

Context Constrained Computation

Robert Atkey

Computer and Information Sciences
University of Strathclyde
robert.atkey@strath.ac.uk

James Wood

Computer and Information Sciences
University of Strathclyde
james.wood.100@strath.ac.uk

1 Introduction

In normal typed λ -calculi, variables may be used multiple times, in multiple contexts, for multiple reasons, as long as the types agree. The disciplines of linear types [Girard 1987] and coeffects [Brunel et al. 2014; Ghica and Smith 2014; Petricek et al. 2014] refine this by tracking variable usage. We might track how many times a variable is used, or if it is used co-, contra-, or invariantly. Such a discipline yields a general framework of “context constrained computing”, where constraints on variables in the context tell us something interesting about the computation being performed.

We will present work in progress on capturing the “intentional” properties of programs via a family of Kripke indexed relational semantics that refines a simple set-theoretic semantics of programs. The value of our approach lies in its generality. We can accommodate the following examples:

1. Linear types that capture properties like “all list manipulating programs are permutations”. This example uses the Kripke-indexing to track the collection of datums currently being manipulated by the program.
2. Monotonicity coeffects that track whether a program uses inputs co-, contra-, or in-variantly (or not at all).
3. Sensitivity typing, tracking the sensitivity of programs in terms of input changes. This forms the core of systems for differential privacy [Reed and Pierce 2010].
4. Information flow typing, in the style of the Dependency Core Calculus [Abadi et al. 1999].

Through discussion at the workshop, we hope to discover more applications of our framework. In future work, we plan to extend our framework with type dependency, and to explore the space of inductive data types and elimination principles possible in the presence of usage information.

The syntax and semantics we present here have been formalised in Agda: <https://github.com/laMudri/quantitative/>. Formalisation of the examples is in progress.

2 Syntax and Typing

We define our type system with respect to a partially ordered semiring R for tracking how variables are used. A *partially ordered semiring* $(R, \leq, 0, +, 1, \cdot)$ is a poset (R, \leq) , commutative monoid $(R, 0, +)$, and monoid $(R, 1, \cdot)$, such that \cdot distributes over 0 and $+$, and $+$ and \cdot are monotonic with respect to \leq . We take $\rho, \pi \in R$.

PL’18, January 01–03, 2018, New York, NY, USA
2018.

Examples 1. The zero-one-many semiring $\{0, 1, \omega\}$ simulates linear typing in our system. 2. Monotonicity typing uses the semiring with carrier $\{0, \uparrow, \downarrow, \equiv\}$, where the multiplicative unit is \uparrow (covariance). The \downarrow represents contra-variance, and \equiv represents invariance. 3. Sensitivity analysis uses the semiring with carrier $\mathbb{R} \cup \{\infty\}$ with \min and $+$ as the addition and multiplication. 4. Information flow analysis uses the semiring with carrier $\mathcal{P}(L)$, where L is some set of labels.

The base language we consider is a bidirectional [Pierce and Turner 2000] simply typed λ calculus with the following types, where ι ranges over some set of base types:

$$S, T ::= S \multimap T \mid !_{\rho} S \mid 1 \mid S \otimes T \mid \top \mid S \& T \mid 0 \mid S \oplus T \mid \iota$$

Bidirectional typing reduces the type annotations required. Since our language is bidirectionally typed, we have two syntactic categories of terms: s ranges over checked terms, and e ranges over synthesising terms. We use t for both.

$$\begin{aligned} s ::= & \lambda x. s \mid \text{bang } s \mid *_{\otimes} \mid (s_0, s_1)_{\otimes} \mid *_{\&} \mid (s_0, s_1)_{\&} \mid \text{inj}_i(s) \mid \underline{e} \\ e ::= & x \mid e s \mid \text{bm}_T(e, \{x\}s) \mid \text{del}_T(e, s) \mid \text{pm}_T(e, \{x, y\}s) \\ & \mid \text{proj}_i(e) \mid \text{ex-falso}_T(e) \mid \text{case}_T(e, \{x\}s_0, \{y\}s_1) \mid s : S \end{aligned}$$

where curly braces and λ denote variable binding and we take $i \in \{0, 1\}$ wherever it appears.

Contexts Γ assign to each variable a type S and a usage $\rho \in R$: $\Gamma = x_1^{\rho_1} : S_1, \dots, x_n^{\rho_n} : S_n$. Contexts whose variables and types match form a left R -semimodule, by pointwise addition and scaling of the usage annotations. The partial order on R is extended pointwise to contexts. Typing judgements for checked and synthesising terms have the same contexts, but either record that a term is checked against a type $(\Gamma \vdash t \ni S)$ or synthesise a type $(\Gamma \vdash e \in T)$.

