

TypeScript is Surprisingly OK for Compilers

Aug 17, 2023

There are two main historical trends when choosing an implementation language for something compiler-shaped.

For more language-centric tasks, like a formal specification, or a toy hobby language, OCaml makes most sense. See, for example, [plzoo](#) or [WebAssembly reference interpreter](#).

For something implementation-centric and production ready, C++ is often chosen: LLVM, clang, v8, HotSpot are all C++.

These days, Rust is a great new addition to the landscape. It is influenced most directly by ML and C++, combines their strengths, and even brings something new of its own to the table, like seamless, safe multithreading. Still, Rust leans heavily towards production readiness side of the spectrum. While some aspects of it, like a “just works” build system, help with prototyping as well, there’s still extra complexity tax due to the necessity to model physical layout of data. The usual advice, when you start building a compiler in Rust, is to avoid pointers and use indexes. Indexes are great! In large codebase, they allow greater decoupling (side tables can stay local to relevant modules), improved performance (an index is u32 and nudges you towards struct-of-arrays layouts), and more flexible computation strategies (indexes are easier to serialize or plug into incremental compilation framework). But they do make programming-in-the-small significantly more annoying, which is a deal-breaker for hobbyist tinkering.

But OCaml is crufty! Is there something better? Today, I realized that TypeScript might actually be OK? It is not really surprising, given how the language works, but it never occurred to me to think about TypeScript as an ML equivalent before.

So, let’s write a tiny-tiny typechecker in TS!

Of course, we start with [deno](#). See [A Love Letter to Deno](#) for more details, but the TL;DR is that deno provides out-of-the-box experience for TypeScript. This is a pain point for OCaml, and something that Rust does better than either OCaml or C++. But deno does this better than Rust! It’s just a single binary, it comes with linting and formatting, there’s no compilation step, and there are built-in task runner and watch mode. A dream setup for quick PLT hacks!

And then there's TypeScript itself, with its sufficiently flexible, yet light-ceremony type system.

Let's start with defining an AST. As we are hacking, we won't bother with making it an IDE-friendly concrete syntax tree, or incremental-friendly "only store relative offsets" tree, and will just tag AST nodes with locations in file:

```
1 | export interface Location {
2 |   file: string;
3 |   line: number;
4 |   column: number;
5 | }
```

Even here, we already see high-level nature of TypeScript — `string` is just a `string`, there's no thinking about `usize` vs `u32` as numbers are just numbers.

Usually, an expression is defined as a sum-type. As we want to tag each expression with a location, that representation would be slightly inconvenient for us, so we split things up a bit:

```
1 | export interface Expr {
2 |   location: Location;
3 |   kind: ExprKind;
4 | }
5 |
6 | export type ExprKind = ExprBool | ExprInt | ... ;
```

One more thing — as we are going for something quick, we'll be storing inferred types directly in the AST nodes. Still, we want to keep raw and type-checked AST separate, so what we are going to do here is to parametrize the `Expr` over associated data it stores. A freshly parsed expression would use `void` as data, and the type checker will set it to `Type`. Here's what we get:

```
1 | export interface Expr<T> {
2 |   location: Location;
3 |   data: T;
4 |   kind: ExprKind<T>;
5 | }
6 |
7 | export type ExprKind<T> =
8 |   | ExprBool<T>
9 |   | ExprInt<T>
10 |  | ExprBinary<T>
11 |  | ExprControl<T>;
```

A definition of `ExprBinary` could look like this:

```

1 export interface ExprBinary<T> {
2   op: BinaryOp;
3   lhs: Expr<T>;
4   rhs: Expr<T>;
5 }
6
7 export enum BinaryOp {
8   Add, Sub, Mul, Div,
9   Eq, Neq,
10  Lt, Gt, Le, Ge,
11 }

```

Note how I don't introduce separate types for, e.g, `AddExpr` and `SubExpr` — all binary expressions have the same shape, so one type is enough!

