

# Phantom Contracts for Better Linking

Daniel Patterson — Northeastern University — dbp@ccs.neu.edu

POPL 2019 Student Research Competition Extended Abstract

**Introduction** There are dozens of programming languages in wide use, and people are designing new ones every day. Increasingly, languages provide programmers with rich types to enforce invariants about their code statically. But an outstanding challenge remains how to integrate components written in new languages into extant systems. Programmers currently rely on unsafe FFIs that defer errors to runtime (if errors are reported at all), failing both to permit a gradual evolution from older languages to a newer ones, and to support graceful integration of domain specific languages.

The essential problem is that, at the point of linking (in low-level bytecode, object code, etc), any static type information that could have been used to report errors has already been discarded. To support linking errors, e.g., that a function is called with the wrong number or type of arguments, we must *preserve* type information. But what should target type systems look like? One option is to add *just enough* types to support what you are compiling. In that case, as soon as you build a compiler from a sufficiently different source language, you will have to extend this target language. Another option is to try to give the target language the richest possible type system: e.g., a dependent high-order separation logic ([5] or similar), increasing the complexity of implementing compilers.

We propose a third option: enrich an existing possibly-typed target language with **phantom contracts**, which are written in a rich but simple operational language, but that *run at type-checking time*. The phantom language can then be used to encode source type information that cannot be captured in the existing target and to enforce flexible linking invariants. Phantom contracts are refinements to target types that operate over ghost state used in type checking. Source types translate to target phantom computations that, after linking and type checking, are erased.

**By Example** We demonstrate the idea by example. We show compilers from two source languages into a target with phantom contracts. We contrast this with the typical approach for type-preserving compilation, using this to show how phantom contracts work and why they are an improvement.

We begin with our source languages. One is  $\text{INDEX}^\pm$ , the simply-typed lambda calculus with integers enriched with an positive/negative index algebra that can be used to track the sign of numbers. The syntax and a selection of the typing rules are in Figure 1; operational semantics are standard, not depending on indexes. Number types  $\eta$  capture our index system, where literal numbers are either  $+0$  or  $-$ , and computations can degrade to the unknown  $?$ . Number type polymorphism allows functions like  $e = \Lambda\alpha.\lambda x:\alpha.x * 2$ , which has type  $\forall\alpha.\alpha * +0$ . Then  $e[+0]2$  would have type  $+0$ , since type application reduces  $+0 * +0$  to  $+0$ .

Our second source language is **BABYDILL**, a simply typed linear lambda calculus based on removing polymorphism from PDILL[13] (and adding a base type). The syntax and a few typing rules are shown in Figure 2 (the reader is encouraged to consult [13] for more details).

**Phantom Contracts** A typical approach to supporting type-preserving compilation from various languages is to incorporate all necessary features directly into the target. We would therefore

$\text{INDEX}^\pm$	$\tau ::= \eta \mid \forall\alpha.\tau \mid \tau \rightarrow \tau$				
	$\eta ::= \alpha \mid +0 \mid - \mid ? \mid \eta + \eta \mid \eta * \eta$				
	$e ::= n \mid x \mid e + e \mid e * e \mid \lambda x:\tau.e \mid e e$				
	$v ::= n \mid \lambda x:\tau.e \mid \Lambda\alpha.e$				
	$\frac{n \geq 0}{H; \Gamma \vdash n : +0}$	$\frac{n < 0}{H; \Gamma \vdash n : -}$	$\frac{x : \tau \in \Gamma}{H; \Gamma \vdash x : \tau}$		
	$\frac{H; \Gamma \vdash e_1 : \eta_1 \quad H; \Gamma \vdash e_2 : \eta_2}{H; \Gamma \vdash e_1 * e_2 : \Downarrow(\eta_1 * \eta_2)}$		$\frac{H, \alpha; \Gamma \vdash e : \tau}{H; \Gamma \vdash \Lambda\alpha.e : \forall\alpha.\tau}$		
				$\Downarrow(+0 + +0) = +0$	
				$\Downarrow(- + +0) = ?$	
				$\Downarrow(? + \eta) = ?$	
				etc...	
	$\frac{H; \Gamma \vdash e : \forall\alpha.\tau}{H; \Gamma \vdash e[\eta] : \Downarrow(\tau[\eta/\alpha])}$			$\Downarrow(\eta_1 + \eta_2) = \eta_1 + \eta_2$	

Figure 1: Syntax and (very) selected static semantics for  $\text{INDEX}^\pm$ .

