Pluggable Types Martin Kellogg

• PA4c1 is due today

- PA4c1 is due today
- I'm out of town for the rest of the week at ICSE

- PA4c1 is due today
- I'm out of town for the rest of the week at ICSE
 I will not hold office hours on Wednesday

- PA4c1 is due today
- I'm out of town for the rest of the week at ICSE
 - I will not hold office hours on Wednesday
 - My PhD student Erfan Arvan will lecture on Wednesday
 - please be nice to him!

- PA4c1 is due today
- I'm out of town for the rest of the week at ICSE
 - I will not hold office hours on Wednesday
 - My PhD student Erfan Arvan will lecture on Wednesday
 please be nice to him!
 - If you want to see a test case before next Monday, you need to ask right after class
 - my flight is at 3pm, so I'm going to go to the airport ~12:30

Agenda

- Motivation
- Limitations of type systems
 - ...or lack thereof (Curry-Howard Correspondence)
- Pluggable Types
 - Formalism
 - Practical example (Nullability in Java)
 - Discussion of state-of-the-art (if time)

Definition: a *type system* is a set of rules that give every program element a *type*, which is an upper bound on the set of possible values that that element can take on at run time

• goal of a type system: prevent errors at run time due to unexpected values

- goal of a type system: **prevent errors** at run time due to unexpected values
 - a type can also encode the set of valid operations on values of that type

- goal of a type system: **prevent errors** at run time due to unexpected values
 - a type can also encode the set of valid operations on values of that type
- many programming languages (including Cool) check types at compile time (*static types*)

- goal of a type system: **prevent errors** at run time due to unexpected values
 - a type can also encode the set of valid operations on values of that type
- many programming languages (including Cool) check types at compile time (*static types*)
 - some languages instead check types at run time (*dynamic types*)

• Traditional static type systems like Cool's are already widely-used by software engineers

- Traditional static type systems like Cool's are already widely-used by software engineers
 - in practice, type systems are one of the few common techniques for improving correctness
 - along with testing and code review

- Traditional static type systems like Cool's are already widely-used by software engineers
 - in practice, type systems are one of the few common techniques for improving correctness
 - along with testing and code review
- However, traditional type systems prevent some but not all run-time errors

- Traditional static type systems like Cool's are already widely-used by software engineers
 - in practice, type systems are one of the few common techniques for improving correctness
 - along with testing and code review
- However, traditional type systems prevent some but not all run-time errors
 - prevented: unsuccessful vtable lookups, missing fields, mixing up incompatible values (e.g., integers and strings), etc.

- Traditional static type systems like Cool's are already widely-used by software engineers
 - in practice, type systems are one of the few common techniques for improving correctness
 - along with testing and code review
- However, traditional type systems prevent some but not all run-time errors
 - prevented: unsuccessful vtable lookups, missing fields, mixing up incompatible values (e.g., integers and strings), etc.
 - not prevented: out-of-bounds array accesses, null-pointer dereferences, memory leaks, mismatched units, etc.

- Many techniques for **verification** of computer programs have been proposed
 - e.g., Floyd-Hoare logic, guarded commands, abstract interpretation, etc. (ask in office hours)

- Many techniques for verification of computer programs have been proposed
 - e.g., Floyd-Hoare logic, guarded commands, abstract interpretation, etc. (ask in office hours)
- Unlike static type systems, these techniques have not been widely-adopted outside of niche industries :(

- Many techniques for verification of computer programs have been proposed
 - e.g., Floyd-Hoare logic, guarded commands, abstract interpretation, etc. (ask in office hours)
- Unlike static type systems, these techniques have not been widely-adopted outside of niche industries :(
 - and computer programs remain full of bugs that a verifier could have prevented...

- Many techniques for verification of computer programs have been proposed
 - e.g., Floyd-Hoare logic, guarded commands, abstract interpretation, etc. (ask in office hours)
- Unlike static type systems, these techniques have not been widely-adopted outside of niche industries :(
 - and computer programs remain full of bugs that a verifier could have prevented...
- One of the biggest obstacles to the adoption of verification technologies in practice is the need to write **specifications**
 - e.g., pre- and post-conditions, loop invariants, etc.

• **Observation**: programmers have already shown that they are willing to write down static types.

- **Observation**: programmers have already shown that they are willing to write down static types.
 - **Insight:** What if we allowed them to write down the specifications that we need for verification as types?

- **Observation**: programmers have already shown that they are willing to write down static types.
 - Insight: What if we allowed them to write down the specifications that we need for verification as types?
 - Pluggable type systems allow us to do just that

- **Observation**: programmers have already shown that they are willing to write down static types.
 - Insight: What if we allowed them to write down the specifications that we need for verification as types?
 - Pluggable type systems allow us to do just that
- More formally, a *pluggable type system* extends a host type system with a set of *type qualifiers* that model a property of interest (usually "the program has no bugs of kind X", for some X)

- **Observation**: programmers have already shown that they are willing to write down static types.
 - **Insight:** What if we allowed them to write down the specifications that we need for verification as types?
 - Pluggable type systems allow us to do just that
- More formally, a *pluggable type system* extends a host type system with a set of *type qualifiers* that model a property of interest (usually "the program has no bugs of kind X", for some X)
 - each type qualifier refines the type in question
 - that is, each qualifier is "more specific" than the base type (i.e., a subtype)

@Positive int x

@Positive int x

Basetype











Pluggable Types: One Slide Summary

 Developers already use static type systems, so they're familiar with the general idea of types => relatively easy to use (compared to other verification techniques)
Pluggable Types: One Slide Summary

- Developers already use static type systems, so they're familiar with the general idea of types => relatively easy to use (compared to other verification techniques)
- Type qualifiers encode property of interest
 - effectively a "second" type system

Pluggable Types: One Slide Summary

- Developers already use static type systems, so they're familiar with the general idea of types => relatively easy to use (compared to other verification techniques)
- Type qualifiers encode property of interest
 o effectively a "second" type system
- Qualified types are like a **Cartesian product** of a type from the pluggable type system and a type from the base type system

Pluggable Types: One Slide Summary

- Developers already use static type systems, so they're familiar with the general idea of types => relatively easy to use (compared to other verification techniques)
- Type qualifiers encode property of interest
 - effectively a "second" type system
- Qualified types are like a **Cartesian product** of a type from the pluggable type system and a type from the base type system
- Typechecking is naturally **modular** = fast
 - but this comes at a cost: programmers need to write types

- Qualified types are like a **Cartesian product** of a type from the pluggable type system and a type from the base type system
 - What does this actually mean?

