

Guidelines for FPGA Gateway Development in LHCb

Alberto Perro^{1,2}, Mitja Vodnik¹, Paolo Durante¹, Guillaume Vouters³, Federico Alessio¹

¹*CERN, Geneva, Switzerland*

²*Aix Marseille Univ, CNRS/IN2P3, CPPM, Marseille, France*

³*Laboratoire d'Annecy de Physique des Particules, Annecy, France*

Abstract

This technical note outlines best practices and methodologies for FPGA development, with a focus on coding standards, verification techniques, and reusable design components tailored for LHCb gateway community. It introduces coding guidelines to ensure consistency, readability, and maintainability in FPGA designs, followed by an exploration of simulation methods and formal verification techniques to guarantee functional correctness and comprehensive coverage. The *colibri* library is presented as a standardized collection of reusable components, enabling portability and reliability across diverse applications. Additionally, the integration of testbenches into Continuous Integration pipelines is discussed, providing automated testing and feedback to maintain code quality throughout the development lifecycle. By combining these approaches, the technical note aims to establish a robust framework for efficient and reliable FPGA development, fostering collaboration and innovation within the LHCb community.

1 Introduction

This technical note provides a comprehensive guide to FPGA development practices intended to support and standardize future developments, focusing on coding standards, verification methodologies, and reusable design components tailored for LHCb gateway development. It begins with coding guidelines for FPGA designs in Section 2, emphasizing standardized practices to ensure readability, maintainability, and collaboration. Section 3 explores simulation techniques, particularly Register Transfer Level (RTL) simulations, as a traditional method for verifying gateway functionality. Formal Verification (FV) is introduced in Section 4, highlighting its advantages and optimization techniques. Section 5 presents the *colibri* library, a standardized collection of reusable components designed to enhance portability and reliability in FPGA designs. Finally, Section 6 discusses the integration of testbenches into Continuous Integration (CI) pipelines, ensuring automated testing and feedback to maintain code quality throughout the development process. Together, these sections provide a structured framework aimed at enabling efficient and reliable FPGA development for future projects in LHCb.

2 Codebase

This section outlines the coding guidelines for FPGA designs in LHCb. Section 2.1 details file structure, formatting, and naming conventions. Architectural design methods aimed at improving readability and debugging are discussed in Section 2.2. Finally, Section 2.3 emphasizes the importance of documentation for ensuring clarity and maintainability.

2.1 VHDL Guidelines

Adhering to standardized coding guidelines is crucial for maintaining readability, consistency, and efficiency in software development. By following conventions for indentation, naming, and commenting, developers can ensure that their code is easier to understand and collaborate on, especially in team environments like LHCb, where multiple programmers contribute to the same project. Consistent coding practices also reduce the likelihood of errors and improve maintainability, enabling future developers to debug or modify the code more effectively. Ultimately, well-structured and standardized code results in higher-quality software that is easier to develop, test, and maintain.

The VHDL Style Guide [1] tool can automatically enforce the following coding style rules using a YAML configuration file, which is available in the [GitLab Repository](#).

File Rules

- File Naming** (mandatory): The name of the file must match the name of the design unit contained in the file followed by the `.vhd` or `.vhd1` extension.
- File Extension** (recommended): The file extension should be `.vhd1`. All tools and OSs support more than 3 letters extension and `.vhd` is already used for virtual hard disk.
- File Content** (recommended): A VHDL File should contain either:

- 39 • An entity and its architecture, the entity must not have multiple architectures.
- 40 • A package declaration, if no body is required.
- 41 • A package declaration and package body.

42 4. **File Header** (mandatory): Each source file has to start with a comment to describe
43 the content of the file and the license. Doxygen syntax should be embraced to
44 automate the generation of documentation. Header template:

```
45     --! @title Round-Robin Arbiter  
46     --! @file arbiter.vhdl  
47     --! @author Alberto Perro (alberto.perro at cern.ch)  
48     --! @date 09-01-2024  
49     --! @version 0.1  
50     --! @copyright CERN 2024  
51  
52  
53     --! This module arbitrates a bus connected to a resource.  
54     --! Inspired by https://github.com/chclau/arbiter_rr  
55     --! Release log:  
56     --! - 0.1 first release  
57     --!  
58     --! **License**: GPLv3  
59
```

60 The ! character at the start of a comment highlights the comment for the automated
61 documenter ([TerosHDL](#)).

62 5. **Line Length** (Mandatory): The recommended line length is 80 characters, the
63 maximum is 120 characters.

64 6. **End of Line** (Mandatory): A single newline \n character should be used following
65 the Unix convention.

66 7. **Language** (Mandatory): Always use the English language in your comments and
67 identifiers.

68 8. **Avoid Special Characters** (Mandatory): Use always the plain 7-bit ASCII
69 encoding. If special chars are necessary, use their correspondent hex encoding.

70 9. **No tabulation** (Mandatory): Horizontal tabulations are not allowed. Indentation
71 must be done with spaces.

72 10. **Last line in a File** (Mandatory): The last line must finish with an end of file and
73 must not be empty.

74 11. **Trailing Spaces** (Mandatory): No trailing spaces are allowed.

75 **Format Rules**

76 1. **Comment style** (Mandatory): Comments should be placed on dedicated lines,
77 although short comments may be added at the end of a line to provide hints.
78 Comments must start with -- followed by a white space, except when a ! immediately
79 follows the --, which serves as an identifier for the documenter.

80 2. **Indentation** (Mandatory): Indentation must be done using two white spaces and
81 applied to:

- 82 • Declarative items,
- 83 • Nested concurrent or sequential statements, and
- 84 • `use` clauses of a library.

85 3. **Spaces** (Mandatory): There must be at least one white space:

- 86 • Before and after `:` in declarations,
- 87 • Before and after `:=` and `<=` in assignments and default values,
- 88 • Before and after comparison operators, and
- 89 • Before and after `=>` in named associations.

90 The number of white spaces can be more than one for alignment purposes. However,
91 there must be **no** white space:

- 92 • Before `,` and `;`,
- 93 • Before `(` when used for conversion, function calls, indexing, or slicing, and
- 94 • Between `process` and `(`.

