Concepts in Programming Languages

Alan Mycroft[1]

Computer Laboratory
University of Cambridge

CST Paper 7: 2018–2019 (Easter Term)

`www.cl.cam.ac.uk/teaching/1819/ConceptsPL/`

Practicalities

- Course web page:

  `www.cl.cam.ac.uk/teaching/1819/ConceptsPL/`

  with *lecture slides*, *exercise sheet* and *reading material*.
  These slides play two roles – both "lecture notes" and
  "presentation material"; not every slide will be lectured in
  detail.

- There are various code examples (particularly for
  JavaScript and Java applets) on the 'materials' tab of the
  course web page.

- One exam question.

- The syllabus and course has changed somewhat from that
  of 2015/16. I would be grateful for comments on any
  remaining 'rough edges', and for views on material which is
  either over- or under-represented.

Main books

- J. C. Mitchell. *Concepts in programming languages*. Cambridge University Press, 2003.

- T. W. Pratt and M. V. Zelkowitz. *Programming Languages: Design and implementation* (3RD EDITION). Prentice Hall, 1999.

- ⋆ M. L. Scott. *Programming language pragmatics* (4TH EDITION). Elsevier, 2016.

- R. Harper. *Practical Foundations for Programming Languages*. Cambridge University Press, 2013.

Context:
so many programming languages

Peter J. Landin: "The Next 700 Programming Languages",
CACM *(published in 1966!)*.

Some programming-language 'family trees' (too big for slide):
http://www.oreilly.com/go/languageposter
http://www.levenez.com/lang/
http://rigaux.org/language-study/diagram.html
http://www.rackspace.com/blog/
infographic-evolution-of-computer-languages/

Plan of this course: pick out interesting programming-language
concepts and major evolutionary trends.

Topics

I. Introduction and motivation.

Part A: Meet the ancestors

II. The first *procedural* language: FORTRAN (1954–58).

III. The first *declarative* language: LISP (1958–62).

IV. *Block-structured* languages: Algol (1958–68), Pascal (1970).

V. *Object-oriented* languages: Simula (1964–67), Smalltalk (1972).

Part B: Types and related ideas

VI. *Types* in programming languages: ML, Java.

VII. *Scripting Languages*: JavaScript.

VIII. *Data abstraction* and *modularity*: SML Modules.

Part C: Distributed concurrency, Java lambdas, Scala, Monads

IX. Languages for *concurrency and parallelism*.

X. Functional-style programming meets object-orientation.

XI. Miscellaneous concepts: Monads, GADTs.

Topic I

Introduction and motivation

Goals

- Critical *thinking* about programming languages.
  - **?** What is a programming language!?
- *Study* programming languages.
  - Be familiar with basic language *concepts*.
  - Appreciate trade-offs in language *design*.
- Trace *history*, appreciate *evolution* and diversity of *ideas*.
- Be prepared for new programming *methods*, *paradigms*.

Why study programming languages?

- ▶ To improve the ability to develop effective algorithms.
- ▶ To improve the use of familiar languages.
- ▶ To increase the vocabulary of useful programming constructs.
- ▶ To allow a better choice of programming language.
- ▶ To make it easier to learn a new language.
- ▶ To make it easier to design a new language.
- ▶ To simulate useful features in languages that lack them.
- ▶ To make better use of language technology wherever it appears.

What makes a good language?

- Clarity, simplicity, and unity.
- Orthogonality.
- Naturalness for the application.
- Support of abstraction.
- Ease of program verification.
- Programming environments.
- Portability of programs.
- Cost of use.
    - Cost of execution.
    - Cost of program translation.
    - Cost of program creation, testing, and use.
    - Cost of program maintenance.

What makes a language successful?

- ▶ Expressive power.
- ▶ Ease of use for the novice.
- ▶ Ease of implementation.
- ▶ Standardisation.
- ▶ Many useful libraries.
- ▶ Excellent compilers (including open-source)
- ▶ Economics, patronage, and inertia.

Note the recent trend of big companies to create/control their own languages: C# (Microsoft), Hack (Facebook), Go (Google), Objective-C/Swift (Apple), Rust (Mozilla) and perhaps even Python (Dropbox hired Guido van Rossum).

? Why are there so many languages?

- ▶ Evolution.
- ▶ Special purposes.
- ▶ No one language is good at expressing all programming styles.
- ▶ Personal preference.

? What makes languages evolve?

- ▶ Changes in hardware or implementation platform
- ▶ Changes in attitudes to safety and risk
- ▶ New ideas from academic or industry

Motivating purpose and language design

A specific purpose or motivating application provides *focus* for language design—what features to include and (harder!) what to leave out. E.g.

- ▶ Lisp: symbolic computation, automated reasoning
- ▶ FP: functional programming, algebraic laws
- ▶ BCPL: compiler writing
- ▶ Simula: simulation
- ▶ C: systems programming [Unix]
- ▶ ML: theorem proving
- ▶ Smalltalk: Dynabook [1970-era tablet computer]
- ▶ Clu, SML Modules: modular programming
- ▶ C++: object orientation
- ▶ Java, JavaScript: Internet applications

Program execution model

Good language design presents *abstract machine*.

- ▶ Fortran: Flat register machine; memory arranged as linear array
- ▶ Lisp: cons cells, read-eval-print loop
- ▶ Algol family: stack of activation records; heap storage
- ▶ BCPL, C: underlying machine + abstractions
- ▶ Simula: Object references
- ▶ FP, ML: functions are basic control structure
- ▶ Smalltalk: objects and methods, communicating by messages
- ▶ Java: Java virtual machine

Classification of programming languages

See `en.wikipedia.org/wiki/Programming_paradigm`
for more detail:

- ► Imperative
  procedural   C, Ada, **Pascal**, **Algol**, **Fortran**, . . .
  object-oriented   **Scala**, C#,Java, **Smalltalk**, **SIMULA**, . . .
  scripting   Perl, Python, PHP, JavaScript, . . .

- ► Declarative
  functional   Haskell, SML, Lisp, Scheme, . . .
  logic   Prolog
  dataflow   Id, Val
  constraint-based   spreadsheets
  template-based   XSLT

Language standardisation

Consider: `int i; i = (1 && 2) + 3 ;`

? Is it valid C code? If so, what's the value of `i`?

? How do we answer such questions!?

> **!** Read the reference manual (ISO C Standard).

> **!** Try it and see!

Other languages may have *informal* standards (defined by a particular implementation but what do we do if the implementation is improved?) or *proprietary* standards.

Language-standards issues

Timeliness. When do we standardise a language?

Conformance. What does it mean for a program to adhere to a standard and for a compiler to compile a standard?

Ambiguity and freedom to optimise – Machine dependence – Undefined behaviour.

A language standard is a treaty setting out the rights and obligations of the programmer and the implementer.

Obsolescence. When does a standard age and how does it get modified?

Deprecated features.

Language standards: unintended mis-specification

- *Function types in Algol 60*, see later.

- In language PL/1 the type `DEC(p,q)` meant a decimal number of p digits (at most 15) with q digits after the decimal point, so `9`, `8`, `3` all had type `DEC(1,0)`.
  Division was defined to so that `8/3` was `DEC(15,14)` with value `2.666 666 666 666 66`.
  But addition was defined so that adding these two was also of type `DEC(15,14)`, which meant that `9 + 8/3` gave `11.666 666 666 666 66`, which didn't fit. This gave either *overflow* or the wrong answer of `1.666 666 666 666 66`.

- A more recent example is C++11's "out of thin air" behaviour, whereby the ISO specification allows the value 42 to appear as the result of a program only involving assignments of 0 and 1.

Argh! Be careful how you specify a language.

Ultra-brief history

1951–55: Experimental use of expression compilers.

1956–60: **Fortran**, COBOL, **Lisp**, **Algol** 60.

1961–65: APL notation, Algol 60 (revised), SNOBOL, CPL.

1966–70: APL, SNOBOL 4, Fortran 66, BASIC, **SIMULA**, Algol 68, Algol-W, BCPL.

1971–75: **Pascal**, PL/1 (Standard), C, Scheme, Prolog.

1976–80: **Smalltalk**, Ada, Fortran 77, ML.

1981–85: Smalltalk-80, Prolog, Ada 83.

1986–90: C++, **SML**, Haskell.

1991–95: Ada 95, TCL, Perl.

1996–2000: Java, JavaScript

2000–05: C#, Python, Ruby, **Scala**.

1990– : Open/MP, MPI, Posix threads, Erlang, X10, MapReduce, Java 8 features.

For more information:

```
en.wikipedia.org/wiki/History_of_programming_
languages
```

$\sim$ *Part A* $\sim$

Meet the ancestors

Santayana 1906: "Those who cannot remember the past are condemned to repeat it."

