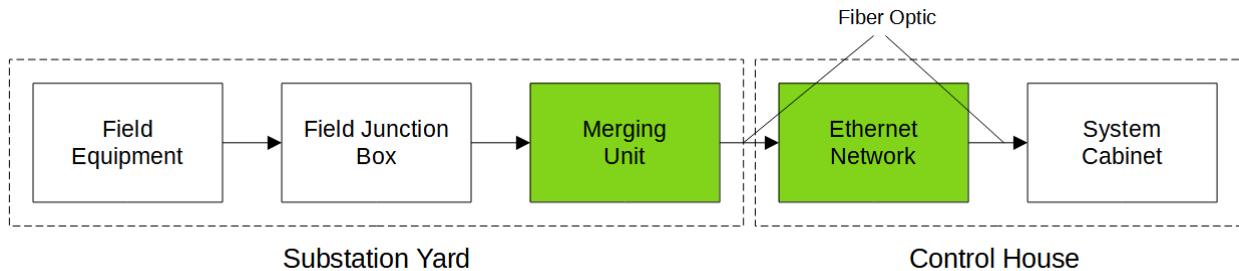
When signals are routed through several junctions both in the substation yard and the control house, the commissioning effort is high due to potential differences seen in a lab environment versus the actual deployment. Within the control house, the same signals can be routed to multiple devices & with the analog interface used in conventional methods, greater amount of testing is required at the site to verify correct operation.

A corollary to this is with documentation in the form of system drawings within the control house which can be particularly complex to formulate & maintain.

### 3 Digital Substation – Effect of Sampled Values on Channel Density

The use of sampled values significantly reduces the need for conventional analog interface wiring between the yard and the control house as well as inside the latter where multiple IEDs need to receive the same signals.

In comparison to fig. 2, the key changes involve the use of merging units which propagate digitized data to the control house via Ethernet. Within the control house, Ethernet switches are used to broadcast the data to relevant IEDs.



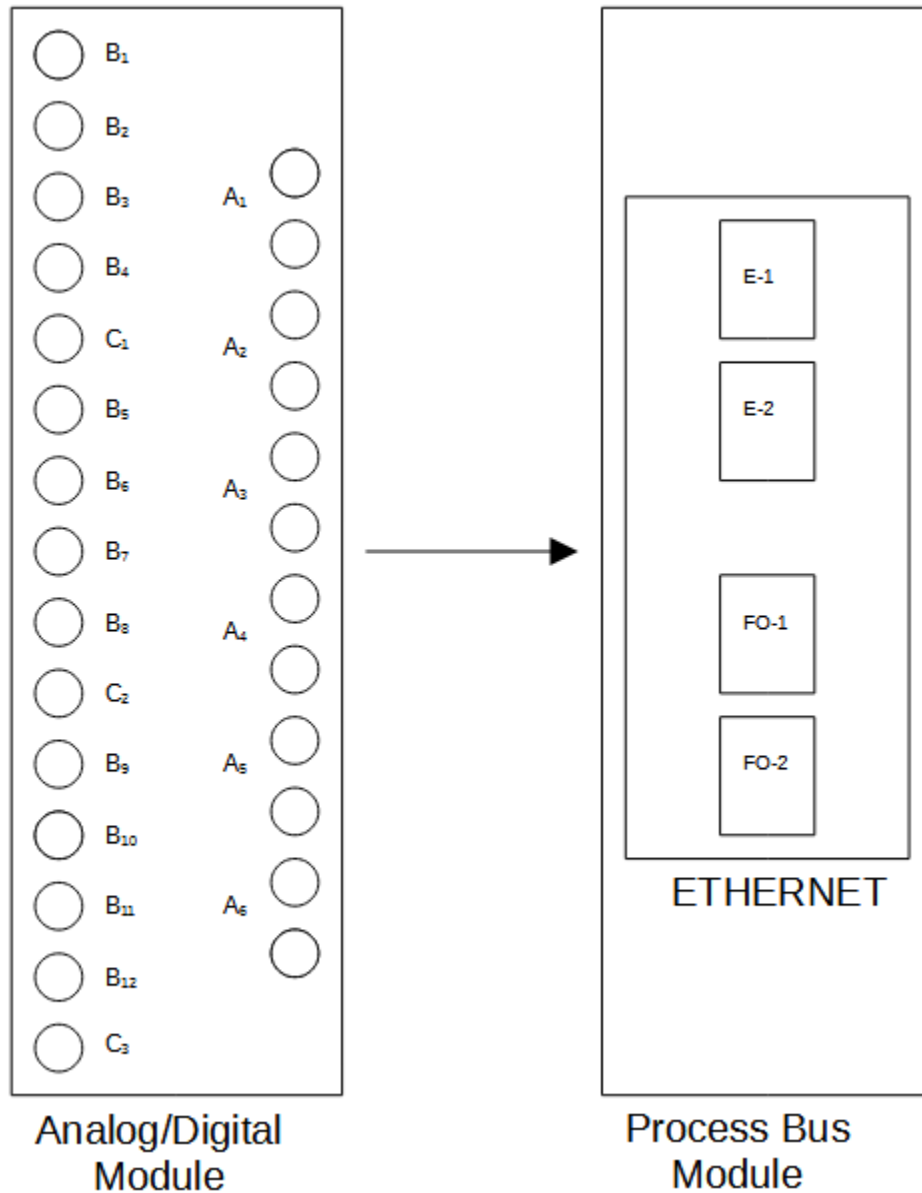
**Fig.4:** Cable routing stages within a digital substation