The typing rules consist of a variable rule, two rules for change of direction, and introduction and elimination rules for each type former. The following rules for variables and function and $!_{\rho}$ introduction and elimination illustrate how usage information is tracked:

$$\begin{array}{c} \frac{\Gamma \leq 0\Gamma_1, x^1 : S, 0\Gamma_2}{\Gamma \vdash x \in S} \qquad \frac{\Gamma, x^1 : S \vdash T \ni s[x]}{\Gamma \vdash S \multimap T \ni \lambda x. s[x]} \\ \frac{\Gamma_1 \vdash e \in S \multimap T \quad \Gamma_2 \vdash S \ni s \quad \Gamma \leq \Gamma_1 + \Gamma_2}{\Gamma \vdash e s \in T} \\ \frac{\Gamma_1 \vdash S \ni s \quad \Gamma \leq \rho \cdot \Gamma_1}{\Gamma \vdash !_{\rho} S \ni \text{bang } s} \qquad \frac{\Gamma_1 \vdash e \in !_{\rho} S \quad \Gamma_2, x^{\rho} : S \vdash T \ni s[x] \quad \Gamma \leq \Gamma_1 + \Gamma_2}{\Gamma \vdash \text{bm}_T(e, \{x\}s[x]) \in T} \end{array}$$

Sub-resourcing, weakening (adding 0-use variables to the context) and substitution are all admissible. In our Agda formalisation, we have constructed our type system in two levels: a non-usage tracked simply-typed λ -calculus, with a usage-tracking system layered above. This emphasises the use of coeffect annotations as an *analysis* of programs, they do not affect the underlying semantics, but comment on it. We introduce our semantic framework in the next section.

3 Semantics

Underlying Semantics We give a standard semantics of types and well typed terms into sets and functions. This semantics ignores the usage information. For types, we have:

$$\begin{aligned} \llbracket l \rrbracket &= A_l \\ \llbracket S \multimap T \rrbracket &= \llbracket S \rrbracket \rightarrow \llbracket T \rrbracket & \llbracket !_{\rho} S \rrbracket &= \llbracket S \rrbracket \\ \llbracket 1 \rrbracket &= \llbracket \top \rrbracket = \{*\} & \llbracket S \otimes T \rrbracket &= \llbracket S \& T \rrbracket = \llbracket S \rrbracket \times \llbracket T \rrbracket \\ \llbracket 0 \rrbracket &= \{\} & \llbracket S \oplus T \rrbracket &= \llbracket S \rrbracket \uplus \llbracket T \rrbracket \end{aligned}$$

Contexts are interpreted as left-nested products. Terms are assigned the usual semantics as functions $\llbracket t \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket S \rrbracket$.

Usage-tracking semantics To derive interesting properties from our type system, we refine the set-theoretic semantics by Kripke-indexed binary relations. This gives a fundamental lemma for our system, that when instantiated in different ways captures the examples in the introduction.

Our framework is parameterised by a category \mathcal{W} of *possible worlds* that track how resources are distributed by programs. To interpret resource separation, we assume that \mathcal{W} has *symmetric promonoidal structure*: profunctors $J : 1 \multimap \mathcal{W}$ and $P : \mathcal{W} \times \mathcal{W} \multimap \mathcal{W}$ such that $P \circ (J \times 1) \cong 1$, $P \circ (1 \times J) \cong 1$, $P \circ (1 \times P) \cong P \circ (P \times 1)$, and $P \cong P \circ (\pi_2 \times \pi_1)$, and the triangle, pentagon, and hexagon laws hold¹.

We now assign to each type T a functor $\llbracket T \rrbracket^R : \mathcal{W}^{op} \rightarrow \text{Rel} \llbracket T \rrbracket$ that captures a notion of \mathcal{W} -indexed “indistinguishability”. To interpret $!_{\rho} S$, we assume we are given a relation transformer $!_A : R^{op} \rightarrow \text{Rel}(A)^{\mathcal{W}^{op}} \rightarrow \text{Rel}(A)^{\mathcal{W}^{op}}$ that satisfies the axioms of a monoidal exponential comonad. The interesting cases are for functions, \otimes -products and the $!_{\rho}$ modality:

$$\begin{aligned} \llbracket S \multimap T \rrbracket^R w (f, f') &= \\ \forall x, y. P(y, w)x \Rightarrow \forall s, s'. \llbracket S \rrbracket^R y(s, s') \Rightarrow \llbracket T \rrbracket^R x(f s, f' s') \\ \llbracket S \otimes T \rrbracket^R w ((s, t), (s', t')) &= \\ \exists x, y. P(x, y)w \wedge \llbracket S \rrbracket^R x(s, s') \wedge \llbracket T \rrbracket^R y(t, t') \\ \llbracket !_{\rho} S \rrbracket^R w (s, s') &= !_{\rho} \llbracket S \rrbracket^R w (s, s') \end{aligned}$$