∼ *Topic II* ∼

FORTRAN: A simple procedural language

Further reading:

▶ *The History of FORTRAN I, II, and III* by J. Backus. In
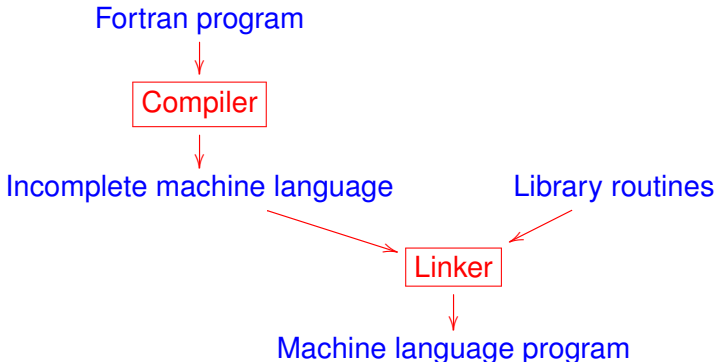 *History of Programming Languages* by R. L. Wexelblat.

  Academic Press, 1981.

FORTRAN = FORmula TRANslator (1957)

▶ Developed (1950s) by an IBM team led by John Backus:

"As far as we were aware, we simply made up the language as we went along. We did not regard *language design* as a difficult problem, merely a simple prelude to the real problem: designing a compiler which could produce efficient programs."

▶ The first high-level programming language to become widely used. At the time the utility of any high-level language was open to question(!), and complaints focused on efficiency of generated code. This heavily influenced the design, orienting it towards execution efficiency.

▶ Standards: 1966, 1977 (FORTRAN 77), 1990 (Fortran 90, spelling change), ..., 2010 (Fortran 2008).

▶ Remains main language for scientific computing.

▶ Easier for a compiler to optimise than C.

Overview: Compilation

Fortran program = main program + subprograms

- ▶ Each is *compiled separately* from all others.
  (Originally no support for cross-module checking, still really true for C and C++.)
- ▶ Translated programs are *linked* into final executable form.

Fortran program

↓

Compiler

↓

Incomplete machine language   Library routines

↘   ↙

Linker

↓

Machine language program

Overview: Data types and storage allocation

- Numerics: Integer, real, complex, double-precision real.

- Boolean. | called logical |

- Arrays. | of fixed declared length |

- Character strings. | of fixed declared length |

- Files.

- Fortran 90 added 'derived data types' (like C structs).

Allocation:

- Originally all *storage was allocated statically* before program execution, even local variables (as early Fortran lacked recursion—machines often lacked the index registers needed for a cheap stack—and we didn't realise how useful stacks and recursion would be!).

- Modern Fortran has recursion and heap-allocated storage.

Overview

Control structures

▶ FORTRAN 66

Relied heavily on statement labels and `GOTO` statements, but did have `DO` (for) loops.

▶ FORTRAN 77

Added some modern control structures (*e.g.*, if-then-else blocks), but `WHILE` loops and recursion had to wait for Fortran 90.

▶ Fortran 2008

Support for concurrency and objects

Example (Fortran 77)

```
      PROGRAM MAIN
        PARAMETER (MaXsIz=99)
        REAL A(mAxSiZ)
 10     READ (5,100,END=999) K
100     FORMAT(I5)
          IF (K.LE.0 .OR. K.GT.MAXSIZ) STOP
          READ *,(A(I),I=1,K)
          PRINT *,(A(I),I=1,K)
          PRINT *,'SUM=',SUM(A,K)
          GO TO 10
999     PRINT *, "All Done"
        STOP
        END
```

Example (continued)

```
C SUMMATION SUBPROGRAM
      FUNCTION SUM(V,N)
        REAL V(N)
        SUM = 0.0
        DO 20 I = 1,N
          SUM = SUM + V(I)
  20      CONTINUE
        RETURN
        END
```

Example

Commentary

- Originally columns and lines were relevant, and blanks and upper/lower case are ignored except in strings. Fortran 90 added free-form and forbade blanks in identifiers (use the `.f90` file extension on Linux).

- Variable names are from 1 to 6 characters long (31 since Fortran 90), letters, digits, underscores only.

- Variables need not be declared: implicit naming convention determines their type, hence the old joke "GOD is REAL (unless declared INTEGER)"; good programming style uses `IMPLICIT NONE` to disable this.

- Programmer-defined constants (`PARAMETER`)

- Arrays: subscript ranges can be declared as (*lwb* : *upb*) with (*size*) meaning (1 : *size*).

- ▶ Data formats for I/O.
- ▶ Historically functions are compiled separately from the main program with no consistency checks. Failure may arise (either at link time or execution time) when subprograms are linked with main program.
  Fortran 90 provides a module system.
- ▶ Function parameters are uniformly transmitted by reference (like C++ '&' types).
  But Fortran 90 provided INTENT(IN) and INTENT(OUT) type qualifiers and Fortran 2003 added pass-by-value for C interoperability.
- ▶ Traditionally all allocation is done statically.
  But Fortran 90 provides dynamic allocation.
- ▶ A value is returned in a Fortran function by assigning a value to the name of a function.

Program consistency checks

- *Static type checking* is used in Fortran, but the checking is traditionally incomplete.

- Many language features, including arguments in subprogram calls and the use of COMMON blocks, were not statically checked (in part because subprograms are compiled independently).

- Constructs that could not be statically checked were often left unchecked at run time (e.g. array bounds).
  (An early preference for speed over ease-of-bug-finding still visible in languages like C.)

- Fortran 90 added a MODULE system with INTERFACEs which enables checking across separately compiled subprograms.

Parameter-passing modes

- Recall the terms *formal parameter* and *actual parameter*.

- Fortran provides *call-by-reference* as historic default, similar to *reference parameters* in C++ (the formal parameter becomes an *alias* to the actual parameter).

- Modern Fortran adds *call-by-value* as in C/C++.

- The language specifies that if a value is assigned to a formal parameter, then the actual parameter must be a variable. This is a traditional source of bugs as it needs cross-module compilation checking:

```
SUBROUTINE SUB(X,Y,Z)
  X = Y
  PRINT *,Z
  END

CALL SUB(-1.0, 1.0, -1.0)
```

We will say more about parameter-passing modes for other languages.

Fortran lives!

- ► Fortran is one of the first languages, and the only early language still in mainstream use (LISP dialects also survive, e.g. Scheme).

- ► Lots of CS people will tell you about all the diseases of Fortran based on Fortran 66, or Fortran 77.

- ► Modern Fortran still admits (most) old code for backwards compatibility, but also has most of the things you expect in a modern language (objects, modules, dynamic allocation, parallel constructs). There's even a proposal for "units of measure" to augment types.
  (Language evolution is preferable to extinction!)

- ► Don't be put off by the syntax—or what ill-informed people say.

～ *Topic III* ～

LISP: functions, recursion, and lists

LISP = LISt Processing (circa 1960)

- Developed in the late 1950s and early 1960s by a team led by John McCarthy at MIT. McCarthy described LISP as a "a scheme for representing the *partial recursive functions* of a certain class of symbolic expressions".

- Motivating problems: Symbolic computation (*symbolic differentiation*), logic (*Advice taker*), experimental programming.

- Software embedding LISP: Emacs (text editor), GTK (Linux graphical toolkit), Sawfish (window manager), GnuCash (accounting software).

- Current dialects: Common Lisp, Scheme, Clojure. Common Lisp is 'most traditional', Clojure is implemented on JVM.

Programming-language phrases

[This classification arose in the Algol 60 design.]

- ▶ *Expressions*. A syntactic entity that may be evaluated to determine its value.
- ▶ *Statement*. A command that alters the state of the machine in some explicit way.
- ▶ *Declaration*. A syntactic entity that introduces a new identifier, often specifying one or more attributes.

Some contributions of LISP

- ▶ LISP is an *expression-based* language.
  LISP introduced the idea of *conditional expressions*.
- ▶ Lists – dynamic storage allocation, *hd* (CAR) and *tl* (CDR).
- ▶ Recursive functions.
- ▶ Garbage collection.
- ▶ Programs as data.
- ▶ Self-definitional interpreter (LISP interpreter explained as a LISP program).

The core of LISP is pure functional, but *impure* (side-effecting) constructs (such as SETQ, RPLACA, RPLACD) were there from the start.

Overview

- ▶ Values in LISP are either *atoms*, e.g. `X`, `FOO`, `NIL`, or *cons cells* which contain two values.
  (Numbers are also atoms, but only literal atoms above can be used as variables below.)

- ▶ A LISP program is just a special case of a LISP value known as an *S-expression*. An S-expression is either an *atom* or a NIL-terminated list of S-expressions (syntactically written in parentheses and separated by spaces), e.g. `(FOO ((1 2) (3)) NIL (4 X 5))`.

- ▶ So right from the start programs are just data, so we can construct a value and then execute it as a program.

