THE PERFORMANCE TRIANGLE IN DIGITAL SUBSTATION ARCHITECTURES

Simon RICHARDS
Alstom Grid - UK
simon.richards@alstom.com

Abraham VARGHESE
Alstom Grid - UK
abraham.varghese@alstom.com

Steve POTTS
Trifford Consulting - UK
steve.trifford@btinternet.com

ABSTRACT

In the traditional domain of teleprotection signalling, there is a recognised triangle of performance. This triangle depends upon the required speed of communication, the expected dependability of the command, and the security of receipt. These three aspects are interlinked, and the correct balance of each is required in order to offer the traditional functionality such as permissive signalling, blocking and intertripping schemes.

This paper relates the traditional concepts to the application of automation schemes within digital substations. In this manner, it is anticipated that protection engineers can relate today's Ethernet technology by analogy with equivalent trusted practices.

INTRODUCTION

Traditionally it had been difficult to get Protection Engineers to engage with telecommunications, and vice-versa for Telecommunications Engineers to engage with Protection. That is, until something went wrong and it sometimes became a blame game with the Protection Engineers accusing the telecommunications of providing an inadequate service and the Telecommunications Engineers claiming that the protection is too demanding. To resolve things requires effective communication between the protection system and the telecommunication system, and this effective communication must extend to the people engaged in the engineering aspects of both domains. The Shannon Weaver model provides a model for communications. For effective communications the source and destination must share a common language. Whilst it might be argued that Protection and Telecommunications Engineers will never speak the same language, there are established terms that are recognised by both.

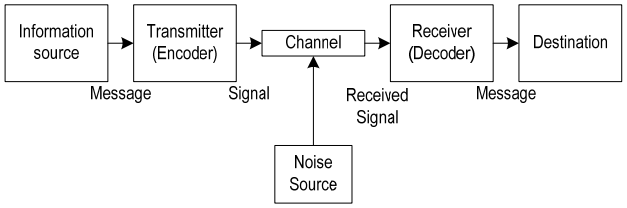

Figure 1: The Shannon-Weaver Model.

TELEPROTECTION

When protection becomes reliant upon telecommunications to meet operational requirements, the term teleprotection is commonly applied. In the language of teleprotection there are three words which describe characteristics that are critical to the overall system performance. They are speed, dependability, and security; and may be expanded as the speed of communication, the expected dependability of the communication, and the security of receipt. These three aspects are interlinked and can be conveniently represented graphically in the form of a recognised triangle, as shown in Fig. 2.

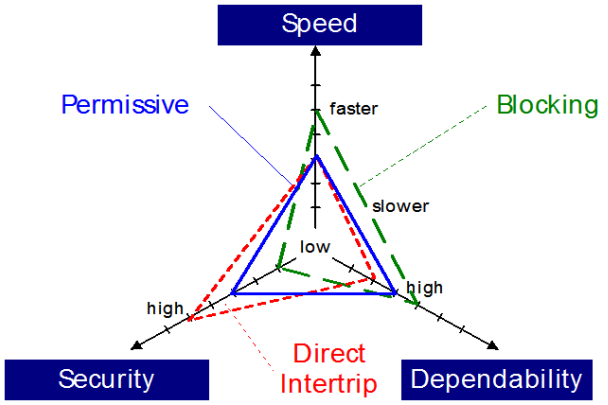

Figure 2: The Performance Triangle

High security means that an intertrip command does not spuriously pick up due to a noisy channel. High dependability is the means by which a blocking or permissive command may easily pass through noise and still be received at the remote line end. Security is assessed by the probability of an unwanted command occurring, and dependability is assessed by the probability of missing a command. Poor security is indicated by a high probability of an unwanted command being received (Puc), therefore a lower Puc figure is generally preferable. Poor dependability is indicated by a high probability of a missing command (Pmc). Generally a lower Pmc figure is also preferable. In practice the correct balance within the triangle is needed to ensure the high performance demanded on a real power system.

COMMUNICATIONS TRIANGLE ANALOGY

Within the Ethernet connected architecture of the digital substation, the same triangle analogy of speed, dependability and security as per Figure 2 still applies.

Speed

Traditional protection schemes involve AC and DC wiring between the relays and primary plant interfaces, i.e. instrument transformers and switchgear respectively. This interface is replaced by Ethernet links in the digital substation, hosting the 'Process Bus' introduced by the IEC 61850 standard. Currents and voltages are sampled in real

time by Merging Units and typically sent to IEDs on 100 Mbit/s Ethernet. The processing performed by IEDs starts with the samples received. The performance of Ethernet, particularly the speed, is essential for the protection scheme; current and voltage samples with sufficient detail (sampling rate) have to reach the relay in the shortest possible time for the protection scheme to perform as well as a conventional scheme.