$\text{BABYDILL}$	$\tau ::= \text{unit} \mid \tau \rightarrow \tau \mid !\tau \mid \tau \& \tau \mid \tau \otimes \tau$				
	$e ::= () \mid x \mid a \mid \lambda a:\tau.e \mid e e' \mid e$				
	$\text{let } !x = e \text{ in } e' \mid \langle e, e' \rangle \mid e.1 \mid e.2$				
	$(e, e) \mid \text{let } (a, a') = e \text{ in } e'$				
	$v ::= () \mid \lambda a:\tau.e \mid e \mid \langle e, e' \rangle \mid (v, v')$				
	$\frac{x : \tau \in \Gamma}{\Gamma; \cdot \vdash x : \tau}$	$\frac{\Gamma; \Delta, a : \tau_1 \vdash e : \tau_2}{\Gamma; \Delta \vdash \lambda a : \tau_1. e : \tau_1 \rightarrow \tau_2}$			
	$\frac{\Gamma; \Delta_1 \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma; \Delta_2 \vdash e_2 : \tau_1 \quad \Delta_1, \Delta_2 = \Delta}{\Gamma; \Delta \vdash e_1 e_2 : \tau_2}$				
	$\frac{\Gamma; \Delta \vdash e_1 : \tau_1 \quad \Gamma; \Delta \vdash e_2 : \tau_2}{\Gamma; \Delta \vdash \langle e_1, e_2 \rangle : \tau_1 \& \tau_2}$	$\frac{\Gamma; \cdot \vdash e : \tau}{\Gamma; \cdot \vdash e : !\tau}$			

Figure 2: Syntax and (very) selected static semantics for **BABYDILL**.

end up with a target with sign indexes and linearity (at least), and we would have some translation into this target. If **BABYDILL** had integers, we could compile literal integers to precise types, but any source integer types would translate to whatever  $?$  translated to, to capture the fact that **BABYDILL** does not track sign. Here the compiler can determine how to translate representations—there also may be cases where we need programmer annotations in order to figure out what code to generate, which is an idea we have explored in previous work called **linking types** [7].

This approach, while theoretically straightforward, has a major downside: it's likely that the target will be specialized to the particular source language(s) chosen, and as a result would have to change in order to support more languages. Over time, it may converge to something rich enough to capture almost everything needed, but at that point the language would be a hodge-podge of features.

Rather than try to incorporate all of the type features of the source languages into the type system of the target, we propose enriching the target language with **phantom contracts**, that on their own have straightforward semantics but can be used by compiler

writers to encode complex invariants. In Figure 3 we present syntax and a small selection of the static semantics for `PHANTOMLC`, which is the simply typed lambda calculus with pairs and recursive types extended with phantom contracts. We highlight a few important details, and then consider our two-language case study.

First, note that our *static* semantics involve *running* our phantom contracts. Our phantom language is a simply typed lambda calculus with *sexpressions*, a first-order store, and *assertions* that can fail (there is no reduction rule for  $\delta^1(\text{assert}, \text{false})$ ). The phantom computations are guaranteed to terminate or fail, and these two possibilities indicate whether the invariants encoded in the phantom contracts were satisfied. This allows compilers to encode their representations of types and the computational fragments that are necessary for checking that the fragments that they receive have the proper phantom representations.

`PHANTOMLC`

$$\begin{aligned} \tau &::= \text{int} \mid \tau \rightarrow \tau' \mid \mu\alpha.\tau \mid \alpha \\ e &::= \hat{e}\{\varphi\} \\ \hat{e} &::= x \mid n \mid e + e \mid e * e \mid \lambda x:\tau.e \mid e e' \\ &\quad (e, e) \mid \text{fst } e \mid \text{snd } e \mid \text{fold } e \\ &\quad \text{unfold } e \mid \text{let } x\{v\}:\tau = e \text{ in } e' \\ v &::= n \mid \text{fold } v \mid \lambda x:\tau.e \mid (v, v) \\ \varphi &::= \text{sexp} \mid \text{ref sexp} \mid \varphi := \varphi \mid !\varphi \mid v \\ &\quad \text{match } \varphi \mid \text{nv}_1.\varphi_1 \mid \text{sv}_2.\varphi_2 \mid \text{bv}_3.\varphi_3 \mid (v_4, v_5).\varphi_4 \\ &\quad \lambda v:\varphi.\tau.\varphi \mid \varphi \varphi' \mid \varphi; \varphi' \mid \delta^{\text{a}}(\text{op}^n, \varphi_1, \dots, \varphi_n) \\ \text{sexp} &::= n \mid s \mid \text{true} \mid \text{false} \mid (\text{sexp}, \text{sexp}') \\ \text{op}^1 &::= \text{assert} \mid \text{not} \mid \text{length} \mid \dots \\ \text{op}^2 &::= \text{sexp} = \mid \text{append} \mid + \mid * \mid \dots \\ \varphi\tau &::= \text{sexp} \mid \text{ref sexp} \mid \varphi\tau \rightarrow \varphi\tau \end{aligned}$$

