

Cost-sensitive programming, verification, and semantics

Dissertation defense

July 29, 2024

Yue Niu

Carnegie Mellon University

Acknowledgements

Thank you for the six years.

Battle of users and designers

Programmers:

- Simplicity
- Familiarity ({...} and ;)
- *Ability* and *can do's*

Researchers:

- Global, mathematical properties
- *Inability* is a feature, not a bug

Balance between ergonomics and safety guarantees.

What is semantics?

Semantics is the process of interpreting PLs into known mathematical structures.

Program verification reduces to mathematical proof.

Most reliable method for establishing properties of not just computer programs, but every discipline amenable to the axiomatic method.

Thesis statement: the internal modal type theory of presheaves furnishes an ergonomic language for ***cost-sensitive*** programming, verification, and semantics.

What is cost?

Computational resource of programs:

- Concrete: time, memory, energy
- Abstract: number of critical operations in an algorithm

A *cost model* is the determination of the resource of interest.

What is a type?

Types can enforce strong safety properties.

Example

Let sort to be the type of functions $f : \text{list}(\mathbb{N}) \rightarrow \text{list}(\mathbb{N})$ such that $f(l)$ is a permutation of l and in ascending order.

Observe that sort contains a *unique* element sort .

Cost as an explicit structure

Profiling: make cost an explicit observable part of the code.

$$\mathbb{N} \times A \rightarrow \mathbb{N} \times B$$

Programming and verification

A new phase distinction

Cost structure vs. meaning as a **phase distinction**.

First introduced in **sterling-harper:2021**, type-theoretic phase distinctions have been used to solve many PL problems:

- Metatheory of type theories and programming languages [**sterling:2021:thesis**]
- Interactive theorem proving [**gratzer-sterling-angiuli-coquand-birkedal:2022**]
- Information flow security [**sterling-harper:2022**]

A new phase distinction

Cost structure vs. meaning as a **phase distinction**.

First introduced in **sterling-harper:2021**, type-theoretic phase distinctions have been used to solve many PL problems:

- Metatheory of type theories and programming languages [**sterling:2021:thesis**]
- Interactive theorem proving [**gratzer-sterling-angiuli-coquand-birkedal:2022**]
- Information flow security [**sterling-harper:2022**]

This thesis: use a phase distinction enable both **cost-sensitive** and **functional** reasoning in a single type theory.

The function-cost phase distinction

We have a cost-profiled function:

$$\mathbb{N} \times A \xrightarrow{f^\bullet} \mathbb{N} \times B$$

The function-cost phase distinction

Ensure that every cost-profiled function f^\bullet lies over a plain function f° :

$$\begin{array}{ccc} \mathbb{N} \times A & \xrightarrow{f^\bullet} & \mathbb{N} \times B \\ \pi_2 \downarrow & & \downarrow \pi_2 \\ A & \xrightarrow{f^\circ} & B \end{array}$$

Commutativity rules out programs disobeying the profiling semantics (*e.g.* badList).

Presheaf semantics of cost

Semantics encapsulated by the category $\widehat{\mathbb{I}}$ of presheaves over the interval $\{\circ \rightarrow \bullet\}$. Think Kripke/possible world semantics.

Presheaf semantics of cost

Semantics encapsulated by the category $\widehat{\mathbb{I}}$ of presheaves over the interval $\{\circ \rightarrow \bullet\}$. Think Kripke/possible world semantics.

- World at \circ = **functional phase**
- World at \bullet = **cost-sensitive phase**
- In cost-sensitive phase, $insertSort \neq mergeSort$.
- In functional phase, $insertSort = mergeSort$.

Presheaf restriction $\bullet \rightarrow \circ$ trivializes/redacts cost structure!