- ▶ LISP is a *dynamically typed* programming language, so heterogeneous lists like the above are fine.

- Programs represented as S-expressions are *evaluated* to give values, treating atoms as variables to be looked up in an environment, and lists as a function (the first element of the list) to be called along with its arguments (the rest of the list).

  Example:

  `(APPEND (QUOTE (FOO 1 Y)) (CONS 3 (CONS 'Z NIL))`

  evaluates to `(FOO 1 Y 3 Z)`.

  Note: the functions `CONS` and `APPEND` behaves as in ML, and the function `QUOTE` returns its argument unevaluated. Numbers and the atoms `NIL` and `T` (also used as booleans) evaluate to themselves.

- To ease typing LISP programs, `(QUOTE e)` can be abbreviated `'e`, see `'Z` above.
  This is done as part of the `READ` function which reads values (which of course can also be programs).

? How does one recognise a LISP program?

```
( defvar x 1 )                    val x = 1 ;
( defun g(z) (+ x z) )            fun g(z) = x + z ;
( defun f(y)                      fun f(y)
    ( + ( g y )                     = g(y) +
        ( let                         let
            ( (x y) )                   val x = y
            (                           in
              g x )                       g(x)
        ) ) )                         end ;
( f (+ x 1) )                     f(x+1) ;
```

! It is full of brackets ("Lots of Irritating Silly Parentheses")!

Core LISP primitives

The following primitives give enough (Turing) power to construct any LISP function.

- ► `CONS`, `CAR`, `CDR`: cons, hd, tl.
- ► `CONSP`, `ATOM`: boolean tests for being a cons cell or atom.
- ► `EQ`, boolean equality test for equal atoms (but beware using it on large numbers which may be boxed, cf. Java boxing).
- ► `QUOTE`: return argument unevaluated.
- ► `COND`, conditional expression: e.g.
  `(COND ((CONSP X) (CDR X)) (T NIL))`.

Note that most LISP functions evaluate their arguments before they are called, but ('special forms') `QUOTE` and `COND` do not.

Along with a top-level form for recursion (LISPs vary):

- ► `DEFUN`, e.g. `(DEFUN F (X Y Z) ⟨body⟩)`

These give Turing power.

Core LISP primitives (2)

Example:

```
(defun subst (x y z)
    (cond  ((atom z)  (cond ((eq z y) x) (T z)))
           (T         (cons (subst x y (car z))
                            (subst x y (cdr z)))))
    )
)
```

Life is simpler with arithmetic and higher-order functions:

- ▶ +, −, < etc. E.g. `(+ X 1)`
- ▶ `LAMBDA`: e.g. `(LAMBDA (X Y) (+ X Y))`
- ▶ `APPLY`: e.g. `(APPLY F '(1 2))`

Note `LAMBDA` is also a special form.

Static and dynamic scope (or binding)

There are two main rules for finding the declaration of an identifier:

- ▶ *Static scope*. A identifier refers to the declaration of that name that is declared in the closest enclosing scope of the program text.
- ▶ *Dynamic scope*. A global identifier refers to the declaration associated with the most recent environment.

Historically, LISP was a dynamically scoped language; [Sethi pp.162] writes: when the initial implementation of Lisp was found to use dynamic scope, its designer, McCarthy[1981], "regarded this difficulty as just a bug".

Static and dynamic scope (example)

```
(defun main () (f 1))
(defun f (x) (g 2 (lambda () x)))
(defun g (x myfn) (apply myfn ()))
```

The question is whether the interpreter looks up the free variable x of the lambda in its static scope (getting 1), or in the (dynamic) scope at the time of the call (getting 2).

Newer dialects of LISP (such as Common Lisp and Scheme) use static scoping for this situation.

However, top-level defvar in Common Lisp can be used to mark a variable as using dynamic binding.

Abstract machines

The terminology *abstract machine* is generally used to refer to an idealised computing device that can execute a specific programming language directly. Systems people use *virtual machine* (as in JVM) for a similar concept.

- ▶ The original Fortran abstract machine can be seen as having only *static* (global, absolute address) *storage* (without a stack as there was no recursion), allocated before execution.

- ▶ The abstract machine corresponding to McCarthy's LISP implementation had a *heap* and *garbage collection*.[2] However, static locations were used to store variables (variables used by a recursive function were saved on entry and restored on exit, leading to dynamic scope).

- ▶ We had to wait for Algol 60 for *stacks* as we know them.

---

[2] He worried whether this term would pass the publication style guides of the day.

Programs as data

- ► One feature that sets LISP apart from many other languages is that it is possible for a program to build a data structure that represents an expression and then evaluates the expression as if it were written as part of the program. This is done with the function EVAL. *But* the environment used is that of the caller of EVAL so problems can arise if the expression being evaluated contains free variables (see 'call by text' below).

- ► McCarthy showed how a *self-definitional* (or *meta-circular*) interpreter for LISP could be written in LISP. See

> www.cl.cam.ac.uk/teaching/current/
> ConceptsPL/jmc.pdf

for Paul Graham's article re-telling McCarthy's construction.

Parameter passing in LISP

▶ Function parameters are transmitted either all *by value* or all *by text* (unevaluated expressions); only built-in functions (such as QUOTE, LAMBDA, COND) should really use pass-by-text. *Why: because we need a special variant of EVAL to evaluate the arguments to COND in the calling environment, and similarly need to capture the free variable of a LAMBDA in the environment of the caller*.

▶ The *actual parameters* in a function call are always expressions, represented as list structures.

▶ Note that pass by text (using either a special form, or explicit programmer use of QUOTE, and with EVAL to get the value of an argument) *resembles* call by name[3] (using LAMBDA to pass an unevaluated expression, and with APPLY to get the value of an argument), but is only equivalent if the EVAL can evaluate the argument in the environment of the function call!

---

[3]See Algol 60 later in the course

Calling by value, by name, by text – dangers of `eval`

Example: Consider the following function definitions

```
(defun CountFrom(n) (CountFrom(+ n 1)))
(defun myEagerIf (x y z) (cond (x y) (T z)))
(defun myNameIf (x y z) (cond (x (apply y ()))
                              (T (apply z ()))))
(defun myTextIf (x y z) (cond (x (eval y)) (T (eval z))))
```

Now suppose the caller has variables `x`, `y` and `z` all bound to 7 and consider:

```
(COND ((eq x z) y) (T (CountFrom 0)))        gives 7
(myEagerIf (eq x z) y (CountFrom 0))         loops
(myNameIf (eq x z)  (lambda () y)
                    (lambda () (CountFrom 0)))   gives 7
(myTextIf (eq x z)  (quote y)
                    (quote (CountFrom 0)))    gives y not 7
```

Note: on Common Lisp implementations this exact behaviour only manifests if we say `(defvar y 7)` before defining `myTextIf`; but all uses of `eval` are risky.

～ *Topic IV* ～

Block-structured procedural languages
Algol, Pascal and Ada

Parameter passing

Note: 'call-by-XXX' and 'pass-by-XXX' are synonymous.

- In *call by value*, the actual parameter is evaluated. The value of the actual parameter is then stored in a new location allocated for the formal parameter.
- In *call by reference*, the formal parameter is an alias to the actual parameter (normally achieved by a 'hidden' pointer).
- In *call by value/result* (IN/OUT parameters in Ada) the formal is allocated a location and initialised as in call-by-value, but its final value is copied back to the actual parameter on procedure exit.
- Algol 60 introduced *call by name*, see later.

Exercise: write code which gives different results under call-by reference and call-by-value/result.

Example Pascal program: The keyword `var` indicates call by reference.

```
program main;
begin
  function f(var x, y: integer): integer;
    begin
      x := 2;
      y := 1;
      if x = 1 then f := 3 else f:= 4
    end;

  var z: integer;
  z := 0;
  writeln( f(z,z) )
end
```

The difference between call-by-value and call-by-reference is important to the programmer in several ways:

- **Side effects.** Assignments inside the function body may have different effects under pass-by-value and pass-by-reference.
- **Aliasing.** Aliasing may occur when two parameters are passed by reference or one parameter passed by reference has the same location as a global variable of the procedure.
- **Efficiency.** Beware of passing arrays or large structures by value.

Examples:

- A parameter in Pascal is normally passed by value. It is passed by reference, however, if the keyword `var` appears before the declaration of the formal parameter.

  ```
  procedure proc(x: Integer; var y: Real);
  ```

- The only parameter-passing method in C is call-by-value; however, the effect of call-by-reference can be achieved using pointers. In C++ true call-by-reference is available using *reference parameters*.

- Ada supports three kinds of parameters:
  1. `in` parameters, corresponding to value parameters;
  2. `out` parameters, corresponding to just the copy-out phase of call-by-value/result; and
  3. `in out` parameters, corresponding to either reference parameters or value/result parameters, at the discretion of the implementation.