These general advantages apply to all devices within the control house but for DFRs which may be required to monitor many analog & digital signals within a single IED, the use of sampled values has a bigger impact on channel density.

#### 3.1 Channel Input Space

By eliminating the need for conventional analog interface wiring and using Ethernet, no pins are required to connect signals physically to the inputs of an IED. Instead, Ethernet via RJ45 and/or fiber optic connectors can take this space. This substantially increases the number of channels an IED can handle,

provided it has the resources (CPU, RAM, etc.) to do so. The diagram below shows the difference between a conventional analog/digital input module and a process bus input module.



**Fig.5:** Difference in inputs between conventional and process bus input modules

The number of channels that can be accommodated by one such module is no longer constrained by physical space for input connectors but is a function of parameters such as processing power, memory, Ethernet sub-system design, etc.

### 3.2 Panel Space

Due to the increase in the number of channels that can be handled per IED, the amount of panel space required to cover the total number of analog and digital channels is reduced as well. A large installation

with multiple DFR units could require several panels and space for other accessories, if applicable. The use of sampled values allows DFRs to pack a large number of channels within a small panel footprint.

### 3.3 Cabling

In a conventional installation involving analog interface to the units, cabling is a labour-intensive process where many factors such as channel polarity, input burden, signal integrity, time synchronization, etc. need to be carefully looked at. In a digital substation that uses sampled values, the digitized data is routed to all IEDs using Ethernet. This marks a shift from analog/mechanical configuration to a software driven approach. Once Ethernet connectivity is achieved with an IED, most of the configuration process is typically in software. Also, IEDs requiring the same inputs merely need to be connected to the same network segment eliminating concerns about input burden, cross talk, etc.

Time sync is another area where the same infrastructure & architecture for sampled value transmission can be utilized to cut down on the need for analog interface-based synchronization to every IED. Time sync protocols such as PTP can help achieve sub-microsecond level sync & use Ethernet.

When combined with protocols such as PRP/HSR, redundancy can also be implemented.

