

Language-based Hardware Communication Safety and Liveness

Aditya Ranjan Jha*
National University of Singapore
Singapore

Jason Zhijingcheng Yu*
National University of Singapore
Singapore

Umang Mathur
National University of Singapore
Singapore

Trevor E. Carlson
National University of Singapore
Singapore

Prateek Saxena
National University of Singapore
Singapore

ABSTRACT

Emerging hardware description languages (HDLs) provide new type annotations to abstract from low-level details [3, 9, 11, 12, 15], enabling opportunities for modular and scalable structural verification at design stage. However, verification of behavioural properties among communicating hardware modules—such as ensuring liveness and communication safety—remain largely unsupported in modern HDLs. In this paper, we present **PACT**, a language extending a recent HDL called Anvil [14] with a novel session type-based system. Our proposed type system explicitly models synchronous clock cycles and introducing new type constructs to capture communication contracts between hardware modules. It offers abstractions to reason about liveness and communication safety between hardware components. Additionally, we give a sound and complete algorithm for verifying them statically. *Our evaluation shows that our type abstractions offer a factor of 2^{70} to 2^{270} in state space reduction over that of concrete RTL designs, when verifying 10 real hardware designs.*

1 INTRODUCTION

Hardware designs typically consist of numerous concurrent modules that communicate with one another. In such communication, a sender shares a value with a receiver in a particular clock cycle. Two properties are highly desirable: (1) *communication safety*—the time the sender intends to send the value and the time the receiver intends to receive the value are synchronized; (2) *liveness*—all hardware modules eventually make progress, and in particular, do not deadlock for message communication.

Traditionally, such properties are verified post design through testing or model checking, which suffers from unsoundness and state explosion. While recent research has started investigating language-based approaches that allow sound design-time hardware verification, the properties studied so far have been limited to local low-level properties surrounding wire and register use, such as timing safety, combinational loop freedom, structural hazard freedom, and so on. Verifying communication safety and liveness, however, requires reasoning about the overall behaviour of modules. In this work, we aim to fill in this gap and propose a language-based approach for verifying communication safety and liveness in hardware. The result of our effort is **PACT**, a hardware description language (HDL) with a type system for verifying such properties.

A natural option is to introduce multiparty session types, which have been used to verify similar properties for software systems, to HDLs. Session types capture relevant behaviour of communication participants. When a collection of session types are compatible, they

$$\begin{aligned} e &::= x \mid v \mid e \text{ bop } e & v &::= \text{true} \mid \text{false} & B &::= \emptyset \mid B, \hat{a} : P \\ a &::= \text{send } m \ e \mid x = \text{recv } m \text{ in } P & \hat{a} &::= \text{send } m \ e \mid x = \text{recv } m \\ P &::= \text{skip} \mid a \mid X \mid \text{cycle} \mid \text{end}_t \mid P \parallel P \mid P; P \mid \text{let } x = e \text{ in } P \\ & \mid \text{if } e \text{ then } P \text{ else } P \mid \text{fair-if}_\eta \ e \text{ then } P \text{ else } P \\ & \mid \text{offer } B \text{ else } P \mid \text{xoffer } B \text{ else } P \mid \mu X. P \end{aligned}$$

Figure 1: Abstract syntax of processes in PACT.

guarantee communication safety and deadlock freedom. However, adapting session types is not straightforward due to a fundamental mismatch between existing session type abstractions and the execution model of synchronous digital hardware. Specifically, such a mismatch requires us to overcome three main challenges: (1) Existing session types [4, 7, 8] typically abstract away cycle-level timing hardware communication; (2) They restrict internal and external choice structures, limiting the expressiveness for communication patterns common to hardware designs; (3) Prior approaches to session types define compatibility syntactically [1, 5] or compromise modularity [2, 10], which is essential for hardware design.

Our work, **PACT**, addresses these challenges first by extending session types with modelling for synchronous hardware behaviour such as synchronous latency and unrestricted multiparty choice. We then devise an automata-based type compatibility algorithm that is sound and complete, overcoming the limitations of existing type compatibility formulations. We implement **PACT** as an extension to Anvil [14], and evaluate its effectiveness. Our experiments across 10 real-world hardware designs show that compared to concrete register-transfer-level (RTL) designs, **PACT** achieves a state space reduction by a factor of 2^{70} to 2^{270} across different types of designs. We intend to open-source **PACT** and potentially merge it into the upstream Anvil once the development has matured.

2 PACT

Following Anvil [14], **PACT** models hardware modules as *processes* which communicate by passing messages across *channels*. Such *message passing* is the only way for processes to communicate.