- Qualified types are like a Cartesian product of a type from the pluggable type system and a type from the base type system
 What does this actually mean?
- Analogy: each pluggable type system is like a **dimension** in a traditional coordinate system:

- Qualified types are like a Cartesian product of a type from the pluggable type system and a type from the base type system
 What does this actually mean?
- Analogy: each pluggable type system is like a **dimension** in a traditional coordinate system:
 - the base type system is the first dimension

- Qualified types are like a Cartesian product of a type from the pluggable type system and a type from the base type system
 What does this actually mean?
- Analogy: each pluggable type system is like a **dimension** in a traditional coordinate system:
 - the base type system is the first dimension
 - each additional qualifier adds one more dimension

- Qualified types are like a Cartesian product of a type from the pluggable type system and a type from the base type system
 What does this actually mean?
- Analogy: each pluggable type system is like a **dimension** in a traditional coordinate system:
 - the base type system is the first dimension
 - each additional qualifier adds one more dimension
 - let's see an example





Object + Object Number Int + Int + Number Real Rational + Rational + Real

Or we can extend it **downwards** with a positive qualifier...



• More formally, in a pluggable type system, every type is a *product type* of a type qualifier and a base type.

- More formally, in a pluggable type system, every type is a *product type* of a type qualifier and a base type.
 - A product type is a logical "and" of two types (more on this in a moment)

- More formally, in a pluggable type system, every type is a *product type* of a type qualifier and a base type.
 - A product type is a logical "and" of two types (more on this in a moment)
- Instead of writing **x** : **T**, we write **x** : **Q T**

- More formally, in a pluggable type system, every type is a *product type* of a type qualifier and a base type.
 - A product type is a logical "and" of two types (more on this in a moment)
- Instead of writing **x** : **T**, we write **x** : **Q T**
 - **T** is the base type, **Q** is the type qualifier from the pluggable type system of interest

- More formally, in a pluggable type system, every type is a *product type* of a type qualifier and a base type.
 - A product type is a logical "and" of two types (more on this in a moment)
- Instead of writing **x** : **T**, we write **x** : **Q T**
 - **T** is the base type, **Q** is the type qualifier from the pluggable type system of interest
- We can extend this formalism arbitrarily, e.g., by writing x : Q₁ Q₂ T
 (As long as Q₁ and Q₂ are from different type systems)

- More formally, in a plugg type of a type qualifier an
 A product type is a lo moment)
 Note that there is not really anything special about T in this formalism - we can view the base type system as just another pluggable type that's "in use by default"
- Instead of writing **x** : **T**, we write **x** : **Q T**
 - **T** is the base type, **Q** is the type qualifier from the pluggable type system of interest
- We can extend this formalism arbitrarily, e.g., by writing x : Q₁ Q₂ T
 (As long as Q₁ and Q₂ are from different type systems)

• I hinted that product types are "like" a logical AND gate

- I hinted that product types are "like" a logical AND gate
- This relationship is one of the cornerstones of the Curry-Howard correspondence between mathematical proofs and computer programs

- I hinted that product types are "like" a logical AND gate
- This relationship is one of the cornerstones of the *Curry-Howard* correspondence between mathematical proofs and computer programs
 - There is an isomorphism between the proof systems and the models of computation: these two families of formalisms can be considered as identical
 - (an *isomorphism* is a one-to-one correspondence between the mathematical structures of two objects)

- I hinted that product types are "like" a logical AND gate
- This relationship is one of the cornerstones of the *Curry-Howard* correspondence between mathematical proofs and computer programs
 - There is an isomorphism between the proof systems and the models of computation: these two families of formalisms can be considered as identical
 - (an *isomorphism* is a one-to-one correspondence between the mathematical structures of two objects)
 - In other words: a proof is a program, and the formula it proves is the type for the program

• The Curry-Howard Correspondence has deep implications:

The Curry-Howard Correspondence has deep implications:
 It is possible to "run" proofs, as we would computer programs

- The Curry-Howard Correspondence has deep implications:
 - It is possible to "run" proofs, as we would computer programs
 - This implies that it's possible to "co-develop" a program and the proof of its correctness

- The Curry-Howard Correspondence has deep implications:
 - It is possible to "run" proofs, as we would computer programs
 - This implies that it's possible to "co-develop" a program and the proof of its correctness
 - Interactive theorem provers like Rocq or Lean operationalize this idea

- The Curry-Howard Correspondence has deep implications:
 - It is possible to "run" proofs, as we would computer programs
 - This implies that it's possible to "co-develop" a program and the proof of its correctness
 - Interactive theorem provers like Rocq or Lean operationalize this idea
 - A type system is just a way to write down a theorem about the program

- The Curry-Howard Correspondence has deep implications:
 - It is possible to "run" proofs, as we would computer programs
 - This implies that it's possible to "co-develop" a program and the proof of its correctness
 - Interactive theorem provers like Rocq or Lean operationalize this idea
 - A type system is just a way to write down a theorem about the program
 - Different type systems let us write down different theorems

- The Curry-Howard Correspondence has deep implications:
 - It is possible to "run" proofs, as we wanted computer programs
 - This implies that it's possible to for more on this topic, ask the proof of its correctness
 For more on this topic, ask in office hours and/or take CS 735 with me in Sp26.
 - Interactive theorem provers like Roug or Lour operationance this idea
 - A type system is just a way to write down a theorem about the program
 - Different type systems let us write down different theorems

- The modifications to most of the type rules to support a pluggable type system are straightforward: just add a qualifier term to each type:
 - **T -> Q T**

- The modifications to most of the type rules to support a pluggable type system are straightforward: just add a qualifier term to each type:
 - **T -> Q T**
 - Effectively, this requires the programmer to write a type qualifier everywhere they write a type

- The modifications to most of the type rules to support a pluggable type system are straightforward: just add a qualifier term to each type:
 - **T -> Q T**
 - Effectively, this requires the programmer to write a type qualifier everywhere they write a type
- The **subtyping** and **least upper bound** rules are bit trickier (but not much).