Contexts $x_1 \stackrel{\rho_1}{:} S_1, \dots, x_n \stackrel{\rho_n}{:} S_n$ are interpreted as if they were $\llbracket (\cdot \cdot \cdot (1 \otimes !_{\rho_1} S_1) \cdot \cdot \cdot \otimes !_{\rho_n} S_n) \rrbracket$. With these definitions, we can prove the following fundamental lemma for our Kripke-indexed relational semantics.

Theorem 3.1 (Fundamental Lemma).

$$\Gamma \vdash t : T \Longrightarrow \llbracket \Gamma \rrbracket^R w (y, y') \Longrightarrow \llbracket T \rrbracket^R w (\llbracket t \rrbracket y, \llbracket t \rrbracket y')$$

¹We don’t need the laws to hold to prove the fundamental lemma.

Example Instantiations The ingredients of our fundamental lemma are perhaps well known (relational interpretations, Kripke-indexing), but the value of our framework lies in the generality of being able to choose \mathcal{W} and its promonoidal structure, and the interpretation the $!_{\rho}$ modality as a relation transformer. Examples include:

Permutation Types With the $\{0, 1, \omega\}$ semiring, we take the category \mathcal{W} to consist of lists of some type of keys, and permutations between them. The relation transformer is defined as: $!_0 R l = \top$, where \top is the total relation, $!_1 R l = R l$ and $R_{\omega} R l = (l = []) \wedge R l$. With suitable types of keys and lists of keys, the fundamental lemma states that all programs are permutations. This result has already been formalised in a one-off type system at <https://github.com/bobatkey/sorting-types>.

Monotonicity Types With R the partially ordered semiring with carrier $\{0, \uparrow, \downarrow, \equiv\}$ ordered $\equiv \leq \uparrow, \downarrow$ and $\uparrow, \downarrow \leq 0$, we take \mathcal{W} to be the one-object, one-arrow category, and define the relation transformer $!$ to be:

$$!_0 R = \top \quad !_\uparrow R = R \quad !_\downarrow R = R^{op} \quad !_\equiv R = R \cap R^{op}$$

If we let our base type be natural numbers with the relational interpretation $R_{\text{nat}}(n, n') \Leftrightarrow n \leq n'$, then the fundamental lemma states that a program of type $x \stackrel{\uparrow}{:} \text{nat} \vdash t : \text{nat}$ is covariant (and similarly for contravariant and invariant).

Sensitivity Analysis With the $R = \mathbb{R} \cup \{\infty\}$ semiring, we let \mathcal{W} be R as well. The relation transformer is given by scaling. With a base type of real numbers with relational interpretation $R_{\text{real}} k(x, x') \Leftrightarrow |x - x'| \leq k$, then the fundamental lemma states that the usage annotations on the input variables tracks the extent to which the program is sensitive to changes in those variables.

Information Flow With the $R = \mathcal{P}(L)$ semiring, we again take $\mathcal{W} = R$, and let the relation transformer to be $!_l R l' = \{\top \text{ when } l \geq l'; R \text{ otherwise}\}$. Then the fundamental lemma yields the same non-interference properties as stated by Abadi et al. for the DCC [Abadi et al. 1999].

Acknowledgments

James Wood is supported by a EPSRC Studentship.

References

- M. Abadi, A. Banerjee, N. Heintze, and J. G. Riecke. 1999. A Core Calculus of Dependency. In *POPL '99*. 147–160.
- A. Brunel, M. Gaboardi, D. Mazza, and S. Zdancewic. 2014. A Core Quantitative Coeffect Calculus. In *ESOP 2014*. 351–370.
- Dan R. Ghica and Alex I. Smith. 2014. Bounded Linear Types in a Resource Semiring. In *ESOP 2014*. 331–350.
- Jean-Yves Girard. 1987. Linear Logic. *Theor. Comp. Sci.* 50 (1987), 1–101.
- Tomas Petricek, Dominic A. Orchard, and Alan Mycroft. 2014. Coeffects: a calculus of context-dependent computation. In *ICFP 2014*. 123–135.
- Benjamin C. Pierce and David N. Turner. 2000. Local type inference. *ACM TOPLAS* 22, 1 (2000), 1–44.
- J. Reed and B. C. Pierce. 2010. Distance Makes the Types Grow Stronger. In *ICFP 2010*, P. Hudak and S. Weirich (Eds.). 157–168.