But we need a tiny adjustment here. Our `Expr` kind is defined as a union type. To match a value of a union type a bit of runtime type information is needed. However, it's one of the core properties of TypeScript that it doesn't add any runtime behaviors. So, if we want to match on expression kinds (and we for sure want!), we need to give a helping hand to the compiler and include a bit of RTTI manually. That would be the tag field:

```

1 export interface ExprBinary<T> {
2   tag: "binary";
3   op: BinaryOp;
4   lhs: Expr<T>;
5   rhs: Expr<T>;
6 }

```

`tag: "binary"` means that the only possible runtime value for `tag` is the string `"binary"`.

Similarly to various binary expressions, boolean literal and int literal expressions have *almost* identical shape. Almost, because the payload (`boolean` or `number`) is different. TypeScript allows us to neatly abstract this over:

```

1 export type ExprBool<T> = ExprLiteral<T, boolean, "bool">;
2 export type ExprInt<T> = ExprLiteral<T, number, "int">;
3
4 export interface ExprLiteral<T, V, Tag> {
5   tag: Tag;
6   value: V;
7 }

```

Finally, for control-flow expressions we only add `if` for now:

```

1 export type ExprControl<T> = ExprIf<T>;
2
3 export interface ExprIf<T> {
4   tag: "if";

```

```

5 |   cond: Expr<T>;
6 |   then_branch: Expr<T>;
7 |   else_branch: Expr<T>;
8 | }

```

This concludes the definition of the ast! Let's move on to the type inference! Start with types:

```

1 | type Type = TypeBool | TypeInt;
2 |
3 | interface TypeBool {
4 |   tag: "Bool";
5 | }
6 | const TypeBool: TypeBool = { tag: "Bool" };
7 |
8 | interface TypeInt {
9 |   tag: "Int";
10 | }
11 | const TypeInt: TypeInt = { tag: "Int" };

```

Our types are really simple, we could have gone with

```
type Type = "Int" | "Bool",
```

but lets do this a bit more enterprisy! We define separate types for integer and boolean types. As these types are singletons, we also provide canonical definitions. And here is another TypeScript-ism. Because TypeScript fully erases types, everything related to types lives in a separate namespace. So you can have a type and a value sharing the same name. Which is exactly what we use to define the singletons!

Finally, we can take advantage of our associated-data parametrized expression and writhe the signature of

```
1 | function infer_types(expr: ast.Expr<void>): ast.Expr<Type>
```

As it says on the tin, inter_types fills in Type information into the void! Let's fill in the details!

```

1 | function infer_types(expr: ast.Expr<void>): ast.Expr<Type> {
2 |   switch (expr.kind.tag) {
3 |     cas
4 |   }
5 | }

```

If at this point we hit Enter, the editor completes:

```

1 | function infer_types(expr: ast.Expr<void>): ast.Expr<Type> {
2 |   switch (expr.kind.tag) {

```

```

3     case "bool":
4     case "int":
5     case "binary":
6     case "if":
7   }
8 }

```

There's one problem though. What we really want to write here is something like

```
const inferred_type = switch(..),
```

but in TypeScript switch is a statement, not an expression. So let's define a generic visitor!

```

1 export type Visitor<T, R> = {
2   bool(kind: ExprBool<T>): R;
3   int(kind: ExprInt<T>): R;
4   binary(kind: ExprBinary<T>): R;
5   if(kind: ExprIf<T>): R;
6 };
7
8 export function visit<T, R>(
9   expr: Expr<T>,
10  v: Visitor<T, R>,
11 ): R {
12   switch (expr.kind.tag) {
13     case "bool": return v.bool(expr.kind);
14     case "int": return v.int(expr.kind);
15     case "binary": return v.binary(expr.kind);
16     case "if": return v.if(expr.kind);
17   }
18 }

```

Armed with the visit, we can write ergonomically match over the expression:

```

1 function infer_types(expr: ast.Expr<void>): ast.Expr<Type> {
2   const ty = visit(expr, {
3     bool: () => TypeBool,
4     int: () => TypeInt,
5     binary: (kind: ast.ExprBinary<void>) => result_type(kind.op),
6     if: (kind: ast.ExprIf<void>) {
7       ...
8     },
9   });
10  ...
11 }
12
13 function result_type(op: ast.BinaryOp): Type {
14   switch (op) { // A tad verbose, but auto-completed!