The DC wiring is also replaced in the digital substation; status information, trips, alarms and controls are carried by IEC 61850 GOOSE messages on Ethernet. Signals that require cross-wiring between IEDs may also be transmitted and received as GOOSE by respective IEDs. Speed is of the essence in GOOSE messaging, as may be seen from the definitions and performance requirements given by the standard. There is the possibility to improve on the overall speed of response of the protection scheme, with the trip signals routed directly to the Merging Unit/digital circuit breaker controller, rather than the traditional hard wired path via the relay output contact followed by the master trip relay.

In summary, achieving high speed within digital substation architectures requires that the issuing of Ethernet messages is timely and deterministic, that the latency of travel to the subscriber over the Ethernet network is short and relatively jitter-free, and that the processing and response to incoming messages is again fast.

Dependability

When using Ethernet for the station bus, this inevitably involves more than two IEDs, hence a network will be configured. With the use of Ethernet comes the possibility of a network at the process bus level too, and the associated flexibility of distributing signals in many ways.

Simple point to point connections are possible, as shown in Figure 3. These have the advantage that they closely imitate AC and DC wiring in traditional protection schemes, and are intuitively appealing. They also avoid the latency associated with networks, and have minimal engineering effort associated with Ethernet. However, there are other aspects to be considered. For instance, the physical layout of the substation may require multiple merging units – one for CT and VT, and another for the circuit breaker interface. In such an arrangement, there are three or more single point failures – the IED, and the Merging Units, which may not be acceptable. Furthermore, the protection relay should be capable of receiving sampled values and GOOSE independently, to maintain point-point connections. The alternative would be to use a switch between the relay and the Merging Units, which then no longer offers a point – point connection.

Using IEC 61850 there are a number of natural opportunities to maximise dependability. First of all, GOOSE messaging has a repetition feature whereby a change of state will not just cause one message to change, but will be repeated in subsequent messages - initially quickly, and then at longer intervals thereafter. This means that even should one or more messages fail to arrive at the intended receiving IED, or should those messages become corrupted, the GOOSE signal will still be received within a few additional milliseconds. This provides a natural defence mechanism against traffic collisions during GOOSE avalanche, and unexpected bit errors introduced by noise or switch performance.

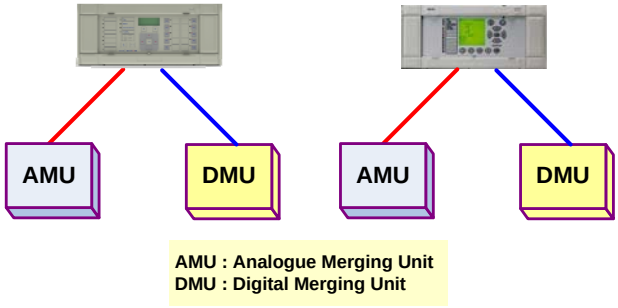

Figure 3: Simple Illustrative Process Bus Architecture

Redundancy is another common means by which the network can be designed to be immune to single device failures, ensuring message transfer dependability.
Sampled value data streams from merging units require time synchronisation, whether to real time, or some other time reference. Depending on the function of the connected IED, and the number of merging units to which it subscribes, the effects of loss of time synchronisation may have an impact on dependability. This is covered in a later section.

Security

In the station bus environment, security is achieved by carefully setting the publishers and the particular GOOSE signals of interest, to which a receiving IED subscribes. The addressing within each message will then ensure that an IED cannot inadvertently respond to any command or status information emanating from a rogue device to which it does not subscribe. The same is true for the process bus, where sampled analogue values will only be taken as valid when received from a correctly addressed logical node.
Security is further enhanced by the use of fibre optic Ethernet connections for any links which run outside of a single local cabinet, such that there are no long, cross-site runs of copper-based communications. This eliminates the risk of induced interference – boosting both security and dependability.

It must be appreciated that the fundamental design of the IED too plays a role. For example, in the case of process bus sampled analogue values, what happens if individual samples fail to arrive on-time for application to protection algorithms? What happens if a number of such instances happen, either sequentially, or at random intervals within say one power frequency cycle? This is akin to a traditional CT wire which has a loose connection, and unless detected might result in a failure to trip an overcurrent function (bad dependability), or a spurious trip of a differential function (bad security).

High-quality protection IEDs may employ techniques such as interpolation of any missing or late samples, based on preceding or following “good” sampled data, until such a point that so much consecutive data is missing, in which

case the affected part of the protection scheme must be blocked, and an alarm raised. A process bus scheme will also be both more dependable and secure if the merging units respond within a delay of maximum 3ms, preferably faster.