2.1 Processes

2.1.1 Syntax. Figure 1 presents the abstract syntax of **PACT** processes. In addition to the basic primitives cycle and message passing primitives in Anvil, **PACT** includes explicit constructs for more expressive communication patterns, particularly for multiparty external choice, and for expressing safety specification.

*Equal Contribution

The exclusive (`xoffer`) and non-exclusive (`offer`) constructs both default to P_{else} if no messages in B can be exchanged. However, if multiple matches exist, `xoffer` non-deterministically executes one, while `offer` executes all in parallel. These unique constructs enable external choices dependent on multiple parties.

Some lightweight constructs help specify aspects of the desired safety properties. For example, consider the branching construct (`fair-if $_{\eta}$ e then P_1 else P_2`) that is *assumed* to be fair, i.e., either neither or both of P_1 and P_2 are executed infinite times. Here, the identifier η allows us to track which branching construct each branching decision corresponds to in the semantics. Likewise, `end $_t$` is used to specify a state of interest, which when reached demarcates progress in the process identified by t . In PACT, the RTL design consists of a set of looping processes P_1, P_2, \dots, P_n provided by the designer. Our safety specification requires that all processes perform infinitely many loop iterations. Therefore, PACT designs are expected to be structured as

$$\mu X.(P_1; \text{end}_1; X) \parallel \mu X.(P_2; \text{end}_2; X) \parallel \dots \parallel \mu X.(P_n; \text{end}_n; X).$$

2.1.2 Semantics. The semantics of PACT processes is defined using a labelled transition system (LTS) (S, Σ, T) , where each transition label $\sigma \in \Sigma$ represents an *action*:

$$\sigma ::= m\checkmark \mid \# \mid \$_t \mid \text{true}_{\eta} \mid \text{false}_{\eta}.$$

An $m\checkmark$ action represents a message exchange over channel m and $\#$ represents elapse of a cycle. A $$_t$ marks reaching of a state that indicates progress, referred to as a *progress state*, e.g., completion of a process iteration, and true_{η} and false_{η} encode choice decisions.

Each state in S is a process (note that a collect of parallel processes can also be expressed as one process). The run of such an LTS from a state s_0 is a finite/infinite sequence $\rho = (s_0 \xrightarrow{\sigma_0} s_1) \cdot (s_1 \xrightarrow{\sigma_1} s_2) \dots$ such that for every i , $(s_i \xrightarrow{\sigma_i} s_{i+1}) \in T$. A run is *complete* if it is either infinite or if it terminates in a state s with no outgoing transition.

2.2 Safety

We define safety based on the LTS. Such a definition needs to capture both communication safety and liveness (under specified fairness assumptions). Informally, the definition requires that all complete run of the LTS from a given start state to satisfy either (1) for all (relevant) process identifiers t , $$_t$ transitions are made infinite times in the run, or (2) for some η , true_{η} or false_{η} transitions are made infinite times, and transitions of the other are made finite times. Note that (1) states progress and (2) states the specified fairness assumptions. Communication safety is implicit: The LTS only allows transitions corresponding to safe communication, and the conditions above imply that the run must be infinite.

2.3 Types

Compared to prior session type formulations, PACT types capture more expressive communication patterns, and use a compatibility checking algorithm that is sound and complete (Section 2.4). The abstract syntax of types is as follows:

$$\begin{aligned} \tau ::= & \alpha \mid \$_t \mid \# \mid \text{skip} \mid \tau + \tau \mid \tau; \tau \mid \tau \parallel \tau \mid X \mid \mu X.\tau \mid \tau \dot{+}_{\eta} \tau \\ & \mid \oplus(\tau, \alpha : \tau, \dots) \mid \wedge(\tau, \alpha : \tau, \dots) \\ \alpha ::= & m! \mid m? \end{aligned}$$

Similar to processes, types include internal choice with fairness assumptions ($\tau \dot{+}_{\eta} \tau$), exclusive and non-exclusive multiparty external choice ($\oplus(\tau, \alpha : \tau, \dots)$ and $\wedge(\tau, \alpha : \tau, \dots)$), as well as the marker for a progress state ($$_t$).

The satisfiability relation between a type and a process is mostly standard, defined inductively over the syntactic structure. We define the operational semantics of types in the same framework (Section 2.1.2), except that the states are here types instead of processes.