Parameter passing
Pass/Call-by-name

The Algol 60 report describes *call-by-name*.

► Such actual parameters are (re-)evaluated every time the formal parameter is used—this evaluation takes place in the scope of the caller (cf. the Lisp discussion).

► This is like beta-reduction in lambda calculus, but can be very hard to understand in the presence of side-effects.

► Lazy functional languages (e.g. Haskell) use this idea, but the absence of side-effects enables re-evaluation to be avoided in favour of caching.

Parameters: positional vs. named

In most languages actual parameters are matched to formals
by position but some languages additionally allow matching by
name and also allow optional parameters, e.g. Ada and to
some extent C++.

```
procedure Proc(Fst: Integer:=0; Snd: Character);

Proc(24,'h');
Proc(Snd => 'h', Fst => 24);
Proc(Snd => 'o');
```

ML can simulate named parameters by passing a record
instead of a tuple.

Algol

- The Algol-like programming languages evolved in parallel with the LISP family of languages, beginning with Algol 58 and Algol 60 in the late 1950s.
- The main characteristics of the Algol family are:
  - the familiar semicolon-separated sequence of statements,
  - block structure,
  - functions and procedures, and
  - static typing.

ALGOL IS DEAD BUT ITS DESCENDANTS LIVE ON!

- Ada, C, C++, Java etc.

Algol innovations

- Use of BNF syntax description.
- Block structure.
- Scope rules for local variables.
- Dynamic lifetimes for variables.
- Nested `if-then-else` expressions and statements.
- Recursive subroutines.
- Call-by-value and call-by-name arguments.
- Explicit type declarations for variables.
- Static typing.
- Arrays with dynamic bounds.

Algol 60

Features

- Simple statement-oriented syntax.
- Block structure.
    - *blocks* contain declarations and exectable statements delimited by begin and end markers.
    - May be nested, declaration visibility: scoping follows lambda calculus (Algol had no objects so no richer O-O visibility from inheritance as well as nesting).
- Recursive functions and stack storage allocation.
- Fewer ad-hoc restrictions than previous languages (*e.g.*, general expressions inside array indices, procedures that could be called with procedure parameters).
- A primitive *static type system*, later improved in Algol 68 and Pascal.

Algol 60

Some trouble-spots

- The Algol 60 type discipline had some shortcomings.
  For instance:
  - The type of a procedure parameter to a procedure does not
    include the types of parameters.

    ```
    procedure myapply(p, x)
      procedure p; integer x;
      begin p(x);
      end;
    ```

  - An array parameter to a procedure is given type array,
    without array bounds.
- Algol 60 was designed around two parameter-passing
  mechanisms, *call-by-name* and *call-by-value*.
  Call-by-name interacts badly with side effects; call-by-value
  is expensive for arrays.

Algol 68

- Algol 68 contributed a *regular, systematic type system*. The types (referred to as *modes* in Algol 68) are either *primitive* (`int`, `real`, `complex`, `bool`, `char`, `string`, `bits`, `bytes`, `semaphore`, `format`, `file`) or *compound* (`array`, `structure`, `procedure`, `set`, `pointer`). Type constructors could be combined without restriction, making the type system more systematic than previous languages.

- Algol 68 used a *stack* for local variables and *heap* storage. Heap data are explicitly allocated, and are reclaimed by *garbage collection*.

- Algol 68 parameter passing is by value, with pass-by-reference accomplished by pointer types. (This is essentially the same design as that adopted in C.)

Pascal (1970)

- Pascal is a *quasi-strong, statically typed* programming language.
  An important contribution of the Pascal *type system* is the rich set of data-structuring concepts: *e.g.* enumerations, subranges, records, variant records, sets, sequential files.

- The Pascal *type system* is more expressive than the Algol 60 one (repairing some of its loopholes), and simpler and more limited than the Algol 68 one (eliminating some of the compilation difficulties).

- Pascal was the first language to propose index checking. The index type (typically a sub-range of integer) of an array is part of its type.

- Pascal lives on (somewhat) as the Delphi language.

Pascal variant records

Variant records have a part common to all records of that type, and a variable part, specific to some subset of the records.

```
type kind = (unary, binary) ;
type                      | datatype
  UBtree = record         |     UBtree = mkUB of
    value: integer ;      |          int * UBaux
    case k: kind of       | and UBaux =
      unary: ^UBtree ;    |     unary of UBtree option
      binary: record      |   | binary of
        left: ^UBtree ;   |       UBtree option *
        right: ^UBtree    |       UBtree option;
      end
end ;
```

We use `UBaux` because ML datatype can only express variants at its top level. Note the use of `option` to encode `NULL`.

Pascal variant records introduced *weaknesses* into its type system.

- ▶ Compilers do not usually check that the value in the tag field is consistent with the state of the record.
- ▶ Tag fields are optional. If omitted, no checking is possible at run time to determine which variant is present when a selection is made of a field in a variant.

C still provides this model with `struct` and `union`. Modern languages provide safe constructs instead (think how a compiler can check for appropriate use):

- ▶ ML provides `datatype` and `case` to express similar ideas. In essence the constructor names provide the discriminator k but this is limited to being the first component of the record.
- ▶ Object-oriented languages provide *subclassing* to capture variants of a class.

See also the 'expression problem' discussion (slide 200).

～ *Topic V* ～

Object-oriented languages: concepts and origins
SIMULA and Smalltalk

Further reading for the interested:

- Alan Kay's "The Early History Of Smalltalk"
  http://worrydream.com/EarlyHistoryOfSmalltalk/

Objects in ML !?

```
exception Empty ;
fun newStack(x0)
  = let val  stack = ref [x0]
    in ref{  push = fn(x)
               => stack := ( x :: !stack )    ,
            pop = fn()
               => case !stack of
                    nil => raise Empty
                  | h::t => ( stack := t; h )
    }end ;


exception Empty
val newStack = fn :
      'a -> {pop:unit -> 'a, push:'a -> unit} ref
```

Objects in ML !?

NB:

- ► **!** The *stack discipline of Algol*[4] for activation records fails!
- ► **?** Is ML an object-oriented language?
  ! Of course not!
  ? Why?

---

[4]The stack discipline is that variables can be allocated on entry to a procedure, and deallocated on exit; this conflicts with returning functions (closures) as values as we know from ML. Algol 60 allowed functions to be passed *to* a procedure but not returned *from* it (nor assigned to variables).

Basic concepts in object-oriented languages

Four main language concepts for object-oriented languages:

1. Dynamic lookup.
2. Abstraction.
3. Subtyping.
4. Inheritance.

Dynamic lookup

▶ *Dynamic lookup*[5] means that when a method of an object is called, the method body to be executed is selected dynamically, at run time, according to the implementation of the object that receives the message (as in Java or C++ *virtual* methods).

▶ For the idea of *multiple dispatch* (not on the course), rather than the Java-style (or single) dispatch, see

http://en.wikipedia.org/wiki/Multiple_dispatch

Abstraction

▶ *Abstraction* means that implementation details are hidden inside a program unit with a specific interface. For objects, the interface usually consists of a set of methods that manipulate hidden data.

---

[5]Also called 'dynamic dispatch' and occasionally 'dynamic binding' (but avoid the latter term as 'dynamic scoping' is quite a different concept).

Subtyping

- *Subtyping* is a relation on types that allows values of one type to be used in place of values of another. Specifically, if an object `a` has all the functionality of another object `b`, then we may use `a` in any context expecting `b`.

- The basic principle associated with subtyping is *substitutivity*: If `A` is a subtype of `B`, then any expression of type `A` may be used without type error in any context that requires an expression of type `B`.

Inheritance

- *Inheritance* is the ability to reuse the definition of one kind of object to define another kind of object.
- The importance of inheritance is that it saves the effort of duplicating (or reading duplicated) code and that, when one class is implemented by inheriting from another, changes to one affect the other. This has a significant impact on code maintenance and modification.

NB: although Java treats subtyping and inheritance as synonyms, it is quite possible to have languages which have one but not the other.
A language might reasonably see `int` as a subtype of `double` but there isn't any easy idea of inheritance here.

Behavioural Subtyping – 'good subclassing'

Consider two classes

```
class MyIntBag {
  protected ArrayList<Integer> xs;
  public void add(Integer x) { xs.add(x); }
  public int size() { return xs.size(); }
  // other methods here...
}
class MyIntSet extends MyIntBag
  @override
  public void add(Integer x) { if (!xs.contains(x))
                                   xs.add(x); }
}
```

Questions: Is `MyIntSet` a subclass of `MyIntBag`? A subtype?

Java says 'yes'. But should it?

