

# Flambda2 Validator

Irene Yoon  
INRIA  
France  
euisun.yoon@inria.fr

Chris Casinghino  
Jane Street  
USA  
ccasinghino@janestreet.com

## Abstract

This talk will describe a validation tool for Flambda2, an optimizing middle-end for OCaml. Its optimizations are centered around inlining (replacing a function call with the function body) and applying simplifications that become possible after inlining. Although such a transformation sounds innocuous, it is one of the most important—and tricky to implement—optimizations in the compiler pipeline. We increase confidence in Flambda 2’s optimizations by providing a relatively small and declarative definition of reduction to reduce optimized and unoptimized versions of the same program to syntactically equivalent terms. The tool is functional and can validate Flambda2’s optimizations for a substantial fraction of the OCaml standard library, our main test suite.

### ACM Reference Format:

Irene Yoon and Chris Casinghino. 2024. Flambda2 Validator. In *Proceedings of The 2024 OCaml Users and Developers Workshop (OCaml Workshop '24)*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 Introduction

This talk will describe a validation tool for Flambda2 [1], an optimizing middle-end for OCaml. Its optimizations are centered around inlining (replacing a function call with the function body) and applying simplifications that become possible after inlining. Although such a transformation sounds innocuous, it is one of the most important—and tricky to implement—optimizations in the compiler pipeline. Both inlining and simplification are easy to get wrong in an effectful language like OCaml, motivating the need for a validation tool.

The core idea of our validation tool is to increase confidence in Flambda 2’s optimizations by providing a relatively small and declarative definition of reduction that is sufficient to reduce optimized and unoptimized versions of the same program to syntactically equivalent terms. We define a lambda-calculus-like simplified core language for the Flambda2 IR, called *Flambda2 Core* ( $F\lambda_2^C$ ). To validate Flambda2’s optimization engine, we translate the original and optimized programs to core, normalize both, and compare them for  $\alpha$ -equivalence. The approach is shown in Figure 1, where the double arrows illustrate components we have implemented, as described further in Section 2.

*OCaml Workshop '24, September, 2024, Milan, Italy*  
2024. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00  
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

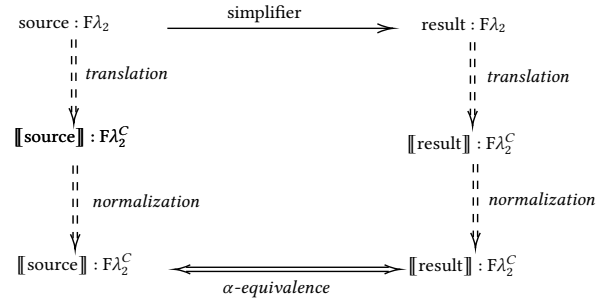


Figure 1. Semantic equivalence for Flambda2

This talk will describe work in progress, and we look forward to feedback and suggestions from the OCaml Workshop community. The tool is functional and can validate Flambda2’s optimizations for a substantial fraction of the OCaml standard library, our main test suite—see Section 3 for details. The goal of providing a declarative reduction relation that captures these optimizations is so far a partial success, we discuss some particular tricky transformations in Section 4.

## 2 Methodology

We highlight the major differences between Flambda2 and  $F\lambda_2^C$  and how we allow full substitution of terms in  $F\lambda_2^C$ .

### 2.1 Full expression subterms

$F\lambda_2^C$  allows for full subexpressions in argument position, so that there can be more aggressive inlining and substitution in the normalization process. Since, for example, substituting an arbitrary let-bound variable for its body is not always sound in the presence of effects, the normalization for  $F\lambda_2^C$  substitutes only when it preserves the behavior of the program, which will be explained further in the talk.

### 2.2 Unifying closure binding of Flambda2

In Flambda2, there are two ways that a binding for a set of closures can be declared. In Flambda2 Core, there is only one syntactic form for this. In  $F\lambda_2^C$ , closure bindings are written  $\text{let } \Phi = \text{closure } f\_id @f; \text{closure } g\_id @g$  where  $\Phi$  is a variable corresponding to the whole set of closures. Each of the function slots can be referred to by using a "slot expression", which is a tuple referring to the set of closures and the function slot. To refer to the function slot that stores  $f\_id$ , we write  $\text{slot } (\Phi, @f)$ . Having a uniform representation

of closures simplifies dealing with primitives that interact with them (e.g., projecting from the closure environment).

### 2.3 Unifying recursive lambda expressions

In Flambda2, there are lambda-like expressions representing code blocks. Code blocks may be mutually recursive; however, one can also have recursive function definitions through local recursive continuations. In  $F\lambda_2^C$ , we unify the representation of recursive lambda expressions so that the normalization scheme can apply uniform treatment for dealing with recursive sets of closures. In  $F\lambda_2^C$ , local continuations definitions cannot be recursive; they are lifted as a recursive form of lambda. Note that the normalization does not prevent divergence, as can be observed in the next section.

## 3 Results

Our primary testsuite for the validator has been the version of the OCaml standard library in the Flambda2 version of the OCaml compiler. We ran the compiler with validation turned on for each file, subject to the limitations described in Section 4. We ran the experiment on a modern AMD Epyc 7002 processor with a time limit of 20 minutes per file. On this testsuite of 67 files, the validator successfully validates 45 files (67%), fails to validate 14 files (21%), and times out on 8 files (12%).