$$\frac{\Gamma; S \vdash \hat{e} : \tau; S' \quad \Gamma \vdash (S', \varphi) \Downarrow (S'', \varphi')}{\Gamma; S \vdash \hat{e}\{\varphi\} : \tau; S''}$$

$$\frac{\Gamma; S_1 \vdash \hat{e}_1 : \tau; S_2 \quad \Gamma \vdash (S_2, \varphi) \Downarrow (S_3, \varphi') \quad \Gamma, x:\tau, v:\varphi'; S_3 \vdash e_2:\tau_2; S_4}{\Gamma; S_1 \vdash \text{let } x\{v\}:\tau = \hat{e}_1\{\varphi\} \text{ in } e_2 : \tau_2; S_4}$$

$$\frac{v:\varphi \in \Gamma \quad \Gamma \vdash (S, \varphi) \Downarrow (S', \text{true})}{\Gamma \vdash (S, v) \Downarrow (S, \varphi) \quad \Gamma \vdash (S, \delta^1(\text{assert}, \varphi)) \Downarrow (S', \text{true})}$$

$$\frac{\Gamma \vdash (S, \varphi_1) \Downarrow (S', s) \quad \Gamma, v:s \vdash (S', \varphi_2) \Downarrow (S'', \varphi_2')}{\Gamma \vdash (S, \text{match } \varphi_1 \dots \mid \text{sv}.\varphi_2 \dots)}$$

Figure 3: Syntax and selected static semantics for `PHANTOMLC`.

**Compilers using Phantom Contracts** To compile `INDEX±`, we consider translating well-typed expressions. This means our compiler will be defined case-wise over the typing rules of the language. We show an illustrative portion of it in Figure 4, where  $(e)^+$  is the compiled term. The key case in the figure, when considering linking, is application. If the argument came from another language, compiler, etc, we could use the same target code, replacing  $e'^+$  with that code. We can see that the rest of the compiler would then be able to enforce the sign restrictions from the source, because they were encoded by the compiler into the phantom contracts and checked at that point.

To compile `BABYDILL`, we again define a translation case-wise over typing derivations. A portion of this translation is shown in Figure 5. Note that the only thing we use our phantom contracts to track is that linearity is not violated, by using reference counts in our phantom store. This presentation is slightly simplified, as the full translation needs to implement a caller-save strategy on the reference counts to support variable shadowing, which we elide for space. One case that is important is lazy (i.e., additive) pairs (of

$$\begin{aligned} (e_1 + e_2 : \eta)^+ &= \text{let } \_e_1\{v_1\} = e_1^+ \text{ in} \\ &\quad \text{let } \_e_2\{v_2\} = e_2^+ \text{ in} \\ &\quad (\_e_1 + \_e_2)\{\langle\eta\rangle\} \\ (x : \tau)^+ &= x\{\delta^1(\text{assert}, \delta^2(\text{sexp} =, \langle\tau\rangle, v_x)); \langle\tau\rangle\} \\ (\lambda x:\tau.e : \tau \rightarrow \tau')^+ &= (\lambda x:\tau^+.\text{let } \_v_x = \mathbf{0}\{\langle\tau\rangle\} \text{ in } e^+) \\ &\quad \{\langle\tau \rightarrow \tau'\rangle\} \\ (e e' : \tau)^+ &= \text{let } f\{v_f\} = e^+ \text{ in} \\ &\quad \text{let } a\{v_a\} = e'^+ \text{ in} \\ &\quad (f a)\{\text{match } v_f \\ &\quad \mid (\_ \rightarrow \_', (v_{\tau'}, v_{\tau'})) \\ &\quad \delta^1(\text{assert}, \delta^2(\text{sexp} =, v_{\tau'}, v_a)); \langle\tau\rangle\} \\ \langle\tau \rightarrow \tau'\rangle &= (\_ \rightarrow \_', \langle\langle\tau\rangle, \langle\tau'\rangle\rangle) \quad ?^+ = \text{int} \\ \langle\langle\mathbf{0}\rangle\rangle &= \_ + \mathbf{0}' \quad +\mathbf{0}^+ = \text{int} \end{aligned}$$

Figure 4: Selections from `INDEX±` to `PHANTOMLC` compiler.

type  $\tau_1 \& \tau_2$ ), for which only one projection is used. In that case, we run the phantom contracts on one side, reset reference counts, and then run the other, so that both sides use all variables. The same technique would also be used to support *if*. Other cases have empty phantom contracts, because we only have to handle cases that affect the reference counts.

$$\begin{aligned} (\lambda a : \tau_1.e : \tau_1 \multimap \tau_2)^+ &= \lambda a : \tau_1^+.\text{let } \_a = \text{ref } 1 \text{ in} \\ &\quad \text{let } \_r\{\_r\} = e^+ \text{ in} \\ &\quad \_r\{\delta^1(\text{assert}, \delta^2(\text{int} =, !a, \mathbf{0})); \_r\} \\ (a : \tau)^+ &= a\{a := \delta^2(-, !a, 1)\} \\ (\Gamma; a_1 : \tau_1 \dots \vdash &= \text{let } v_1\{v_1\} = e_1^+ \text{ in let } \_ = \mathbf{0}\{\delta^1(\text{assert}, \\ (e_1, e_2) : \tau_1 \& \tau_2)^+ &\quad \delta^2(\text{int} =, !a_1, \mathbf{0})\} \dots; a_1 := 1 \dots) \text{ in } (v_1, e_2^+) \end{aligned}$$

Figure 5: Selections from `BABYDILL` to `PHANTOMLC` compiler.