2.4 Type Compatibility

We devise an automata-based sound and complete algorithm for deciding compatibility between types. The high-level idea of the algorithm is straightforward. We start by showing that starting with any type, the set of reachable types in the type semantics is finite (under the restriction that all recursions are at tail positions). This allows us to construct a *finite* LTS which captures the semantics of $\tau_{\text{system}} = \tau_1 \parallel \dots \parallel \tau_n$. We then observe that the safety properties (Section 2.2) can be encoded as Rabin conditions for ω -automata. This can be achieved through checking for emptiness of the product Rabin automaton. Since types abstract away many irrelevant details in concrete designs, we expect the LTS thus constructed to be smaller in state space. We show this empirically in Section 5.

3 LANGUAGE OVERVIEW

PACT is a language extension to Anvil [14]. We first illustrate the programming model with a simple process description in Anvil:

```
chan arb_worker_ch {
  left req : (logic@#1),
  left done : (logic@#1)
}
proc worker(ep[2] : right arb_w_ch) {
  reg cyc : logic[32];
  loop {
    send ep[0].req(1'd1) >> cycle 1 >>
    if ((*cyc)%3==0) {
      send ep[1].req(1'd1) >> cycle 1 >>
      send ep[0].done(1'd1); send ep[1].done(1'd1)
    } else {
      send ep[0].done(1'd1)
    } >> cycle 1
  } // loop that tracks cycle count (cyc) (skipped for brevity)
}
```

The process `worker` takes two endpoints of channel type `arb_w_ch`. In the channel definition, message labels associated with the `right` endpoint correspond to messages this process receives, while the remaining labels correspond to messages it sends [14].

Operationally, the process first sends a `req` message to the first endpoint and waits one cycle. It then checks whether `cyc` is divisible by three. If so, it sends a `req` message to the second endpoint, waits one cycle, and then sends `done` messages to both endpoints in parallel. Otherwise, it sends a `done` message only to the first endpoint. The loop then repeats after one cycle. Since PACT extends Anvil, every Anvil process is valid in PACT. The behavior of the above process can be checked against the PACT *interface type*:

```
interface worker_itfc(ep[2] : right w_arb_ch) =
  ep[0].req; #1;
  (ep[0].done + (ep[1].req; #1; (#ep[0].done || #ep[1].done))); #1
```

This interface abstracts the communication behavior of the process. It specifies that the process first sends `ep[0].req`, waits one

cycle, and then either sends $ep[0].done$ or performs a transaction on the second endpoint consisting of $ep[1].req$, a one-cycle delay, and parallel completion messages on both endpoints. The behavior then repeats after one cycle. All PACT interfaces are implicitly tail-recursive since processes in Anvil/PACT are tail-recursive.

In this syntax, a message identifier such as $ep[0].req$ denotes a blocking message exchange and is syntactic sugar for the abstract type¹ $\mu X. \oplus(\#, X, ep[0].req! : skip)$. In contrast, $\#ep.req$ denotes a non-blocking exchange and corresponds directly to the underlying abstract type, while the remaining operators correspond to constructs of the abstract type language. In PACT, a process may take an instantiation of an interface type as an argument instead of channel endpoints, and the compiler checks that the implementation conforms to the specified interface. It is straightforward to see that the above process conforms to this interface.

Next consider the other endpoint of the channel, implemented by an arbiter that arbitrates between the two workers. Since the arbiter accepts requests from either worker and completes the corresponding transaction, its interface can be written as follows:

```
interface arbiter_itfc(ep[2] : left w_arb_ch) =
  ^ (ep[0].req : #1; ep[0].done, ep[1].req : #1; ep[1].done); #1
```

This interface uses the exclusive external choice operator \wedge (corresponding to \oplus in the abstract syntax) to indicate that the arbiter will engage in a transaction with exactly one of the endpoints. Intuitively, the arbiter waits until a request arrives from either worker and then completes the corresponding request–done transaction. The program type is syntactic sugar for $\mu X. \oplus(\#, X, ep[0].req? : \#, ep[0].done?, ep[1].req? : \#, ep[1].done?)$. A PACT implementation that satisfies this type :

```
reg prio : logic; // process signature skipped for brevity
loop {
  xoffer with *prio {
    x = recv ep[0].req => (
      let dn = recv ep[0].done >> set prio := 1'd1
    ),
    ep[1].req => (
      let dn = recv ep[1].done >> set prio := 1'd0
    )
  }
}
```

The `xoffer` construct waits for requests from both endpoints and resolves them via exclusive offer. PACT allows designers to enforce fair arbitration by explicitly specifying priority with `xoffer`. The annotation with `*prio` in the example specifies that the priority between branches is determined by the content of the register `prio`. If such an annotation is omitted, the first branch takes precedence.