## 4 Challenges for High Density Sampled Values Deployment

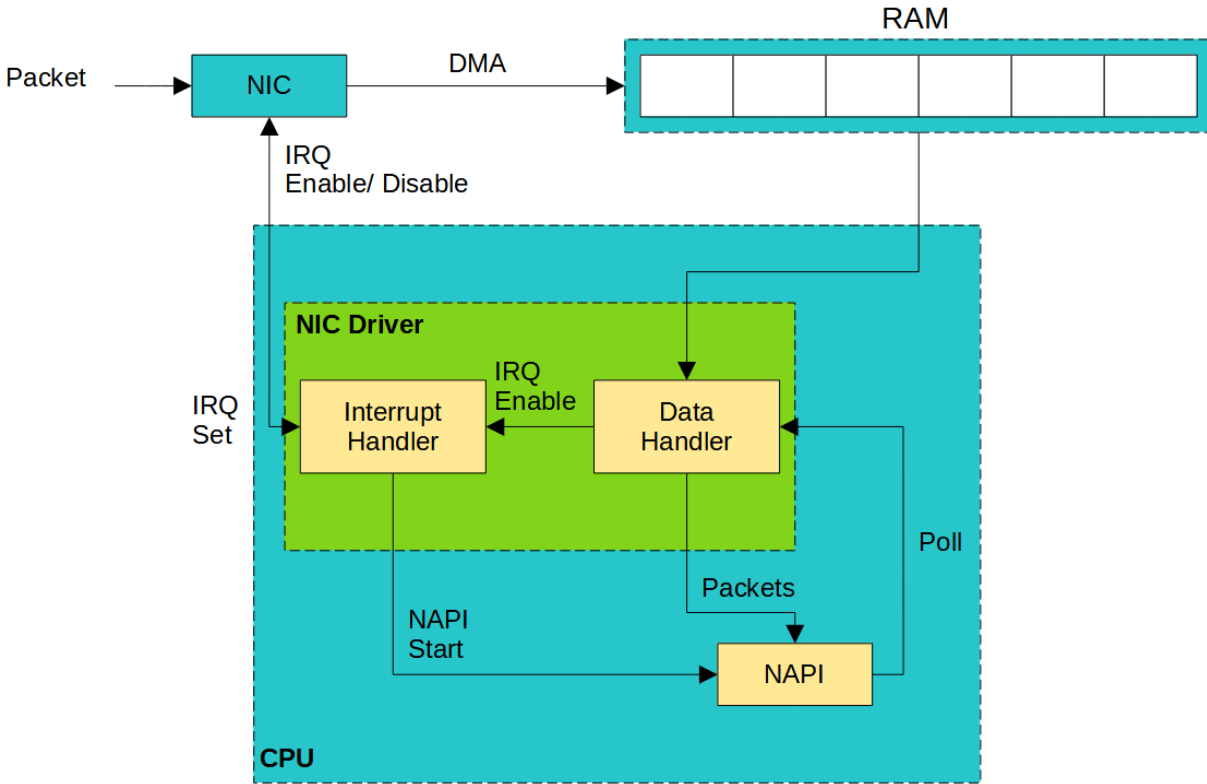
While the benefits of sampled values can be well understood in previous sections, there're challenges for achieving high density installations. Density is affected by the ability of an IED or its sub-modules to subscribe to several sampled value streams and the number of channels in each stream. The limitations can be both hardware and software dependent. The sheer amount of processing power is a straightforward constraint but other areas with an IED, particularly the Ethernet subsystem, can play a major role.

### 4.1 Conventional Ethernet Architecture

Ethernet sub-system design varies between manufacturers at several levels covering both hardware and software. Operating systems also contribute to the architecture for packet transmission/reception via Ethernet. In general, an Ethernet sub-system comprises of the following major parts –

1. Network Interface Card (NIC) – Hardware responsible for conversion between physical media involving electrical or fiber optic signals and a digital interface.
2. CPU – Hardware & software for exchange and processing of Ethernet packets.
3. Network driver – Software to service interrupts & exchange data between the NIC & the CPU. Typically provided by the NIC manufacturer for specific operating systems.
4. Driver to user space interface – Software interface for exchanging packets from an OS level to the relevant user application. This would also comprise of OS driven optimization mechanisms for efficient delivery of packets such as NAPI in Linux or NDIS PacketDirect in Windows.