```

```

15     case ast.BinaryOp.Add: case ast.BinaryOp.Sub:
16     case ast.BinaryOp.Mul: case ast.BinaryOp.Div:
17         return TypeInt
18
19     case ast.BinaryOp.Eq: case ast.BinaryOp.Neq:
20         return TypeBool
21
22     case ast.BinaryOp.Lt: case ast.BinaryOp.Gt:
23     case ast.BinaryOp.Le: case ast.BinaryOp.Ge:
24         return TypeBool
25     }
26 }

```

Before we go further, let's generalize this visiting pattern a bit! Recall that our expressions are parametrized by the type of associated data, and type-checker-shaped transformations are essentially an

$$\text{Expr}\langle U \rangle \rightarrow \text{Expr}\langle V \rangle$$

transformation.

Let's make this generic!

```

1 | export function transform<U, V>(expr: Expr<U>, v: Visitor<V, V>): Expr<V>

```

Transform maps an expression carrying T into an expression carrying V by applying an f visitor. Importantly, it's $\text{Visitor}\langle V, V \rangle$, rather than a $\text{Visitor}\langle U, V \rangle$. This is counter-intuitive, but correct — we run transformation bottom up, transforming the leaves first. So, when the time comes to visit an interior node, all subexpression will have been transformed!

The body of `transform` is wordy, but regular, rectangular, and auto-completes itself:

```

1 | export function transform<U, V>(expr: Expr<U>, v: Visitor<V, V>): Expr<V>
2 |     switch (expr.kind.tag) {
3 |         case "bool":
4 |             return {
5 |                 location: expr.location,
6 |                 data: v.bool(expr.kind),
7 |                 kind: expr.kind, 1
8 |             };
9 |         case "int":
10 |            return {
11 |                location: expr.location,
12 |                data: v.int(expr.kind),
13 |                kind: expr.kind,
14 |            };
15 |         case "binary": {
16 |             const kind: ExprBinary<V> = { 2

```

```

17     tag: "binary",
18     op: expr.kind.op,
19     lhs: transform(expr.kind.lhs, v),
20     rhs: transform(expr.kind.rhs, v),
21   };
22   return {
23     location: expr.location,
24     data: v.binary(kind), 2
25     kind: kind,
26   };
27 }
28 case "if": {
29   const kind: ExprIf<V> = {
30     tag: "if",
31     cond: transform(expr.kind.cond, v),
32     then_branch: transform(expr.kind.then_branch, v),
33     else_branch: transform(expr.kind.else_branch, v),
34   };
35   return {
36     location: expr.location,
37     data: v.if(kind),
38     kind: kind,
39   };
40 }
41 }
42 }

```

1 Note how here `expr.kind` is both `Expr<U>` and `Expr<V>` — literals don't depend on this type parameter, and TypeScript is smart enough to figure this out without us manually re-assembling the same value with a different type.

2 This is where that magic with `Visitor<V, V>` happens.

The code is pretty regular here though! So at this point we might actually recall that TypeScript is a dynamically-typed language, and write a generic traversal using `Object.keys`, *while keeping the static function signature in-place*. I don't think we need to do it here, but there's comfort in knowing that it's possible!