Employing cybersecurity features in IEDs will tend to ensure that deliberate, or inadvertent corruption of the scheme configuration will not occur. This naturally will improve not just physical security, but all aspects of performance within the triangle, as all IEDs in the scheme should remain as per the configuration finalised in the factory, and site acceptance tests. This is particularly significant where the process bus is used, or where trip logic is digitised too. For example, if a technician were to mistakenly rename a logical node of a merging unit, this could be akin to routing CT and VT circuits to the wrong relay in a traditional wired scheme.

ETHERNET NETWORK TOPOLOGY

The architecture chosen for the station and process bus networks influences all three aspects of the performance triangle. The topology of a network is the way in which devices are connected together. Figure 4 shows the two simplest topologies which are the star and the ring, so-called because they resemble stars and rings when drawn.

In the star topology, a physical connection runs from each device on the network to a central location, usually an Ethernet switch. In the ring topology, a physical connection is daisy-chained around the devices in the form of a ring.

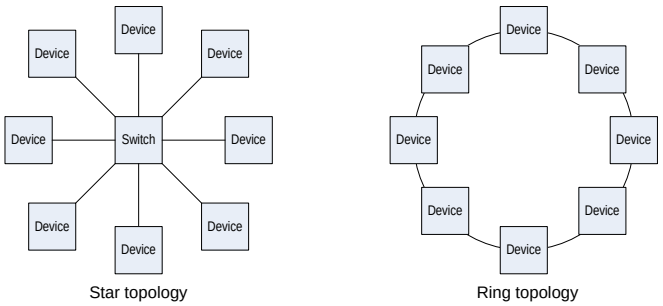

Figure 4: Simple star and ring network topologies

Principles of Redundancy in Communications Networks

In the context of communication networks "Redundancy is any resource that would not be needed if there were no failures". Redundancy is therefore a provision of transparent backup. It is required where failure cannot be tolerated, such as in critical applications like primary substation automation.

Figure 5 shows how redundancy can be incorporated into a star network. A connection from each node goes to a different switch, providing an alternative path. This topology is called Dual Homing Star.

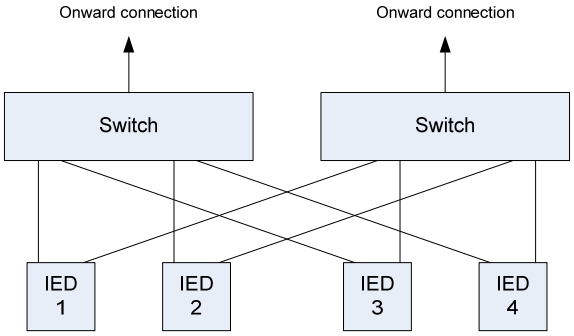

Figure 5: Redundant connections in star topology

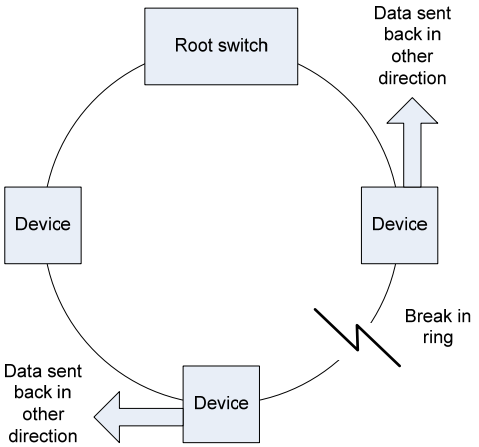

Figure 6: Redundancy in a ring topology

Figure 6 shows simple redundancy with a ring topology. The links are from device to device in a ring, the idea being that if the link from one direction fails, the link in the other direction can be used to effect transactions. In the event that there is a break at one point of the ring, appropriate redundancy protocols can automatically readjust the ring such that the data is sent back in the opposite direction, ensuring it will get to its destination.

Techniques for Redundancy in Substations

IEC 62439-3 clause 4 describes a Parallel Redundancy Protocol (PRP), and IEC 62439-3 clause 5 describes High-availability Seamless Redundancy (HSR). Systems based on PRP and HSR are appropriately suited to the needs of substation automation networks, providing true static redundancy (zero recovery/switchover delay). This is also known as 'bumpless' redundancy.

IEC 62439-3 PRP

PRP is capable of providing bumpless redundancy for real-time systems, and hence becomes the reference standard for star-topology networks in the substation environment. A PRP compatible device has two ports operating in parallel, each port being connected to a separate LAN (Local Area Network) segment. IEC62439-3 assigns the term DANP (Doubly Attached Node running PRP) to such devices.
When a node sends a frame of data, the frame is duplicated on both ports and thus on both LAN segments, providing a

redundant path for the data frame in the event of failure of one of the segments. When both LAN segments are operational, as is the normal case, each port receives identical frames and these identical frames need to be carefully handled. The handling brings overheads, but bumpless redundancy is assured.

