Virtual Protection Relay: Myth or Reality?

Jose Ruiz
Solutions
Doble Engineering
Chattanooga, TN, U.S.

Montie Smith
Global Energy Team
Dell Technologies
Dallas, TX, U.S.

***Abstract*—Not long after the advent of the large-scale power grid over 100 years ago, the need for circuit protection was soon identified and developed. While grid protection began with fuses and later shifted to electromechanical relays, for the last four decades the utility industry has used the most recent iteration of this technology: the microprocessor-based protection relay. Though protection relays have advanced in the last 40 years, operationally they have changed relatively little. However, as we move to a more digital power grid to support the ongoing transition to new forms of clean energy, a new and radically different protection technology is poised to take over: the Virtual Protection Relay (VPR).**

Virtualization is a mature technology that is deployed in IT environments in nearly every industry, allowing a single high-performance server to seamlessly split and share resources to multiple virtual machines running segregated workloads. However, the advantages of this technology are driving development of virtual versions of traditional operational technology (OT) workloads such as substation protection. Communications standards such as IEC 61850 are driving digital OT control networks, rather than the analog systems of the past.

This paper focuses on introducing the reader, with little or no knowledge of VPR, to this new protection technology. Topics such as introduction to virtual protection relaying, hardware requirements for adopting this technology, lessons learned during a computer server setup, and test results comparison of a protection element with VPR and traditional protection relays will be covered.

Keywords—virtual protection relay; VPR; hypervisor; protection relay; GOOSE; SV

I. INTRODUCTION TO VIRTUAL PROTECTION

A. History of protection relaying

Shortly after the development of the large-scale power grid over a century ago, the need for circuit protection became apparent. Initially this protection was provided by fuses, later followed by the earliest prototypes of the electromechanical relay (EMR) around 1905. Early EMR relied on the operation of physical switches, which moved to break an overloaded circuit based on the physical movement of a plate pulled by the magnetic field generated by a relay coil. These relays were bulky and required individual wiring for each circuit meant to be protected, limiting their flexibility and scalability.

Over the last 100 years, there have been many technological developments in protection, with the latest being the development and subsequent adoption of the microprocessor-based relay (MPR) over the last four decades.

MPRs were (as the name suggests) the first to include some level of onboard processing and computerization in the protection system. While the MPR is designed to emulate the principle by which EMRs operate, it does so through the operation of software algorithms. The MPR introduced several advantages over EMR, including increased reaction speed, ability to log data on fault events, and more complex logic that could be applied to one or multiple circuits. They were also smaller and far more scalable. Since their introduction, MPRs have continued to improve their features, such as through the introduction of more complex protection schemes/algorithms and a greater ability to be integrated into a utility's DMS/SCADA system. The complexity of protection systems has continued to increase with the IEC 61850 standard (which allows for the digitalization of analog signals into data streams for protection interpretation) becoming more widespread.

[1] takes the reader through the technological evolution of this marvelous world of protection relaying until the MPRs era.

As we move to a more digitalized power grid to support the ongoing transition to new forms of clean energy, a new and radically different approach to protection technology is poised to take over: the Virtual Protection Relay (VPR).

B. Transition to virtualized protection

Virtualization is a widespread technology utilized in the information technology (IT) space, and it is a standard practice in enterprise architectures in nearly every industry [2]. Virtualization is the process by which a piece of software or operating system (OS) is run from within a virtual machine (VM), rather than installed on a dedicated piece of physical hardware.

From within the VM, the OS, and all software running within the OS, has no awareness that it is not running on a physical computer. This allows multiple VMs to be run on a single server, rather than needing multiple servers to achieve the same number of workloads [3]. This provides several benefits, such as greater resource utilization for the same number of processes, and segregation of workloads to aid management and avoid cascading failures. Fig. 1 below visualizes a comparison of a traditional versus virtualized system architecture.

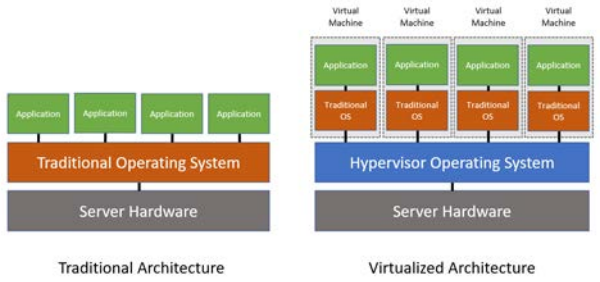

Fig. 1. Traditional vs Virtual system architecture.