Algorithms vs. functions

$$\begin{array}{ccc}
 \mathbb{N} \times \text{list}(\mathbb{N}) & \xrightarrow{\text{isort}} & \mathbb{N} \times \text{list}(\mathbb{N}) \\
 \pi_2 \downarrow & & \downarrow \pi_2 \\
 \text{list}(\mathbb{N}) & \xrightarrow{\text{sort}} & \text{list}(\mathbb{N})
 \end{array}$$

Distinguished in the cost-sensitive phase, identified in the functional phase.

$$\begin{array}{ccc}
 \mathbb{N} \times \text{list}(\mathbb{N}) & \xrightarrow{\text{msort}} & \mathbb{N} \times \text{list}(\mathbb{N}) \\
 \pi_2 \downarrow & & \downarrow \pi_2 \\
 \text{list}(\mathbb{N}) & \xrightarrow{\text{sort}} & \text{list}(\mathbb{N})
 \end{array}$$

Modal types

In general, a function in $\widehat{\mathbb{I}}$ looks like a square:

$$\begin{array}{ccc} A^\bullet & \xrightarrow{f^\bullet} & B^\bullet \\ \pi_A \downarrow & & \downarrow \pi_B \\ A^\circ & \xrightarrow{f^\circ} & B^\circ \end{array}$$

Combining modalities

A cost-sensitive function with nontrivial functional behavior given by amalgamating function and cost modal types:

$$\begin{array}{ccc}
 \mathbb{N} \times A^\circ & \xrightarrow{f^\bullet} & \mathbb{N} \times B^\circ \\
 \pi_2 \downarrow & & \downarrow \pi_2 \\
 A^\circ & \xrightarrow{f^\circ} & B^\circ
 \end{array}$$

Square encodes possible dependencies: in $\mathbb{N} \times B^\circ$, \mathbb{N} can depend on both components of $\mathbb{N} \times A^\circ$, but B° can only depend on A° .

Internal characterization of modal types

Proposition

A type A is function-modal when $(\mathbb{N} \rightarrow A) \cong A$.

Proposition

A type A is cost-modal when $(\mathbb{N} \rightarrow A) \cong 1$.

In other words, a function-modal type “thinks” the functional phase holds and a cost-modal type “thinks” the functional phase is false.

Semantics

Cost semantics

calf supports *denotational/equational* cost analysis for *total functions*.

The rest of the thesis addresses the following equations:

1. How does this relate to *operational cost semantics*?
2. How to deal with *partial functions*?

Cost semantics

calf supports *denotational/equational* cost analysis for *total functions*.

The rest of the thesis addresses the following equations:

1. How does this relate to *operational cost semantics*?
2. How to deal with *partial functions*?

Extend **calf** with **general recursion** and prove an internal **computational adequacy** property.

Operational semantics

PL researchers often define the semantics of PLs in terms of an **operational semantics** given by a system of transitions between program states [plotkin_81_structural].

$$\frac{}{\text{bind}(\text{ret}(a), f) \mapsto f(a)}$$

$$\frac{}{\text{ret}(a) \text{ val}}$$

Operational semantics

PL researchers often define the semantics of PLs in terms of an **operational semantics** given by a system of transitions between program states [plotkin_81_structural].

$$\frac{}{\text{bind}(\text{ret}(a), f) \mapsto f(a)} \qquad \frac{}{\text{ret}(a) \text{ val}}$$

Meaning not by what a program *is*, but by what it *does*.

Operational cost semantics

Can extend operational semantics to account for cost.

$$\frac{}{\text{step}(e) \mapsto 1, e}$$

Operational cost semantics

Can extend operational semantics to account for cost.

$$\frac{}{\text{step}(e) \mapsto 1, e}$$

How relate operational and denotational cost semantics?

Internal program semantics

Study operational semantics as an object theory *within calf!*

Internal program semantics

Study operational semantics as an object theory *within* **calf**!

Fixing an internal PL \mathcal{L} in **calf**, a denotational model of \mathcal{L} is just a function $\llbracket - \rrbracket$.

Internal program semantics

Study operational semantics as an object theory *within* **calF**!

Fixing an internal PL \mathcal{L} in **calF**, a denotational model of \mathcal{L} is just a function $\llbracket - \rrbracket$.

The image of $\llbracket - \rrbracket$ carves out cost-profiled functions that are operationally definable.