95 4. **Context clauses** (Mandatory): Context clauses must be organized into groups,
96 starting with a library clause followed by `use` clauses and a blank line. The library
97 clause should be omitted for `std` and `work`. Each `library` clause must contain only
98 one library name, while each `use` clause must contain only one selected name. A `use`
99 clause should make an entire package visible, written as the library name followed by
100 the package name and `all`. The `std.textio` clause must be the last `use` clause in
101 the `ieee` group, if present. The first group should be for the `ieee` library, followed
102 by vendor libraries, project libraries, and finally `use` clauses for the `work` library.

103 5. **Use of context keyword** (Mandatory): Due to incompatibility with Xilinx Vivado,
104 the use of the `context` keyword is not allowed in files used for synthesis. However,
105 simulation-only files such as testbenches may use it.

106 6. **Entity declaration layout** (Mandatory): An entity declaration must follow this
107 layout: `entity`, white space, entity name, white space, `is`, and a new line. If there
108 are generic interfaces, they must be declared one per line, aligned to the group if
109 separated by comments. The `(` should be on the same line as the `generic` keyword,
110 followed by a new line, and the `)` should be on a new line with correct indentation,
111 followed by `;` and a new line. These rules also apply to port declarations.

112 7. **Structures** (Mandatory): Each structure requiring an `end` keyword (e.g., `if`, `for`,
113 `while`, `entity`, `procedure`) must be closed with the `end` keyword followed by the
114 name of the structure and, if present, the label of the structure. Loops are an
115 exception, as they are closed with the `loop` label. For example:

```
116     for i in v a l range loop  
117     ...  
118     end loop;
```

121
122
123
124
125
126
127

```
procedure pack
(signal arg : std_logic) is
begin
...
end procedure pack;
```

128
129
130
131
132
133
134
135
136
137

8. **Instantiation** (Mandatory): For component or entity instantiation, generics and ports must be associated by name, following the order of declaration, one per line, with arrows (=>) aligned. The label and instantiated unit must appear on the first line, followed by generics (if present) on a separate line, and then ports on another separate line.
9. **Process label** (Mandatory): Each process must have a label that clearly explains its purpose, with a brief description provided in a comment just before the declaration.
10. **Parenthesis** (Mandatory): Parentheses must be used in expressions to make the evaluation order explicit, and conditions in **if** and **while** statements must be enclosed in parentheses.

138 **Identifiers**

139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158

1. **Keywords case** (Mandatory): Keywords case (VHDL reserved identifiers) must be written in lower case.
2. **Identifiers case** (Mandatory): Identifiers should be written in lower case. The only exceptions are generics and constants, which should be written in upper case. Identifiers should use underscores for multi-word naming (e.g., **this_is_multi_word**); no **camelCase** or **PascalCase** is allowed.
3. **References** (Mandatory): References must use the same case as the object they are referring to.
4. **Architecture names** (Mandatory): Only two architecture names are allowed: **rtl** should be used for synthesizable code, while **sim** must be used for simulation-only code.
5. **Constants names** (Mandatory): Constants declarations should be in UPPER case with a **c_** prefix.
6. **Generics names** (Mandatory): Generic identifiers should also be in UPPER case, with the **g_** prefix.
7. **Ports names** (Mandatory): Port names must be in lower case and should be followed by a suffix specifying the direction of the port:
 - **_i** for inputs,
 - **_o** for outputs, and
 - **_io** for bidirectional ports.

- 159 8. **Variables names** (Mandatory): Variables must have the `v_` prefix.
- 160 9. **Types names** (Mandatory): User-declared types and subtypes must have the `_t`
161 suffix.
- 162 10. **Signals name** (Mandatory): Clock signals should have the `clk_` prefix and the
163 direction suffix. In case of multiple clock domains, the domain name and the
164 frequency (if known) should be specified. Reset signals should have the `reset_`
165 prefix and the direction suffix. In case of multiple clocks, there should be one reset
166 per clock, and the reset name should report the corresponding domain name. List
167 of suffixes:
- 168 • Active-low signals should have a `_n` suffix,
 - 169 • Asynchronous signals should have a `_a` suffix,
 - 170 • Delayed signals should have a `_d` suffix. If delays are concatenated, use an
171 index (e.g., `_d1`, `_d2`, ...) or a vector.
- 172 The order of suffixes should follow alphabetical order, except for delays or port
173 directions, which should be the last ones.
- 174 11. **Process labels** (Mandatory): Process labels should have the `proc_` prefix.
- 175 12. **Generate label** (Mandatory): Generate labels should have the `gen_` prefix.

176 Language Rules

- 177 1. **VHDL Version** (Mandatory): Synthesizable units must follow the VHDL-2008
178 standard. The only exception is rule Format.5.
- 179 2. **IEEE Packages** (Mandatory): The only IEEE packages allowed are:
- 180 • `std_logic_1164`,
 - 181 • `numeric_std`,
 - 182 • `numeric_bit`,
 - 183 • `math_real`,
 - 184 • `math_complex`,
 - 185 • `std_logic_misc` (only the reduce functions), and
 - 186 • `std_logic_textio`.
- 187 In particular, `std_logic_unsigned`, `std_logic_signed`, and `std_logic_arith` are
188 **not** allowed. Use `numeric_std` instead.
- 189 3. **Attribute declarations** (Mandatory): It is not allowed to declare custom attributes,
190 except for tool-specific attributes.
- 191 4. **Enumerations** (Mandatory): All user-defined enumerations must use names for
192 literals. Characters are not allowed.

- 193 5. **Buffer, linkage, guarded signals** (Mandatory): bus, register, buffer, and
194 linkage are not allowed.
- 195 6. **Block statements** (Mandatory): Block statements can be used to group concurrent
196 statements. No ports, generics, and implicit guards are allowed.
- 197 7. **Configurations** (Mandatory): Configurations are not allowed.