Multiple VMs can be stored and run on a single host server that is running a hypervisor. A hypervisor is a specific type of software that allows for the creation, operation, and management of virtual machines, coordinating their resource needs with the underlying hardware [2]. The hypervisor dynamically splits the resources of the physical host system: CPU, memory, storage, networking, etc., among each virtual machine as required by the virtual OS.

One primary benefit of a virtualized architecture is the segregation of individual workloads running on a single piece of hardware. Rather than individual pieces of software running within a single OS on a single computer, software can operate within multiple VMs that are then run on a single server allowing for the individualized management and operation of each service, without influencing other VMs on the host. Additionally, this presents a more cost-effective operating model: as a single server running N workloads at 100% capacity is more efficient than multiple servers each running a single workload at 100/N % capacity.

While virtualization offers numerous benefits in general, the technology delivers many benefits when applied to substation applications specifically. In substations, often the same functionality will be seen repeated across multiple pieces of hardware. One such example is in the substation protection systems. A substation may have between 10 to 30 protection relays, each broadly speaking performing the same functionality. There exists an opportunity, utilizing virtualization, to simplify the protection architecture of a substation from many protection relays to just a few. A cluster of 2 to 3 servers has more than enough operating capacity to perform the functionality of 30 or more protection relays, in a smaller and simpler to manage form factor. This also allows for simpler upgrades and life cycle management, with fewer outages.

C. *Benefits of virtual protection*

There are multiple benefits of virtual protection, which can be summarized as follows:

- Simpler management of protection system (one vs many).
- Smaller protection footprint.
- Fewer outages associate with life cycle management.
- Reduced operations and maintenance costs.
- Improved protection performance.
- Standard platform to add additional virtualized services in the future.
- Less copper cabling and simpler wiring complexity

II. VIRTUAL PROTECTION RELAY HARDWARE

Virtualization is by its very nature meant to be general purpose, assuming minimum specifications are met. An individual VM is hardware agnostic if the underlying hardware meets the minimum requirements for its operation. All aspects of the underlying hardware are abstracted by the hypervisor when resources are presented to the VM.

However, introduction of virtualized protection into a substation requires additional consideration imposed by the specific requirements of the substation environment. Standards such as IEC 61850-3 [4] or IEEE 1613 [5] impose stricter environmental requirements for computer equipment deployed in a substation than those housed in a datacenter. Utilities must consider this when selecting the hardware platform to host their virtual protection deployments. The server hardware used in testing for this paper met all minimum required specifications for the VPR software testing, including meeting IEC 61850-3 standards, allowing it to be deployed in rugged environments such as utility substations.

Some considerations to be aware of when selecting a server platform for virtualized protection include the base hardware specifications related to processor, networking, and storage, as well as environmental/regulatory standards that must be met for deployment in a substation environment.

A. *Processor Specifications*

Ensure that the processor selected is compatible with virtualization. For example, Intel Virtualization Technology (VT-x) or AMD AMD-V are the settings within a server basic input/output system (BIOS) that enables splitting physical cores for virtualized workloads. The BIOS is the base level system that governs the physical hardware settings.

Some functions in substations, such as protection, are time critical, and need to react as fast as possible. The processor chosen for a virtual protection system should be equipped with real-time computing features.

Protection functionality is a modest workload (4 CPU compute cores were required for the VPR software evaluated in this paper), but overall core count chosen by the utility will be driven by the needs of additional workloads intended for virtualization alongside virtual protection. Additionally, a utility may elect to run multiple separate protection related VMs on the same hardware, requiring additional core count.

B. Memory Specifications

Random Access Memory (RAM), or simply "memory," is the data storage pool used to store data being used in active processes in a computer. Data center servers are generally equipped with Error Correction Code (ECC) memory, which is resistant to data corruption. This type of memory is recommended for hardware running critical applications, such as virtual protection.

Overall, memory requirements are dependent on the nature and number of VMs deployed. The VPR software under evaluation in this paper requires 8 GB of memory. Additional VMs deployed to the substation will have their own memory requirements. A safe minimum recommendation for a small virtualization deployment in the substation would be 64 GB.

