SAND2023-03283A

Development of a General, Modular, Reprogrammable Information Barrier for Arms Control Applications

Jay K. Brotz, J. Kyle Polack, Rachel Helguero, Michael Hamel, Thomas Weber, Peter Marleau

Sandia National Laboratories, Albuquerque, NM

Abstract

Next-generation treaty verification will require a wide range of enabling capabilities to support future transparency in monitoring systems. Trusted capabilities must be developed to support inspections, monitoring, and confirmation of warheads while protecting sensitive information such as warhead design and facility operations. Prior trusted system development efforts have often incorporated information barriers (IBs), which help address host certification concerns (safety and security of equipment used in host facilities) but make authentication (inspector process to gain trust in monitoring equipment) challenging. For all but the simplest of measurements, some type of complex processing device is required to operate the system and process acquired data. This complexity makes authentication and certification of measurement systems and collected data challenging and time consuming. To meet that challenge, we present our development of a general-purpose, extrinsic information barrier that will protect against the release of sensitive information collected from sensor measurements while providing inspectors confidence in the measurement results. Our design will use a field programmable gate array (FPGA) to operate the system, which provides the requisite flexibility for all required operations while reducing extraneous functionality found in a microprocessor that could potentially be exploited. By developing hardware that can serve as an information barrier for a range of different verification measurements, we will eliminate the need to develop authentication and certification procedures for multiple independent systems. Additionally, because it is reprogrammable, we will be able to leverage the same IB hardware for systems with different sensors and facilitate the re-use of firmware portions that are more generally applicable (such as waveform processing), further streamlining authentication and certification procedures.

Introduction and Motivation

Next-generation arms control will require a wide range of enabling capabilities to support future transparency in monitoring and verification. Such agreements would require on-site inspection, monitoring, and verification of sensitive materials within a nation's nuclear weapons stockpile that *does not compromise sensitive nuclear weapons design information or significantly impact operations at national security sites* (confirmed through host certification of equipment), while at the same time *providing for credible verification* (confirmed through inspector authentication of equipment). If warhead confirmation measurements are required for verification of potential future arms control treaties, trusted systems will be required to ensure that critical design information is protected. We define a trusted system as a verification tool that the host country has certified for use in their facilities or to measure their sensitive assets and that an inspecting country has authenticated so they trust the system's output and conclusions.

Trusted systems proposed for treaty verification often use an information barrier (IB) designed to perform one specific analysis task for one specific instrument, communicating only an affirmative

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. Approved for unlimited release. SAND2023-03283A

or negative response to the inspector. This response is based on analysis of data collected by the instrument that is highly likely to be sensitive and therefore not shareable with the inspector. These responses may indicate attributes of a warhead, such as whether or not a minimum mass of fissile material is present or whether or not a declared warhead signature matches that of trusted reference item. For all but the simplest of measurements, some type of complex processing device is required to control the system and process acquired data. This complexity makes authentication and certification (A&C) of measurement systems and collected data challenging and time consuming. As a range of capabilities may be required to support future treaties, stakeholders would benefit from the development of a modular information barrier that could be used in conjunction with a variety of sensors to support numerous use cases, or measurement modalities, while reducing the variety of complex equipment requiring authentication. Thus, as new phenomena are developed to support warhead or material confirmation, our information barrier can be used to streamline the maturation of the technique for use in treaty monitoring and verification. Examples of desired sensors may include gamma-ray spectrometers [1], neutron detectors [2], neutron multiplicity counters, or combinations of these instruments. Examples of analyses include template matching (to a trusted reference) [3,4] and attribute measurement [1,5,6].