198 Synthesis Rules

- 199 1. **Type of top-level ports** (Mandatory): The type of the ports in a top-level
200 entity must be `std_logic`, `std_logic_vector`, `unsigned`, or a user-defined bounded
201 composite type (array or record) composed of these types.
- 202 2. **Type of top-level generics** (Mandatory): The type of generics in a top-level entity
203 must be `string`, `integer`, `boolean`, `std_logic`, `std_logic_vector`, an enumerated
204 type, or a user-defined bounded composite type of these types.
- 205 3. **Wrapper of top-level units** (Recommended): Bus signals should be grouped into
206 records to reduce the number of connections. There can be two versions of the
207 top-level unit: a wrapped one and a non-wrapped one. The name of the wrapped
208 unit is the name of the normal one, but with the `_x` prefix. Each bus should have
209 two records: one for the input signals and one for the output signals. The records
210 should be declared in a package. The default top-level entity should be the wrapped
211 version, while the unwrapped version should only unwrap the signals.
- 212 4. **Process for a register** (Recommended): Use only registers triggered on the rising
213 edge of the clock and with a synchronous reset.

```
214     proc_reg : process (clk) is  
215     begin  
216         if rising_edge(clk) then  
217             if (reset = '1') then  
218                 q <= '0';  
219             else  
220                 q <= d;  
221             end if;  
222         end if;  
223     end process proc_reg;  
224  
225
```

- 226 5. **Asynchronous reset** (Mandatory): If reset is asynchronous, it must be syn-
227 chronously deasserted.
- 228 6. **Register reset** (Mandatory): All registers **must** be initialized during reset.
- 229 7. **Signal attributes** (Mandatory): Do not use time-related signal attributes for
230 synthesis (`event`, `active`, `delayed`). Use the function `rising_edge(clk)` or
231 `falling_edge(clk)`, if necessary.

- 232 8. **Index direction** (Mandatory): Subtypes of `std_logic_vector`, `signed`, and
233 `unsigned` must always use the `downto` direction. `integer` and array indexes can
234 use the `to` direction.
- 235 9. **Array size** (Mandatory): Array length should be 2 or more if the size is not
236 computed.
- 237 10. **Clock and reset ports** (Mandatory): The main clock should be the first port of
238 the list, and the main reset must be the second port.
- 239 11. **Usage of clocks** (Mandatory): Logic on clocks is not allowed.
- 240 12. **FSM code style** (Recommended): If all outputs are a function of the current
241 state, the FSM can be written using only one process (Moore). In other cases, the
242 FSM can be written with two processes (Mealy): one handling the register, and the
243 second computing states and outputs.
- 244 13. **Unused Ports map** (Mandatory): Unused ports should be assigned to `open` signals.

245 2.2 Architecture

246 This section is based on the work of Jiri Gaisler [2].

247 The *dataflow* style is a traditional approach for synthesizable VHDL models, relying
248 on numerous concurrent statements and small processes connected through signals. While
249 effective, this style can be difficult to read and understand, as execution depends on signal
250 changes rather than the order of written statements. Engineers often need to create block
251 diagrams to understand dependencies, making maintenance challenging, especially in
252 large designs like memory controllers. Additionally, dataflow VHDL has a low abstraction
253 level, relying on basic constructs like multiplexers and bitwise operators, which makes
254 recognizing complex algorithms difficult. Simulation performance is also inefficient, as
255 signal assignments are slower than variable assignments due to event scheduling overhead.

256 The *two-process* coding method is proposed as an alternative for synchronous single-
257 clock VHDL designs. This method improves readability, abstraction, debugging, and
258 simulation efficiency by:

- 259 • Using record types for ports and signals.
- 260 • Limiting each entity to two processes: one for combinational logic and one for
261 sequential logic.
- 262 • Employing high-level sequential statements for algorithm encoding.

263 Unlike traditional VHDL, where concurrent statements complicate readability as
264 designs grow, the two-process method structures logic sequentially, similar to standard pro-
265 gramming languages like C. This approach simplifies analysis, debugging, and maintenance
266 while maintaining a single model for both synthesis and simulation.

267 Listing 1 demonstrates the Run Length Encoder from *colibri* written using this
268 architectural style.

Listing 1: Run Length Encoder

```

269 architecture rtl of rle_decode is
270
271
272 constant c_WORD_MAX : positive := 2 ** g_COUNT_WIDTH - 1;
273
274 type sig_t is record
275     word_cnt : natural range 0 to c_WORD_MAX;
276     buf      : snk_data_i'subtype;
277     odata    : src_data_o'subtype;
278     valid    : std_logic;
279     ready    : std_logic;
280 end record sig_t;
281
282 signal rreg, rcmb : sig_t;
283
284 begin
285
286     rreg <= rcmb when rising_edge(clk_i);
287
288     proc_encode : process (all) is
289
290         variable v_int : sig_t;
291
292     begin
293
294         v_int := rreg;
295
296         -- acknowledge a read
297         if (src_ready_i = '1' and rreg.valid = '1') then
298             v_int.valid := '0';
299         end if;
300
301         if (v_int.valid = '0' and rreg.word_cnt /= 0) then
302             v_int.valid := '1';
303             v_int.odata := rreg.buf(g_WORD_WIDTH - 1 downto 0);
304             v_int.word_cnt := rreg.word_cnt - 1;
305         end if;
306
307         v_int.ready := not v_int.valid;
308
309         -- if input is valid and no word is buffered
310         if (snk_valid_i = '1' and v_int.valid = '0') then
311             v_int.buf := snk_data_i;
312             v_int.word_cnt := to_integer(unsigned(snk_data_i(snk_data_i
313                 'left downto g_WORD_WIDTH)));
314             -- forward the current word
315             v_int.valid := '1';
316             v_int.odata := snk_data_i(g_WORD_WIDTH - 1 downto 0);
317             v_int.word_cnt := v_int.word_cnt - 1;
318         end if;

```

```

319
320     if (reset_i) then
321         v_int.word_cnt := 0;
322         v_int.buf      := (others => '0');
323         v_int.valid    := '0';
324         v_int.ready    := '0';
325     end if;
326
327     rcmb <= v_int;
328
329 end process proc_encode;
330
331 src_data_o <= rreg.odata;
332 src_valid_o <= rreg.valid;
333 snk_ready_o <= rcmb.ready;
334
335 end architecture rtl;
336

```

337 While not mandatory, adopting this architectural style is highly recommended, partic-
338 ularly for large designs where debugging can become complex.

339 2.3 Documentation

340 Clear and comprehensive documentation is vital for improving code readability, main-
341 tainability, and collaboration among developers. Most documentation can be included
342 directly in comments, ensuring that explanations remain close to the relevant code sections.
343 Structured comments help future developers understand the design and functionality,
344 simplifying debugging and modifications.