Internal program semantics

Study operational semantics as an object theory *within* **calf**!

Fixing an internal PL \mathcal{L} in **calf**, a denotational model of \mathcal{L} is just a function $\llbracket - \rrbracket$.

The image of $\llbracket - \rrbracket$ carves out cost-profiled functions that are operationally definable.

We have equational cost bounds on such functions (*e.g.* *isort*).

Internal program semantics

Study operational semantics as an object theory *within* **calF**!

Fixing an internal PL \mathcal{L} in **calF**, a denotational model of \mathcal{L} is just a function $\llbracket - \rrbracket$.

The image of $\llbracket - \rrbracket$ carves out cost-profiled functions that are operationally definable.

We have equational cost bounds on such functions (*e.g.* *isort*).

Such equational cost bound in **calF** also holds wrt operational semantics of \mathcal{L} when *computational adequacy* holds.

Cost-sensitive computational adequacy

Suppose we have a program of ground type $e : F1$.

Definition

Classic computational adequacy [plotkin:1977]: $\llbracket e \rrbracket$ defined if and only if e terminates operationally.

Definition

Cost-sensitive Plotkin adequacy: $\llbracket e \rrbracket = c$ if and only if e terminates in c steps operationally.

Recursion in type theory

We aim to prove computational adequacy for **PCF**, a language with general recursive functions.

But type theory is a theory of *total functions* — divergence not in the image of $\llbracket - \rrbracket$!

Introduce partiality by using some form of **domains**.

Recursion in type theory

1. *Internal domain theory* [dejong:2023:thesis:1]

Recursion in type theory

1. *Internal domain theory* [**dejong:2023:thesis:1**]
 - (+) Mathematically canonical solution.

Recursion in type theory

1. *Internal domain theory* [**dejong:2023:thesis:1**]
 - (+) Mathematically canonical solution.
 - (-) Mismatch between domain-theoretic functions and type-theoretic functions.

Recursion in type theory

1. *Internal domain theory* [**dejong:2023:thesis:1**]
 - (+) Mathematically canonical solution.
 - (-) Mismatch between domain-theoretic functions and type-theoretic functions.
2. *Synthetic guarded domain theory (SGDT)* [**sgdt:2011; paviotti:2016**]

Recursion in type theory

1. *Internal domain theory* [**dejong:2023:thesis:1**]
 - (+) Mathematically canonical solution.
 - (-) Mismatch between domain-theoretic functions and type-theoretic functions.
2. *Synthetic guarded domain theory (SGDT)* [**sgdt:2011; paviotti:2016**]
 - (+) Can work naively with guarded domains — *every* (guarded) type-theoretic endofunction has a (guarded) fixed-point.

Recursion in type theory

1. *Internal domain theory* [**dejong:2023:thesis:1**]
 - (+) Mathematically canonical solution.
 - (-) Mismatch between domain-theoretic functions and type-theoretic functions.
2. *Synthetic guarded domain theory (SGDT)* [**sgdt:2011**; **paviotti:2016**]
 - (+) Can work naively with guarded domains — *every* (guarded) type-theoretic endofunction has a (guarded) fixed-point.
 - (-) Fixed-point equations do not hold exactly; correct notion of equality is *weak bisimulation*.

Recursion in type theory

1. *Internal domain theory* [**dejong:2023:thesis:1**]
 - (+) Mathematically canonical solution.
 - (-) Mismatch between domain-theoretic functions and type-theoretic functions.
2. *Synthetic guarded domain theory (SGDT)* [**sgdt:2011**; **paviotti:2016**]
 - (+) Can work naively with guarded domains — *every* (guarded) type-theoretic endofunction has a (guarded) fixed-point.
 - (-) Fixed-point equations do not hold exactly; correct notion of equality is *weak bisimulation*.
3. *Synthetic domain theory (SDT)* [**hyland:1991**; **reus-streicher:1999**]

Recursion in type theory

1. *Internal domain theory* [dejong:2023:thesis:1]
 - (+) Mathematically canonical solution.
 - (-) Mismatch between domain-theoretic functions and type-theoretic functions.
2. *Synthetic guarded domain theory (SGDT)* [sgdt:2011; paviotti:2016]
 - (+) Can work naively with guarded domains — *every* (guarded) type-theoretic endofunction has a (guarded) fixed-point.
 - (-) Fixed-point equations do not hold exactly; correct notion of equality is *weak bisimulation*.
3. *Synthetic domain theory (SDT)* [hyland:1991; reus-streicher:1999]
 - (+) Like internal domain theory, fixed-points hold on the nose.