The following example describes the typical flow of incoming packets within Linux through the NAPI (New API) subsystem –



**Fig.6:** Incoming packet flow through Linux using the NAPI system

1. When a packet is received by the NIC for the first time, it exits idle mode & copies the packet to RAM. This is typically done using DMA for low latency & CPU offloading, & it also raises an interrupt to the CPU.
2. The NIC driver receives this interrupt through its interrupt handler routine & notifies the NAPI subsystem. The driver disables further interrupts from the NIC.
3. The NAPI subsystem exits idle mode and starts polling for packets from the NIC driver.
4. If the driver sees that there're packets available in RAM, it'll pass those to the NAPI subsystem. This is an infinite polling process that runs until there're no packets left. The NIC does not generate any interrupts during this entire period. This is critical to avoid wasting CPU resources in handling interrupts. If every packet required an interrupt, a busy network could result in a large number of interrupts which can starve the CPU.
5. Once all packets have been received, the driver exits polling mode & re-enables interrupts from the NIC.

The key aspect of this system that allows high bandwidth data transfers over Ethernet with minimal wastage of CPU resources is the NAPI subsystem which relies on selective polling rather than a blind interrupt-based mechanism. A single interrupt can be sufficient to cover a transfer of several gigabytes of data.

#### 4.2 Uniqueness of Sampled Value Packets

The general idea incorporated within NAPI on Linux (and similar mechanisms on other operating systems) was focused around the efficient use of Ethernet for high volume/bandwidth data transfers. As an example, a large file transfer could involve several gigabytes of data being exchanged. The key aspect here is that

the gigabytes of data within the file is available at once and can be broken down into chunks & transmitted with no delays in between. This means that each packet is fully used and there're little to no delays between packets until the entire file is transferred.

Sampled value packets are unique in the sense that they're a continuous stream of data but uniformly spaced in time. Consider a 60 Hz system with a sample rate of 80 samples per cycle & 1 ASDU. This yields 4800 samples and effective packets per second. The packet-to-packet delay is uniform at 208 microseconds.

<b>Sample Count</b>	0	1	2	3	4	5	6	7
<b>Delay (us)</b>	208	208	208	208	208	208	208	208
<b>Packet Size (bytes)</b>	131	131	131	131	131	131	131	131

**Table.1:** 80 spc, 60 Hz samples – 4800 samples or packets per second uniformly spaced at 208 us

Assuming each packet is ~132 bytes, which is less than the typical MTU size of 1500 bytes, the bandwidth required is just ~630 kiB/s – far below modern Ethernet capabilities. When multiple SV streams are present on the network, they would all maintain a similar timing pattern. However, the key difference between such a stream of data and bulk data transfers is the consistent delay between packets. For a system such as NAPI to function effectively, it's required that large amounts of data be available simultaneously with no delays in between. The introduction of a 208 us delay in the above example means that the Ethernet sub-system can no longer make effective use of the single interrupt for full transfer approach. Instead, the behaviour would look like this –

1. First packet is received by the NIC which copies it using DMA to RAM & raises an interrupt.
2. CPU catches the interrupt and lets the NIC driver service it.
3. Driver disables further interrupts and notifies NAPI which exits idle mode and starts the infinite poll process.
4. NAPI polls the driver for packets and receives SV packet/s for a time instant (or in other words, for one sample number).
5. NAPI pushes the packet to the userspace and polls the driver again for more packets.
6. Due to the 208 us delay before the next SV packet is received, the poll returns empty handed.
7. The driver re-enables interrupts from the NIC thinking there're no packets to be received.
8. At the next 208 us mark, a new SV packet is received which causes the entire process above to be repeated.