345 In VHDL, the `--!` comment keyword can be used to generate automatic documentation
346 with tools like Doxygen [3], providing a systematic way to describe modules, ports, and
347 functionality. Additionally, tools like WaveDrom [4] enable easy waveform visualization,
348 clarifying timing and signal behavior. Integrated development environments such as
349 TerosHDL [5] can further automate documentation generation.

350 3 Simulations

351 Simulation is a widely used and traditional method for verifying gateway functionality.
352 This section focuses on Register Transfer Level (RTL) simulations, which provide cycle-
353 accurate behavioral information while remaining independent of implementation details.
354 A typical testbench, illustrated in Figure 1, consists of three main components: the Device
355 Under Test (DUT), which is the component being verified; a stimuli generator, which
356 provides inputs to the DUT; and a checker, which compares the DUT's outputs against
357 expected values. The designer is responsible for developing test cases and providing a
358 quantitative measure of test coverage.

359 To ensure effective verification, the following generic guidelines are recommended when
360 designing a verification plan:

- 361 • **Test Minimal Components:** Focus on verifying the smallest functional units
362 possible. Testing isolated sub-components simplifies test case development and

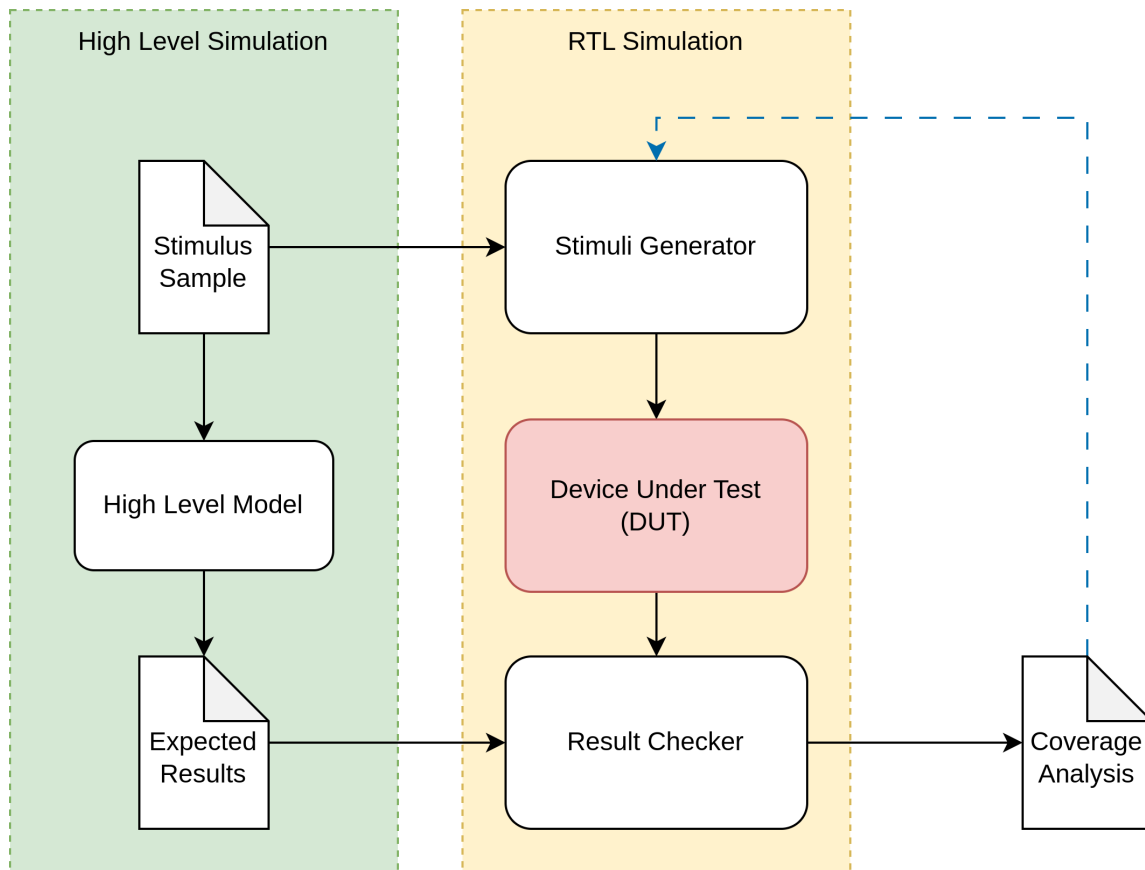


Figure 1: An example diagram of a functional verification testbench. The primary components of the simulation are highlighted in yellow. Stimuli generators and result checkers are often implemented in HDL code, while high-level simulation is typically implemented using languages like SystemC or C.

363 reduces complexity. Rigorous verification of these smaller units facilitates the
 364 subsequent verification of the complete design. Additionally, organizing tests based
 365 on specifications improves bug tracking and collaboration by making the tests easier
 366 to understand.

367 • **Cover Normal Operational Modes:** Ensure that the design performs as expected
 368 in all operational modes under normal conditions, as specified in the functional
 369 requirements. For larger or parameterizable designs, test all combinations of normal
 370 modes to verify the correct interaction between modules. The system should always
 371 return to an idle state after the tests.

372 • **Test Exception Conditions:** Validate the system's behavior under abnormal or
 373 exceptional conditions. This includes testing the system's recovery capabilities when
 374 operating outside normal modes. Tests should cover illegal conditions and protocol
 375 violations, ensuring the system can gracefully handle unexpected scenarios.

376 **Functional Coverage** Coverage is a quantitative metric used to measure the complete-
 377 ness of the verification effort. It is typically expressed as the percentage of verification
 378 objectives met. Two primary types of coverage metrics are used:

- 379 • **Code Coverage:** This metric, derived directly from the code, indicates whether
380 and how frequently specific lines of code are exercised during testing. While useful,
381 it does not provide insights into the context or purpose of the exercised code.
- 382 • **Functional Coverage:** Defined by the designer, this metric ties verification to the
383 design’s intended functionality. It measures whether specific scenarios, corner cases,
384 design requirements, and other critical conditions of the DUT have been observed,
385 validated, and tested.