- The modifications to most of the type rules to support a pluggable type system are straightforward: just add a qualifier term to each type:
 - **T -> Q T**
 - Effectively, this requires the programmer to write a type qualifier everywhere they write a type
- The **subtyping** and **least upper bound** rules are bit trickier (but not much).
 - We need to account for all possible cases

Adding Pluggable Types to Cool: Subtyping
• Subtyping can generally be split into two parts:

- Subtyping can generally be split into two parts:
 - the base types part, which just defers to the old subtyping rule

- Subtyping can generally be split into two parts:
 - the base types part, which just defers to the old subtyping rule
 - the type qualifier part

- Subtyping can generally be split into two parts:
 - the base types part, which just defers to the old subtyping rule
 - the type qualifier part
 - for this part, assume that the type qualifiers are arranged into a lattice (sound familiar?)

- Subtyping can generally be split into two parts:
 - the base types part, which just defers to the old subtyping rule
 - the type qualifier part
 - for this part, assume that the type qualifiers are arranged into a lattice (sound familiar?)

Nullable T | Nonnull T

- Subtyping can generally be split into two parts:
 - the base types part, which just defers to the old subtyping rule
 - the type qualifier part
 - for this part, assume that the type qualifiers are arranged into a lattice (sound familiar?)

Nullable T | Nonnull T

Any Int | + Int

- Subtyping can generally be split into two parts:
 - the base types part, which just defers to the old subtyping rule
 - the type qualifier part
 - for this part, assume that the type qualifiers are arranged into a lattice (sound familiar?)
 - then, we can just read subtyping information from the lattice

Nullable T | Nonnull T

Any Int | + Int

- Subtyping can generally be split into two parts:
 - the base types part, which just defers to the old subtyping rule
 - the type qualifier part

If this looks similar to **abstract interpretation**, that's because it is (Cousot proved it in POPL 1997). Pluggable types can be viewed as an alternative syntax for abstract interpretation. the type to a <mark>lattice</mark> typing ce Nullable T | Nonnull T

Any Int | + Int

- Least upper bounds can also be split between the base types and the qualifiers:
 - $\circ \quad \mathsf{LUB}(\mathbf{Q}_{1} \mathbf{T}_{1}, \mathbf{Q}_{2} \mathbf{T}_{2}) = \mathsf{LUB}(\mathbf{Q}_{1}, \mathbf{Q}_{2}) \, \mathsf{LUB}(\mathbf{T}_{1}, \mathbf{T}_{2})$

- Least upper bounds can also be split between the base types and the qualifiers:
 - $\circ \quad LUB(\mathbf{Q}_1 \mathbf{T}_1, \mathbf{Q}_2 \mathbf{T}_2) = LUB(\mathbf{Q}_1, \mathbf{Q}_2) LUB(\mathbf{T}_1, \mathbf{T}_2)$
- Some pluggable type systems might have interactions between qualifiers and base types.

- Least upper bounds can also be split between the base types and the qualifiers:
 - $\circ \quad LUB(\mathbf{Q}_1 \mathbf{T}_1, \mathbf{Q}_2 \mathbf{T}_2) = LUB(\mathbf{Q}_1, \mathbf{Q}_2) LUB(\mathbf{T}_1, \mathbf{T}_2)$
- Some pluggable type systems might have interactions between qualifiers and base types.
 - e.g., it might not make sense to have a "positive string"

- Least upper bounds can also be split between the base types and the qualifiers:
 - $\circ \quad LUB(\mathbf{Q}_1 \mathbf{T}_1, \mathbf{Q}_2 \mathbf{T}_2) = LUB(\mathbf{Q}_1, \mathbf{Q}_2) LUB(\mathbf{T}_1, \mathbf{T}_2)$
- Some pluggable type systems might have interactions between qualifiers and base types.
 - e.g., it might not make sense to have a "positive string"
 - in practice, this is fine just treat like bottom

• While it's plausible to require programmers to write a type qualifier at each type use, in practice that's inconvenient

- While it's plausible to require programmers to write a type qualifier at each type use, in practice that's inconvenient
 - instead, we'd like to pick sensible defaults for unqualified types

- While it's plausible to require programmers to write a type qualifier at each type use, in practice that's inconvenient
 - instead, we'd like to pick sensible defaults for unqualified types
- For example, consider a nullability pluggable type system with two qualifiers: Nullable and NonNull.

- While it's plausible to require programmers to write a type qualifier at each type use, in practice that's inconvenient
 - instead, we'd like to pick sensible defaults for unqualified types
- For example, consider a nullability pluggable type system with two qualifiers: Nullable and NonNull.
 - if we see an unqualified type **T**, should we treat it as:
 - Nullable T?
 - or **NonNull T**?

- While it's plausible to require programmers to write a type qualifier at each type use, in practice that's inconvenient
 - instead, we'd like to pick sensible defaults for unqualified types
- For example, consider a nullability pluggable type system with two qualifiers: Nullable and NonNull.
 - if we see an unqualified type **T**, should we treat it as:
 - Nullable T?
 - or **NonNull T**?
 - interestingly, **either choice is correct**! (Why?)

• Practical pluggable type systems use a specific defaulting scheme called CLIMB-to-top

- Practical pluggable type systems use a specific defaulting scheme called *CLIMB-to-top*
 - this approach uses a simple dataflow analysis to propagate types within a method body

- Practical pluggable type systems use a specific defaulting scheme called CLIMB-to-top
 - this approach uses a simple dataflow analysis to propagate types within a method body
 - also saves user effort!