C. Networking Specifications

Due to the critical nature of substation applications, it is recommended to incorporate parallel redundancy protocol (PRP) into deployed hardware. This networking protocol allows for the seamless duplication of network traffic across two separate networks, protecting against the failure of any single network component. In the event of network failure, the traffic will still reach its destination across the secondary network. This is a "hot-hot" topology, as all traffic is continuously sent across both networks. PRP is based in the IEC 62439-3 standard. [6] provides a deeper understanding for redundancy networks.

If a given device does not support PRP natively or through an add-in card, a PRP redbox may be used to protect this device's traffic from network failure. A PRP redbox is a device that acts as an interface between redundant PRP networks and devices which do not support PRP, protecting their network traffic.

A common time synchronization source is crucial in a digital protection deployment. The multiple merging units (MUs) feeding the VPR must be time synchronized. A MU converts the conventional voltage and current signals from the instrument transformers in the substation to digital signals: sampled values (SV). The server where the VPR will be installed must be capable to support the IEEE C37.238 or IEC 61850-9-3 standards.

D. Environmental Specifications

Substation control buildings are rarely environmentally controlled. For this reason, any virtualization server deployed to a substation should be rated for harsh environments, such as extreme heat and cold ambient temperature. The specific ranges required vary by industry standard, but -10 °C to 55 °C is a common minimum range of operation to consider. Some manufacturers may have offerings beyond this range for more extreme environments.

Some servers used for rugged applications have been designed to operate without fans. These fanless designs offer several tradeoffs. While they do have reduced moving parts, they often have lower-power processors, limiting the maximum capabilities deployable on these systems. From 4 to 6 core CPUs are common on fanless systems. Other servers that do incorporate fans can incorporate higher capability processors due to the addition of forced air cooling. Up to 24 core processors are available for these systems.

E. Power Specifications

Substation A substation server should utilize redundant power supply units (PSU) across multiple circuits, if possible, to protect from system failure to a single failed power supply or circuit. Each single power supply should be rated to handle the entire system load in the event of redundant power supply failure. There should be no interruption or pickup time between primary and secondary PSU transition in the event of a failure in either power supply.

Direct current (DC) power is often preferred for substation control equipment. 48 Vdc and 125 Vdc are the most common voltages found for this equipment. Most rugged servers appropriate for substations may be configured for either alternate current (AC) or DC power. Consider the type of power available onsite when choosing a server platform.

F. Hypervisors

There are many hypervisors available for consideration for a virtualized protection deployment. A utility may consider consulting with their existing IT department to learn which hypervisor is currently used in-house, simplifying transfer of institutional knowledge on hypervisor operation.

The VPR software evaluated in this paper currently supports VMware ESXi and Linux KVM. Consider learning which hypervisor operating systems are supported or recommended by the specific software you are deploying to the substation.

III. VIRTUAL PROTECTION RELAY FIRST IMPRESSIONS

Working with a VPR based protection system differs from traditional protection relays, whether the operator is familiar with EMRs or MPRs. It does share broad overlap in functionality with MPRs, but the actual setup and operation is significantly different.

In traditional relays familiar to most utility operators, visual indicators and tactile buttons (such as LEDs and front panel push buttons) are common. VPR systems generally lack built-in visual and tactile indicators, as the relay itself has been virtualized onto general-purpose computer hardware. Subsequently, a human machine interface (HMI) becomes essential equipment in the substation to allow the user to interface via a web browser with the control system, when a VPR system is put in effect. Since the VPR application has become decoupled from specific vendor hardware, awareness of how to navigate the control system via HMI is critical.

Most traditional protection relays are operational out of the box. EMRs and most MPRs only require connection to the analog inputs and binary inputs and outputs. Relay settings are accessible through the relay itself or with computer software in the case of the most modern microprocessor-based relays. VPR relay settings are only accessible through a computer, this is assuming the VPR device has been delivered with the protection relay application installed in. Otherwise, the protection relay software will have to be installed in the hardware by the user or an original equipment manufacturer (OEM).

Due to the fully digital nature of a VPR system, the testing methodology must change as well. Special testing equipment is not necessary to test VPR since all signals are digital and can be monitored by industry standard software over the network. The most challenging new aspect for relay engineers will be the mental transition from a physical relay to a digital relay that cannot be seen or touched in physical space. This is similar conceptually to the idea of a SCADA system, which is often run in a virtualized environment in the utility's on-premises datacenter.