386 To design an effective functional coverage testbench, methodologies such as *Constrained*
387 *Random Testing* can be employed. Instead of manually crafting each test case, pseudo-
388 random stimuli generators can automatically create a set of test cases based on general
389 specifications. The testbench is equipped with checkers to verify that coverage goals
390 are met. A high-level diagram of this methodology is shown in Figure 2. Verification
391 frameworks, discussed in Section 3.1, can simplify this process by providing integrated
392 utilities and dedicated components. An example of this methodology can be found in the
393 FIFO testbench available in the *colibri* library [6].

394 Constrained random stimuli can also be produced outside of the VHDL testbench.
395 Stimuli data can be provided from Monte-Carlo simulations and external synthetic data
396 generators. For example, Python frameworks such as *Hypothesis* [7] can be used to
397 generate constrained random data based on specification constraints. If this approach is
398 used, the developers should also provide an external checker and detail the data formats
399 for both the stimuli input and the DUT output.

400 **Transaction Level Modeling** FPGA designs often involve interactions between compo-
401 nents over interfaces, with signals divided into data and control paths. These interactions,
402 defined by a *protocol*, consist of specific *transactions* that dictate component responses
403 based on control signal combinations.

404 Abstracting these interactions into interfaces enables higher-level modeling, known
405 as Transaction Level Modeling (TLM). TLM separates communication details from the
406 implementation specifics of functional units, simplifying the representation of complex
407 systems.

408 Most FPGA designs use two primary categories of data communication mechanisms:

- 409 • **Streaming Interfaces:** These are point-to-point links where the transmitter
410 is referred to as the source (or master) and the receiver as the sink (or slave).
411 Transactions on a streaming interface are defined by a basic *handshake* mechanism:
412 the source drives the data along with a *valid* signal, indicating that the data is
413 ready to be read. The sink responds with a *ready* signal to indicate its readiness to
414 receive the data. Additional control signals can transform the communication from
415 a cycle-based approach to a packet-based one, where a single packet is transmitted
416 over multiple clock cycles.

417 Standard streaming interfaces include AXI-Stream (used by AMD Xilinx) and
418 Avalon Stream (used by Altera). An example of a streaming transaction is shown
419 in Figure 3a.

- 420 • **Memory-Mapped Interfaces:** These interfaces are used for bus communication
421 and consist of a bidirectional data path, an address path, and control signals. A host

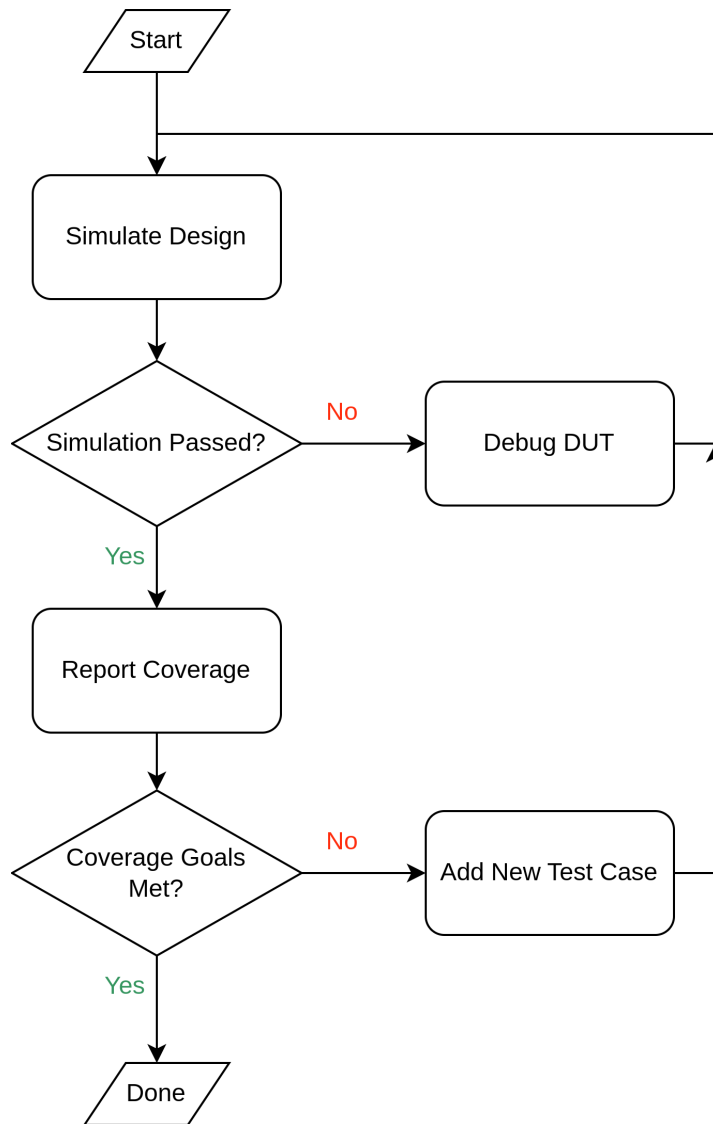


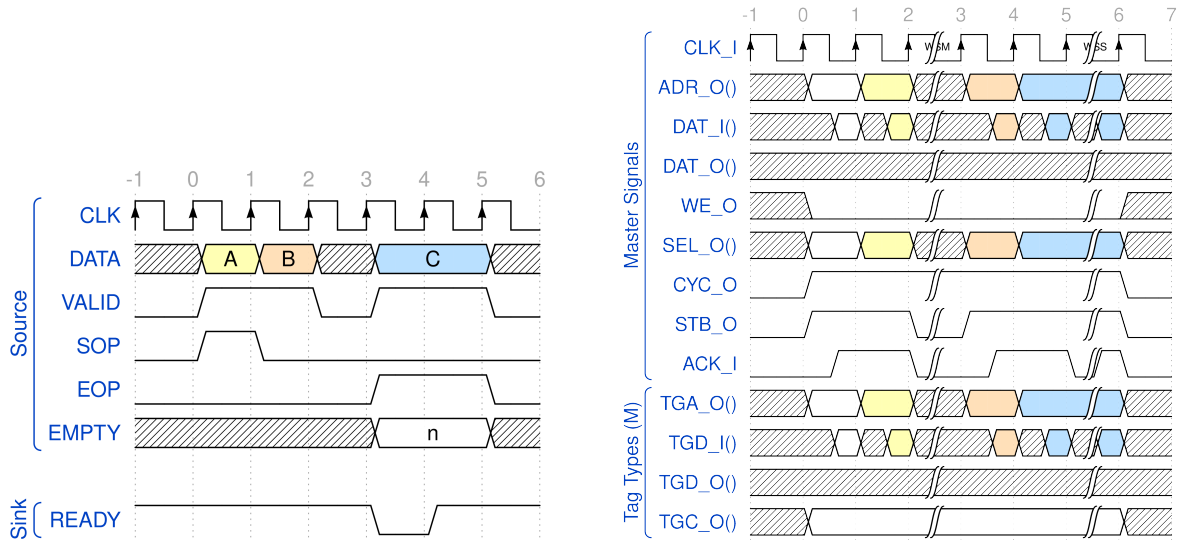
Figure 2: Flow chart of the constrained random test methodology.