- Practical pluggable type systems use a specific defaulting scheme called **CLIMB-to-top**
 - this approach uses a simple dataflow analysis to propagate types within a method body
 - also saves user effort!
 - it defaults any location where dataflow discovers a qualifier to that qualifier, regardless of the "usual" defaulting rules

- Practical pluggable type systems use a specific defaulting scheme called **CLIMB-to-top**
 - this approach uses a simple dataflow analysis to propagate types within a method body
 - also saves user effort!
 - it defaults any location where dataflow discovers a qualifier to that qualifier, regardless of the "usual" defaulting rules
 - e.g., in Object x = null, x will be Nullable even if the usual default is NonNull

- Practical pluggable type systems use a specific defaulting scheme called **CLIMB-to-top**
 - this approach uses a simple dataflow analysis to propagate types within a method body
 - also saves user effort!
 - it defaults any location where dataflow discovers a qualifier to that qualifier, regardless of the "usual" defaulting rules
 - e.g., in Object x = null, x will be Nullable even if the usual default is NonNull
 - there are a number of other subtleties which I'll elide for time

Pluggable Types: Custom Type Rules

• Many pluggable type systems exist to enforce some kind of custom type rule (which might replace an existing rule)

Pluggable Types: Custom Type Rules

- Many pluggable type systems exist to enforce some kind of custom type rule (which might replace an existing rule)
 - for example, a nullability type system might have a custom rule that forbids dereferencing a type with a nullable qualifier:

$$\frac{\Gamma + \mathbf{e}_1 : \text{NonNull T} \quad \Gamma + \mathbf{e}_2 : \mathbf{QT}}{\Gamma + \mathbf{e}_1 \cdot \mathbf{e}_2 : \mathbf{QT}} \text{[Deref]}$$

Pluggable Types: Custom Type Rules

- Many pluggable type systems exist to enforce some kind of custom type rule (which might replace an existing rule)
 - for example, a nullability type system might have a custom rule that forbids dereferencing a type with a nullable qualifier:

$$\frac{\mathbf{\nabla} \mathbf{e}_{1} : \mathsf{NonNull T} \quad \mathbf{\nabla} \mathbf{e}_{2} : \mathbf{QT}}{\mathbf{\nabla} \mathbf{e}_{1} \cdot \mathbf{e}_{2} : \mathbf{QT}}$$
[Deref]

 These rules can be arbitrary, but generally the type system designer chooses them to make some soundness proof work

- Null dereferences (@NonNull)
 - >200 errors in Google Collections, javac, ...

- Null dereferences (@NonNull)
 - >200 errors in Google Collections, javac, ...
- Equality tests (@Interned)
 - >200 problems in Xerces, Lucene, ...

- Null dereferences (@NonNull)
 - >200 errors in Google Collections, javac, ...
- Equality tests (@Interned)
 - >200 problems in Xerces, Lucene, ...
- Concurrency / locking (@GuardedBy)
 - >500 errors in BitcoinJ, Derby, Guava, Tomcat, ...

- Null dereferences (@NonNull)
 - >200 errors in Google Collections, javac, ...
- Equality tests (@Interned)
 - >200 problems in Xerces, Lucene, ...
- Concurrency / locking (@GuardedBy)
 - > >500 errors in BitcoinJ, Derby, Guava, Tomcat, ...
- Array bounds (@IndexFor, @NonNegative, etc.)
 - defects in Google Guava, JFreeChart, ...

- Null dereferences (@NonNull)
 - >200 errors in Google Collections, javac, ...
- Equality tests (@Interned)
 - >200 problems in Xerces, Lucene, ...
- Concurrency / locking (@GuardedBy)
 - > >500 errors in BitcoinJ, Derby, Guava, Tomcat, ...
- Array bounds (@IndexFor, @NonNegative, etc.)
 - defects in Google Guava, JFreeChart, ...
- Resource leaks (@MustCall)
 - o defects in Zookeeper, Hadoop, HBase, ...

- Null dereferences (@NonNull)
 - >200 errors in Google Collections, javac, ...
- Equality tests (@Interned)
 - >200 problems in Xer These last two were
- Concurrency / locking (@ part of my PhD work!
 >500 errors in Bitcoin, _____, ___, ____, __, ___, ___, ___, __, ___, ___, ___, _
- Array bounds (@IndexFor, @NonNegative, etc.)
 - defects in Google Guava, JFreeChart, ...
- Resource leaks (@MustCall)
 - defects in Zookeeper, Hadoop, HBase, ...

Trivia Break: History of Mathematics

This Ancient Greek mathematician, physicist, engineer, astronomer, and inventor from the ancient city of Syracuse in Sicily is regarded as the greatest mathematician of ancient history. He anticipated modern calculus and analysis by applying the concept of the infinitely small and the method of exhaustion to derive and rigorously prove many geometrical theorems. These include the area of a circle, the surface area and volume of a sphere, the area of an ellipse, the area under a parabola, the volume of a segment of a paraboloid of revolution, the volume of a segment of a hyperboloid of revolution, and the area of a spiral.
Let's consider one specific example in detail: nullability in Java
 Tony Hoare's "billion-dollar mistake" etc

- Let's consider one specific example in detail: nullability in Java
 Tony Hoare's "billion-dollar mistake" etc
- While my examples in this section will all concern nullability, many of the **principles** at play are applicable to any pluggable typechecker

- Let's consider one specific example in detail: nullability in Java
 Tony Hoare's "billion-dollar mistake" etc
- While my examples in this section will all concern nullability, many of the **principles** at play are applicable to any pluggable typechecker
 - I'll try to point these out as we come to them, but as a hint about exam questions - asking you about how one of these principles might apply to a different type system would be fair game

• Where is the defect in the following code?

- Where is the defect in the following code?
 - Whose fault: library implementer or client?

Where is the defect in the following code?
 Whose fault: library implementer or client?

```
String op(Data in) {
    return "transform: " + in.getF();
}
```



Where is the defect in the following code?
 Whose fault: library implementer or client?



• Where is the defect in the following code?

```
• Whose fault: library implementer or client?
```

```
String op(Data in) {
   return "transform: " + in.getF();
}
```

```
String s = op(null);
```

Cannot decide without a **specification**!