Now implementing type inference should be a breeze! We need some way to emit type errors though. With TypeScript, it would be trivial to accumulate errors into an array as a side-effect, but let's actually represent type errors as instances of a specific type, `TypeError` (pun intended):

```

1 | type Type = TypeBool | TypeInt | TypeError;
2 |
3 | interface TypeError {

```

```
4 tag: "Error";
5 location: ast.Location;
6 message: string;
7 }
```

To check ifs and binary expressions, we would also need an utility for comparing types:

```
1 function type_equal(lhs: Type, rhs: Type): boolean {
2   if (lhs.tag == "Error" || rhs.tag == "Error") return true;
3   return lhs.tag == rhs.tag;
4 }
```

We make the Error type equal to any other type to prevent cascading failures. With all that machinery in place, our type checker is finally:

```
1 function infer_types(expr: ast.Expr<void>): ast.Expr<Type> {
2   return ast.transform(expr, {
3     bool: (): Type => TypeBool,
4     int: (): Type => TypeInt,
5
6     binary: (kind: ast.ExprBinary<Type>, location: ast.Location): Type =>
7       if (!type_equal(kind.lhs.data, kind.rhs.data)) {
8         return {
9           tag: "Error",
10          location,
11          message: "binary expression operands have different types",
12        };
13      }
14     return result_type(kind.op);
15  },
16
17  if: (kind: ast.ExprIf<Type>, location: ast.Location): Type => {
18    if (!type_equal(kind.cond.data, TypeBool)) {
19      return {
20        tag: "Error",
21        location,
22        message: "if condition is not a boolean",
23      };
24    }
25    if (!type_equal(kind.then_branch.data, kind.else_branch.data)) {
26      return {
27        tag: "Error",
28        location,
29        message: "if branches have different types",
30      };
31    }
32    return kind.then_branch.data;
33  },
34  });
```



```

35 }
36
37 function result_type(op: ast.BinaryOp): Type {
38     ...
39 }

```

Astute reader will notice that our visitor functions now take an extra `ast.Location` argument. TypeScript allows using this argument only in cases where it is needed, cutting down verbosity.

And that's all for today! The end result is pretty neat and concise. It took some typing to get there, but TypeScript autocompletion really helps with that! What's more important, there was very little fighting with the language, and the result feels quite natural and directly corresponds to the shape of the problem.

I am not entirely sure in the conclusion just yet, but I think I'll be using TypeScript as my tool of choice for various small language hacks. It is surprisingly productive due to the confluence of three aspects:

- deno is a perfect scripting runtime! Small, hermetic, powerful, and optimized for effective development workflows.
- TypeScript tooling is great — the IDE is helpful and productive (and deno makes sure that it also requires zero configuration)
- The language is powerful both at runtime and at compile time. You can get pretty fancy with types, but you can also just escape to dynamic world if you need some very high-order code.

Just kidding, here's one more cute thing. Let's say that we want to have lots of syntactic sugar, and also want type-safe desugaring. We could tweak our setup a bit for that: instead of `Expr` and `ExprKind` being parametrized over associated data, we circularly parametrize `Expr` by the whole `ExprKind` and vice verse:

```

1 interface Expr<K> {
2     location: Location,
3     kind: K,
4 }
5
6 interface ExprBinary<E> {
7     op: BinaryOp,
8     lhs: E,
9     rhs: E,
10 }

```

This allows expressing desugaring in a type-safe manner!

```
1 // Fundamental, primitive expressions.
2 type ExprKindCore<E> =
3     ExprInt<E> | ExprBinary<E> | ExprIf<E>
4
5 // Expressions which are either themselves primitive,
6 // or can be desugared to primitives.
7 type ExprKindSugar<E> = ExprKindCore<E>
8     | ExprCond<E> | ExprUnless<E>
9
10 type ExprCore = Expr<ExprKindCore<ExprCore>>;
11 type ExprSugar = Expr<ExprKindSugar<ExprSugar>>;
12
13 // Desugaring works by reducing the set of expression kinds.
14 function desugar(expr: ExprSugar): ExprCore
15
16 // A desugaring steps takes a (potentially sugar) expression,
17 // whose subexpression are already desugared,
18 // and produces an equivalent core expression.
19 function desugar_one(
20     expr: ExprKindSugar<ExprCore>,
21 ): ExprKindCore<ExprCore>
```

 [fix typo](#)  [rss](#)  [matklad](#)