422 orchestrates transactions across multiple peripherals. Unlike streaming interfaces,
 423 memory-mapped interfaces include an address path to specify the physical location
 424 of a register, memory, or resource. The use of a bus allows multiple peripherals to
 425 share the same signals, saving resources.

426 Basic memory-mapped interfaces support read and write operations, where each
 427 operation specifies the address and control signals. Extended features, such as
 428 pipelined burst transfers or variable latency support, can add complexity.

429 Common memory-mapped protocols include AXI4, AXI4-Lite, Avalon Memory-
 430 Mapped, and Wishbone Bus. An example of a memory-mapped transaction is shown
 431 in Figure 3b.

432 Verification frameworks often provide functional models for these standards, allowing
 433 developers to import models into testbenches and connect them to the appropriate
 434 interfaces. These frameworks integrate seamlessly with constrained random testing,



(a) Example of an Avalon Streaming packet transaction. SOP and EOP signal the start and the end of the packet, respectively. EMPTY indicates how many symbols are valid in the last data word. (b) Example of a Wishbone Block Read transaction. CYC_O high signals the transaction is happening. The operation spans over five clock cycles.

Figure 3: Example waveforms of interfaces.

435 further simplifying and enhancing the verification process.

436 3.1 Verification Frameworks

437 VHDL-based designs rely on several open-source verification frameworks. While these
 438 frameworks offer similar functionalities, they differ in implementation focus and method-
 439 ology. Designers can choose the most suitable framework for their application or combine
 440 multiple frameworks within the same testbench.

441 **UVVM** [8] is a free and open-source framework for developing structured, reusable
 442 testbenches. Supported by industry and research institutions like the European Space
 443 Agency, UVVM is fully implemented in VHDL-2008, making it compatible with a wide
 444 range of simulators. Its standout feature is the VHDL Verification Components (VVCs),
 445 which extend the TLM concept by autonomously handling transaction scheduling and
 446 component implementation. VVCs support complex protocols, including packet-based
 447 streams and out-of-order communication. However, their complexity results in a steeper
 448 learning curve, particularly for customization.

449 **OSVVM** [9] is another free and open-source framework. It offers features such as
 450 verification components, functional coverage, and constrained random tools. Its strengths
 451 include intuitive pseudo-random generators and powerful scoreboards. However, its
 452 Tcl-based scripting interface may feel outdated compared to modern alternatives.

453 **VUnit** [10] provides a modern approach to HDL verification, with a Python-based
 454 scripting interface. Its test runner automates test discovery, file scanning, and compilation

455 order management, simplifying test design and organization. Python scripting enhances
 456 portability and testbench development.

457 A comparison of these frameworks is summarized in Table 1.

	UVVM	OSVVM	VUnit
Checkers and Loggers	yes	yes	yes
Verification Components	excellent	good	basic
Randomized Coverage	basic	excellent	not available
Scoreboards	good	excellent	not available
Scripting	not available	poor (Tcl)	excellent (Python)

Table 1: Comparison of key features among the three major open-source VHDL verification frameworks.

458 3.2 Testbench Design

459 Testbenches should be all designed using the VUnit python framework which allows
 460 portability between simulators and simplifies dependencies management. For simple
 461 components and VHDL functions, a single testbench file is sufficient. These testbenches
 462 generate randomized stimuli and verify results against expected values, leveraging utility
 463 and random functions provided by OSVVM and UVVM.

464 For more complex components, a modular testbench structure is recommended, follow-
 465 ing UVVM’s guidelines. This structure consists of three core elements:

- 466 • **Test Package:** Contains common constants, functions, and procedures used across
 467 the testbench. It also defines types and subtypes for custom interfaces, such as clock
 468 periods and data widths.
- 469 • **Test Harness:** A non-synthesizable entity responsible for instantiating and in-
 470 terconnecting all components used in the tests, including the DUT and auxiliary
 471 entities. The harness entity is designed without ports, ensuring all signal sources
 472 and sinks are contained within the entity.
- 473 • **Test Bench:** Hosts the test procedure itself. It instantiates the test harness and
 474 controls its behavior through a single sequencer process. Commands are dispatched
 475 via the framework’s internal communication network to the various TLM entities in
 476 the test harness.

477 This modular architecture, shown in Figure 4, ensures a clear separation of implemen-
 478 tation details from the test procedure. High-level TLM commands are managed in the
 479 sequencer process, while low-level signal handling resides in a separate file. This approach
 480 enhances test clarity and facilitates reusability.

481 4 Formal Verification

482 Formal Verification (FV) is a mathematical approach used to analyze the entire space of
 483 possible behaviors of a design. It can be applied to both software and hardware problems
 484 and is useful at various stages of the development process.

