

High-level error messages for modules through diffing

Gabriel RADANNE

Inria

gabriel.radanne@inria.fr

Florian ANGELETTI

Inria

florian.angeletti@inria.fr

Modules are one of the most complex features of ML languages. This complexity is reflected in error messages. Whenever two module types are mismatched, it is hard to identify and report the exact source of the error. Consequently, typecheckers often resort to printing the whole module types, and hope that the human user will navigate the sea of definitions.

We propose to improve module error messages by coupling classical typechecking with a diffing algorithm. The typechecker deals with the gritty details of the ML module system whereas the diffing algorithm summarizes the error through a higher level view. The large literature on diffing algorithms allows us to pick and choose the exact algorithm adapted for signatures, functors applications, submodules, etc.

1 Motivation

Typical uses of module are quite simple. Few people use directly higher-order functor with anonymous module arguments and labyrinth of module type definitions. Quite often, the names of functor arguments and parameters even match:

```
module type XT = sig type x end
module type YT = sig type y end
module F(X: XT)(Y:YT) = struct ... end
module X = struct type x = A end
module Y = struct type y = B end
module Result = F(X)(Y)
```

Furthermore, module type errors often happens during code refactoring, where a handful of changes are common:

- Adding a new item
- Removing an old item
- Changing the type of an item

For instance, if we refactor the definition of the functor F above and remove the first argument but forget to update the definition of the Result module, the previously working code now yields an error

```
module F(Y:sig type y end) = struct ... end
module Result = F(X)(Y)
```

Error: Signature mismatch:

Modules do not match: sig type x = X.x = A end is not included in YT
The type y is required but not provided

However, we are not interested in the possibly lengthy mismatch between the module types XT and YT. We are more interested in the fact that there is an extra functor argument. This combination of hard to decipher module type errors and a small class of common high-level errors makes a good argument for trying to give user a higher-level view in module-level type error messages. We propose to use the tree-like shape of module types to leverage *diffing algorithm* on trees and lists.

Diffing algorithms have a long history in many domains. For linear texts, the Longest Common Subsequence problem [3] is used for code versioning and wikis. A more general version using edit distances [4] is commonly used for spellchecking and bioinformatics. More recently, diffing for trees [5] has found a large application in Web programming for UIs [6]. This diversity gives us a very fertile ground to pick diffing algorithms adapted to the exact mismatch at hand, from edit distances for functor applications to tree diffs for signatures with submodules.

We now demonstrate these ideas on functor applications, which we implemented in the OCaml compiler (<https://github.com/ocaml/ocaml/pull/9331>). In the talk, we will present the larger context, its application to signatures, and some technical details.

2 Optimising edit-distance for functors

As we pointed out before, functors applications are often the source of complex type errors. This is partially due to the contrast between the higher-level view of a functor multi-application

```
module R = F(X_1)...(X_n)
```

and the left-to-right biased view of the typechecker. If type-checking this functor application reports an error at position k , the standard way to report an error would be simply to report the mismatch between expected type of the k th-argument and the k th-parameter. But by doing so, we are losing the context of the functor multi-application.

By considering the whole list of arguments, we can find the smallest patched argument lists that make the multi-application typechecks. A patched argument is here either:

1. An accepted argument from the original argument list
2. A deleted argument that we throw away from the original argument list
3. An additional argument pulled from the expected parameter list
4. A mismatched argument from the original argument list

It is then natural to assign a positive weight to those patched arguments to represent how far away they are from the arguments that the user has written. If only accepted arguments have a zero weight, then we are guaranteed that there is a patched argument list with minimal weight. For instance, in our example

```
module F(Y:sig type y end) = struct ... end
module Result = F(X)(Y)
```

The possible patched arguments would be

1. [Delete(X); Accept(Y)]
2. [Change(X);Delete(Y)]
3. [Add(YT);Delete(X);Delete(Y)]

The smallest change is clearly the first one. Consequently, rather than reporting the mismatch between XT and YT, we can simply report X as an extra-argument:

Error: The functor application is ill-typed.

These arguments:

X Y

do not match these parameters:

functor (Y : YT) -> ...

1. The following extra argument is provided X : sig type x = X.x = A end
2. Module Y matches the expected module type YT

It is important to inform users that we consider the functor multi-application as a whole and how we think the code should be corrected. In our experience, these error messages scale to fairly complex uses (as can be seen in the patch linked previously) with high-order, dependent or even variadic functors, using the full panel of advanced ML features. This work also applies to inclusion tests between two functor signatures.

Our unoptimized implementation uses a variant of the Wagner–Fischer algorithm [1] with a complexity of $O(\max(|arguments|, |parameters|)^2)$ module comparisons in the worse case. This is sufficient in practice, since functors applications are usually short. More efficient algorithms [2] could also be used.

So far, we only considered the Levenshtein distance, which includes insertions, deletions, and substitutions (i.e., mismatches). We could easily extend our implementation to transpositions but it is unclear how to present such a mix of operations to the user in a clear manner. One case that could be worth investigating is when a simple permutation, without any other operation, is sufficient to make the application typechecks. We could then simply inform the user to reorder the arguments in the right order.

3 Conclusion

We propose a methodology to improve the user experience of error message for modules in ML languages by providing high-level reports. This is done by combining the classical inclusion check algorithm with diffing algorithms, leading to much more focused error messages. In the talk, we will present the application to both functors arguments and signature, explore technical details such as the handling of variadic functors, and compare with existing error messages.

References

- [1] Robert A. Wagner and Michael J. Fischer. “The String-to-String Correction Problem”. In: *J. ACM* 21.1 (1974), pp. 168–173. DOI: [10.1145/321796.321811](https://doi.org/10.1145/321796.321811). URL: <https://doi.org/10.1145/321796.321811>.
- [2] Gad M. Landau, Eugene W. Myers, and Jeanette P. Schmidt. “Incremental String Comparison”. In: *SIAM J. Comput.* 27.2 (1998), pp. 557–582. DOI: [10.1137/S0097539794264810](https://doi.org/10.1137/S0097539794264810). URL: <https://doi.org/10.1137/S0097539794264810>.

- [3] Lasse Bergroth, Harri Hakonen, and Timo Raita. “A Survey of Longest Common Subsequence Algorithms”. In: *Seventh International Symposium on String Processing and Information Retrieval, SPIRE 2000, A Coruña, Spain, September 27-29, 2000*. Ed. by Pablo de la Fuente. IEEE Computer Society, 2000, pp. 39–48. DOI: [10.1109/SPIRE.2000.878178](https://doi.org/10.1109/SPIRE.2000.878178). URL: <https://doi.org/10.1109/SPIRE.2000.878178>.
- [4] Gonzalo Navarro. “A guided tour to approximate string matching”. In: *ACM Comput. Surv.* 33.1 (2001), pp. 31–88. DOI: [10.1145/375360.375365](https://doi.org/10.1145/375360.375365). URL: <https://doi.org/10.1145/375360.375365>.
- [5] Philip Bille. “A survey on tree edit distance and related problems”. In: *Theor. Comput. Sci.* 337.1-3 (2005), pp. 217–239. DOI: [10.1016/j.tcs.2004.12.030](https://doi.org/10.1016/j.tcs.2004.12.030). URL: <https://doi.org/10.1016/j.tcs.2004.12.030>.
- [6] The ReactJS team. *Reconciliation in ReactJS*. 2020. URL: <https://reactjs.org/docs/reconciliation.html> (visited on 2020).