- Specification option 1:
 - Non-null parameter

```
String op(@NonNull Data in) {
   return "transform: " + in.getF();
}
```

```
String s = op(null);
```

- Specification option 1:
 - Non-null parameter

```
String op(@NonNull Data in) {
   return "transform: " + in.getF();
}
String s = op(null);
   with this specification, defect is here!
```

- Specification option 2:
 - Nullable parameter

```
String op(@Nullable Data in) {
   return "transform: " + in.getF();
}
```

```
String s = op(null);
```

- Specification option 2:
 - Nullable parameter

```
String op(@Nullable Data in) {
   return "transform: " + in.getF();
}
with this specification, defect
is here, instead!
```

• Two **type qualifiers** (= abstract values):

- Two **type qualifiers** (= abstract values):
 - **@Nullable** = might be null

- Two **type qualifiers** (= abstract values):
 - **@Nullable** = might be null
 - **@NonNull** = definitely not null

- Two **type qualifiers** (= abstract values):
 - **@Nullable** = might be null
 - @NonNull = definitely not null

Lattice: @Nullable | @NonNull

• Default is **@NonNull**

- Two **type qualifiers** (= abstract values):
 - **@Nullable** = might be null
 - **@NonNull** = definitely not null
- Default is **@NonNull**
 - requires fewer annotations



- Two **type qualifiers** (= abstract values):
 - **@Nullable** = might be null
 - **@NonNull** = definitely not null
- Default is **@NonNull**
 - requires fewer annotations
 - makes the dangerous case explicit



- Two **type qualifiers** (= abstract values):
 - **@Nullable** = might be null
 - **@NonNull** = definitely not null
- Default is **@NonNull**
 - requires fewer annotations
 - makes the dangerous case explicit
 - note: opposite of Java's implicit default!



- Two **type qualifiers** (= abstract values):
 - **@Nullable** = might be null
 - **@NonNull** = definitely not null
- Default is **@NonNull**
 - requires fewer annotations
 - makes the dangerous case explicit
 - note: opposite of Java's implicit default!
- Nearly no annotations required in method bodies



- Two **type qualifiers** (= abstract values):
 - **@Nullable** = might be null
 - **@NonNull** = definitely not null
- Default is **@NonNull**
 - requires fewer annotations
 - makes the dangerous case explicit
 - note: opposite of Java's implicit default!
- Nearly no annotations required in method bodies
 - dataflow analysis handles most cases (CLIMB-to-top)



- Two **type qualifiers** (= abstract values):
 - **@Nullable** = might be null
 - **@NonNull** = definitely not null
- Default is **@NonNull**
 - requires fewer annotations
 - makes the dangerous case explicit
 - note: opposite of Java's implicit default!
- Nearly no annotations required in method bodies
 - dataflow analysis handles most cases (CLIMB-to-top)
 - o ...except for e.g., generics: List<@Nullable String>



- **Dataflow analysis** performs refinement within method bodies:
- if (myField != null) {
 myField.hashCode();
 }

- **Dataflow analysis** performs refinement within method bodies:
- if (myField != null) {
 myField.hashCode();

}

• No need to e.g., declare a new local variable

- **Dataflow analysis** performs refinement within method bodies:
- if (myField != null) {
 myField.hashCode();

- No need to e.g., declare a new local variable
- This is the same "*refinement*" that we saw when we talked about abstract interpretation earlier

Nullability: Side-Effects

- if (myField != null) {
 method1();
 myField.hashCode();
 }
- What about the code above? Should it typecheck?

Nullability: Side-Effects

- if (myField != null) {
 method1();
 myField.hashCode();
 }
- What about the code above? Should it typecheck?
 - No! method1() could side-effect myField and set it to null

Nullability: Side-Effects

if (myField != null) {
 method1();
 myField.hashCode();
}

- What about the code above? Should it typecheck?
 - No! method1() could side-effect myField and set it to null
 - Three ways to tell the typechecker this is ok:
 - **@SideEffectFree** void method1() { ... }
 - @MonotonicNonNull myField;
 - @EnsuresNonNull("myField") method1() {...}

• Three annotations for methods:

- Three annotations for methods:
 - **@SideEffectFree**
 - Does not modify externally-visible state

- Three annotations for methods:
 - **@SideEffectFree**
 - Does not modify externally-visible state
 - **@Deterministic**
 - If called with == args again, gives == result

- Three annotations for methods:
 - **@SideEffectFree**
 - Does not modify externally-visible state
 - **@Deterministic**
 - If called with == args again, gives == result
 - **@Pure**
 - Both side-effect-free and deterministic
- Three annotations for methods:
 - **@SideEffectFree**
 - Does not modify externally-visible state
 - **@Deterministic**
 - If called with == args again, gives == result
 - **@Pure**
 - Both side-effect-free and deterministic
- These are *method annotations*, not type qualifiers

- Three annotations for methods:
 - **@SideEffectFree**
 - Does not modify externally-visible state
 - **@Deterministic**
 - If called with == args again, gives == result
 - **@Pure**
 - Both side-effect-free and deterministic
- These are *method annotations*, not type qualifiers
 - Checker Framework trusts rather than checks them by default

- Three annotations for methods:
 - @SideEffectFree
 - Does not modify externally-visible state
 - **@Deterministic**
 - If called with == args again, gives == result
 - **@Pure**
 - Both side-effect-free and deterministic
- These are *method annotations*, not type qualifiers
 - Checker Framework trusts rather than checks them by default
 - checking these with few false positives is an open problem

- Three annotations for methods:
 - @SideEffectFree
 - Does not modify externally-visible state
 - **@Deterministic**
 - If called with == args again, gives == result
 - **@Pure**
 - Both side-effect-fr
- These are *method annotat*
 - Checker Framework t
 - checking these with

These can be used with any Checker Framework checker, not just the nullability checker

- Consider the following code:
- if (method2() != null) {
 method2().hashCode();

- Consider the following code:
- if (method2() != null) {
 method2().hashCode();

}

• This code is safe if and only if method2() returns the same value every time we call it

- Consider the following code:
- if (method2() != null) {
 method2().hashCode();

- This code is safe if and only if method2() returns the same value every time we call it
 - **@Deterministic** exactly expresses this property

- Consider the following code:
- if (method2() != null) {
 method2().hashCode();

}

- This code is safe if and only if method2() returns the same value every time we call it
 - **@Deterministic** exactly expresses this property
 - Why might a method not be deterministic?