As before, we can see what linking we can support. With no shared base types, the default interactions are pretty limited, but again if we extended `BABYDILL` with integers, we could extend our compiler to support them. But in this case, the compilation is more interesting, because what we would compile them to is the particular choice of phantom contracts used by the `INDEX±` compiler, rather than to a built-in feature of the target language. The resulting interoperability would be the same, and indeed, that's the goal: phantom contracts are not about making fundamentally easier the task of building a given type-preserving compiler, but rather to prevent the issue of having to migrate the entire ecosystem because some feature needs to be added to the target language. With such a system, type-preserving compilers could be developed independently, and only at the point of interface do common types and calling conventions need to be established, similarly to how they have to be established with existing FFIs. Individual compiler writers could experiment with different representations, benefiting from not having to change the underlying target language infrastructure, and getting baseline interoperability to the extent that they provide phantom interoperability wrappers.

**Related Work** The name should hint this derives inspiration both from phantom types ([1]) and contracts (e.g., soft contract verification [6]) / gradual typing, where the notion of encoding typing constraints operationally is familiar (e.g., AGT [3]), though the fact that our operational semantics is at type-checking time makes things quite different. But we also draw much inspiration from refinement type systems (e.g., [2, 8, 12], where the last made an explicit connection to static higher-order contracts) and from ghost state, both from separation logic (e.g., [4]) and F\*'s ghost monad [11] (and earlier ghost refinements [10]). We also see connections to work translating typing rules into rewrite systems (e.g., [9]).

## References

- [1] Matthew Fluet and Riccardo Pucella. 2002. Phantom Types and Subtyping. In Proceedings of the Second IFIP International Conference on Theoretical Computer Science (TCS '02).
- [2] Tim Freeman. 1994. Refinement Types for ML. PhD Thesis.
- [3] Ron Garcia, Alison Clark, Éric Tanter. 2016. Abstracting Gradual Typing. In 43rd ACM Symposium on Principles of Programming Languages (POPL '16).
- [4] Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-Order Ghost State. In 21st ACM International Conference on Functional Programming (ICFP '16).
- [5] Aleksander Nanevski, Amal Ahmed, Greg Morrisett, and Lars Birkedal. 2007. Abstract Predicates and Mutable ADTs in Hoare Type Theory. In 16th European Symposium on Programming (ESOP '07).
- [6] Phc C. Nguyen, Sam Tobin-Hochstadt, and David Van Horn. 2014. Soft contract verification. In 19th ACM International Conference on Functional Programming (ICFP '14).
- [7] Daniel Patteron and Amal Ahmed. 2017. Linking Types for Multi-Language Software: Have Your Cake and Eat It Too. In SNAPL: Summit on Advances in Programming Languages (SNAPL'17), May 2017.
- [8] Patrick Rondon, Ming Kawaguchi, and Ranjit Jhala. 2008. Liquid Types. In 30th ACM Conference on Programming Language Design and Implementation (PLDI '08).
- [9] Aaron Stump, Garrin Kimmell, Roba El Haj Omar. 2011. Type Preservation as a Confluence Problem. In Proceedings of the 22nd International Conference on Rewriting Techniques and Applications (RTA '11).
- [10] Nikhil Swamy, Juan Chen, Cdric Fournet, Pierre-Yves Strub, Karthikeyan Bhargavan, Jean Yang. Secure Distributed Programming with Value-dependent Types. 2011. In 16th ACM SIGPLAN International Conference on Functional Programming (ICFP '11).
- [11] Nikhil Swamy, Ctlin Hricu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cdric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoue, Santiago Zanella-Bguelin. 2016. Dependent Types and Multi-Monadic Effects in F\*. In 43rd ACM Symposium on Principles of Programming Languages (POPL '16).
- [12] Niki Vazou, Patrick M. Rondon, and Ranjit Jhala. 2013. Abstract Refinement Types. In 22nd European Symposium on Programming (ESOP '13).
- [13] Jianzhou Zhao, Qi Zhang, and Steve Zdancewic. 2010. Relational Parametricity for a Polymorphic Linear Lambda Calculus. In Proceedings of the 8th Asian conference on Programming languages and systems (APLAS'10).