

## Research Statement

### Higher-Order Contracts for JavaScript

JavaScript is an untyped and dynamic programming language with objects and first-class functions. While it is most well-known as the client-side scripting language for web sites, it is also increasingly used for non-browser development, such as developing server-side applications with *Node.js*, for game development, to implement platform-independent mobile applications, or as an intermediate language for other languages to target, such as *TypeScript* or *Dart*. Hence, it does not come as a surprise that JavaScript is in the focus of many research works, out of the need to create better development tools for JavaScript programmers and to launch new language features.

Unfortunately, JavaScript itself has no real security awareness: there is no namespace or encapsulation management, there is a global scope for functions and variables, all scripts have the same authority, and everything can be modified, from the fields and methods of an object over its prototype property to the scope chain of a function closure. As a consequence, JavaScript code is prone to injection attacks, library code can read and manipulate everything reachable from the global scope, and third-party code can get access to sensitive data. Furthermore, side effects may cause unexpected behavior so that program understanding and maintenance become difficult.

To overcome this limitations we propose using contracts with run-time monitoring. Software contracts were introduced with Meyer's *Design by Contract*<sup>TM</sup> methodology which stipulates invariants for objects as well as Hoare-like pre- and postconditions for functions. Since Meyer's work, the contract idea has taken off and attracted a plethora of follow-up works that ranges from contract monitoring of higher-order functional languages over semantic investigations and studies on blame assignment to extensions in various directions: polymorphic contracts, behavioral and temporal contracts, etc.

Contract monitoring has become a prominent mechanism to provide strong guarantees for programs in dynamically typed languages while preserving their flexibility and expressiveness. Hence, the first higher-order contract systems were devised for Scheme and Racket, but other dynamic languages like JavaScript, Python, PHP, Ruby, and Lua have followed suit.

My prior and current research focusses on the design and implementation of *TreatJS* [7], a language embedded, higher-order contract system for JavaScript which enforces contracts by run-time monitoring. Beyond the standard abstractions for higher-order contracts (flat contract, function contracts, dependent contracts), *TreatJS* comes with the following unique features:

- *TreatJS* provides a contract constructor that constructs and composes contracts at run-time using contract abstraction. A contract constructor may contain arbitrary JavaScript code and it may encapsulate a local state. Contract constructors are the building blocks for dependent contracts, parameterized contracts, and recursive contracts.

- *TreatJS* provides intersection and union operators for contracts. The specification of these operators is inspired by their type-theoretic counterparts. They enable developers to specify independent properties in independent contracts and to combine them using intersection and union.
- *TreatJS*'s blame assignment for higher-order contracts with intersection and union [6] has a number of novel aspects. First, it uses constraints to create a structure for computing positive and negative blame according to the semantics of subject and context satisfaction, respectively. Second, it applies a compatibility check to distinguish contracts from different sides of an intersection or union, and third it provides three general monitoring semantics that handle the visibility of contracts inside of predicate code.
- *TreatJS* gives noninterference a high priority. Its implementation employs a membrane-based sandbox to keep the predicate code apart from the normal program execution and it encapsulates objects that are passed through the membrane to enforce write protection and to withhold external bindings from functions. Contracts are guaranteed not to exert side effects on a contract abiding program execution.
- *TreatJS* is implemented as a library in JavaScript. It enables a developer to specify all aspects of a contract using the full JavaScript language. Proxies implement delayed contract checking of function and object contracts and guarantee full interposition for the full JavaScript language, including the with-statement and eval.
- The implementation of *TreatJS* illustrates the need for a different proxy constructor that is better suited for the implementation of contract wrappers. One issue with the current contract implementation arises because a contract wrapper is different (not pointer-equal) from the target object so that an equality test between wrapper and target returns false instead of true. Thus, *TreatJS* comes with an implementation of a transparent object proxy [4] which ensures transparent operations with all JavaScript programs.

However, to bring efficient and user-friendly contracts to JavaScript, there are still some issues and directions open for future work. The following sections address some open research topics I plan to continue in the short and medium term.

---

## Static Contract Simplification

Writing formal and precise specifications in form of contracts sounds appealing, but it comes with a cost: Dynamic contract monitoring degrades the execution time of the underlying program [9, 7]. All existing contract systems like Racket's contract framework [3, Chapter 7], Disney's JavaScript contract system *contracts.js* [2], *JSTest2* [5], or *TreatJS* [7] for JavaScript report a considerable slowdown when extending programs with contracts.

These costs arise because every contract extends a program with additional code that checks the contract while the program executes. Moreover, developers may add contracts to frequently used functions and objects on hot-paths in a program. In particular, predicates may repeatedly check the same values and different predicates may check redundant parts.

In contrast, static contract checking [11] avoids runtime costs by removing contracts after inspection. However, static contract checking is not suitable for a language like JavaScript.

The dynamic nature of JavaScript requires dynamic contract monitoring. Completely static techniques would lead to a huge number of false positives.

My ongoing work on *Static Contract Simplification* attacks this issue with compile-time program transformation. It adapts ideas from previous work on hybrid contract checking [10] and static contract verification [8] to evaluate as much of a contract as possible and to collapse the remaining parts to a smaller contract that is more efficient to check at run-time. To this end we unroll contracts through the program code, we detect and remove redundant parts, we check predicates where possible, and we lift the remaining fragments to the enclosing module boundary. Finally we combine the remaining fragments to new contracts which only contain parts that must be checked at run-time. Such a simplification can be done even without knowing the concrete execution of a program.