 Another way to prove that a value is persistent across side-effects is the @MonotonicNonNull type annotation, written on a field type

- Another way to prove that a value is persistent across side-effects is the @MonotonicNonNull type annotation, written on a field type
 - Might be null or non-null

- Another way to prove that a value is persistent across side-effects is the @MonotonicNonNull type annotation, written on a field type
 - Might be null or non-null
 - May only be (re-)assigned a non-null value

- Another way to prove that a value is persistent across side-effects is the @MonotonicNonNull type annotation, written on a field type
 - Might be null or non-null
 - May only be (re-)assigned a non-null value
- Purpose: avoid re-checking

- Another way to prove that a value is persistent across side-effects is the @MonotonicNonNull type annotation, written on a field type
 - Might be null or non-null
 - May only be (re-)assigned a non-null value
- Purpose: avoid re-checking
- Once non-null, always non-null
 - Example: Singleton pattern

- Another way to prove that a value is persistent across side-effects is the @MonotonicNonNull type annotation, written on a field type
 - Might be null or non-null
 - May only be (re-)assigned a non-null value
- Purpose: avoid re-checking
- Once non-null, always non-null
 - Example: Singleton pattern
- Fields with this type don't need to be "unrefined" at possibly-side-effecting method calls

• More method annotations that express facts that a method requires or guarantees about program elements other than parameters/return values (e.g., in-scope fields)

- More method annotations that express facts that a method requires or guarantees about program elements other than parameters/return values (e.g., in-scope fields)
 - **@RequiresNonNull (pre-condition)**

- More method annotations that express facts that a method requires or guarantees about program elements other than parameters/return values (e.g., in-scope fields)
 - **@RequiresNonNull (pre-condition)**
 - **@EnsuresNonNull (post-condition)**

- More method annotations that express facts that a method requires or guarantees about program elements other than parameters/return values (e.g., in-scope fields)
 - **@RequiresNonNull (pre-condition)**
 - **@EnsuresNonNull (**post-condition)
 - **@EnsuresNonNullIf (conditional post-condition)**

- More method annotations that express facts that a method requires or guarantees about program elements other than parameters/return values (e.g., in-scope fields)
 - **@RequiresNonNull (pre-condition)**
 - **@EnsuresNonNull (post-condition)**
 - **@EnsuresNonNullIf (conditional post-condition)**
- Example:

@EnsuresNonNullIf(expression="#1", result=true)
public boolean equals(@Nullable Object obj) { ... }

Nullability: Qualifier Polymorphism

Nullability: Qualifier Polymorphism

```
/** Interns a String, and handles null. */
@PolyNull String intern(@PolyNull String a) {
   if (a == null) {
      return null;
    }
   return a.intern();
}
```

Nullability: Qualifier Polymorphism

```
/** Interns a String, and handles null. */
@PolyNull String intern(@PolyNull String a) {
   if (a == null) {
      return null;
    }
   return a.intern();
}
```

- Like defining two overloaded methods:
 - O @NonNull String intern(@NonNull String a) {...}
 - O @Nullable String intern(@Nullable String a) {...}

• It's possible for an @NonNull field to contain null at run time! How?

- It's possible for an @NonNull field to contain null at run time! How?
- Consider the following code:

 @NonNull String name;
 MyClass() { // constructor
 ... this.name.hashCode() ...
 }

- It's possible for an @NonNull field to contain null at run time! How?
- Consider the following code:

 @NonNull String name;
 MyClass() { // constructor
 ... this.name.hashCode() ...
 }
- Problem: Java initializes all non-primitive fields to null during object initialization, before running the constructor

- It's possible for an @NonNull field to contain null at run time! How?
- Consider the following code:

 @NonNull String name;
 MyClass() { // constructor
 ... this.name.hashCode() ...
 }
- Problem: Java initializes all non-primitive fields to null during object initialization, before running the constructor
 Does Cool do this? How do you know?

- It's possible for an @NonNull field to contain null at run time! How?
- Consider the following code:

 @NonNull String name;
 MyClass() { // constructor
 ... this.name.hashCode() ...
 }
- Problem: Java initializes all non-primitive fields to null during object initialization, before running the constructor
 - Does Cool do this? How do you know?
 - (new operational semantics rule)

• So, another type system tracks the initialization status of each object:

- So, another type system tracks the initialization status of each object:
 - **@Initialized** = constructor has completed

- So, another type system tracks the initialization status of each object:
 - **@Initialized** = constructor has completed
 - **@UnderInitialization(Frame.class)**
 - Frame's constructor is currently executing

- So, another type system tracks the initialization status of each object:
 - **@Initialized** = constructor has completed
 - **@UnderInitialization(Frame.class)**
 - Frame's constructor is currently executing
 - **@UnknownInitialization(Frame.class)**
 - Might be initialized or under initialization

- So, another type system tracks the initialization status of each object:
 - **@Initialized** = constructor has completed
 - **@UnderInitialization(Frame.class)**
 - Frame's constructor is currently executing
 - **@UnknownInitialization(Frame.class)**
 - Might be initialized or under initialization
- Type rule for initialization: cannot dereference fields (even nonnull fields!) until the object is **@Initialized**

- So, another type system tracks the initialization status of each object:
 - **@Initialized** = constructor has completed
 - **@UnderInitialization(Frame.class)**
 - Frame's constructor is currently executing
 - **@UnknownInitialization(Frame.class)**
 - Might be initialized or under initialization
- Type rule for initialization: cannot dereference fields (even nonnull fields!) until the object is **@Initialized**
 - What's the type hierarchy/lattice for these?
• Nullability checker in the CF also has a type system for tracking which values are known to be keys in which maps

- Nullability checker in the CF also has a type system for tracking which values are known to be keys in which maps
 - Map.get() can return null in general, but map operations are so common that we need to handle the common case

- Nullability checker in the CF also has a type system for tracking which values are known to be keys in which maps
 - Map.get() can return null in general, but map operations are so common that we need to handle the common case
 - @KeyFor("m") means "this value is definitely a valid key in map m, so calling Map.get() on this value will return something nonnull"

- Nullability checker in the CF also has a type system for tracking which values are known to be keys in which maps
 - Map.get() can return null in general, but map operations are so common that we need to handle the common case
 - @KeyFor("m") means "this value is definitely a valid key in map m, so calling Map.get() on this value will return something nonnull"
- These sort of auxiliary type systems are a **common pattern** in practical typecheckers

- Nullability checker in the CF also has a type system for tracking which values are known to be keys in which maps
 - Map.get() can return null in general, but map operations are so common that we need to handle the common case
 - @KeyFor("m") means "this value is definitely a valid key in map m, so calling Map.get() on this value will return something nonnull"
- These sort of auxiliary type systems are a **common pattern** in practical typecheckers
 - My first project in grad school had 7!