Our experience to date suggests that the remaining validation failures are likely caused by missing reduction rules. The most common cause has been missing rules for particular OCaml primitives, but in some cases we have needed to add more complex rules to account for Flambda2’s optimizations, and we discuss some interesting examples in Section 4.

The performance of the validator is bimodal. Of the 59 files on which it does not time out, 43 complete in under a second, and only 5 require more than a minute. The speed is not strongly correlated with program size—some large files validate quickly and some smaller files time out. Based on this performance profile, we suspect certain constructs or patterns may cause the validator to loop or term size to explode—further investigation is future work.

At present the validator comprises 7215 lines of code. This includes the definition of  $F\lambda_2^C$ , translation from Flambda2 to  $F\lambda_2^C$ , normalization of  $F\lambda_2^C$ , and equivalence checking of  $F\lambda_2^C$  terms. In the version of the compiler we are using for testing, Flambda2 itself comprises 95457 lines of code (for a fair comparison, we included the translation into Flambda2 from the previous IR, but not the translation out of Flambda2 to the next IR, and did not include tests in either case).

At less than 10% the size of Flambda2, we believe it is reasonable to consider the validator a more declarative specification of correctness of the middle-end, despite some of the complications discussed in the next section.

## 4 Limitations

The goal of this project is to increase confidence in Flambda2’s optimizations by providing a relatively small and declarative definition of reduction that is sufficient to reduce optimized and unoptimized versions of the same program to syntactically equivalent terms. While the results from the previous section show partial success, we have also encountered some optimizations performed by Flambda2 that are challenging to validate with this technique. This section provides examples of optimizations we are not yet validating successfully, or where validation required a more complicated notion of reduction.

### 4.1 Cross-module inlining

Flambda2 can perform inlining not just within a compilation unit, but also between compilation units when the compiled version of a library is available. This cross-module inlining does not create any fundamental theoretical obstacle for our approach—we could in principle fully inline every function whose definition is available, even if from another module. However, this would be problematic for the performance of the validator, as term sizes would grow massively.

A more nuanced approach would be to *replicate* Flambda2’s cross-module inlining decisions, and only inline a function body from another module if this was also done during the actual optimization pass. Conveniently, Flambda2 maintains an “inlining tree” data structure that records these inlining decisions. However, we have not yet implemented this heuristic, and our present results are collected with cross-module inlining disabled.

### 4.2 Function argument introduction/elimination

We have encountered two transformations applied by Flambda2 that modify the number of arguments to functions. If these functions are not eliminated by reduction, this defeats our simple syntactic equivalence check.

The first such example concerns invariant arguments to recursive continuations. Translation from OCaml’s Lambda IR into Flambda2 creates recursive continuations both to model recursive functions themselves and as the compilation of loops. In the case that such a recursive continuation has an *invariant* argument—an argument that does not change from recursive call to recursive call—and the continuation is applied only once, Flambda2 will eliminate the argument from the continuation and simply substitute in the value from the application site.

To handle this, we have included a “reduction” rule in our system that applies a similar transformation to eliminate arguments. This rule must do a deep analysis of the body of the continuation to detect invariant arguments. As such, this rule reduces the extent to which our reduction relation serves as a simple declarative specification of correctness for Flambda2.

221 The second example concerns *common subexpression elim-*  
 222 *ination* (CSE), where Flambda2 attempts to eliminate dupli-  
 223 cated computations of pure expressions. If Flambda2 achieved  
 224 this goal by creating a simple let binding of the common  
 225 expression, our straightforward reduction rule for let bind-  
 226 ings would validate this optimization. However, as Flambda2  
 227 uses continuations for control flow, it instead propagates the  
 228 subexpression by adding an extra argument to the continu-  
 229 ation called by the code where the it is initially calculated.  
 230 This continuation therefore has an extra argument in the  
 231 optimized version.

232 It is possible the validator could accomodate this optimiza-  
 233 tion by applying a similar rule as we used to handle the  
 234 invariant arguments, where we observe this continuation  
 235 is always called with the same value and simply eliminate  
 236  
 237  
 238  
 239  
 240  
 241  
 242  
 243  
 244  
 245  
 246  
 247  
 248  
 249  
 250  
 251  
 252  
 253  
 254  
 255  
 256  
 257  
 258  
 259  
 260  
 261  
 262  
 263  
 264  
 265  
 266  
 267  
 268  
 269  
 270  
 271  
 272  
 273  
 274  
 275

the argument and substitute in the value. However, we have  
 not yet implemented such an optimization, and have instead  
 temporarily disabled the CSE optimization when running  
 validation.

## 5 Acknowledgements

This work was done in collaboration with Xavier Clerc, Mark  
 Shinwell, and Leo White at Jane Street. The authors also  
 received guidance from the Flambda 2 team at OCamlPro,  
 particularly Guillaume Bury, Pierre Chambart, Nathana elle  
 Courant, and Vincent Lavi ron.

## References

- [1] Flambda2 Team. 2023. Efficient OCaml compilation with Flambda 2.  
 OCaml Workshop.