Behavioural Subtyping – 'good subclassing' (2)

It shouldn't really be a subtype, because it violates *behavioural subtyping* – members of a subtype should have the same behaviour as the members of the supertype. Consider:

```
int foo(MyBag b) { int n = b.size;
                   b.add(42);
                   return b.size - n; }
```

For every `MyBag` this returns 1. However if I pass it a `MySet` already containing 42, then it returns 0.

So `MySet` shouldn't be subtype of `MyBag` as its values behave differently, e.g. results of `foo`. So properties of `MyBag` which I've proved to hold may no longer hold! We say that `MySet` is not a *behavioural subtype* of `MyBag`. (It is legally a subclass in Java, but arguably reflects 'bad programmer design'.)

[Liskov and Wing's paper "A Behavioral Notion of Subtyping" (1994) gives more details.]

History of objects
SIMULA and Smalltalk

- ▶ Objects were invented in the design of SIMULA and refined in the evolution of Smalltalk.
- ▶ SIMULA: The first object-oriented language.
  - ▶ Extremely influential as the first language with classes, objects, dynamic lookup, subtyping, and inheritance. Based on Algol 60.
  - ▶ Originally designed for the purpose of *simulation* by Dahl and Nygaard at the Norwegian Computing Center, Oslo,
- ▶ Smalltalk: A dynamically typed object-oriented language. Many object-oriented ideas originated or were popularised by the Smalltalk group, which built on Alan Kay's then-futuristic idea of the Dynabook (Wikipedia shows Kay's 1972 sketch of a modern tablet computer).

- A generic event-based simulation program (pseudo-code):

```
Q := make_queue(initial_event);
repeat
  select event e from Q
  simulate event e
  place all events generated by e on Q
until Q is empty
```

  naturally requires:
    - A data structure that may contain a variety of kinds
      of events.   ⤳ subtyping
    - The selection of the simulation operation according to
      the kind of event being processed.   ⤳ dynamic lookup
    - Ways in which to structure the implementation of
      related kinds of events.   ⤳ inheritance

SIMULA: Object-oriented features

- ▶ *Objects*: A SIMULA object is an activation record produced by call to a class.
- ▶ *Classes*: A SIMULA class is a procedure that returns a pointer to its activation record. The body of a class may initialise the objects it creates.
- ▶ *Dynamic lookup*: Operations on an object are selected from the activation record of that object.
- ▶ *Abstraction*: Hiding was not provided in SIMULA 67; it was added later and inspired the C++ and Java designs.
- ▶ *Subtyping*: Objects are typed according to the classes that create them. Subtyping is determined by class hierarchy.
- ▶ *Inheritance*: A SIMULA class could be defined, by *class prefixing*, to extend an already-defined class including the ability to override parts of the class in a subclass.

SIMULA: Sample code

```
CLASS POINT(X,Y); REAL X, Y;
  COMMENT***CARTESIAN REPRESENTATION
BEGIN
  BOOLEAN PROCEDURE EQUALS(P); REF(POINT) P;
    IF P =/= NONE THEN
      EQUALS := ABS(X-P.X) + ABS(Y-P.Y) < 0.00001;
  REAL PROCEDURE DISTANCE(P); REF(POINT) P;
    IF P == NONE THEN ERROR ELSE
      DISTANCE := SQRT( (X-P.X)**2 + (Y-P.Y)**2 );
END***POINT***
```

```
CLASS LINE(A,B,C); REAL A,B,C;
  COMMENT***Ax+By+C=0 REPRESENTATION
BEGIN
  BOOLEAN PROCEDURE PARALLELTO(L); REF(LINE) L;
    IF L =/= NONE THEN
      PARALLELTO := ABS( A*L.B - B*L.A ) < 0.00001;
  REF(POINT) PROCEDURE MEETS(L); REF(LINE) L;
    BEGIN REAL T;
      IF L =/= NONE and ~PARALLELTO(L) THEN
        BEGIN
          ...
          MEETS :- NEW POINT(...,...);
        END;
    END;***MEETS***
  COMMENT*** INITIALISATION CODE (CONSTRUCTOR)
  REAL D;
  D := SQRT( A**2 + B**2 )
  IF D = 0.0 THEN ERROR ELSE
    BEGIN
      A := A/D;  B := B/D;  C := C/D;
    END;
END***LINE***
```

[Squint and it's almost Java!]

SIMULA: Subclasses and inheritance

SIMULA syntax: `POINT CLASS COLOUREDPOINT`.
To create a `COLOUREDPOINT` object, first create a `POINT` object
(activation record) and then add the fields of `COLOUREDPOINT`.

Example:

```
POINT CLASS COLOUREDPOINT(C); COLOUR C;   << note arg
BEGIN
  BOOLEAN PROCEDURE EQUALS(Q); REF(COLOUREDPOINT) Q;
    ...;
END***COLOUREDPOINT***

REF(POINT) P;  REF(COLOUREDPOINT) CP;
P :- NEW POINT(1.0,2.5);
CP :- NEW COLOUREDPOINT(2.5,1.0,RED);   << note args
```

NB: SIMULA 67 did not hide fields. Thus anyone can change
the colour of the point referenced by `CP`:

```
    CP.C := BLUE;
```

SIMULA: Object types and subtypes

- All instances of a class are given the same *type*. The name of this type is the same as the name of the class.
- The class names (types of objects) are arranged in a *subtype* hierarchy corresponding exactly to the subclass hierarchy.
- The Algol-60-based type system included explicit REF types to objects.

Subtyping Examples – essentially like Java:

1. ```
   CLASS A;    A CLASS B;

   REF(A) a;   REF(B) b;

   a :- b; COMMENT***legal since B is
                   ***a subclass of A
   ...

   b :- a; COMMENT***also legal, but checked at
                   ***run time to make sure that
                   ***a points to a B object, so
                   ***as to avoid a type error
   ```

2. ```
   inspect a
      when B do b :- a
      otherwise ...
   ```

Smalltalk

- ▶ Extended and refined the object metaphor.
  - ▶ Used some ideas from SIMULA; but it was a completely new language, with new terminology and an original syntax.
  - ▶ Abstraction via private *instance variables* (data associated with an object) and public *methods* (code for performing operations).
  - ▶ Everything is an object; even a class. All operations are messages to objects. Dynamically typed.
  - ▶ Objects and classes were shown useful organising concepts for building an entire programming environment and system. Like Lisp, easy to build a self-definitional interpreter.
- ▶ Very influential, one can regard it as an object-oriented analogue of LISP: "Smalltalk is to Simula (or Java) as Lisp is to Algol".

Smalltalk Example

Most implementations of Smalltalk are based around an IDE environment ("click here to add a method to a class"). The example below uses GNU Smalltalk which is terminal-based with `st>` as the prompt.

```
st> Integer extend [
     myfact [
      self=0 ifTrue: [^1] ifFalse: [^((self-1) myfact) * self]
    ] ]
st> 5 myfact
120
```

- ▶ Send an `extend` message to (class) `Integer` containing `myfact` and its definition.
- ▶ The body of `myfact` sends the two-named-parameter message (`[^1]`, `[^((self-1) myfact) * self]`) to the boolean resulting from `self=0`, which has a method to evaluate one or the other.
- ▶ '^' means return and `(self-1) myfact` sends the message `myfact` to the Integer given by `self-1`.

Reflection, live coding and IDEs

Above, I focused on Smalltalk's API, e.g. the ability to dynamically add methods to a class.

Note that objects and classes are often shared (via *reflection*) between the interpreter and the executing program, so swapping the 'add' and 'multiply' methods in the Integer class may have rather wider effects than you expect!

While the API is of interest to implementers, often a user interface will be menu-driven using an IDE:

- ▶ click on a class, click on a method, adjust its body, etc.
- ▶ the reflective structure above gives us a way to control the interpreter and IDE behaviour by adjusting existing classes.
- ▶ this is also known as 'live coding', and gives quite a different feel to a system than (say) the concrete syntax for Smalltalk.

Remark: *Sonic Pi* is a live-coding scripting language for musical performance.

∼ *Part B* ∼

Types and related ideas

Safety, static and dynamic types, forms of polymorphism,
modules

$\sim$ *Topic VI* $\sim$

Types in programming languages

Additional Reference:

- **Sections 4.9 and 8.6** of *Programming languages: Concepts & constructs* by R. Sethi (2ND EDITION).

  Addison-Wesley, 1996.

Types in programming

- A *type* is a collection of computational entities that share some common property.
- Three main uses of types in programming languages:
  1. naming and organising concepts,
  2. making sure that bit sequences in computer memory are interpreted consistently,
  3. providing information to the compiler about data manipulated by the program.
- Using types to organise a program makes it easier for someone to read, understand, and maintain the program. Types can serve an important purpose in documenting the design and intent of the program.
- Type information in programs can be used for many kinds of optimisations.

Type systems

- A *type system* for a language is a set of rules for associating a type with phrases in the language.