---

## Native Contract Proxies

My work on Transparent Object Proxies for JavaScript [4] already examined the issue with transparency in various use cases of JavaScript proxies and we showed that a significant number of object comparisons would fail when gradually adding contracts to a program. Therefore, we propose an alternative design for transparent proxies that is better suited for implementing a contract system like *TreatJS*. However, the presented transparent proxy is a straightforward extension of the already existing opaque proxy, i.e. it provides the same features and it enables to override the same handler traps as the opaque counterpart.

But this power comes with some danger. Proxies may redefine the semantics of the underlying target object arbitrarily and thus they prevent certain optimizations in a compiler. Examples from our work show that the sole introduction of simple forwarding proxies degrades the execution time of a JavaScript program dramatically.

However, we already showed that it is sufficient to restrict contract proxies to projections [4]. A native *observer proxy* that implements a projection could be more efficient as it does not change the semantics of the underlying target object. JIT compilers would still be able to optimize a program as usual. To make dynamic contract monitoring efficient, special contract proxies are essential to improve the run-time costs of contract monitoring.

---

## Realm-aware Pure Functions

*TreatJS* uses normal JavaScript functions as predicates. To guarantee noninterference with the actual program execution it executes predicate code in a sandbox. But this sandboxing impacts the execution of the underlying program and it complicates the writing of contracts because each needed reference must be imported into the sandbox.

In JavaScript, determining the effects of a function is nearly impossible as even a simple property access might be the call of a side-effecting getter function or the call of a handler trap which causes an undesired behaviour. Because of this flexibility, JavaScript would benefit from a new function constructor that implements a *pure function*.

A pure function is a function that only maps its input into an output without causing any

observable side effects. A pure function can inspect its input and it can evaluate pure expressions, including function calls of pure functions and the access to a property that is bound to a getter function (if this getter is also a pure function). Moreover, JavaScript proxies would also benefit as handler traps can be restricted to pure functions to implement an observer proxy that does not change the semantics of the target object.

To grant effects on certain objects, pure functions could be realm-aware, i.e. they are bound to a certain realm in which effects are permitted. One example for such a realm could be the constraint set of a contract monitor, which might be influenced by the evaluation of a handler trap.

---

## Looking Forward

In the short term I plan to continue my work on static contract simplification. So far I implemented a prototype using PLT Redex, but there is no implementation that works with JavaScript contracts. To this end I plan to develop a compile-time program transformation that simplifies contract definitions to residual contracts that are collectively cheaper to check at runtime. This transformation (which is somehow related to hybrid contract checking) combines techniques from static and dynamic contract checking and can also simplify contracts that cannot be verified entirely at compile time.

In a larger perspective I plan to continue my work on contract systems and to pursue my overall goal to bring efficient and user-friendly contracts to JavaScript. Two concrete projects in this line are 1. the development of more efficient contract wrappers that implement a projection, and 2. the design and implementation of a pure function construct for JavaScript.

## Reference

- [1] John Tang Boyland, editor. *29th European Conference on Object-Oriented Programming, ECOOP 2015, July 5-10, 2015, Prague, Czech Republic*, volume 37 of *LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
- [2] Tim Disney. `contracts.js`. <https://github.com/disnet/contracts.js>, April 2013.
- [3] Matthew Flatt, Robert Bruce Findler, and PLT. *The Racket Guide*. <http://docs.racket-lang.org/guide/index.html>, February 2017. Version 6.8.
- [4] Matthias Keil, Sankha Narayan Guria, Andreas Schlegel, Manuel Geffken, and Peter Thiemann. Transparent object proxies in javascript. In Boyland [1], pages 149–173.
- [5] Matthias Keil and Peter Thiemann. Efficient dynamic access analysis using javascript proxies. In Antony L. Hosking, Patrick Th. Eugster, and Carl Friedrich Bolz, editors, *DLS'13, Proceedings of the 9th Symposium on Dynamic Languages, part of SPLASH 2013, Indianapolis, IN, USA, October 26-31, 2013*, pages 49–60. ACM, 2013.
- [6] Matthias Keil and Peter Thiemann. Blame assignment for higher-order contracts with intersection and union. In Kathleen Fisher and John H. Reppy, editors, *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015, Vancouver, BC, Canada, September 1-3, 2015*, pages 375–386. ACM, 2015.
- [7] Matthias Keil and Peter Thiemann. `Treatjs`: Higher-order contracts for javascripts. In Boyland [1], pages 28–51.
- [8] Phuc C. Nguyen, Sam Tobin-Hochstadt, and David Van Horn. Soft contract verification. In Johan Jeuring and Manuel M. T. Chakravarty, editors, *Proceedings of the 19th ACM SIGPLAN international conference on Functional programming, Gothenburg, Sweden, September 1-3, 2014*, pages 139–152. ACM, 2014.
- [9] Asumu Takikawa, Daniel Feltey, Ben Greenman, Max S. New, Jan Vitek, and Matthias Felleisen. Is sound gradual typing dead? In Rastislav Bodík and Rupak Majumdar, editors, *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 456–468. ACM, 2016.
- [10] Dana N. Xu. Hybrid contract checking via symbolic simplification. In Oleg Kiselyov and Simon J. Thompson, editors, *Proceedings of the ACM SIGPLAN 2012 Workshop on Partial Evaluation and Program Manipulation, PEPM 2012, Philadelphia, Pennsylvania, USA, January 23-24, 2012*, pages 107–116. ACM, 2012.
- [11] Dana N. Xu, Simon L. Peyton Jones, and Koen Claessen. Static contract checking for haskell. In Zhong Shao and Benjamin C. Pierce, editors, *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, pages 41–52. ACM, 2009.