We present here the development of a general-purpose, extrinsic information barrier that protects against the release of sensitive information collected from sensor measurements. An extrinsic information barrier is a component that is separable from the detector and that can be added onto it in order to provide trustable conclusions without revealing the sensitive measurement data or any other sensitive information about the measurement. This arrangement is in contrast to intrinsic information barriers, in which the detection system itself is designed to minimize (or eliminate altogether) the collection of sensitive data. The reprogrammable aspect of the information barrier allows it to be easily adapted for a variety of use cases. This adaptability is advantageous, as the design could serve as a platform from which verification tools for hypothetical future treaties are built, regardless of the needed sensor or analysis technique. This work builds upon numerous prior efforts such as the Information Barrier Experimental (IBX), created by faculty and students at Princeton University [7], the UKNI IB created as a collaboration between the UK and Norway [5], and the Sandia-developed Trusted Radiation Identification System (TRIS) and Trusted Radiation Attribute Demonstration System (TRADS) [1,3]. Our design relies on a field programmable gate array (FPGA) for data processing, which provides the requisite flexibility for needed functionality, but presents authentication challenges that will need to be addressed in our design. We hypothesize that the limited, but still adequate, amount of functionality provided by an FPGA may be advantageous for A&C compared to that of a microprocessor, which would contain more extraneous functionality. Some of the A&C challenges have been considered in prior work focused on FPGA authentication [8]. By developing hardware that can be used for a range of different verification equipment, we will eliminate the need to develop A&C procedures for multiple independent systems. Additionally, leveraging the same IB hardware for systems with different sensors will facilitate the re-use of firmware portions that are more generally applicable (such as pulse processing), further streamlining A&C procedures.

Design Principles

Confirmation measurements of a treaty-accountable item (TAI) have been proposed for many verification scenarios. Measurements can be used to determine a variety of TAI characteristics using a variety of signals and physical phenomena, often (but not exclusively) based on ionizing radiation. For these confirmation measurements, the purpose of any information barrier is to provide sufficient confidence in the result of the measurement to an inspecting party while providing sufficient information protection for the host. Part of this confidence on both sides comes from knowing that the system, including the information barrier, is designed to correctly perform the intended function and is not capable of performing any other functions. The other part of confidence comes from knowing that the actual physical system used in a measurement is a faithful realization of the trusted design, containing no more, or no less than the agreed upon functionality. A measurement system with an information barrier must allow host certification to provide confidence in information protection (as well as safety and security to meet facility regulations) and inspector authentication to provide confidence in measurement results. Both certification and authentication rely on inspections at a number of points in the lifecycle of the verification equipment. The primary design principles for the modular reprogrammable information barrier, therefore, are flexibility, inspectability, and modularity.

Flexibility

There are many types of TAI confirmation measurements that could involve sensitive data. For example, passive radiation signatures are typically sensitive, meaning that count rates and energy spectra are not shareable with an inspecting party. These use cases, while potentially useful for confirming that an object is (or in some cases, that an object is not) a nuclear warhead or warhead component, would not typically be allowable without an information barrier. We aim to develop an information barrier that can be employed to facilitate what would be otherwise unallowable measurements. This goal means that our IB needs to be able to handle data input from a variety of sensors and run a variety of algorithms to determine measurement conclusions using those data. For our initial effort, we have chosen four use cases to develop and test: low-resolution gamma spectroscopy for template matching, high-resolution gamma spectroscopy for attribute measurement, neutron multiplicity counting (either for attribute measurements or template measurements), and multi-modal measurement involving the simultaneous collection and analysis of gamma spectroscopy and neutron multiplicity counting data. In addition, the system will be designed to allow additional sensor interfaces and run additional algorithms in the future.