IV. FIRST STEPS ON SETTING UP A VIRTUAL PROTECTION RELAY

In large utility companies, there is generally a clear definition of roles between the communication and protection teams. However, the adoption of IEC 61850 standard-based protection applications will likely change the status quo of separation between IT and OT teams, leading to protection relay engineers and technicians requiring new familiarity with concepts common in networking and IT.

Given that re-training protection engineers to be IT technicians would be a difficult undertaking, the more prudent course of action taken by many utilities may be to work in close collaboration with their internal IT team. This may be a better approach for easier implementation of this technology. Protection relay engineers could continue to focus on what has been done for decades and allow the IT personnel to assist in server related issues that might arise during the deployment, commissioning, and maintenance of a VPR.

Setting up a VPR could be summarized as follows:

A. Server set up - basic configuration

The first step to setting up a successful VPR deployment is the configuration of the underlaying physical hardware. Specific setup may differ from vendor to vendor, but the most critical steps are those which allow for the seamless installation of a hypervisor OS in step two.

Consideration must be made for cases where the server hardware is equipped with an IPMI (intelligent platform management interface) port. This type of port, common in datacenter hardware, allows for the remote management of the server, including power cycling and installing operating systems remotely. Careful consideration should be taken to ensure this port is isolated on a management network if use by the utility is planned or disabled otherwise.

Station bus as defined in the IEC 61850 standard is intended for control and GOOSE messages communication. Sampled values should not be present on this network. GOOSE and SV are multicast messages at the level 2 of the Open Systems Interconnection (OSI) model [7].

The links between the power grid and the VPR are the network interface cards (NIC). It is important to properly identify them before the VPR application software installation process. It could become a tedious task without a tool in the server to identify them. If a tool is available, plugging in a cable on each port at the time and verifying it with the tool is a convenient method. The tool should be able to display the medium access control (MAC) address of the port. Once each port MAC address is identified, it will be easier to identify the ports that will be allocated in the VPR. A MAC address is unique for each NIC. Ideally, it would be easier if the MAC address for each NIC is available at the time of purchasing the server.

Formatting the disks and choosing the preferred redundancy disk mode is a balance between redundancy and capacity. Given VPR workloads generally do not consume large amounts of data, while also operating as a critical workload, the recommendation would be to optimize for system redundancy over overall capacity. Redundant disc operating modes such as RAID 10, RAID 5, and RAID 6 should be considered to protect the VPR server from failing from single or multiple drive failure.

B. *Hypervisor OS installation*

Each of the listed steps requires more or less effort depending on the amount of knowledge has and experience of the individual. A cost-benefit analysis should be considered prior to selecting a hypervisor to deploy in your VRP configuration. Open-source hypervisor solutions tend to be attractive in terms of cost, but may be costly in the overall process without technical support. On the other hand, commercially available hypervisors have a cost, but it comes with a technical support team that can save time during the overall process set up. Additionally, open-source hypervisors may lack some features available on mature enterprise options.

Commercially available hypervisor solutions offer step-by-step instructions on how to install a full system, on the other hand, open-source solutions will require an in depth understanding of the platform to define the necessary steps to be applied on the deployment of the hypervisor in a server.

C. *VPR application software installation*

Step by step installation instructions are usually provided by the protection relay manufacturer for commercial and open-source hypervisor applications.

V. Overcurrent Protection Element Testing

An instantaneous overcurrent (ANSI/IEEE 50P) protection element was used in this test, because it is a common protection element, available in the different relay technologies chosen for this research and the power industry.

The 50P pickup value was set for 10 A secondary or 2,000 A primary.

The relays, depending on the technology that it was built-in, had contact outputs and GOOSE messages for tripping purposes.

Fig. 2 shows the initial set up for the devices.

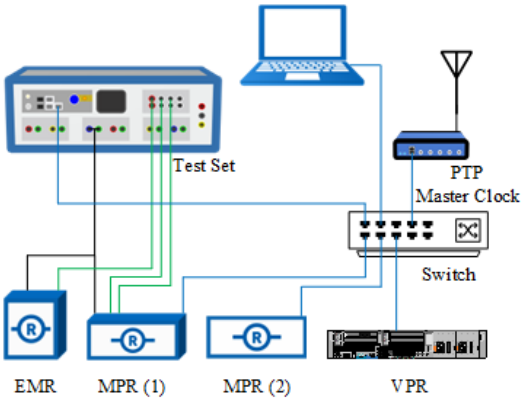

Fig. 2. Relays test set up.

