

=====

Polymorphism and Generics

Recall from chapter 3:

ad hoc polymorphism: fancy name for overloading

subtype polymorphism in OO languages allows code to do the "right thing" when a ref of parent type refers to an object of child type (implemented with vtables (to be discussed in chapter 9))

parametric polymorphism

type is a parameter of the code, implicitly or explicitly

implicit (true)

language implementation figures out what code requires of object at compile-time, as in ML or Haskell
at run-time, as in Lisp, Smalltalk, or Python
lets you apply operation to object only if object has everything the code requires

explicit (generics)

programmer specifies the type parameters explicitly
mostly what I want to talk about today

Generics found in Clu, Ada, Modula-3, C++, Eiffel, Java 5, C# 2.0, ...

C++ calls its generics **templates**.

Allow you, for example, to create a single stack abstraction, and instantiate it for stacks of integers, stacks of strings, stacks of employee records, ...

```
template <class T>
class stack {
    T[100] contents;
    int tos = 0; // first unused location
public:
    T pop();
    void push(T);
    ...
}
stack<double> s;
```

I could, of course, do

```
class stack {
    void* [100] contents;
    int tos = 0;
public:
    void* pop();
    void push(void *);
    ...
}
```

But then I have to use type casts all over the place. Inconvenient and, in C++, unsafe.

Lots of containers (stacks, queues, sets, strings, mappings, ...) in the C++ standard library.

Similarly rich libraries exist for Java and C#.

(Also for Python and Ruby, but those use implicit parametric polymorphism with run-time checking.)

Some languages (e.g. Ada and C++) allow things other than types to be passed as template arguments:

```
template <class T, int N>
class stack {
    T[N] contents;
    int tos = 0;
public:
    T pop();
    void push(T);
    ...
}
stack<double, 100> s;
```

Implementation

C# generics do run-time instantiation (**reification**). When you say `stack<foo>`, the run-time system invokes the JIT compiler and generates the appropriate code. Doesn't box native types if it doesn't need to -- more efficient.

Java doesn't do run-time instantiation. Internally everything is `stack<Object>`. You avoid the casts in the source code, but you have to pay for boxing of native types. And since the designers were unwilling (for backward compatibility reasons) to modify the VM, you're stuck with the casts in the generated code (automatically inserted by the compiler) -- even though the compiler knows they're going to succeed -- because the JVM won't accept the byte code otherwise: it will think it's unsafe. Also, because everything is an `Object` internally, reflection doesn't work.

The Java implementation strategy is known as **erasure** -- the type parameters are simply erased by the compiler. One more disadvantage: you can't say `new T()`, where `T` is generic parameter, because Java doesn't know what to create.

C++ does compile-time instantiation; more below.

C# may do compile-time instantiation when it can, as an optimization.

Constraints

The problem:

If I'm writing a sorting routine, how do I insist that the elements in the to-be-sorted list support a `less_than()` method or `<` operator?

If I'm writing a hash table, how do I insist that the keys support a `hash()` method?

If I'm writing output formatting code, how do I insist that objects support a `to_string()` method?

Related question:

Do I (can I) type-check the generic code, independent of any particular instantiation, or do I type-check the instantiations independently?

Tradeoffs are nicely illustrated by comparing Java, C#, and C++.

C++ is very flexible: every instantiation is independently type-checked. Constraints are implicit: if we try to instantiate a template for a type that doesn't support needed operations, the instance won't type-check. This has led, historically, to **really** messy error messages.

Most other languages type-check the generic itself, so you don't get any instantiation-specific error messages. To support this, they require that the operations supported by generic parameter types be explicitly listed. Java and C# leverage interfaces for this purpose.

C++20 adds **concepts**, which provide an (optional, for backward compatibility) superset of the functionality found in Java and C#.

Java example:

```
public static <T implements Comparable<T>> void sort(T A[]) {
    ...
    if (A[i].compareTo(A[j]) >= 0) ...
}
Integer[] myArray = new Integer[50];
sort(myArray);
```

Note that Java puts the type parameters right in front of the return type of the function, rather than in a preceding "template" clause.

Comparable is a standard library interface that implements Comparable<T>.

Comparable<T>.

C# syntax is similar:

```
static void sort<T>(T[] A) where T : IComparable {
    if (A[i].CompareTo(A[j]) >= 0) ...
}
int[] myArray = new int[50];
sort(myArray);
```

C# puts the type parameter between the function name and the parameter list and the constraints after the parameter list. Java won't let you do this with a generic parameter, but C# is happy to; it creates a custom version of `sort` for ints.

(pre-20) C++ doesn't require that constraints be explicit.

```
template <class T> void sort(T[] A) {
    if (A[i].CompareTo(A[j]) >= 0) ...
}
int[] myArray = new int[50];
sort(myArray);
```

(C++ can't figure out the size of an array, so you have to pass it in. Alternatively you could make it another generic parameter.)

This works fine, but it breaks if I try

```
class ObjComp implements Comparator<Object> {
    public Boolean ordered(Object a, Object b) {
        return a < b;
    }
}
Sorter<Integer> s = new Sorter<Integer>(new ObjComp());
s.sort(myArray);
```

The call to `new ObjComp()` generates a type clash message, because we're passing a `Comparator<Object>` rather than a `Comparator<Integer>`.

Because we're passing a `Comparator<Object>` rather than a `Comparator<Integer>`, the compiler can't figure out what to do with the `Object` parameter. It expects `Object` to be used in a context where `Object` is a `Comparable`, but it's not.

Typically happens in the case where `T` objects are returned from `Sorter<T>` methods, but never passed into them as parameters.

For example, I can probably pass a `Comparator<Object>` object to anybody who expects a `Comparator<Object>` object:

```
- If the generator gives them a Object instead, they're happy.
- If the generator gives them a String instead, they're happy.
- If the generator gives them a Boolean instead, they're happy.
- If the generator gives them a Number instead, they're happy.
- If the generator gives them a Object instead, they're happy.
```

Conversely (and more commonly), there are times when a `Object` object can be used in a context that expects `Object`.

When this happens, we say `Object` is **contravariant** in `T`.

For example, if you have a `Printer<P>` object to anybody who expects a `Printer<P>` object:

```
- If the generator gives them a Object instead, they're happy.
- If the generator gives them a String instead, they're happy.
- If the generator gives them a Boolean instead, they're happy.
- If the generator gives them a Object instead, they're happy.
```

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.

Java gives you more flexibility (and arguably more confusion) by allowing you to annotate `uses` of the generic -- in effect, saying "I promise (and the compiler should verify) that I never use methods that pass objects to the generic."

If you're always going to pass objects that pass objects to the generic, then you can use `uses` to indicate covariance and contravariance respectively. This restricts what can be done with `T` inside the generic.