Inspectability

Our design must facilitate inspection to be useful in a treaty verification context. For an instrument containing an extrinsic information barrier, both the host and the inspector must be able to confirm that the measurement equipment and the IB are built as expected with no additional functionality or interfaces. While the host and inspector have different aims, the inspections used to verify the equipment will often be the same. Inspectability appears in MRPIB in two ways: concerning inspections of the design and concerning inspections of the realized (as-built) system. The design of both hardware and firmware should be inspectable so that each party can confirm that, if built to the

agreed-upon design, the system will exhibit the expected functionality and only the expected functionality. Inspectable design is that which is easy to read and understand with clear interfaces, inputs and outputs, behavior, and intent. The realized system should be inspectable so that each party can confirm that the system that will be used in an inspection matches the previously agreed upon design and contains no additional features. Inspectable hardware is easy to see (visually or with technological aid using magnification or non-visible methods such as x-ray radiography or thermal imaging) and test functionality. Inspectable firmware is easy to extract in a trusted way (i.e., with the knowledge that the extracted firmware is actually loaded onto the hardware) and to compare to a trusted reference.

Modularity

In order to provide a flexible and inspectable system, our design is modular. Both hardware and firmware designs are subdivided into small portions that are well documented. The firmware will need to be different for each use case employed, but these modular blocks allow the reuse of functions in firmware between use cases and reduce the time needed to develop new use-case firmware in the future. Modularity also enables inspectability by reducing the complex firmware to easily understandable and testable (i.e., confirmable) blocks. We envision a library of basic firmware blocks that can be used in a variety of designs in the future.

Conceptual Design

The MRPIB meets the goals of flexibility, inspectability, and modularity with a split hardware design that includes both a “Red Side” (for sensitive processing) and a “Black Side” (for non-sensitive processing) for each party, as described in the System Architecture section below. In addition, it meets these goals with firmware developed for a field-programmable gate array (FPGA) written in the VHDL[i] language that can be used to configure the functionality of the information barrier immediately prior to measurement, which is described in the Firmware Design section below.

System Architecture

As shown in Figure 1, the system architecture of MRPIB includes a Red Side and two Black Sides: one for each party. The Red Side collects, stores, and operates on sensitive data. The Black Sides display non-sensitive information to the host and inspector and allow a limited set of commands and other input (such as hash challenges) to be send to the Red Side.

