

Robust Hyperproperty Preservation for Secure Compilation

(Extended Abstract)

Deepak Garg¹ Cătălin Hrițcu² Marco Patrignani³ Marco Stronati² David Swasey¹

¹MPI-SWS ²Inria Paris ³CISPA

Abstract

We map the space of soundness criteria for secure compilation based on the preservation of hyperproperties in arbitrary adversarial contexts, which we call robust hyperproperty preservation. For this, we study the preservation of several classes of hyperproperties and for each class we propose an equivalent "property-free" characterization of secure compilation that is generally better tailored for proofs. Even the strongest of our soundness criteria, the robust preservation of all hyperproperties, seems achievable for simple transformations and provable using context back-translation techniques previously developed for showing fully abstract compilation. While proving the robust preservation of hyperproperties that are not safety requires such powerful context back-translation techniques, for preserving safety hyperproperties robustly, translating each finite trace prefix back to a source context seems to suffice.

Extended Abstract

Secure compilation is an emerging field that puts together advances in programming languages, verification, compilers, and security enforcement mechanisms to devise secure compiler chains that eliminate many of today's devastating low-level vulnerabilities. One class of low-level vulnerabilities arises when code written in a safe language is compiled and interacts with unsafe code written in a lower-level language, e.g., when linking with libraries. While currently all the guarantees of the source code are generally lost in such cases, we would like to devise secure compilers that protect some of the security guarantees established in the source language even against adversarial low-level contexts.

What is a good soundness criterion for a compiler that attains this? Fully abstract compilation [1] is a criterion that provides one potential answer to this question: a fully abstract compiler preserves (and reflects) the observational equivalence of partial source programs. In more detail, a compiler is fully abstract when any two partial source programs that are observationally indistinguishable by all compatible adversarial source contexts get compiled to two target-level programs that are indistinguishable by all adversarial target-level contexts. While fully abstract compilation has received significant attention in the literature, the indistinguishability of partial programs in all contexts is not the only security property one might be interested in preserving. In this work we set out to explore a much larger space of security properties that can be preserved even against adversarial target-level contexts.

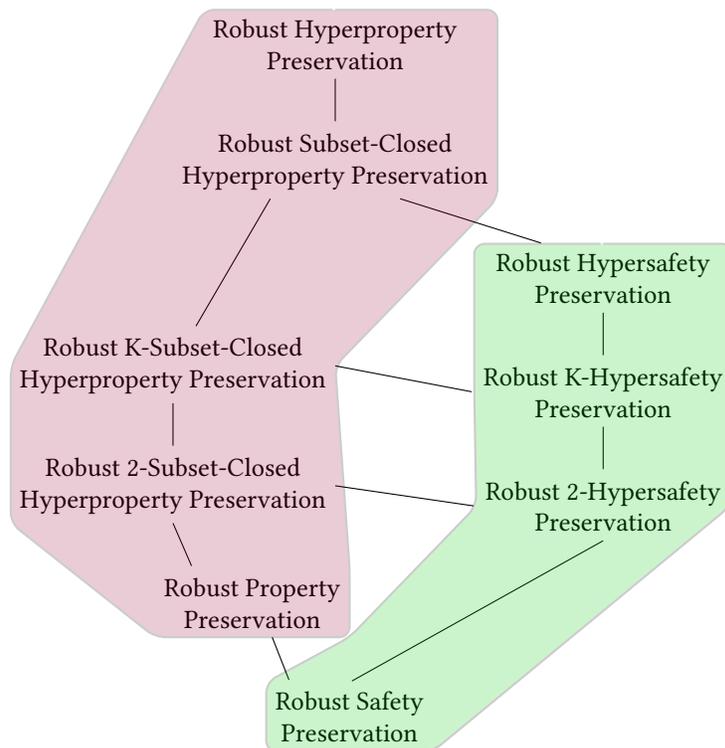


Figure 1. Different notions of robust hyperproperty preservation. Notions higher in the figure are stronger.

Specifically, we look at preserving classes of *hyperproperties* despite adversarial contexts. Hyperproperties [2] are a generalization of trace properties that can express important security policies such as noninterference. While trace properties are formally expressed as sets of (potentially infinite) traces, hyperproperties are sets of sets of traces. Concretely, these traces are built over events such as inputs from and outputs to the environment [8]. We say that a complete program P satisfies a hyperproperty H when the set of traces of P is a member of H , or formally $\{t \mid P \rightsquigarrow t\} \in H$, where $P \rightsquigarrow t$ indicates that program P emits trace t . We say that a partial program P *robustly satisfies* [7] a hyperproperty H when P linked with any (adversarial) context satisfies H . Armed with this notion of robust satisfaction of hyperproperties, we define secure compilation as preserving the robust satisfaction of a class of hyperproperties \mathcal{H} , so if a partial source program P robustly satisfies a hyperproperty $H \in \mathcal{H}$ (wrt. all source contexts) then its compilation $P\downarrow$ must also robustly satisfy H (wrt. all target-level contexts).