The VPR device is an IEC 61850-3 compliant server with an Intel Xeon Gold 6312U CPU processor running at 2.4 GHz, and has 256 GB of memory RAM. The processor has virtualization technology enabled, and the VPR software is running as a virtual machine within a type-1 hypervisor. The server meets all other specifications laid out in this paper for a recommended VPR system. The Ethernet switch is managed with support for IEEE 1588. It has gigabit ports.

The protection relays from different manufacturers used for this test are as follow:

- Electromechanical relay with a regular contact output (EMR)
- Microprocessor-based relay (MPR (1))
  - Contact output
  - High-speed contact output
  - GOOSE message
- Microprocessor-based relay with a GOOSE message (MPR (2))
- Virtual protection relay with a GOOSE message (VPR)

The power system simulator or test set used for this test simulated current using a conventional current source and sampled values (SV) at the same time. The EMR and MPR (1) are connected to the conventional current source as an input. The MPR (2) and VPR are connected to the SV as a source input. The MPR (2), VPR, and the set were time synchronized to a precision time protocol (PTP) master clock.

The test set sensed the relays operation by checking the contact outputs and generic object-oriented system event (GOOSE) messages status: open, False or close, True for the contact outputs or GOOSE messages respectively.

MPR (1) is a special case because the same relay was configured with three different outputs for indicating the 50P protection element operation. In this case, it has a regular contact output, a high-speed contact output (transistor type), and a GOOSE message.

It is important to highlight that the GOOSE message on both MPRs was configured in such a way that the payload of the data packet was minimized by adding only the necessary data attributes.

The simulation consisted of a pre-fault, fault, and a post-fault state.

A timer was set up in the test set to sense the physical and digital output of the 4 relays. Timers started when the fault state initiated and stopped when the contact output closed or the trip indication contained in the GOOSE message became True.

The test was repeated twenty times to calculate an average operating time.

Four testing scenarios were considered for this paper:

- Fault current 1% above the pickup value.
- Fault current 10% above the pickup value.
- Fault current 100% above the pickup value.
- Fault current 10% above the pickup value with Ethernet traffic conditions in the switch.

The idea behind the four testing scenarios is to check the VPR operation response, when compared to known relay technologies.

A. Fault current 1% above the pickup value

A current simulation of 10.01 A secondary or 2,020 A primary was injected in the relays.

Table I shows the test results for each test run along with the minimum, maximum and average values for all 20 test runs.

The EMR is the slowest one to operate of all follow by the MPR (2). In the case of the MPR (1), the high-speed contact output is a little bit faster than the GOOSE message output from the same relay, and approximately 3 ms faster than the regular contact output. The VPR under analysis in this paper shows the fastest operation time for this test case.

TABLE I. TEST RESULTS FOR A FAULT CURRENT 1% ABOVE PICKUP VALUE IN MILLISECONDS

Test #	EMR	MPR (1) CO	MPR (1) HS CO	MPR (1) GOOSE	MPR (2) GOOSE	VPR GOOSE
1	68.0	26.2	23.6	24.6	35.6	19.9
2	69.2	26.2	23.6	24.6	36.5	19.9
3	72.0	26.2	23.6	24.8	37.5	19.9
4	70.4	26.3	23.6	24.7	35.4	19.9
5	75.9	26.3	23.6	25.0	37.5	19.9
6	73.4	26.2	23.6	24.8	36.3	19.9
7	70.4	26.2	23.6	24.8	35.6	19.9
8	72.9	26.2	23.6	24.8	37.1	19.9
9	72.8	26.2	23.6	24.6	35.6	19.9
10	75.6	26.2	23.6	24.7	36.1	19.9
11	72.2	26.3	23.6	24.6	35.5	19.9
12	71.7	26.2	23.6	24.6	35.5	19.9
13	72.0	26.2	23.6	24.6	35.6	19.9
14	66.5	26.2	23.6	24.7	38.2	20.0
15	69.2	26.3	23.6	24.6	35.5	19.9
16	66.7	26.2	23.6	24.6	38.4	19.9
17	69.8	26.2	23.6	24.7	37.2	19.9
18	71.2	26.2	23.6	24.7	37.2	19.9
19	70.5	26.2	23.6	24.6	35.6	19.9
20	72.7	26.2	23.6	24.6	35.6	19.9
Min. Value	66.5	26.2	23.6	24.6	35.4	19.9
Max. Value	75.9	26.3	23.6	25.0	38.4	20.0
Average	71.2	26.2	23.6	24.7	36.4	19.9

B. Fault current 10% above the pickup value

A current simulation of 10.1 A secondary or 2,200 A primary was injected in the relays. Table II shows the test results for each test run along with the minimum, maximum and average values for all 20 test runs.