Finally, consider a closed system consisting of two workers and two arbiters, where each worker connects to both arbiters and each arbiter connects to both workers. To check compatibility, PACT translates interface types into automata derived from their operational semantics. The synchronous cross-product of these automata exposes a deadlock state in which both workers simultaneously attempt to issue requests to the opposite arbiter while each arbiter waits to complete a transaction with its current endpoint.

¹Note that we omit the direction (! or ?) of the message passing operation in the concrete syntax as it can be inferred from the endpoint and channel definitions.

Table 1: Comparison of state space: RTL vs type ($\log_2 |N|$)

Component	Participants	RTL State Space	Type State Space	Compatibility
FIFO	3	81	9.7	✓
Stream FIFO	3	80	8.8	✓
Spill Register	3	70	8.8	✓
RR Arbiter	6	176	18.5	✓
TileLink Crossbar	10	184	34.5	✓
ACE Cache Controller	11	241	39.9	✓
AXI Lite Demux (Read)	8	155	27.4	✓
AXI Lite Demux (Write)	10	144	38.5	✓
AXI Lite Mux (Read)	9	226	33.2	✓
AXI Lite Mux (Write)	11	311	41.7	✓

4 SOUNDNESS

At a high level, we prove soundness—if τ_i satisfies P_i and $\tau_{\text{system}} = \tau_1 \parallel \dots \parallel \tau_n$ is compatible, then $P_{\text{system}} = P_1 \parallel \dots \parallel P_n$ is safe—by showing that (1) satisfiability between any τ and P implies that the semantics of P refines that of τ , thus safety of τ transfers to P , and (2) compatibility of τ implies its safety.

5 EVALUATION

We evaluate expressivity using open-source IPs, comparing the state space of our behavioural type models against SystemVerilog implementations (via yosys-abc latch counts). Table 1 demonstrates orders-of-magnitude reductions in state bits ($\log_2 |N|$).

6 DISCUSSION AND FUTURE WORK

We briefly discuss some key aspects of PACT and future directions. **Practicality.** For practical convenience, our PACT implementation supports syntactic sugar such as quantification (\forall) as a basic form of metaprogramming in the interface types. It also provides type inference, allowing designers to start from a high-level specification and refine it according to the desired behavior.

Real-World Applications. The two main aspects of our work—type satisfiability and type compatibility, are useful in different settings. Type satisfiability is particularly useful when specifying the abstract behavior of a component, such as in type-state programming for drivers [13] which enables formal guarantees at design time for hardware-software co-design. Type compatibility, on the other hand, is useful when complex designs are divided among multiple designers and the overall system must be checked for liveness before implementation. Once the abstract behaviors agree, each designer only needs to ensure that their implementation satisfies the corresponding interface specification. Such a layered approach also opens opportunities for generalizing the method to hybrid session types [6].

Fairness. Although fairness is undecidable in general, it is important for improving the expressivity and scalability of the type system, particularly in cases involving unrestricted or unbalanced choices. Therefore, tractable verification of fairness as a language feature for common cases, such as through verifying monotonicity of priority changes, is an interesting direction for future work.

7 ACKNOWLEDGEMENT

We thank the anonymous reviewers for their constructive feedback and insightful comments. This research is funded, in part, by Singapore Ministry of Education Tier 2 grant MOE-T2EP20124-0007.