Recursion in type theory

1. *Internal domain theory* [dejong:2023:thesis:1]
 - (+) Mathematically canonical solution.
 - (-) Mismatch between domain-theoretic functions and type-theoretic functions.
2. *Synthetic guarded domain theory (SGDT)* [sgdt:2011; paviotti:2016]
 - (+) Can work naively with guarded domains — *every* (guarded) type-theoretic endofunction has a (guarded) fixed-point.
 - (-) Fixed-point equations do not hold exactly; correct notion of equality is *weak bisimulation*.
3. *Synthetic domain theory (SDT)* [hyland:1991; reus-streicher:1999]
 - (+) Like internal domain theory, fixed-points hold on the nose.
 - (+) Like SGDT, every type-theoretic function is domain-theoretic (*i.e.* continuous), justifying one to work naively in type theory

Synthetic domain theory

Identify a class of *predomains* amongst ordinary types that behave like *e.g.* ω -cpo from classic domain theory.

Concretely, models of SDT are given by toposes equipped with a distinguished object Σ called the *dominance*.

Everything derived by imposing structures on Σ (*cf.* the line object R in synthetic differential geometry).

Axioms of synthetic domain theory

The dominance serves as the classifier of the *termination support* of partial maps.

Definition

A *dominance* is a collection of propositions closed under \perp , \top and dependent sums.

A Σ -*subset* is one whose characteristic map is valued in Σ (think *computational* predicate).

Lifting structure


The dominance Σ induces a *lifting structure* $L(A) = \sum_{\phi:\Sigma} \cdot \phi \rightarrow A$:
partial maps $A \xrightarrow{\Sigma} D \rightarrow B$ as total maps $A \rightarrow L(B)$.

Lifting induces an incidence relation $\omega \hookrightarrow \bar{\omega}$ including the
initial lift algebra ω into the final lift coalgebra $\bar{\omega}$.

Think of $\omega \hookrightarrow \bar{\omega}$ as a *figure shape* that we use to state the
completeness properties of predomains.


Complete types

A type is *complete* when it has the unique extension property along $\omega \hookrightarrow \bar{\omega}$:



Complete types


A type is *complete* when it has the unique extension property along $\omega \hookrightarrow \bar{\omega}$:



Unique extension properties as closure under *generalized limit processes* [**streicher:1997:repletion**].

Complete types

A type is *complete* when it has the unique extension property along $\omega \hookrightarrow \overline{\omega}$:



Unique extension properties as closure under *generalized limit processes* [streicher:1997:repletion].

For example, ω -cpo's defined by the unique extension property relative to $\{0 \sqsubseteq 1 \sqsubseteq \dots\} \hookrightarrow \{0 \sqsubseteq 1 \sqsubseteq \dots \sqsubseteq \infty\}$.

Candidates for predomains

To close complete types under lifting, one may consider well-complete types [**longley-simpson:1997**].

Definition

A type A is *well-complete* when $\mathcal{L}A$ is complete.

Well-complete types are closed under lifting and form a reflective exponential ideal — all the structures necessary for everyday denotational semantics.

An ergonomic interface for SDT

Well-complete types are adequate to obtain all results in this thesis.

However, notions like unique extension properties are not familiar to many *users* of denotational semantics.

In the thesis, work relative to an *order-theoretic interface* of SDT that closes the gap between classic and synthetic domains.

An order-theoretic formulation of SDT

Observe that we did not use the *information/approximation* order to define predomains.

An order-theoretic formulation of SDT

Observe that we did not use the *information/approximation* order to define predomains.