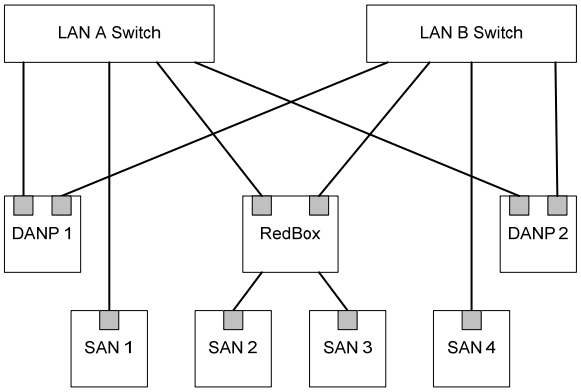

Figure 6: Example PRP redundant network

IEC 62439-3 HSR

HSR is capable of providing bumpless redundancy for real-time systems and becomes the reference standard for ring-topology networks in the substation environment. HSR works on the premise that each device connected in the ring is a doubly attached node running HSR. IEC62439 assigns the term DANH (Doubly Attached Node running HSR) to such devices. As per PRP, singly attached nodes are connected via a so-called Redbox.

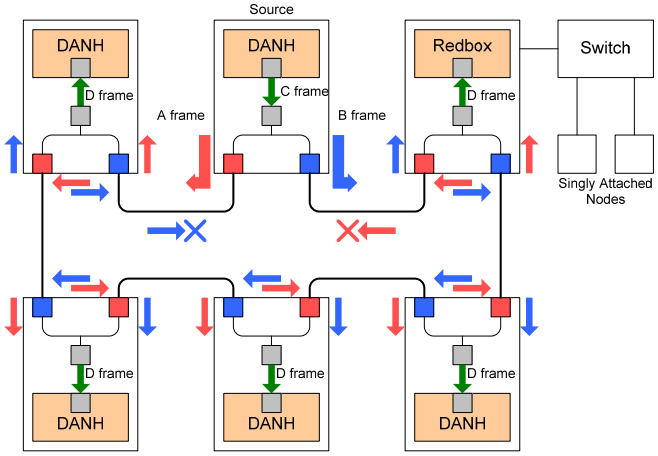

Figure 7: HSR for multicast traffic

Figure 7 shows a simple HSR network, where a doubly attached node is sending a multicast frame (that is a frame that is intended for multiple recipients on the network). The frame (C frame) is duplicated, and each duplicate frame is tagged with the destination MAC address and the sequence number. The frames differ only in their sequence number, which is used to identify one copy from another. For convenience, the duplicate frames are labelled the A frame and B frame. Each frame is sent to the network via a separate port. The destination DANH receives two identical frames from each port, removes the HSR tag of the first frame received and passes this to its upper layers. This now becomes the D frame. The duplicate frame is discarded. The nodes forward frames from one port to another unless the particular node is the node that originally injected it into the ring. With unicast frames (frames that are intended for a single destination), there is just one destination and the frames are sent to that destination alone. All non-recipient devices simply pass the frames on.

In summary, the use of redundancy primarily improves dependability, as it offers a degree of immunity to single point failures. Likewise, the new standards on bumpless redundancy offer this redundancy without sacrifice of speed.

MERGING UNIT TIME SYNCHRONISATION

The outputs of merging units are time stamped, such that each sampled value is given a sequential tag which indicates it's ordinal position since the start of the last second. This usually requires the merging units to be accurately time-synchronised. Synchronisation can be achieved thanks to the global positioning satellite system. Synchronising signals are typically delivered over either fibre-optic links in the form of one-pulse-per-second (1pps) signals or over Ethernet according to IEEE 1588.

In instances where a protection IED derives all of its sampled value data from a single merging unit, any loss of the time synchronising input to that merging unit can be configured not to impact the protection scheme. This is because the global real time may be irrelevant, and fallback to use the merging unit's own local time is sufficient. In the event of an antenna failure, it may be that the GPS clock source is lost, and that time synchronising device can also provide a substitute local reference time instead.

There are instances where protection IEDs subscribe to multiple merging units in a substation for their primary function, for example in the case of transformer or busbar differential. In such cases, loss of GPS is not critical, provided that a common local clock reference is available, even if only from a single free-running device. In instances where the protection function runs between substations, such as line differential, accurate time alignment of sampled data from each (all) ends is essential. Particularly in cases such as this, redundant clock sources may be employed at each substation, or if not possible, fallback protection elements such as distance zones may be configured.

CONCLUSION

The introduction of the digital substation and the underlying technologies confronts protection engineering practices with a whole new set of challenges. Thinking about system requirements for speed, dependability and security offers a solid direction of what might be achieved, and what if any compromises need to be reached.