The EMR is the slowest one to operate of all followed by the MPR (1) with its regular contact output. In the case of the MPR (1), the high-speed contact output is a little bit faster than the GOOSE message output from the same relay, and approximately 3 ms faster than the regular contact output. The operation time for the MPR (1) and MPR (2) with their GOOSE message output are similar. The VPR under analysis in this paper shows the fastest operation time for this test case.

TABLE II. TEST RESULTS FOR A FAULT CURRENT 10% ABOVE PICKUP VALUE IN MILLISECONDS

Test #	EMR	MPR (1) CO	MPR (1) HS CO	MPR (1) GOOSE	MPR (2) GOOSE	VPR GOOSE
1	34.4	23.7	21.0	23.3	23.4	16.6
2	34.5	23.7	21.0	23.4	23.5	17.3
3	39.0	23.6	21.0	22.1	21.9	16.6
4	42.7	23.7	21.0	22.1	22.2	16.6
5	38.7	23.6	21.0	22.2	22.1	16.6
6	34.4	23.6	21.0	22.1	22.0	16.6
7	39.0	23.6	21.0	22.3	22.2	16.6
8	41.7	23.6	21.0	22.5	22.1	16.6
9	39.3	23.6	21.0	22.4	22.3	16.6
10	34.6	23.6	21.0	22.3	22.2	16.6
11	34.5	23.7	21.0	22.0	22.1	16.6
12	39.0	23.6	21.0	22.1	22.0	16.6
13	41.1	23.6	21.0	22.1	22.3	16.6
14	34.6	23.7	21.1	22.1	22.2	16.6
15	34.4	23.6	21.0	22.0	22.3	16.6
16	43.6	23.6	21.0	22.1	22.4	16.6
17	34.5	23.7	21.0	22.2	22.1	16.6
18	34.4	23.7	21.0	22.2	22.1	16.6
19	43.7	23.6	21.0	22.1	22.2	16.6
20	44.0	23.6	21.0	22.1	22.2	16.6
Min. Value	34.4	23.6	21.0	22.0	21.9	16.6
Max. Value	44.0	23.7	21.1	23.4	23.5	17.3
Average	38.1	23.6	21.0	22.3	22.3	16.6

C. Fault current 100% above the pickup value

A current simulation of 20 A secondary or 4,000 A primary was injected in the relays. Table III shows the test results for each test run along with the minimum, maximum and average values for all 20 test runs.

The MPR (1) with its regular contact output is the slowest one to operate of all follows by its own GOOSE message output and the EMR. The VPR under analysis in this paper is slower when compared to the MPR (2), which is the fastest one to operate in this test case.

TABLE III. TEST RESULTS FOR A FAULT CURRENT 100% ABOVE PICKUP VALUE IN MILLISECONDS

Test #	EMR	MPR (1) CO	MPR (1) HS CO	MPR (1) GOOSE	MPR (2) GOOSE	VPR GOOSE
1	13.6	16.1	13.5	14.5	9.1	12.3
2	21.4	16.1	13.5	14.6	9.0	9.9
3	13.6	16.0	13.4	14.4	9.4	12.5
4	13.6	16.0	13.4	14.4	9.1	9.9
5	13.6	16.0	13.4	16.4	9.0	9.9
6	13.6	16.0	13.4	14.4	9.0	9.9
7	22.7	16.0	13.3	15.3	9.0	12.4
8	13.6	16.0	13.3	15.7	9.3	9.9
9	13.6	16.0	13.3	14.5	8.9	9.9
10	13.6	16.1	13.4	16.2	8.9	9.9
11	13.6	16.0	13.4	14.6	9.0	9.9
12	13.6	16.0	13.3	14.5	9.0	9.9
13	13.6	16.0	13.3	14.5	9.4	9.9
14	13.6	16.0	13.3	14.4	9.2	9.8
15	13.6	16.0	13.3	15.0	9.9	10.1
16	13.6	16.0	13.4	14.4	8.9	9.9
17	13.6	16.0	13.3	15.3	10.2	10.4
18	13.6	16.0	13.3	14.4	9.0	9.9
19	23.2	16.0	13.3	15.2	9.4	9.8
20	13.6	16.0	13.3	14.5	9.2	9.9
Min. Value	13.6	16.0	13.3	14.4	8.9	9.8
Max. Value	23.2	16.1	13.5	16.4	10.2	12.5
Average	14.9	16.0	13.4	14.9	9.2	10.3