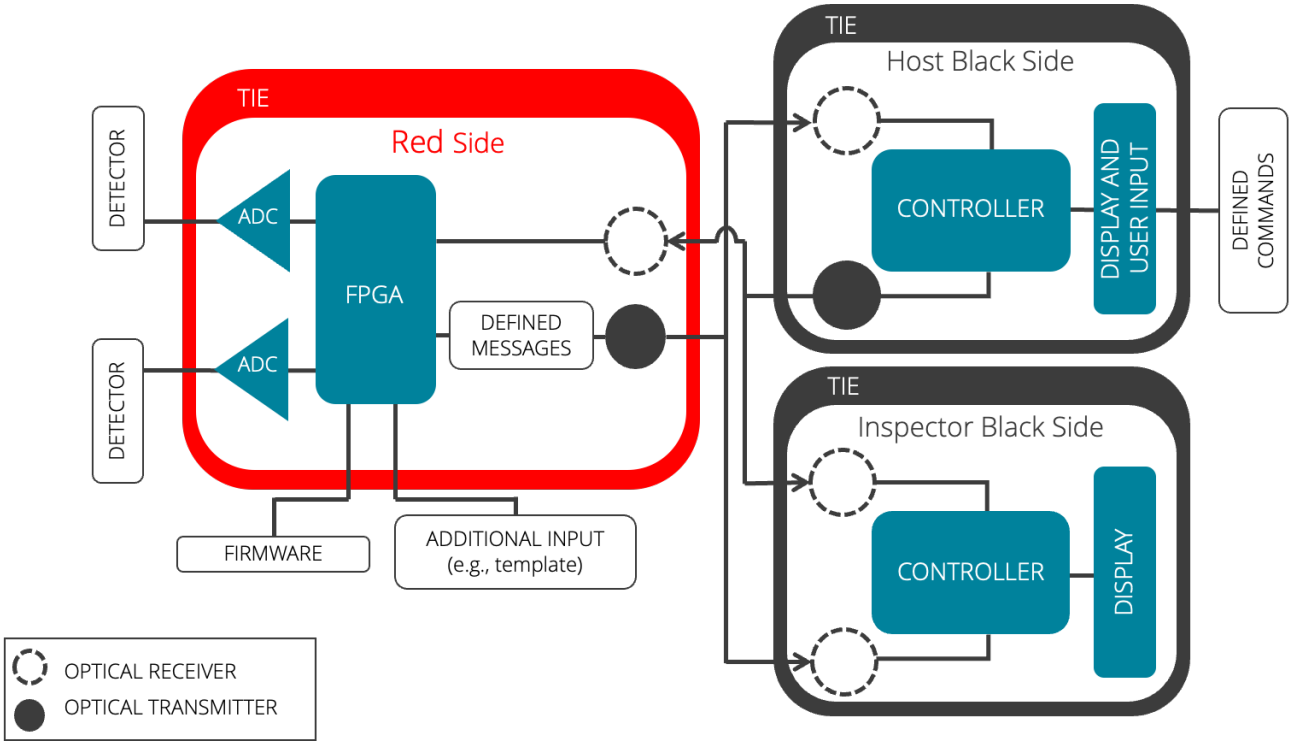

Figure 1: MRIPB System Architecture

The Red Side hardware is provided by the host and the Red Side firmware is provided by the inspector, allowing each side to trust the most critical part of the system for their own A&C goals. The host is most concerned about sensitive information being given, in some way, to the inspectors. By providing the physical hardware for the Red Side, the host can have increased confidence that there are no side channels for transmitting information from the FPGA or bypassing the FPGA altogether. The inspector is most concerned about a correct measurement and measurement result being reported. By providing the firmware, the inspector can have increased confidence that the algorithm analyzing the incoming data is correct and is sending the correct messages to the Black Sides. This is a new concept that we present with a hypothesis that it provides more confidence overall to the host and inspecting parties. As a contrast, the TRIS concept for provision includes the inspector providing all of the Red Side and the host providing all of the (single) Black Side [3].

The Red Side consists of detector interfaces (we will implement two in the initial prototype, though a single interface or more than two interfaces could be accommodated by the architecture) with a fast analog-to-digital converter (ADC) on each input, an FPGA for processing the detector data, an optical transmitter and receiver for communicating to the Black Sides, and only the additional electronics needed to support these components. In addition, there are two more interfaces for templates of reference items (or other additional inputs needed to support future use cases) and firmware loading. Given that the design relies on inspector provision of firmware, an interface is needed to securely load that firmware from an external digital media source. No non-volatile memory exists on the Red Side, so that when the system is powered down, no sensitive data should remain. The Red Side hardware is inside a tamper-indicating enclosure (TIE).

There is one Black Side provided by each party. The interfaces are limited to optical inputs and outputs from and to the Red Side, with limited defined messages permitted, a user interface display, and a mechanism for input on only the host Black Side for system control (and potentially for inputting a hash challenge sequence or something similar). Both Black Sides could be housed within TIEs.

Concept of Operations

Prior to use, the MRPIB design is authenticated and certified. The inspector authenticates the Red Side hardware design by examining drawings and other design information, while the host certifies the Red Side firmware by examining the VHDL files. These design verifications have the same goal: confirm that the functionality that both parties have agreed upon is reflected in the designs and that no additional functionality or interfaces are present. Following design verification, both sides may send a built unit (physical hardware in the case of authentication and a synthesized bitfile for the target FPGA on the physical hardware in the case of certification) that will allow each side to confirm through examination or testing that the implementation still exhibits the agreed functionality with no additional functionality or interfaces. In addition, both sides can examine the design and implementation of the other party's Black Side.