The above may happen for every SV packet or even packets from multiple SV streams cause all IEDs are time synchronized & transmit SV packets every 208 us. This defeats the efficiency measures of high bandwidth data transfer mechanisms like NAPI & the system ends up receiving a huge number of interrupts. The processing associated with each interrupt can be spread throughout the OS at different levels & thereby causes a lot of unnecessary CPU resource utilization.

### 4.3 Sampled Value Reception - CPU Consumption Experiment

An experiment was conducted to analyse the above behaviour on a BeagleBone Black board when several SV streams are active. The setup consisted of –

1. BeagleBone Black Board running Linux
2. 2x SV publishers with three 9-2LE streams each.

Both SV publishers were setup to send 4800 packets per stream (80 spc at 60 Hz) and 3 streams each. This results in a total of  $4800 \times 6 = 28,800$  packets per second. Each packet is ~130 bytes in size which means a bandwidth of 3.4 MiB/s. This bandwidth is well below the 100/1000M speeds typically seen with Ethernet designs but the uncommon packet to packet timing makes it difficult for systems to effectively utilize bulk transfer optimizations. As such, the CPU consumption on the BeagleBone Black board was seen to spike up significantly –

```

CPU[|||||] 60.2% Tasks: 14, 0 thr, 52 kthr; 1 running
Mem[|||||] 23.8M/491M Load average: 0.19 0.27 0.27
Swp[|||||] 0K/0K Uptime: 02:59:08
Network: rx: 3.28MiB/s tx: 1KiB/s (28071/16 packets)

CPU usage bar: [low/normal/kernel/irq/soft-irq/steal/guest/io-wait used%]

```

**Fig.7:** CPU consumption with 6 IEC 61850 9-2LE streams

It can be noticed that the CPU consumption is entirely owed to soft-irqs which are in turn driven by the hard interrupts from the NIC. As explained in section 4.1 & 4.2, the timing of the sampled value packets causes the system to spend a lot of cycles servicing interrupts to extract packets from the Ethernet HW rather than on the actual application specific processing of packets. Packets may also be lost due to this excessive CPU consumption & this implies practical limits due to the Ethernet subsystem design.

## 5 Conclusion

IEC 61850 sampled values present a unique opportunity to reduce the footprint of substations by achieving high channel density installations. It also addresses many challenges associated with conventional analog interface wiring. However, new challenges need to be taken into consideration on the digital side, particularly with Ethernet architectures and their efficient use for such applications.

1. IEC 61850 9-2 LE proposes the use of 4 current and 4 voltage channels per sampled value stream. This results in a typical packet size of ~130 bytes for protection rates such as 80 spc. These small packets are transmitted by merging units at periodic intervals at the order of few hundreds of microseconds. To minimize the number of SV packets on the network and the burden on subscribing devices, the IEC 61869-9 provision of supporting higher channel counts per stream should be used, wherever possible. This reduces the sheer number of packets required to transmit voltage & current samples for a large number of channels and allows many Ethernet architectures to make effective utilization of bulk transfer techniques.
2. IEDs such as DFRs which are typically required to support high channel counts stand to benefit from careful optimizations in their hardware & software architectures, particularly around Ethernet. Typical Ethernet designs are aimed at efficiency with bulk transfer of data where high bandwidth is required. The uniqueness of the timing of sampled value packets means that the systems might need to be tuned/modified for such precision timed packets.
3. While packing significant processing power in a small space is a relatively well understood concept, sampled values can present bottlenecks in other areas such as networking. Inability of networking hardware & software to reliably capture high volumes of SV packets would pose a constraint & can also cause inconsistent behaviour in IEDs.
4. Careful consideration is required in the selection of switches and routers meant for high density SV applications. It's essential that such switches have the ability to process large number of small packets, often indicated through the packets-per-second (pps) specification.

5. Isolated testing of IEDs in a lab environment can yield different results compared to the field. It's beneficial to perform RTDS style full system testing where all equipment including Ethernet switches and other networking equipment are considered.