- Terms strong and weak refer to the effectiveness with which a type system prevents errors. A type system is *strong* if it accepts only *safe* phrases. In other words, phrases that are accepted by a strong type system are guaranteed to evaluate without type error. A type system is *weak* if it is not strong.

- Perhaps the biggest language development since the days of Fortran, Algol, Simula and LISP has been how type systems have evolved to become more expressive (and perhaps harder to understand)—e.g. Java generics and variance later in this lecture.

Type safety

A programming language is *type safe* if no program is allowed to violate its type distinctions.

Safety	Example language	Explanation
Not safe	C, C++	Type casts, pointer arithmetic
Almost safe	Pascal	Explicit deallocation; dangling pointers
Safe	LISP, SML, Smalltalk, Java	Type checking

Type checking

A *type error* occurs when a computational entity is used in a manner that is inconsistent with the concept it represents. *Type checking* is used to prevent some or all type errors, ensuring that the operations in a program are applied properly. Some questions to be asked about type checking in a language:

- Is the type system *strong* or *weak*?
- Is the checking done *statically* or *dynamically*?
- How *expressive* is the type system; that is, amongst safe programs, how many does it accept?

Static and dynamic type checking

Run-time type checking:

- ► Compiler generates code, typically adding a 'tag' field to data representations, so types can be checked at run time.
- ► Examples: LISP, Smalltalk.
  (We will look at dynamically typed languages more later.)

Compile-time type checking:

- ► Compiler checks the program text for potential type errors and rejects code which does not conform (perhaps including code which would execute without error).
- ► Examples: SML, Java.
- ► Pros: faster code, finds errors earlier (safety-critical?).
- ► Cons: may restrict programming style.

NB: It is arguable that object-oriented languages use a mixture of compile-time and run-time type checking, see the next slide.

Java Downcasts

Consider the following Java program:

```
class A { ... };                  A a;
class B extends A { ... };        B b;
```

▶ Variable `a` has Java type `A` whose valid values are all those of *class* `A` along with those of all classes subtyping class `A` (here just class `B`).

▶ Subtyping determines when a variable of one type can be used as another (here used by assignment):

```
a = b;        √(upcast)
a = (A)b;     √(explicit upcast)
b = a;        ×(implicit downcast—illegal Java)
b = (B)a;     √(but needs run-time type-check)
```

▶ Mixed static and dynamic type checking!

See also the later discussion of subtype polymorphism.

Type equality

When type checking we often need to know when two types are *equal*. Two variants of this are *structural equality* and *nominal equality*.

Let *t* be a type expression (e.g. `int * bool` in ML) and make two type definitions

```
type n1 = t; type n2 = t;
```

- ▶ Type names `n1` and `n2` are structurally equal.
- ▶ Type names `n1` and `n2` are not nominally equal.
  Under nominal equality a name is only equal to itself.

We extend these definitions to *type expressions* using structural equivalence for all type constructors not involving names.

Examples:

- **Type equality in C/C++.** In C, type equality is structural for `typedef` names, but nominal for `struct`s and `union`s; note that in

  ```
  struct { int a; } x;    struct { int a; } y;
  ```

  there are two different (anonymously named) `struct`s so `x` and `y` have unequal types (and may not be assigned to one another).

- **Type equality in ML.** ML works very similarly to C/C++, structural equality except for datatype names which are only equivalent to themselves.

- **Type equality in Pascal/Modula-2.** Type equality was left ambiguous in Pascal. Its successor, Modula-2, avoided ambiguity by defining two types to be compatible if

  1. they are the same name, or
  2. they are `s` and `t`, and `s = t` is a type declaration, or
  3. one is a subrange of the other, or
  4. both are subranges of the same basic type.

Type declarations

We can classify type definitions `type n = ` *t* similarly:

Transparent. An alternative name is given to a type that can also be expressed without this name.

Opaque. A new type is introduced into the program that is not equal to any other type.

In implementation terms, type equality is just tree equality, except when we get to a type name in one or both types, when we either (transparently) look inside the corresponding definition giving a structural system, or choose insist that both nodes should be identical types names giving a nominal system.

Type compatibility and subtyping

- ► Type equality is symmetric, but we might also be interested in the possibly non-symmetric notion of *type compatibility* (e.g. can this argument be passed to this function, or be assigned to this variable).

- ► This is useful for subtyping, e.g. given Java `A a; B b; a = b;` which is valid only if `B` is a subtype of (or equal to) `A`.

Similarly we might want type definitions to have one-way transparency. Consider

```
type age = int; type weight = int;
var x : age,  y : weight,  z : int;
```

We might want to allow implicit casts of `age` to `int` but not `int` to `age`, and certainly not `x := y;`.

Polymorphism

Polymorphism [Greek: "having multiple forms"] refers to constructs that can take on different types as needed. There are three main forms in contemporary programming languages:

- **Parametric (or generic) polymorphism.** A function may be applied to any arguments whose types match a type expression involving type variables.
  Subcases: ML has *implicit* polymorphism, other languages have *explicit* polymorphism where the user must specify the instantiation (e.g. C++ templates, and the type system of "System F").

- **Subtype polymorphism.** A function expecting a given class may be applied to a subclass instead. E.g. Java, passing a `String` to a function expecting an `Object`. (Note we will return to *bounded subtype polymorphism* later.)

- ▶ **Ad-hoc polymorphism or overloading.** Two or more implementations with different types are referred to by the same name. E.g. Java, also addition is overloaded in SML (which is why `fn x => x+x` does not type-check).
  (Remark 1: Haskell's *type classes* enable rich overloading specifications. These allow functions be to implicitly applied to a range of types specified by a Haskell *type constraint*.)
  (Remark 2: the C++ rules on how to select the 'right' variant of an overloaded function are arcane.)

Although we've discussed these for function application, it's important to note that Java generics and ML parameterised datatypes (e.g. `Map<Key,Val>` and `'a list`) use the same idea for type constructors.

Type inference

- *Type inference* is the process of determining the types of phrases based on the constructs that appear in them.
- An important language innovation.
- A cool algorithm.
- Gives some idea of how other static analysis algorithms work.

Type inference in ML – idea

Idea: give every expression a new type variable and then emit constraints $\boxed{\alpha \approx \beta}$ whenever two types have to be equal. These constraints can then be solved with Prolog-style unification.

For more detail see Part II course: "Types".

Typing rule (variable):

$$\frac{}{\Gamma \vdash x : \tau} \quad \text{if } x : \tau \text{ in } \Gamma$$

Inference rule:

$$\frac{}{\Gamma \vdash x : \gamma} \quad \boxed{\gamma \approx \alpha} \text{ if } x : \alpha \text{ in } \Gamma$$

Typing rule (application):

$$\frac{\Gamma \vdash f : \sigma \to \tau \qquad \Gamma \vdash e : \sigma}{\Gamma \vdash f(e) : \tau}$$

Inference rule:

$$\frac{\Gamma \vdash f : \alpha \qquad \Gamma \vdash e : \beta}{\Gamma \vdash f(e) : \gamma} \qquad \boxed{\alpha \approx \beta \to \gamma}$$

Typing rule (lamda):

$$\frac{\Gamma, x : \sigma \vdash e : \tau}{\Gamma \vdash (\texttt{fn } x => e) : \sigma \to \tau}$$

Inference rule:

$$\frac{\Gamma, x : \alpha \vdash e : \beta}{\Gamma \vdash (\texttt{fn } x => e) : \gamma} \qquad \boxed{\gamma \approx \alpha \to \beta}$$

Example:

$$\frac{\dfrac{\surd}{f:\alpha_1, x:\alpha_3 \vdash f:\alpha_5} \qquad \dfrac{\dfrac{\surd}{f:\alpha_1, x:\alpha_3 \vdash f:\alpha_7} \qquad \dfrac{\surd}{f:\alpha_1, x:\alpha_3 \vdash x:\alpha_8}}{f:\alpha_1, x:\alpha_3 \vdash f(x):\alpha_6}}{\dfrac{f:\alpha_1, x:\alpha_3 \vdash f(f(x)):\alpha_4}{\dfrac{f:\alpha_1 \vdash \mathtt{fn}\ x => f(f(x)):\alpha_2}{\vdash \mathtt{fn}\ f => \mathtt{fn}\ x => f(f(x)):\alpha_0}}}$$

$$\boxed{\begin{aligned} \alpha_0 &\approx \alpha_1 -> \alpha_2\ , \quad \alpha_2 \approx \alpha_3 -> \alpha_4\ , \quad \alpha_5 \approx \alpha_6 -> \alpha_4\ , \quad \alpha_5 \approx \alpha_1\ , \\ \alpha_7 &\approx \alpha_8 -> \alpha_6\ , \quad \alpha_7 \approx \alpha_1\ , \quad \alpha_8 \approx \alpha_3 \end{aligned}}$$

Solution: $\alpha_0 = (\alpha_3 -> \alpha_3) -> \alpha_3 -> \alpha_3$

let-polymorphism

► The 'obvious' way to type-check `let val` x `=` e `in` e' `end` is to treat it as `(fn` x `=>` e' `)(` e `)`.