Can equip types with a *derived synthetic order* \sqsubseteq .

An order-theoretic formulation of SDT

Observe that we did not use the *information/approximation* order to define predomains.

Can equip types with a *derived synthetic order* \sqsubseteq .

The order \sqsubseteq its closure properties (ω -joins \bigvee) furnish a convenient and familiar *interface* to users of SDT.

An order-theoretic formulation of SDT

Observe that we did not use the *information/approximation* order to define predomains.

Can equip types with a *derived synthetic order* \sqsubseteq .

The order \sqsubseteq its closure properties (ω -joins \bigvee) furnish a convenient and familiar *interface* to users of SDT.

Think of \sqsubseteq and \bigvee as the “ $\{ \dots \}$ ” and “ $;$ ” of SDT.

Predomains and domains

We have a topos \mathcal{E} equipped with a class of predomains Predom :

- Predom is Cartesian closed and (internally) (co)complete
- Domains (predomains with least elements) support fixed-points of endomaps.

Predomains and domains

We have a topos \mathcal{E} equipped with a class of predomains Predom :

- Predom is Cartesian closed and (internally) (co)complete
- Domains (predomains with least elements) support fixed-points of endomaps.

Predomains $\sim \omega$ -cpos and domains \sim pointed ω -cpos.

Complete relative to ω -chains instead of \mathbb{N} -chains.

Sufficient for the ordinary denotational semantics of **PCF**.

SDT model of the phase distinction

To incorporate cost structure as a phase distinction, assume a distinguished proposition $\phi : \Sigma$ — ensures the cost modality $\llbracket \vee \rrbracket$ — produces a predomain.

Definition

An *SDT model of the phase distinction* is a topos \mathcal{E} equipped with a complete dominance Σ and a Σ -proposition $\phi : \Sigma$.

In this SDT we define a model of *cost-sensitive* version of **PCF**.

Cost-sensitive PCF

We work with a version of **PCF** equipped with the cost interface **FA**:

interface Cost

$F : \text{type} \rightarrow \text{type}$

$\text{ret} : A \rightarrow \text{FA}$

$\text{bind} : \text{FA} \rightarrow (A \rightarrow \text{FB}) \rightarrow \text{FB}$

$\text{step} : \text{FA} \rightarrow \text{FA}$

$\text{fix} : (\text{FA} \rightarrow \text{FA}) \rightarrow \text{FA}$

Note that we now support *partial* cost-sensitive functions.

Denotational cost semantics

The Cartesian closed structure of **PCF** is interpreted in a canonical way into predomains.

For cost structure, we lift the interpretation of FA in **calf**:

$$\llbracket FA \rrbracket = L(\mathbb{C} \times \llbracket A \rrbracket)$$

Assume \mathbb{C} *cost-modal*, i.e. $(\mathbb{1} \rightarrow \mathbb{C}) \cong 1$. Recover ordinary semantics in the functional phase: $\mathbb{1} \rightarrow (\llbracket FA \rrbracket = L\llbracket A \rrbracket)$.

Operational cost semantics

Recall that we may define the operational cost semantics of the cost effect:

$$\frac{}{\text{step}(e) \mapsto \mathbf{1}, e}$$

We can extend this to a big-step semantics $e \Downarrow^c v$.

Operational cost semantics

Writing $\text{Tm}(A)$ for terms of **PCF**, define a recursive function $\text{eval} : \text{Tm}(\text{FA}) \rightarrow \text{Tm}(\text{FA}) \rightarrow L(\mathbb{C})$:

```

fun eval( $e, v$ ) =
  case out( $e$ ) of
    inl  $\cdot * \hookrightarrow (e = v, \lambda u : (e = v), 0)$ 
    inr  $\cdot (c, e') \hookrightarrow (c \boxplus \text{eval}(e', v))$ 
  
```

- We have $\text{out}(e, v) = \text{inr}(c, e')$ if and only if $e \mapsto c, e'$.
- We write $c \boxplus e$ to increment the cost of a computation e by c units.