D. Fault current 10% above the pickup value with Ethernet traffic conditions in the switch

The intention of this test is to check the impact of Ethernet traffic on the measured operation time of the relays with GOOSE messages outputs. Only one case scenario was chosen out of the three fault currents. Perhaps, a more in-depth analysis of the impact of Ethernet traffic in the GOOSE messages output could be the subject for another paper.

A MU was added to the initial setup and 7 sampled values streams published by this MU. The test set also published 6 more sampled values streams besides the one used for testing the relays. A total of 14 sampled value streams were in the Ethernet network. All SV streams were 9-2 LE. The VPR and Ethernet switch have gigabit Ethernet ports. The MPR (1) and MPR (2) have 100 Mbps Ethernet ports.

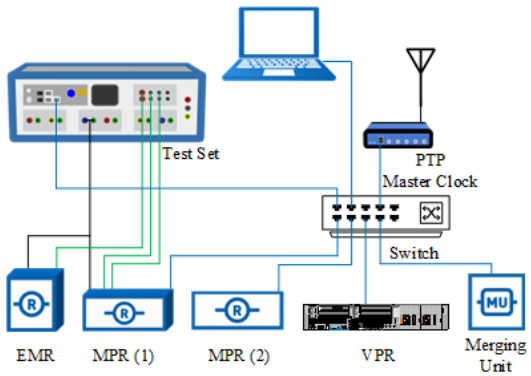

Fig. 3. Relays test setup with added traffic to the Ethernet network.

A current simulation of 10.1 A secondary or 2,200 A primary was injected in the relays. Table IV shows the test results for each test run along with the minimum, maximum and average values for all 20 test runs.

The EMR is the slowest one to operate of all followed by the MPR (1) with its regular contact output. The MPR (1) with its high-speed contact output has a similar operation time as the MPR (2), but faster than its own GOOSE message output.

The VPR under analysis in this paper shows the fastest operation time for this test case.

TABLE IV. TEST RESULTS FOR A FAULT CURRENT 10% ABOVE PICKUP VALUE WITH TRAFFIC CONDITIONS IN MILLISECONDS

Test #	EMR	MPR (1) CO	MPR (1) HS CO	MPR (1) GOOSE	MPR (2) GOOSE	VPR GOOSE
2	34.7	25.0	22.4	23.6	22.2	16.6
3	39.8	25.0	22.4	23.7	22.1	16.6
4	42.5	25.0	22.4	23.7	22.2	16.7
5	41.7	25.0	22.4	23.7	23.6	17.2
6	39.5	25.0	22.4	24.2	24.1	16.6
7	34.8	25.0	22.4	24.3	24.2	16.6
8	40.8	25.0	22.4	23.5	22.3	16.6
9	42.0	25.0	22.4	23.5	21.9	16.6
10	40.7	25.0	22.4	23.5	22.4	16.6
11	38.3	25.0	22.4	23.5	22.1	16.6
12	41.7	25.0	22.4	23.5	22.5	16.6
13	41.2	25.0	22.4	23.5	22.4	16.6
14	41.0	25.0	22.4	23.5	22.2	16.6
15	41.1	25.0	22.4	23.5	22.0	16.6
16	41.6	25.0	22.4	23.5	22.4	16.6
17	43.5	25.0	22.4	23.5	22.1	16.6
18	39.3	25.0	22.4	23.5	22.2	16.6
19	39.5	25.0	22.4	23.5	22.0	16.6
20	41.8	25.0	22.4	23.6	22.2	16.6
21	38.9	25.0	22.4	23.6	22.2	16.6
Min. Value	34.7	25.0	22.4	23.5	21.9	16.6
Max. Value	43.5	25.0	22.4	24.3	24.2	17.2
Average	40.2	25.0	22.4	23.6	22.5	16.6

E. Testing results summary

Table V summarizes all results from each of the four series of tests.