► But Milner invented a more generous way to type let-expressions (involving *type schemes*—types qualified with $\forall$ which are renamed with new type variables at every use).

► For instance

```
let val f = fn x => x in f(f) end
```

type checks, whilst

```
(fn f => f(f)) (fn x => x)
```

does not.

► Exercise: invent ML expressions e and e' above so that both forms type-check but have different types.

Surprises/issues in ML typing

The mutable type `'a ref` essentially has three operators

```
ref : 'a -> 'a ref
(!) : 'a ref -> 'a
(:=) : 'a ref * 'a -> unit
```

Seems harmless. But think about:

```
val x = ref [];  (* x : ('a list) ref *)
x := 3 :: (!x);
x := true :: (!x);
print x;
```

We expect it to type-check, but it doesn't and trying to execute the code shows us it shouldn't type-check!

► ML type checking needs tweaks around the corners when dealing with non-pure functional code. See also the exception example on the next slide.

► This is related to the issues of *variance* in languages mixing subtyping with generics (e.g. Java).

Polymorphic exceptions

Example: Depth-first search for finitely-branching trees.

```
datatype
  'a FBtree = node of 'a * 'a FBtree list ;
fun dfs P (t: 'a FBtree)
  = let
      exception Ok of 'a;
      fun auxdfs( node(n,F) )
        = if P n then raise Ok n
          else foldl (fn(t,_) => auxdfs t) NONE F ;
    in
      auxdfs t handle Ok n => SOME n
    end ;
```

Type-checks to give:

```
val dfs = fn : ('a -> bool) -> 'a FBtree -> 'a option
```

This use of a polymorphic exception is OK.

But what about the following nonsense:

```
exception Poly of 'a ;    (*** ILLEGAL!!! ***)
(raise Poly true) handle Poly x => x+1 ;
```

When a *polymorphic exception* is declared, SML ensures that it is used with only one type (and not instantiated at multiple types). A similar rule (the 'value restriction') is applied to the declaration

```
val x = ref [];
```

thus forbidding the code on Slide 101.

▶ This is related to the issue of *variance* in languages like Java to which we now turn.

Interaction of subtyping and generics—variance

In Java, we have that `String` is a subtype of `Object`.

- ▶ But should `String[]` be a subtype of `Object[]`?
- ▶ And should `ArrayList<String>` be a subtype of `ArrayList<Object>`?
- ▶ What about `Function<Object,String>` being a subtype of `Function<String,Object>`?

Given generic `G` we say it is

- ▶ *covariant* if `G<String>` is a subtype of `G<Object>`.
- ▶ *contravariant* if `G<Object>` is a subtype of `G<String>`.
- ▶ *invariant* or *non-variant* if neither hold.
- ▶ variance is a per-argument property for generics taking multiple arguments .

But what are the rules?

Java arrays are covariant

The Java language decrees so. Hence the following code
type-checks.

```
String[] s = new String[10];
Object[] o;
o = s;                // decreed to be subtype
o[5] = "OK so far";
o[4] = new Integer(42);  // whoops!
```

However, it surely can't run! Indeed it raises exception
`ArrayStoreException` at the final line. Why?

▶ The last line would be unsound, so all *writes* into a Java
  array need to check that the item stored is a subtype of the
  array they are stored into. The type checker can't help.

▶ Note that there is no problem with *reads*.

▶ this is like the ML polymorphic `ref` and exception issue.

Java generics are invariant (by default)

The Java language decrees so. Hence the following code now fails to type-check.

```
ArrayList<String> s = new ArrayList(10);⁶
ArrayList<Object> o;
o = s;                    // fails to type-check
o.set(5,"OK so far");     // type-checks OK
o.set(4, new Integer(42)); // type-checks OK
```

So generics are safer than arrays. But covariance and contravariance can be useful.

▶ What if I have an immutable array, so that writes to it are banned by the type checker, then surely it's OK for it to be covariant?

---

⁵Legal note: it doesn't matter here, but to exactly match the previous array-using code I should populate the ArrayList with 10 NULLs. Real code would of course populate both arrays and ArrayLists with non-NULL values.

Java variance specifications

In Java we can have safe covariant generics using syntax like:

```
ArrayList<String> s = new ArrayList(10);
ArrayList<? extends Object> o;
o = s;                  // now type checks again
```

But what about reading and writing to `o`?

```
s.set(2,"Hello");
System.out.println((String)o.get(2)+"World"); //fine
o.set(4,"seems OK");    //faulted at compile-time
```

The trade is that the covariant `ArrayList o` cannot have its elements written to, in exchange for covariance.

► Java has *use-site variance* specifications: we can declare variance at every use of a generic.

► By contrast Scala has *declaration-site variance* which many find simpler (see later).

Java variance specifications (2)

Yes, there is a contravariant specification too (which allows writes but not reads):

```
ArrayList<? super String> ss;
```

So `ss` can be assigned values of type `ArrayList<String>` and `ArrayList<Object>` only.

For more information (beyond the current course) see:
en.wikipedia.org/wiki/Covariance_and_
contravariance_(computer_science)

~ *Topic VII* ~

Scripting Languages and Dynamic Typing

Scripting languages

"*A scripting language is a programming language that supports scripts; programs written for a special run-time environment that automate the execution of tasks that could alternatively be executed one-by-one by a human operator. Scripting languages are often interpreted (rather than compiled). Primitives are usually the elementary tasks or API calls, and the language allows them to be combined into more complex programs.*" [Wikipedia]

From this definition it's clear that many (but not all) scripting languages will have a rather ad-hoc set of features – and therefore tend to cause computer scientists to have strong views about their design (indeed it's arguable that we don't really teach any scripting languages on the Tripos).

Scripting languages (2)

- A *script* is just a program written for a scripting language, indeed the usage seems to be 'respectability driven': we write a Python *program* (respectable language) but a shell *script* or a Perl *script*.

- The definition is a bit usage-dependent: ML is a 'proper' stand-alone language, but was originally was the scripting language for creating proof trees in the Edinburgh LCF system (history: ML was the *meta-language* for manipulating proofs about the *object language* PPλ).

- Similarly, we've seen Lisp as a stand-alone language, but it's also a scripting language for the *Emacs* editor.

- But it's hard to see Java or C as a scripting language. Let's turn to why.

Scripting languages (3)

- *scripting language* means essentially "language with a REPL (read-evaluate-print loop)" as interface, or "language which can run interactively" (otherwise it's not so good for abbreviating a series of manual tasks, such API calls, into a single manual task).
- interactive convenience means that scripting languages generally are dynamically typed or use type inference, and execution is either interpretation or JIT compilation.

Incidentally, CSS and HTML are generally called *mark-up languages* to reflect their weaker (non-Turing powerful) expressivity.

Dynamically typed languages

We previously said:

- ▶ Using types to organise a program makes it easier for someone to read, understand, and maintain the program. Types can serve an important purpose in documenting the design and intent of the program.

So why would anyone want to lose these advantages?

- ▶ And why is JavaScript one of the most-popular programming languages on surveys like RedMonk?

Perhaps there is a modern-politics metaphor here, about elites using Java and ordinary programmers using JavaScript to be free of the shackles of types?

Questions on what we teach vs. real life

- ▶ why the recent rise in popularity of dynamically typed languages when they are slower and can contain type errors?
- ▶ why C/C++ still used when its type system is unsafe?
- ▶ why language support for concurrency is so 'patchwork' given x86 multi-core processors have existed since 2005?

We start by looking at a recent survey (by RedMonk, considering GitHub projects and StackOverflow questions) on programming language popularity.

RedMonk language rankings 2016

Rank	RedMonk (2016)	Rank	RedMonk (2016)
1	JavaScript	11	Shell
2	Java	12	R
3	PHP	13	Perl
4	Python	14	Scala
5=	C#	15	Go
5=	C++	16	Haskell
5=	Ruby	17	Swift
8	CSS	18	MATLAB
9	C	19	Visual Basic
10	Objective-C	20=	Clojure
		20=	Groovy

Note 1. CSS (Cascading Style Sheets) not really a language.
Note 2. Swift is Apple's language to improve on Objective C.
Note 3. Don't trust such surveys too much.

What are all these languages?

- ▶ Java, C, C++, C#, Objective-C, Scala, Go, Swift, Haskell
  These are small variants on what we already teach in Tripos; all are statically typed languages.

- ▶ JavaScript, PHP, Python, Ruby, R, Perl, MATLAB, Clojure, Groovy
  These are all dynamically typed languages (but notably Groovy is mixed as it tries to contains Java as a proper subset).