Prior to an on-site inspection using the MRPIB connected to a detector, the inspector verifies the Red Side hardware and the host Black Side, and the host verifies the Red Side firmware and the inspector Black Side. These activities include confirming that the hardware and firmware presented for use match the trusted hardware and firmware that was already inspected. Then the inspector may apply chain of custody (CoC) measures, such as seals or unique identifiers, to the Red Side hardware in order to trust that it will not be modified for the duration of the inspection. The inspector loads the firmware onto the hardware in a way that gives the host confidence that the trusted firmware is actually configuring the FPGA. The Black Sides are connected to the Red Side by the optical connections, and the system is ready for use.

During use, to give the host increased confidence that the inspector cannot gain any sensitive information, the inspector Black Side can be shielded from the view of the inspectors until the host has reviewed the user interface display. This step may not be necessary if the host has sufficient confidence that the firmware cannot pass sensitive data to the Black Sides.

Following the on-site inspection, the host retains all equipment, including the MRPIB, to ensure that any data remanence cannot be recovered by the inspector.

Use Cases

While the MRPIB is designed for the flexibility to be used on many use cases, we have chosen four examples to implement in the first phase of this work. Those four use cases are:

1. Template matching with low-resolution gamma spectroscopy
2. Attribute measurement with high-resolution gamma spectroscopy
3. Attribute measurement or template matching with neutron multiplicity counting

4. Multi-modal attribute measurement using both gamma spectroscopy and neutron multiplicity counting.

Many attribute measurement systems rely on multiple measurements of different physical signatures running different analysis algorithms. One benefit of an FPGA-based information barrier is the inherent parallelization of FPGAs; they are capable of performing many simultaneous tasks on one or more inputs.

Firmware Design

The firmware will be modular, breaking down complex computations into simple functional blocks, to provide both flexibility to implement many use cases with a set of common blocks and making the firmware more easily understandable and therefore more inspectable. The Red Side interfaces to a variety of radiation detectors that have pulse outputs and digitizes the pulse waveforms so that all pulse processing is done in the digital domain. Thus, the firmware blocks can be broken into five major tasks: pulse processing, calibration, histogramming, analysis, and utility functions. Each task is shown in Figure 2 through Figure 6 with a collection of potential blocks (the blocks shown could change by the completion of initial development). Additional blocks can and will be created for future use cases, with the goal of reusing blocks that have already been written.