• Practical checkers still have **false positives**

- Practical checkers still have **false positives**
 - Users need ways to suppress them easily, or they won't use the tools

- Practical checkers still have **false positives**
 - Users need ways to suppress them easily, or they won't use the tools
- Two main ways:

- Practical checkers still have **false positives**
 - Users need ways to suppress them easily, or they won't use the tools
- Two main ways:
 - Casts. A cast in "regular Java" (without a pluggable typechecker) is already used to suppress a warning from the regular type system
 - Can also write a type qualifier in the casted-to type

- Practical checkers still have **false positives**
 - Users need ways to suppress them easily, or they won't use the tools
- Two main ways:
 - Casts. A cast in "regular Java" (without a pluggable typechecker) is already used to suppress a warning from the regular type system
 - Can also write a type qualifier in the casted-to type
 - Explicit warning suppressions via @SuppressWarnings
 - More flexible and easier to read

Nullability: Effectiveness

Nullability: Effectiveness

	Null pointer errors		False	Annotations
	Found	Missed	warnings	written
Checker Framework	9	0	4	35
FindBugs	0	9	1	0
Jlint	0	9	8	0
PMD	0	9	0	0
Eclipse, in 2017	0	9	8	0
Intellij (@NotNull default), in 2017	0	9	1	0
	3	6	1	925 + 8

Checking the Lookup program for file system searching (4kLOC)



• Extant type systems are **pretty good**

- Extant type systems are **pretty good**
 - Nullness Checker I've just shown is in wide use

- Extant type systems are **pretty good**
 - Nullness Checker I've just shown is in wide use
 - Uber's NullAway is a big competitor
 - "No false negatives in practice"

- Extant type systems are **pretty good**
 - Nullness Checker I've just shown is in wide use
 - Uber's NullAway is a big competitor
 - "No false negatives in practice"
 - Type systems for other problems (e.g., resource leaks) becoming more popular

- Extant type systems are **pretty good**
 - Nullness Checker I've just shown is in wide use
 - Uber's NullAway is a big competitor
 - "No false negatives in practice"
 - Type systems for other problems (e.g., resource leaks) becoming more popular
- Big problem: getting people to use them

- Extant type systems are **pretty good**
 - Nullness Checker I've just shown is in wide use
 - Uber's NullAway is a big competitor
 - "No false negatives in practice"
 - Type systems for other problems (e.g., resource leaks) becoming more popular
- Big problem: getting people to use them
 - Annotation burden for legacy code is the most serious issue

- Extant type systems are **pretty good**
 - Nullness Checker I've just shown is in wide use
 - Uber's NullAway is a big competitor
 - "No false negatives in practice"
 - Type systems for other problems (e.g., resource leaks) becoming more popular
- Big problem: getting people to use them
 - Annotation burden for legacy code is the most serious issue
 - Now I'm going to talk about some ongoing research work

Solution: Type Inference

• Traditional type inference: **constraint solving**

Solution: Type Inference

- Traditional type inference: constraint solving
 - Problem: need a new constraint system for each type system

Solution: Type Inference

- Traditional type inference: constraint solving
 - **Problem**: need a new constraint system for each type system
- Recent work has proposed 3 alternative type inference techniques:
 - Checker Framework Whole-Program Inference (ASE 2023)
 - NullAway Annotator (FSE 2023)
 - NullGTN (arxiv 2024)

• Most pluggable typecheckers already implement local type inference within method bodies

- Most pluggable typecheckers already implement local type inference within method bodies
 - Reduces user effort: no annotations on local variables

- Most pluggable typecheckers already implement **local type inference** within method bodies
 - Reduces user effort: no annotations on local variables
 - Implemented as intra-procedural dataflow analysis

- Most pluggable typecheckers already implement **local type inference** within method bodies
 - Reduces user effort: no annotations on local variables
 - Implemented as intra-procedural dataflow analysis
 - Typically implemented at the **framework level**

- Most pluggable typecheckers already implement **local type inference** within method bodies
 - Reduces user effort: no annotations on local variables
 - Implemented as intra-procedural dataflow analysis
 - Typically implemented at the framework level
- Key idea: iteratively run local inference and propagate results

- Most pluggable typecheckers already implement local type inference within method bodies
 - Reduces user effort: no annotations on local variables
 - Implemented as intra-procedural dataflow analysis
 - Typically implemented at the framework level
- Key idea: iteratively run local inference and propagate results
 - Advantage: works with any typechecker built on a framework "for free" (no per-typechecker code required)

```
class DataHandler {
  Data data = null;
```

```
void lazyInit() {
    if (data == null)
        data = new Data(...);
}
String serialize() {
    lazyInit();
    return data.toString(); // safe
```

```
class DataHandler {
  +@Nullable Data data = null;
```

_ local dataflow proves that null can flow into this field

```
void lazyInit() {
    if (data == null)
        data = new Data(...);
}
String serialize() {
    lazyInit();
    return data.toString(); // safe
```

```
class DataHandler {
  +@Nullable Data data = null;
  +@EnsuresNonNull(expr={"this.data"})
  void lazyInit() {
                                         CF also supports postcondition
    if (data == null)
                                         annotations (WPI advantage!)
      data = new Data(...);
  String serialize() {
    lazyInit();
    return data.toString(); // safe
```

```
class DataHandler {
  +@Nullable Data data = null;
  +@EnsuresNonNull(expr={"this.data"})
  void lazyInit() {
                                          CF also supports postcondition
    if (data == null)
                                          annotations (WPI advantage!),
      data = new Data(...);
                                          enabling verification of this code
  String serialize() {
    lazyInit();
    return data.toString(); // safe
```