All four relays' operation time was reduced as the simulated fault current was increased from the set pickup value. This is expected for a 50P protection element.

The VPR is the fastest one in the first two test cases: 1% and 10% fault current above pickup value. The MPR (2) is the fastest one in the 100% fault current above pickup test case. The possible explanation for this 1 ms difference could be related to the specific algorithm parameters implemented on each relay.

In the case of the 10% fault current above pickup value with Ethernet traffic and without it, there is no significant operation time difference between the MPR (2) and VPR. The added traffic does not seem to have an impact on the relay's operation time. The MPR (1) with

a GOOSE message output shows approximately 1 ms difference when traffic is added to the network. But notice that the EMR and MPR (1) with the contact outputs and a conventional current input also show a time difference, to which the traffic in the network must not have an influence.

In summary, testing shows that the VPR operates faster or similarly to the known relays' technology: electromechanical and microprocessor-based relays.

TABLE V. AVERAGE TEST RESULTS SUMMARY IN MILLISECONDS

Fault Above PU Setting	EMR	MPR (1) CO	MPR (1) HS CO	MPR (1) GOOSE	MPR (2) GOOSE	VPR GOOSE
1%	71.2	26.2	23.6	24.7	36.4	19.9
10%	38.1	23.6	21.0	22.3	22.3	16.6
100%	14.9	16.0	13.4	14.9	9.2	10.3
10% with traffic	40.2	25.0	22.4	23.6	22.5	16.6

VI. CONCLUSION

Similar to the initially-slow adoption of the microprocessor-based relay, virtual protection will likely face resistance to widespread adoption at the outset. Until there is a greater body of successful use cases and institutional experience with the technology, there will of course be some hesitation. This paper intends to be one such body of knowledge, demonstrating that while the hardware and deployment may be different, the performance of a VPR based system exceeds that of traditional mechanical and microprocessor-based relay protection. Many early adopters of this technology may feel compelled to deploy VPR in limited pilot projects before widespread deployment. Additionally, the skills needed for the successful deployment of a VPR system require knowledge of both the IT and OT domains, converging the two technologies that have traditionally been managed separately.

VII. REFERENCES

[1] B. Lundqvist, "100 years of relay protection, the Swedish ABB relay history," ABB Automation Products.

[2] IBM, "What is virtualization?," IBM, [Online]. Available: https://www.ibm.com/topics/virtualization.

[3] Red Hat, "What is Virtualization?," 2 March 2018. [Online]. Available: https://www.redhat.com/en/topics/virtualization/what-is-virtualization. [Accessed 20 February 2024].

[4] "Communication networks and systems for power utility automation - Part 3: General requirements," IEC 61850-3, 2013.

[5] "IEEE Standard Environmental and Testing Requirements for Communications Networking Devices in Electric Power Substations," IEEE 1613, 2003.

[6] "Communication networks and systems for power utility automation – Part 90-4: Network engineering guidelines," IEC/TR 61850-90-4 , 2013.

[7] "Communication networks and systems for power utility automation – Part 8-1: Specific communication service mapping (SCSM) – Mappings to MMS (ISO 9506-1 and ISO 9506-2) and to ISO/IEC 8802-3," IEC 61850-8-1 , 2011.

VIII. BIOGRAPHIES

Jose Ruiz is a Principal Technical Application Engineer for Doble Engineering based in the United States of America. He previously worked as a protection application engineer with ABB supporting pre-sales and post-sales of transmission and sub transmission level protection relays. He received his M.S. degree from the University of Tennessee at Chattanooga in 2012. During his graduate study, he learned and tested microprocessor-based relays using the IEC 61850 communication standard with different vendors. Since then, Jose has shared his knowledge in this subject through conference presentations, trainings and assisting with the development of new products related to this matter. He is an active senior member of the IEEE PSRCC.

Montie Smith is a Business Development Executive for Dell Technologies supporting Dell's energy and utility business in North and South America. He joined Dell in 2022 as an Energy Solutions Specialist, working to develop OT solutions for the energy industry. He previously worked as an applications engineer with ABB supporting the sale of distribution switchgear and recloser products in the Southern United States. He received his Bachelors (2016) and Masters (2018) in Electrical Engineering from the University of Tennessee. During his time as a graduate research assistant, he studied power electronics and researched TSN communications and their applicability in smart grids and microgrids.