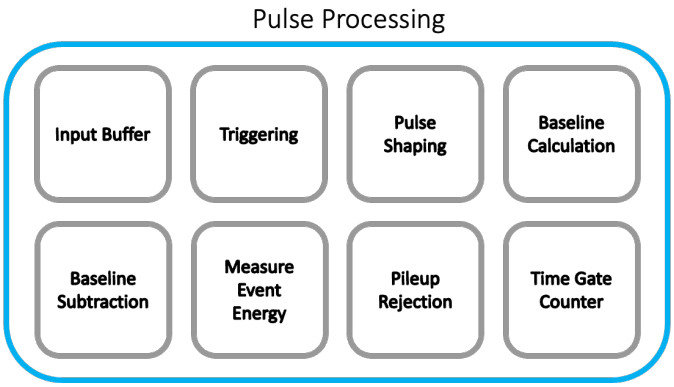

Figure 2: Pulse Processing Firmware Blocks

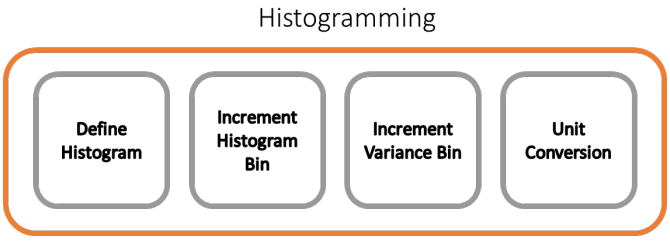

Figure 3: Histogramming Firmware Blocks

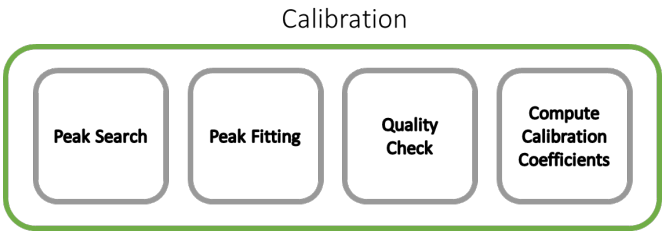

Figure 4: Calibration Firmware Blocks

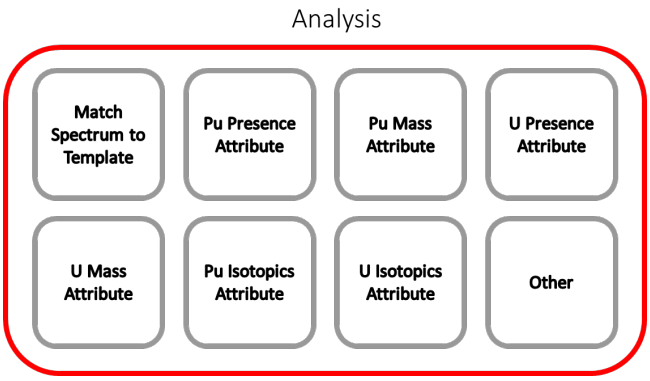

Figure 5: Analysis Firmware Blocks

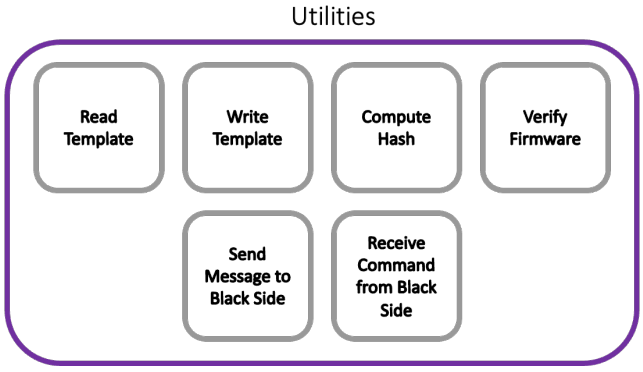

Figure 6: Utility Firmware Blocks

These blocks will be combined as needed to perform a particular use case. For example, the template matching using low-resolution gamma spectroscopy use case will use the blocks displayed in Figure 7. As seen in the diagram, while there are many functions needed to fully execute the use case, many of the blocks are re-used even within this use case, limiting the required firmware development.

We envision a library of VHDL functional blocks that are flexible enough through the use of global variables (called generics in VHDL) to be used in a number of ways in a variety of use cases. Blocks from this library can be combined to meet the needs of future use cases, requiring only the creation of blocks unique to that use case and a top-level VHDL file that combines the blocks into a single circuit.