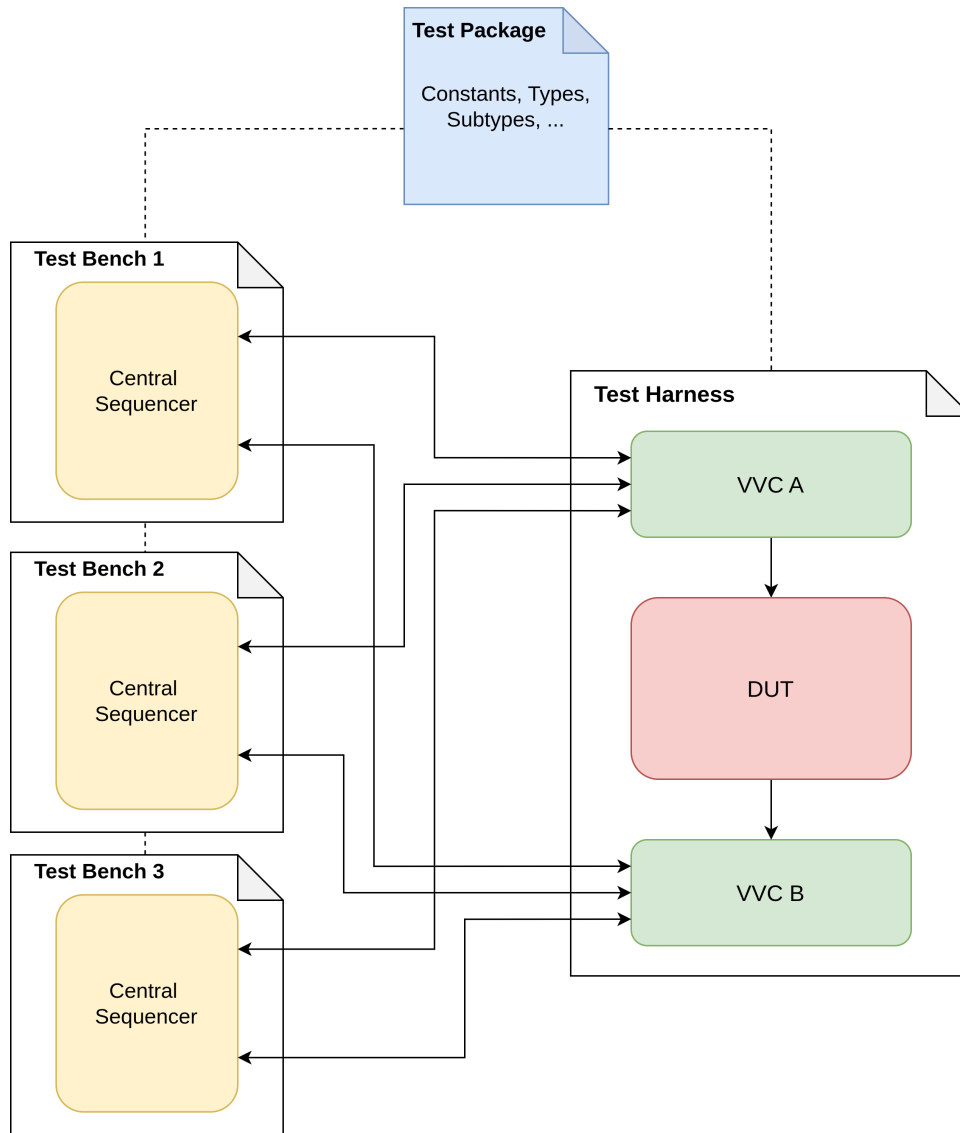


Figure 4: Modular testbench structure. This implementation allows the reuse of the test package and the harness across multiple testbenches. For instance, VVC A stimulates the DUT by generating the low-level signals from the sequencer commands. VVC B receives the output signals of the DUT and transmits them to the sequencer for checking.

485 Like simulation, FV employs tools to ensure that a design satisfies the properties
 486 and requirements of its specification. However, the methodology differs fundamentally:
 487 while simulation evaluates individual test cases, FV explores the entire design space at
 488 once. This does not mean that FV executes all possible simulations; instead, it uses
 489 mathematical techniques to efficiently compute and simplify the behavior of the design.

490 FPGA designs are particularly well-suited to FV because they are deterministic digital
 491 systems defined by Boolean logic. Although Boolean Satisfiability Problems (SAT) are
 492 NP-complete—meaning the time required to solve them grows exponentially with input
 493 size—clever strategies can significantly reduce problem complexity, making FV a practical
 494 approach. Figure 5 illustrates the difference in coverage between simulation and FV.




Figure 5: Coverage level of formal verification compared to simulation: simulation covers only specific points, while formal verification aims to cover the entire design space. In practice, full coverage is achievable only in certain areas.

495 **Bounded Model Checking** Bounded Model Checking (BMC) is the most widely
 496 used FV technique. In BMC, the design is represented as a finite-state system, and its
 497 properties are expressed in temporal logic. The tool translates the design into a logical
 498 expression that models all possible state transitions from the initial state. A finite bound
 499 of steps is specified by the user, limiting the scope of verification to a manageable range.

500 A counterexample is found when a combination of variables falsifies the logical expres-
 501 sion. If no counterexample exists within the specified bound, the tool guarantees that the
 502 design satisfies the property for all states within that range. It is the user’s responsibility
 503 to determine whether the chosen bound is sufficient for verification closure.

504 **Property Specification Languages** To perform FV, the properties of the design
 505 must be specified using a Property Specification Language (PSL). PSL enables precise
 506 definitions of expected behavior and verification methodology. Properties are categorized
 507 as:

- 508 • **Assertions:** Define the allowed behavior of the DUT.
- 509 • **Assumptions:** Specify constraints on the inputs to the DUT.

510 PSL assertions are more expressive than those used in simulation, as they can incor-
 511 porate temporal logic to describe how conditions evolve over time. IEEE PSL [11] is a
 512 generic property specification language that supports multiple HDL languages, including
 513 VHDL. For example, a complete FV testbench for the FIFO in the *colibri* library can be
 514 implemented in fewer than 20 lines of PSL, whereas an equivalent constrained random
 515 simulation requires at least ten times as much code without guaranteeing comprehensive
 516 coverage.

517 **Optimization Techniques** In computational theory, FV is an *NP-hard* problem,
 518 meaning that the time and memory required to compute FV tests grow exponentially
 519 with the size of the design. However, several techniques can reduce problem complexity
 520 and improve tool efficiency:

- 521 • **Reducing Problem Size:** FPGA designs often involve wide data paths and large
522 memories. If no constraints are provided, the tool will attempt to model all possible
523 combinations of these bits, leading to an explosion of states. Reducing data widths
524 and memory sizes to a few bits typically preserves the design’s functionality while
525 allowing the tool to catch most bugs. If reducing data widths is not feasible, limiting
526 the search space to specific data patterns (e.g., most bits being 0) can also help.
- 527 • **Decomposing Tests:** Complex properties can be broken down into simpler,
528 independent ones. Similarly, independent behaviors of the DUT can be tested
529 separately. This decomposition reduces the complexity of individual tasks, enabling
530 tools to manage them more efficiently and significantly reducing computation time
531 and memory usage.