REFERENCES

- [1] Marco Carbone, Sonia Marin, and Carsten Schürmann. 2023. A Logical Interpretation of Asynchronous Multiparty Compatibility. In *Logic-Based Program Synthesis and Transformation: 33rd International Symposium, LOPSTR 2023, Cascais, Portugal, October 23-24, 2023, Proceedings* (Cascais, Portugal). Springer-Verlag, Berlin, Heidelberg, 99–117. https://doi.org/10.1007/978-3-031-45784-5_7
- [2] Giuseppe Castagna, Mariangiola Dezani-Ciancaglini, and Luca Padovani. 2011. On Global Types and Multi-party Sessions. In *Formal Techniques for Distributed Systems*, Roberto Bruni and Juergen Dingel (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–28.
- [3] Michael Christensen, Timothy Sherwood, Jonathan Balkind, and Ben Hardekopf. 2021. Wire sorts: a language abstraction for safe hardware composition. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (Virtual, Canada) (PLDI)*. Association for Computing Machinery, New York, NY, USA, 175–189. <https://doi.org/10.1145/3453483.3454037>
- [4] Pierre-Malo Deniérou and Nobuko Yoshida. 2012. Multiparty session types meet communicating automata. In *Proceedings of the 21st European Conference on Programming Languages and Systems (Tallinn, Estonia) (ESOP'12)*. Springer-Verlag, Berlin, Heidelberg, 194–213. https://doi.org/10.1007/978-3-642-28869-2_10
- [5] Pierre-Malo Deniérou and Nobuko Yoshida. 2013. Multiparty compatibility in communicating automata: characterisation and synthesis of global session types. In *Proceedings of the 40th International Conference on Automata, Languages, and Programming - Volume Part II (Riga, Latvia) (ICALP'13)*. Springer-Verlag, Berlin, Heidelberg, 174–186. https://doi.org/10.1007/978-3-642-39212-2_18
- [6] Lorenzo Gheri and Nobuko Yoshida. 2023. Hybrid Multiparty Session Types: Compositionality for Protocol Specification through Endpoint Projection. *Proc. ACM Program. Lang.* 7, OOPSLA1, Article 79 (April 2023), 31 pages. <https://doi.org/10.1145/3586031>
- [7] Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2016. Multiparty Asynchronous Session Types. *J. ACM* 63, 1, Article 9 (March 2016), 67 pages. <https://doi.org/10.1145/2827695>
- [8] Hans Hüttel, Ivan Lanese, Vasco T. Vasconcelos, Luis Caires, Marco Carbone, Pierre-Malo Deniérou, Dimitris Mostrous, Luca Padovani, António Ravara, Emilio Tuosto, Hugo Torres Vieira, and Gianluigi Zavattaro. 2016. Foundations of Session Types and Behavioural Contracts. *ACM Comput. Surv.* 49, 1, Article 3 (2016), 36 pages. <https://doi.org/10.1145/2873052>
- [9] Minseong Jang, Jungin Rhee, Woojin Lee, Shuangshuang Zhao, and Jeehoon Kang. 2024. Modular Hardware Design of Pipelined Circuits with Hazards. *Proceedings of the ACM on Programming Languages* 8, PLDI (June 2024), 28–51. <https://doi.org/10.1145/3656378>
- [10] Elaine Li, Felix Stutz, Thomas Wies, and Damien Zufferey. 2023. Complete Multiparty Session Type Projection with Automata. In *Computer Aided Verification: 35th International Conference, CAV 2023, Paris, France, July 17–22, 2023, Proceedings, Part III* (Paris, France). Springer-Verlag, Berlin, Heidelberg, 350–373. https://doi.org/10.1007/978-3-031-37709-9_17
- [11] Xun Li, Mohit Tiwari, Jason K. Oberg, Vineeth Kashyap, Frederic T. Chong, Timothy Sherwood, and Ben Hardekopf. 2011. Caisson: A Hardware Description Language for Secure Information Flow. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM, San Jose California USA, 109–120. <https://doi.org/10.1145/1993498.1993512>
- [12] Rachit Nigam, Pedro Henrique Azevedo de Amorim, and Adrian Sampson. 2023. Modular Hardware Design with Timeline Types. *Proc. ACM Program. Lang.* 7, PLDI, Article 120 (June 2023), 25 pages. <https://doi.org/10.1145/3591234>
- [13] Tyler Potyondy, Anthony Tarbinian, Leon Schuermann, Eric Mugnier, Adin Ackerman, Amit Levy, and Pat Pannuto. 2026. Rage against the state Machine: Type-States Hardware Peripherals for increased driver correctness. In *Proceedings of the 31st ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '26)*. Association for Computing Machinery, Pittsburgh, PA, USA. Preprint available at <https://tylerpotyondy.com/assets/pdf/asplos26-potyondy.pdf>.
- [14] Jason Zhijiangcheng Yu, Aditya Ranjan Jha, Umang Mathur, Trevor E Carlson, and Prateek Saxena. 2026. Anvil: A General-Purpose Timing-Safe Hardware Description Language. In *Proceedings of the 31st ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '26)*. Association for Computing Machinery, Pittsburgh, PA, USA. <https://doi.org/10.1145/3779212.3790125> To appear. Preprint available at arXiv:2503.19447.
- [15] Danfeng Zhang, Yao Wang, G. Edward Suh, and Andrew C. Myers. 2015. A Hardware Design Language for Timing-Sensitive Information-Flow Security. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '15)*. Association for Computing Machinery, New York, NY, USA, 503–516. <https://doi.org/10.1145/2694344.2694372>