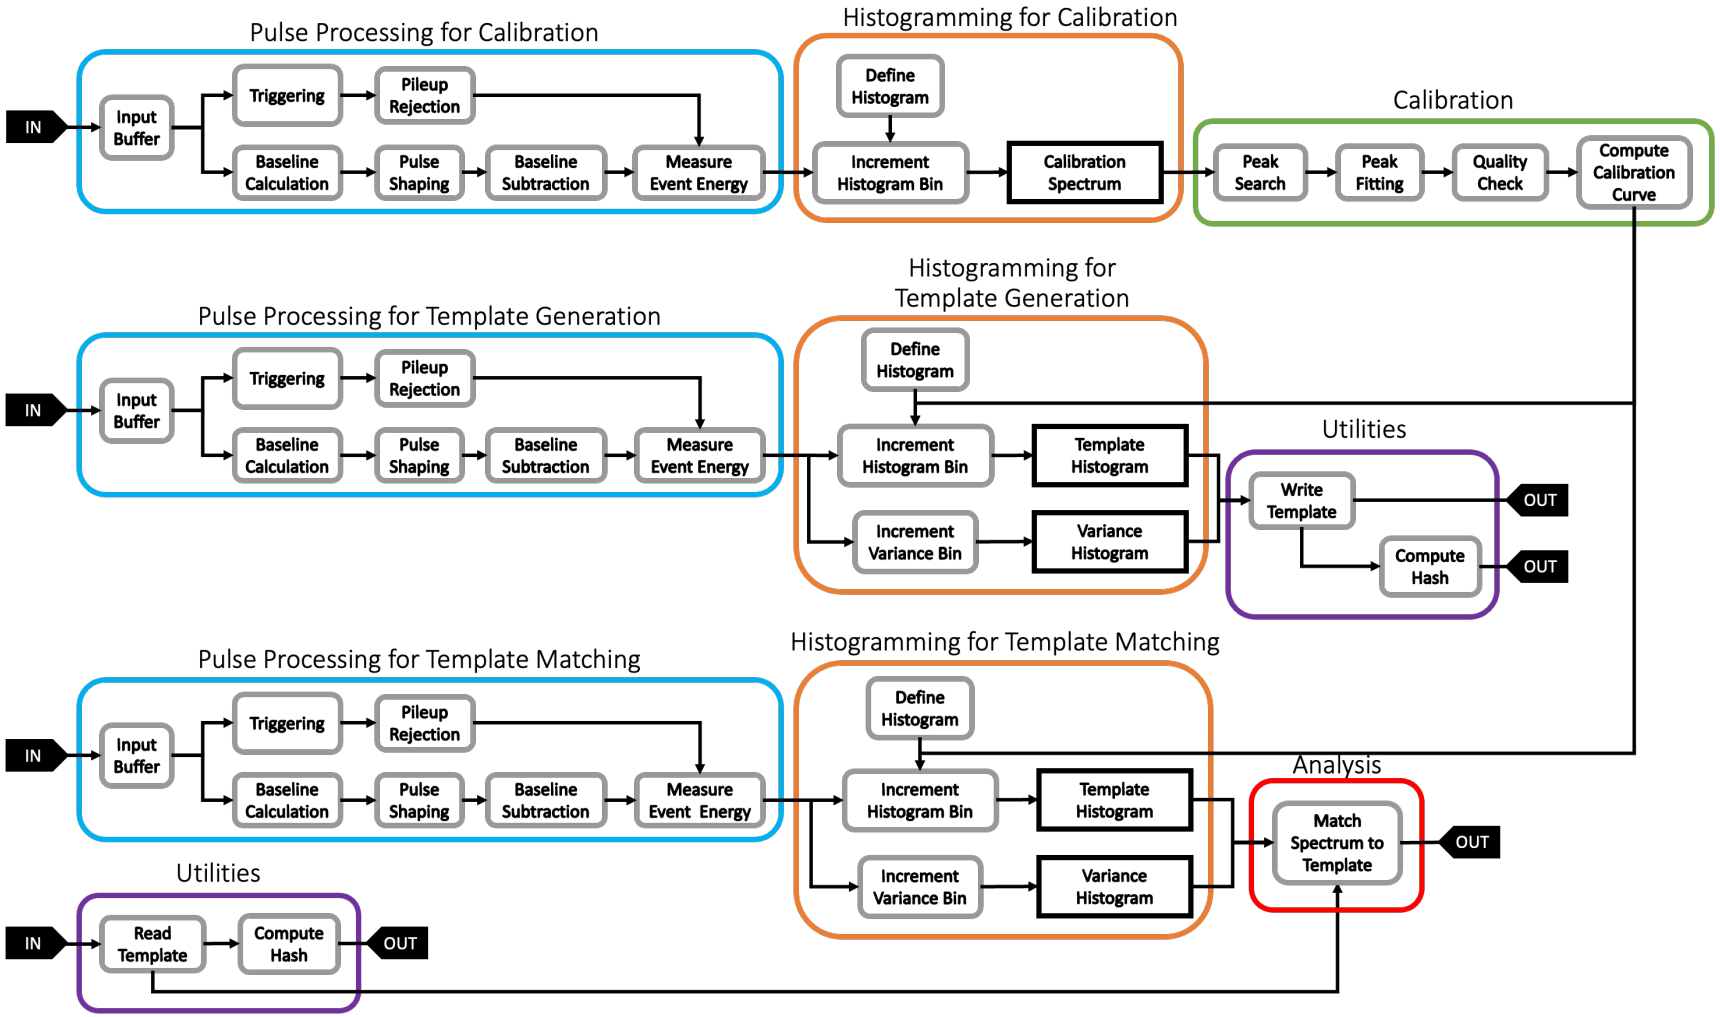

Figure 7: Firmware Block Diagram for Template Matching Use Case

Conclusion

The Modular Reprogrammable Information Barrier provides a solution for ensuring host and inspector confidence for a variety of warhead confirmation technologies and use cases. It allows the host to certify the measurement system, generating sufficient confidence that sensitive information is not revealed to inspectors. It allows the inspector to authenticate the measurement system, generating sufficient confidence in the measurement analysis and conclusion of the confirmation algorithm. We envision this system reducing the time needed to develop a mature treaty verification system when such a system is needed for a future treaty, after further development by scientists with expertise working on warhead confirmation detection methodologies.

Acknowledgements

We recognize the generous support of the U.S. Department of Energy National Nuclear Security Administration (NNSA) Defense Nuclear Nonproliferation Research and Development office.

References

[1] D. J. Mitchell and K. M. Tolk, "Trusted Radiation Attribute Demonstration System," Proceedings of the Institute of Nuclear Materials Management 41st Annual Meeting, New Orleans, 2000.

[2] R. I. Ewing and K. W. Marlow, "A fast-neutron detector used in verification for the INF Treaty," Nuclear Instruments and Methods in Physics Research A, vol. 299, pp. 559-561, 1990.