We study the preservation of robust satisfaction for various classes of hyperproperties, many of which are mentioned in Figure 1, and which include all hyperproperties, subset-closed hyperproperties, safety hyperproperties, trace properties, and safety properties. For each such class we propose an equivalent “property-free” characterization of secure compilation that is generally better suited for proofs. For instance, we prove that preserving all hyperproperties robustly can be equivalently stated as the following criterion we call *hyper-robust compilation* (where C are contexts and $C[P]$ is the linking of a context C with a program P):

$$\forall P. \forall C_T. \exists C_S. \forall t. C_T[P \downarrow] \rightsquigarrow t \iff C_S[P] \rightsquigarrow t$$

This requires that, given a program P , each target context C_T can be mapped to a source context C_S in a way that perfectly preserves the set of traces produced when linking with P and $P \downarrow$ respectively. On the other hand, preserving all trace properties robustly is equivalent to the following *robust compilation* criterion:

$$\forall P. \forall C_T. \forall t. \exists C_S. C_T[P \downarrow] \rightsquigarrow t \Rightarrow C_S[P] \rightsquigarrow t$$

Compared to the previous definition, the $\exists C_S$ and $\forall t$ quantifiers in this definition are swapped and the implication is in just one direction: Each (bad) trace in the target can be emulated using a different source context C_S . The intuition is that if the compiled program is able to produce a trace, that same trace must also be produceable in the source. Swapping the quantifiers is crucial for transitioning from hyperproperties (sets of traces) to properties (traces). The $\forall t. \exists C_S$ quantifiers of robust compilation let us pick a different context C_S for each trace t . The reversed ordering $\exists C_S. \forall t$ of robust hyperproperty preservation instead requires us pick the context C_S before the traces. As a final example, preserving only safety properties robustly [4] is equivalent to the following *robustly safe compilation* criterion:

$$\forall P. \forall C_T. \forall t. C_T[P \downarrow] \rightsquigarrow t \Rightarrow \forall m \leq t. \exists C_S t'. C_S[P] \rightsquigarrow t' \wedge m \leq t'$$

Here only the (bad) finite prefixes m of a potentially infinite trace t in the target need to be back-simulated in the source. Safety properties are concerned with (bad) prefixes that must not happen for the property to hold. If a safety property holds in the source and some prefix broke this property in the target, then the same prefix would also exist in the source, contradicting the fact that the property holds in the source.

Even the strongest of our secure compilation criteria, hyper-robust compilation, which is, as explained above, equivalent to the robust preservation of all hyperproperties, seems achievable. We plan to demonstrate this by adapting a recent fully abstract translation of a simply typed λ -calculus into the untyped λ -calculus [3]. For this to be interesting, we first extend the two λ -calculi with a notion of trace by adding inputs from and outputs to the environment. For achieving hyper-robust compilation we also extend the source language with

recursive types. This allows us to encode the untyped λ -calculus values using recursive, product, and sum types, allowing for a precise back-translation of contexts. We expect that the logical relation proof technique of Devriese et al. [3] can be adapted to prove the hyper-robust compilation of their translation. Moreover, if we drop recursive types from the source we expect to still be able to use the approximate back-translation of Devriese et al. [3] to show robust hypersafety preservation, a weaker security criterion.

While preserving hyperproperties that are not safety seems to require powerful context back-translation techniques, for preserving safety hyperproperties translating each finite trace prefix individually back to a source context is also possible. This could potentially be simpler as it can benefit from proof techniques that are based on trace semantics [5], which was also used in the context of full abstraction proofs [6, 10]. In Figure 1 we mark in green the secure compilation criteria for which mapping finite trace prefixes is possible, and in light purple the ones for which it is not.

Finally, the property-free characterizations of all classes of hyperproperties from Figure 1 have quantifier alternation of the form $\forall P. \forall C_T \dots \exists C_S \dots$, so a (constructive) proof that a compiler satisfies such a characterization can define C_S as a function of the source program P and the target context C_T . While the dependence of C_S on C_T is essential in most cases, the dependence of C_S on P is necessary only when the target language allows the context to make observations that the source does not allow. For example, the target language may have reflection but the source language may not have it. However, in many cases, this kind of an abstraction mismatch does not exist and, in fact, many existing proof techniques [9], including the aforementioned context back-translation techniques, construct C_S only from C_T , independent of the source program P . This begs the question of what kinds of properties are actually preserved by a compiler that satisfies a *stronger* criterion of the form $\forall C_T. \exists C_S. \forall P \dots$. For example, what properties of source programs are preserved by a compiler that satisfies the following stronger variant of hyper-robust compilation?

$$\forall C_T. \exists C_S. \forall P. \forall t. C_T[P \downarrow] \rightsquigarrow t \iff C_S[P] \rightsquigarrow t$$

We conjecture that such strong soundness criteria correspond to the robust preservation of *relational* properties of programs. In particular, the criterion listed above implies (the interesting direction of) full abstraction, which is the robust preservation of a specific relational property, namely, observational equivalence. Investigating which of the criteria of Figure 1 imply or are implied by fully abstract compilation is interesting future work.

References

- [1] M. Abadi. Protection in programming-language translations. *Secure Internet Programming*. 1999.
- [2] M. R. Clarkson and F. B. Schneider. Hyperproperties. *Journal of Computer Security*, 18(6):1157–1210, 2010.
- [3] D. Devriese, M. Patrignani, and F. Piessens. Fully-abstract compilation by approximate back-translation. *POPL*, 2016.
- [4] A. D. Gordon and A. Jeffrey. Types and effects for asymmetric cryptographic protocols. *JCS*, 12(3-4):435–483, 2004.
- [5] A. Jeffrey and J. Rathke. Java Jr: Fully abstract trace semantics for a core Java language. *ESOP*. 2005.
- [6] Y. Juglaret, C. Hritcu, A. Azevedo de Amorim, B. Eng, and B. C. Pierce. Beyond good and evil: Formalizing the security guarantees of compartmentalizing compilation. *CSF*, 2016.
- [7] O. Kupferman and M. Y. Vardi. Robust satisfaction. *CONCUR*. 1999.
- [8] X. Leroy. A formally verified compiler back-end. *JAR*, 43(4):363–446, 2009.
- [9] M. S. New, W. J. Bowman, and A. Ahmed. Fully abstract compilation via universal embedding. *ICFP*. 2016.
- [10] M. Patrignani and D. Clarke. Fully abstract trace semantics for protected module architectures. *Computer Languages, Systems & Structures*, 42:22–45, 2015.