Operational cost semantics

Define $e \Downarrow^c v$ if and only if $\text{eval}(e, v) = c$.

Observe that $e \Downarrow^c v$ is a Σ -proposition — we will use this in the proof of computational adequacy.

Defining $e \Downarrow^c v$ as the reflexive-transitive closure of $e \mapsto c, e'$ need *not* yield a Σ -proposition without additional assumptions.

Logical relation for computational adequacy

Following **plotkin:1977**, define a family of relations

$\triangleleft_A \subseteq \llbracket A \rrbracket \times \text{Tm}(A)$ between the semantics and syntax of **PCF**.

Interpretation of \triangleleft_{FA} :

$$\begin{aligned}
 e \triangleleft_{FA} e' = & \\
 & \forall [f : \llbracket A \rightarrow F1 \rrbracket, f' : \text{Tm}(A \rightarrow F1)] \\
 & (\forall [a \triangleleft_A a'] f(a) \ll e'(a')) \rightarrow \\
 & \text{bind}(e, f) \ll \text{bind}(e', f')
 \end{aligned}$$

For $e : \llbracket F1 \rrbracket$ and $e' : \text{Tm}(F1)$, define $(e \ll e') \stackrel{\text{def}}{=} (e \sqsubseteq \text{eval}(e', *)).$

Computational adequacy

Observe that $(- \ll e') \subseteq L(\mathbb{C})$ is an *admissible* subset, allowing us to prove the fundamental lemma by fixed-point induction.

Theorem

Given $\Gamma \vdash e : A$, we have $\Gamma \vdash \llbracket e \rrbracket \triangleleft_A e$.

Cost-sensitive computational adequacy follows directly from the fundamental lemma:

Theorem

Given $e : F1$, we have that $\llbracket e \rrbracket = c$ if and only if $e \Downarrow^c *$.

Model construction

Model of cost-sensitive SDT

Define a *relative* model of SDT following **sterling-harper:2022**, *i.e.* a model fibred over the presheaf model of **calf**.

Model of cost-sensitive SDT

Define a *relative* model of SDT following **sterling-harper:2022**, *i.e.* a model fibred over the presheaf model of **calf**.

Isolate a (small) category \mathcal{C} of *internal dcpos* in $\widehat{\mathbb{I}}$.

- Presheaves on \mathcal{C} is *almost* a model of SDT.
- Restrict to sheaves on \mathcal{C} for the extensive coverage: preserves \emptyset and $+$.

Model of cost-sensitive SDT

Define a *relative* model of SDT following **sterling-harper:2022**, *i.e.* a model fibred over the presheaf model of **calF**.

Isolate a (small) category \mathcal{C} of *internal dcpos* in $\widehat{\mathbb{I}}$.

- Presheaves on \mathcal{C} is *almost* a model of SDT.
- Restrict to sheaves on \mathcal{C} for the extensive coverage: preserves \emptyset and $+$.

Theorem

The category of (internal) sheaves on \mathcal{C} furnishes a model of SDT such that the functional phase proposition $\heartsuit : \mathcal{C}$ is preserved by the Yoneda embedding.

Conclusions

The big picture (programming and verification)

The PL community has been converging towards type theory as the premier (meta)language for doing formalized mathematics and programming.

The big picture (programming and verification)

The PL community has been converging towards type theory as the premier (meta)language for doing formalized mathematics and programming.

From a computer science perspective, cost analysis and cost-sensitive verification has been a glaring gap.

The big picture (programming and verification)

The PL community has been converging towards type theory as the premier (meta)language for doing formalized mathematics and programming.

From a computer science perspective, cost analysis and cost-sensitive verification has been a glaring gap.

This thesis addressess the question of cost in a framework that is both ergonomic in practice (battle tested in Agda) and theoretically sound.

The big picture (semantics)

Relate denotational and operational cost semantics by viewing PLs as *internal* theories.

The big picture (semantics)

Relate denotational and operational cost semantics by viewing PLs as *internal* theories.

The results suggest an *internal* extraction procedure for producing real executable programs that meet their proven cost bounds.

Thanks for listening!

References I