[3] P. B. Merkle, T. M. Weber, J. D. Strother, J. Etzkin, A. J. Flynn, J. C. Bartberger, W. C. Amai and L. F. Anderson, "Next Generation Trusted Radiation Identification System," Proceedings of the INMM 51st Annual Meeting, Baltimore, 2010.

[4] P. Marleau, R. Krentz-Wee and P. Schuster, "Proof of concept demonstration of CONFIDANTE (CONfirmation using a Fast-neutron Imaging Detector with Anti-image Null-positive Time Encoding)," Proceedings of the INMM 59th Annual Meeting, Baltimore, 2018.

[5] D. Keir, D. M. Chambers, S. Høibråten, S. Backe, S. Allen and H. E. Torkildsen, "UK-Norway Initiative: Research into Information Barriers to Allow Warhead Attribute Verification Without Release of Sensitive or Proliferative Information," Proceedings of the INMM 51st Annual Meeting, Baltimore, 2010.

[6] G. White, “Review of Prior U.S. Attribute Measurement Systems,” Proceedings of the INMM 53rd Annual Meeting, 2012.

[7] M. Kütt, M. Göttsche, and A. Glaser, “Information Barrier Experimental: Toward a Trusted and Open-source Computing Platform for Nuclear Warhead Verification”, Measurement, Volume 114, 2018, pp. 185-190.

[8] J. Brotz, R. Hymel, N. Grant, and N. Evans, “Framework for Evaluating Authentication Methods for Treaty-Related Processor Systems,” Proceedings of the INMM 57th Annual Meeting, 2016.

[i] VHDL stands for VHSIC Hardware Description Language, where VHSIC stands for Very High Speed Integrated Circuits Program – the U.S. Department of Defense program in the 1980s that created the language.