532 **Formal Verification Tools** An open-source toolkit is available for performing FV,
533 consisting of the following components:

- 534 • **Yosys:** An HDL synthesizer.
- 535 • **GHDL:** A VHDL interpreter.
- 536 • **SymbiYosys (SBY):** A formal verification frontend.

537 These tools are often packaged together in pre-configured Docker images, such as
538 [hdlc/formal](#), simplifying setup and usage.

539 An FV testbench typically consists of three essential files:

- 540 • **DUT:** The design under test.
- 541 • **PSL Specification:** A file (usually with a `.psl` extension) containing the verifica-
542 tion unit, where assertions and assumptions are defined. Additional glue HDL code
543 can be included to encode more complex behaviors.
- 544 • **SBY Configuration:** A file specifying the configuration of the FV toolkit and
545 listing the files required to synthesize the design.

546 Many examples of FV testbenches can be found in the *colibri* library, demonstrating
547 how to apply these tools effectively.

548 5 Common Components Library

549 Gateware components, often referred to as Intellectual Property (IP), are essential for
550 developing FPGA designs. Many components used in designs are generic and implement
551 reusable functionalities. These generic components are often provided directly by vendors
552 under licensing agreements. However, relying on vendor-provided components can lock
553 developers into a specific ecosystem, limiting flexibility and choice.

554 To avoid vendor lock-in, many designers choose to develop custom implementations
555 tailored to their specific needs. While this approach provides independence from proprietary
556 solutions, it often lacks thorough verification and proper documentation, leading to
557 challenges such as reduced portability and collaboration.

558 The *colibri* library [6] was developed to address these limitations, providing a stan-
559 dardized components library for CERN developers. The library ensures functionality and
560 reliability through rigorous verification and validation, employing the techniques discussed
561 in previous sections.

562 **Library Content** The *colibri* library consolidates commonly used components for DAQ
563 applications, drawing inspiration from the selection used in the current PCIe40 gateway.
564 The components are designed to be generic and parameterizable, accommodating a wide
565 range of user applications. As of this writing, *colibri* contains 57 components, organized
566 into the following categories:

- 567 • **Common:** Entities and packages implementing frequently used functions, such as
568 synchronizers, buffers, and debouncers.
- 569 • **Communication:** Entities designed for communication tasks, including scramblers
570 and gearboxes.
- 571 • **Encoders:** Components for standard data encoding and decoding, such as Run-
572 Length Encoding.
- 573 • **File I/O:** Functions and packages that simplify file operations during simulation.
- 574 • **Interfaces:** Components for translating and converting between different interfaces,
575 such as Avalon and AXI.
- 576 • **Input/Output:** Basic low-level protocols (I²C, SPI, UART) for interacting with
577 external sensors and microcontrollers.
- 578 • **Memory:** Implementations of FIFOs for single and dual clocks, along with RAM
579 entities.
- 580 • **Packet:** Common operations for packet-based streaming interfaces.
- 581 • **Pipes:** Components for managing streaming interfaces, such as arbiters, broadcast-
582 ers, and routers.
- 583 • **Protocols:** High-level protocols like Ethernet and Aurora.

584 The *colibri* library is hosted and maintained on GitLab, provided by the CERN IT
585 department. Developers are encouraged to provide feedback through issues and contribute
586 via merge requests.

587 6 Continuous Integration

588 Testbenches developed using the methods described above can be seamlessly integrated into
589 the Continuous Integration (CI) system in GitLab. This integration provides immediate
590 feedback after each commit, ensuring code quality and reliability. Open-source VHDL
591 simulators such as GHDL [12] and NVC [13] are recommended, as they do not require
592 licenses.

593 A suggested layout for the CI pipeline is the one used by *colibri*. For both simulation
594 and FV, GitLab pipelines capture detailed test artifacts, such as logs and waveforms, in
595 the event of test failures. These artifacts can be downloaded and analyzed for debugging
596 without the need to rerun the tests.

597 The pipelines run on a distributed computing cluster provided by CERN IT, which is
598 accessible to all FPGA developers at CERN. Each pipeline stage executes within a Docker
599 container preloaded with the necessary tools and frameworks. Pipelines are triggered by
600 new commits to the `devel` and `master` branches, as well as merge requests. To optimize
601 resource usage, tests are executed selectively, only when relevant source or simulation files
602 have been modified.

603 References

- 604 [1] J. Leary, *jeremiah-c-leary/vhdl-style-guide*, [https://github.com/jeremiah-c-](https://github.com/jeremiah-c-leary/vhdl-style-guide)
605 [leary/vhdl-style-guide](https://github.com/jeremiah-c-leary/vhdl-style-guide).
- 606 [2] Jiri Gaisler, *A structured VHDL design method*, in *Fault-tolerant Microprocessors*
607 *for Space Applications*.
- 608 [3] *Doxygen*, <https://www.doxygen.nl/>.
- 609 [4] A. Chapyzenka and J. Probell, *Rendering Beautiful Waveforms from Plain Text*,
610 2016.
- 611 [5] Teros Technology, *terosHDL*, [https://github.com/TerosTechnology/vscode-](https://github.com/TerosTechnology/vscode-terosHDL)
612 [terosHDL](https://github.com/TerosTechnology/vscode-terosHDL).
- 613 [6] *colibri*, <https://gitlab.cern.ch/colibri/colibri>.
- 614 [7] D. MacIver, Z. Hatfield-Dodds, and M. Contributors, *Hypothesis: A new approach to*
615 *property-based testing*, *Journal of Open Source Software* **4** (2019) 1891.
- 616 [8] UVVM, *UVVM/UVVM*, <https://github.com/UVVM/UVVM>.
- 617 [9] OSVVM, *Open Source VHDL Verification Methodology*, <https://osvvm.org/>.
- 618 [10] *VUnit/vunit*, <https://github.com/VUnit/vunit>.
- 619 [11] *IEEE Standard for Property Specification Language (PSL)*, *IEEE Std 1850-2005*
620 *(2005)* 1, Conference Name: IEEE Std 1850-2005.
- 621 [12] *ghdl/ghdl*, <https://github.com/ghdl/ghdl>.
- 622 [13] N. Gasson, *nickg/nvc*, <https://github.com/nickg/nvc>.