NullAway Annotator (FSE 2023)

• Key idea: use warnings from the checker as a fitness function for annotations
NullAway Annotator (FSE 2023)

- Key idea: use warnings from the checker as a fitness function for annotations
- Iterative, bounded-depth search for annotation set that minimizes checker warnings
 - With some optimizations to reduce the search space

NullAway Annotator (FSE 2023)

- Key idea: use warnings from the checker as a fitness function for annotations
- Iterative, bounded-depth search for annotation set that minimizes checker warnings
 - With some optimizations to reduce the search space
- Only implementation is for NullAway (FSE '19, developed at Uber)

NullAway Annotator (FSE 2023)

- Key idea: use warnings from the checker as a fitness function for annotations
- Iterative, bounded-depth search for annotation set that minimizes checker warnings
 - With some optimizations to reduce the search space
- Only implementation is for NullAway (FSE '19, developed at Uber)
 Only infers @Nullable annotations

(same example from a few slides ago...)

NullAway Annotator: Example

```
class DataHandler {
```

```
Data data = null;
```

```
void lazyInit() {
    if (data == null)
        data = new Data(...);
}
String serialize() {
    lazyInit();
    return data.toString(); // safe
}
```

(same example from a few slides ago...

NullAway Annotator: Example

```
class DataHandler {
```

Data data = null;

```
void lazyInit() {
```

```
if (data == null)
```

```
data = new Data(...);
```

```
String serialize() {
```

```
lazyInit();
```

```
return data.toString(); // safe
```

Annotator considers adding an @Nullable annotation, but does not: doing so would not decrease the error count

(same example from a few slides ago...

NullAway Annotator: Example

```
class DataHandler {
```

```
Data data = null;
```

```
void lazyInit() {
```

```
if (data == null)
```

```
data = new Data(...);
```

Annotator considers adding an @Nullable annotation, but does not: doing so would not decrease the error count (new error here)

```
String serialize() {
```

lazyInit();

return data.toString(); // safe

NullAway Annotator: Another Example

```
class MyActivity {
 Name name;
 Address addr;
 MyActivity() {
    name = null; addr = null; }
  @Initializer
  void onStart() { name = defaultName(); }
  void doFirst() { name.showFirst(); }
  void doLast() {
    name.showLast();
    if (addr != null) addr.show();
```

```
Nima Kariminour, Justin Pham, Lazaro Clapp, and Manu Sridharan, Practical Inference of Nullability Types. In ESEC/ESE 2021
```

NullAway Annotator: Another Example

```
class MyActivity {
                                 Annotator does not add @Nullable to
  Name name; 🔶
                                 name, because it would add two new
  Address addr;
                                 errors (correct answer!)
 MvActivity() {
    name = null; addr = null; }
  @Initializer
  void onStart() { name = defaultName(); }
  void doFirst() { name.showFirst(); }
  void doLast() {
    name.showLast();
    if (addr != null) addr.show();
```

NullAway Annotator: Another Example

```
class MyActivity {
  Name name;
  +@Nullable Address addr;
                                      However, it does add
 MyActivity() {
                                      @Nullable to addr (always
    name = null; addr = null; }
                                      null-checked before use)
  @Initializer
  void onStart() { name = defaultName(); }
  void doFirst() { name.showFirst(); }
  void doLast() {
    name.showLast();
    if (addr != null) addr.show();
```

• Graph-based deep learning model

- Graph-based deep learning model
 - Inspired by recent success of similar ML-based techniques for inferring Python and TypeScript type annotations

- Graph-based deep learning model
 - Inspired by recent success of similar ML-based techniques for inferring Python and TypeScript type annotations
- Key idea: place annotations like a human would

- Graph-based deep learning model
 - Inspired by recent success of similar ML-based techniques for inferring Python and TypeScript type annotations
- Key idea: place annotations like a human would
- Trained on ~32k classes with at least one @Nullable annotation from GitHub
 - Data from many sources: checkers, documentation, etc.

- Graph-based deep learning model
 - Inspired by recent success of similar ML-based techniques for inferring Python and TypeScript type annotations
- Key idea: place annotations like a human would
- Trained on ~32k classes with at least one @Nullable annotation from GitHub
 - Data from many sources: checkers, documentation, etc.
- Only infers @Nullable

NullGTN: Example

```
public int sumLengths(
    int [] v1,
    int [] v2) {
        int result = 0;
        result += (v1 == null) ? 0 : v1.length;
        result += (v2 == null) ? 0 : v2.length;
        return result;
```

NullGTN: Example

public int sumLengths(int +@Nullable [] v1, int +@Nullable [] v2) { int result = 0; result += (v1 == null) ? 0 : v1.length; result += (v2 == null) ? 0 : v2.length; return result;

NullGTN can infer

Key Differences

Key Differences

- CF WPI *implementation* is the only one that supports multiple different pluggable type systems
 - Others claim they should generalize, but it's not evaluated

Key Difference Implication: comparison has to be focused on nullability type systems, for now

- CF WPI *implementation* is the only one that supports multiple different pluggable type systems
 - Others claim they should generalize, but it's not evaluated

Key Differences

- CF WPI *implementation* is the only one that supports multiple different pluggable type systems
 - Others claim they should generalize, but it's not evaluated
- Only NullGTN can possibly annotate **entrypoint parameters**
 - (Assuming no test cases)
 - In WPI's evaluation, this was the largest cause of missed human-written annotations (11%)

Pluggable Types in Practice

- <u>checkerframework.org</u> has a few dozen checkers
 - also practical to build your own!

Pluggable Types in Practice

- <u>checkerframework.org</u> has a few dozen checkers
 - also practical to build your own!
- Specialized checkers are now being built at tech companies
 e.g., NullAway at Uber

Pluggable Types in Practice

- <u>checkerframework.org</u> has a few dozen checkers
 - also practical to build your own!
- Specialized checkers are now being built at tech companies
 e.g., NullAway at Uber
- If we can get type inference working, maybe these will be in widespread use in a few more years :)