- ▶ Most of the dynamically typed languages have a principal use for scripting.

- ▶ Some static languages (Scala, Haskell, ML) are also used for scripting (helpful to have type inference and lightweight top-level phrase syntax).

Let's look at some of them.

JavaScript

- ▶ Originally called Mocha, shares little heritage with Java (apart from curly braces), but renamed 'for advertising reasons'.

- ▶ Designed and implemented by Brendan Eich in 10 days (see fuller history on the web).

- ▶ Dynamically typed, *prototype-based* object system (there was a design requirement not to use Java-like classes).

- ▶ Has both object-oriented and functional language features (including higher-order functions)

- ▶ Implemented within browsers (Java *applets* did this first, but security and commercial reasons make these almost impossible to use nowadays).

- ▶ *callback*-style approach to scheduling within browsers.

Browsers: Java Applets and JavaScript

There are two ways to execute code in a browser:

- [Java Applets] compile a Java application to JVM code and store it on a web site. A web page references it via the `<applet>` tag; the JVM code is run in a browser sandbox. In principle the best solution, but historically beset with security holes, and unloved by Microsoft and Apple for commercial reasons. Effectively dead in 2019.

- [JavaScript] the browser contains a JavaScript interpreter. JavaScript programs are embedded in source form in a web page with the `<script>` tag, and interpreted as the page is rendered (now sometimes JIT compiled). Interaction between browser and program facilitated with the DOM model. Questionable confidentiality guarantees.

https://en.wikipedia.org/wiki/Java_applet
https://en.wikipedia.org/wiki/JavaScript

The JavaScript language

Three language inspirations [quoting Douglas Crockford]:

- ▶ Java (for its syntax)
- ▶ Scheme (a Lisp dialect for its lexical closures)
- ▶ Self (a Smalltalk dialect with prototypes instead of classes)

What's 'wrong' with classes? Java has:

- ▶ Classes to define the basic qualities and behaviours of objects.
- ▶ Object instances as particular manifestations of a class.

So has Smalltalk.

Problems with classes

- When designing a system with the waterfall design method, then classes and objects are natural.
- In more flexible design styles (spiral and agile), then we might want to change the classes during program evolution.
- In Smalltalk this can be achieved by changing the structure of a class, or the class an object belongs to, by assignment.
- But it's not quite so simple, in that users inheriting a class may be surprised by the behaviour change of the superclass (the *fragile base class* problem)
- One partial solution is to use *prototypes* . . .

Prototypes instead of classes

- In JavaScript (and Self), there are no classes.
- When a constructor function is defined, it acquires a prototype property (the JS for 'field' and 'method') initially an empty object.
- When a constructor is called, created objects inherit from the prototype (i.e. field and method lookup look first in the object and then in the prototype).
- You can add properties to the prototype object, or indeed replace the it entirely with another object.
- This gives a prototype inheritance chain similar to a class inheritance chain.
- The Self community argue that this is a more flexible style of programming, but many JavaScript systems start by defining prototypes which play the role of classes, and copying from these plays the role of *new*.

Prototypes example

[Taken from the 'prototype' section of
https://www.w3schools.com/js/]

```
function Person(first, last, age, eye) {
    // Constructor: caller needs to use 'new'
    this.firstName = first;
    this.lastName = last;
    this.age = age;
    this.eyeColor = eye;
}
Person.prototype.nationality = "English";
var myFather = new Person("John", "Doe", 50, "blue");
```

After this, then `myFather.nationality` gives `"English"`.

Note the difference between updating a field in the prototype
and replacing the prototype with a new object.

Duck typing

If we have classes, we can use Java-style type-testing:

```
if (x instanceof Duck} { ... }
```

This is a form of *nominal* property. But what do we do without classes? One answer is *Duck Typing*:

> *"In other words, don't check whether it IS-a duck: check whether it QUACKS-like-a duck, WALKS-like-a duck, etc, etc, depending on exactly what subset of duck-like behaviour you need to play your language-games with." [Alex Martelli, comp.lang.python (2000)]*

More precisely, we take a *structural* rather than *nominal* view of being a duck: if an object has sufficiently many methods and fields (perhaps `walks()` and `quacks()` to be considered a Duck) then it *is* a Duck.
[Beware: JavaScript also has `instanceof` but this is merely shorthand for a test on an object's prototype.]

The JavaScript DOM model

JavaScript is most often used in browsers, so needs to be tightly connected with web pages. This is done by:

► The DOM (Document Object Model), a W3C standard.

► The browser converts the HTML into a tree of objects representing the HTML page structures. This is made available as the JavaScript variable *document*.

► Executing JavaScript code reads and modifies *document* which the browser renders on-screen.

► Example *button-click2* on 'materials' tab on course web page
(view the source as well as clicking).

JavaScript interaction within the browser

- ► Uses event-based GUI programming, based on callbacks, just like the Swing library for Java.

WebAssembly – turning full circle?

Because JavaScript is the only language universally supported by browsers (we're not counting plug-ins because of lack of ubiquitous support and security risks), we find that various compilers have developed JavaScript-targeting back-ends.

- ▶ E.g. js_of_caml, ghcjs, Scala.js.
- ▶ Try out OCaml in a browser: `http://ocsigen.org/js_of_ocaml/2.8.3/files/toplevel/index.html`
- ▶ But JavaScript is a terrible compiler target language (a Java Applet compiled to JVM would be more effective).
- ▶ Upcoming answer: WebAssembly. Machine-agnostic, safety-checked assembly language JIT'ted in the browser.
- ▶ Isn't this just a variant of applets?

Some other scripting Languages

- ▶ Python: popular for scripting scientific package API calls.
- ▶ Ruby: popularised by Heinemeier Hansson's *Rails* framework (server-side web-application programming and easy database integration).
- ▶ Clojure: Lisp meets the JVM.
- ▶ Typescript: Microsoft's typed version of JavaScript.
- ▶ R: a statistical scripting package.
- ▶ Lua: original for scripting for game engines.
- ▶ Matlab.

Gradual type systems

Dynamically typed languages are great (at least for small programs):

▶ *"You can write your whole prototype application in Python while a Java programmer is still writing his/her class hierarchy"*

Statically typed languages are better for large programs, due to the documentation and cross-module static checks, provided by types.

▶ Imagine the discussions at Facebook when management started to confront the issue that they had one million lines of PHP in their system.

So what you might want is:

▶ a language in which small programs can be written without types, but the IDE encourages you to add types faster than software rust overtakes your system – *gradual types*.

Gradual type systems (2)

"Gradual typing is a type system in which some variables and expressions may be given types and the correctness of the typing is checked at compile-time (which is static typing) and some expressions may be left untyped and eventual type errors are reported at run-time (which is dynamic typing). Gradual typing allows software developers to choose either type paradigm as appropriate, from within a single language. In many cases gradual typing is added to an existing dynamic language, creating a derived language allowing but not requiring static typing to be used. In some cases a language uses gradual typing from the start." [Wikipedia]

Example languages:

- ▶ Microsoft's TypeScript
- ▶ Facebook's Hack (for PHP)
- ▶ Cython and mypy (for Python)
- ▶ C#'s *dynamic* type.

Sometimes gradual types are called *optional types*.

Object-orientation: static vs dynamic, compiled vs interpreted

- ► Java and Scala are probably sweet-spots for ahead-of-time compilation for object-oriented languages, due to their static type systems.

- ► The dynamic features of Self and JavaScript (and indeed Smalltalk) mean that these languages need to be interpreted or JIT'ed.

- ► You might (non-examinably) think about what makes a nice type system possible and what is nice about a type system to enable ahead-of-time compilation.

$\sim$ *Topic VIII* $\sim$

Data abstraction and modularity
SML Modules

Additional Reference:

▶ **Chapter 7** of *ML for the working programmer* (2ND EDITION) by L. C. Paulson. CUP, 1996.

May also be useful:

▶ *Purely Functional Data Structures* by Chris Okasaki, Cambridge University Press. [Clever functional data structures, implemented in Haskell and SML Modules.]

▶ Previous versions of this course (available on the teaching pages) had more information on SML Modules.

▶ http://www.standardml.org/

Need for modules

► We want control name visibility in large systems.

► Name visibility is important because it expresses large-scale architecture (which components can use which other components). Excessive visibility increases the attack surface for malevolent use.

► Traditional Java does quite well on class-level visibility (private, protected, public) – but less well on package-level visibility (no formal way to control who might use a package-level export, or malevolent uses of reflection).

► Java 9 has just added a module system.

► We focus on the SML module system (seminal and well-designed).

There are also issues (only briefly mentioned) concerning *versioning*, when modules evolve over time and multiple versions are available, e.g. Linux packages or Windows DLLs.

The Core and Modules languages

SML consists of two sub-languages:

- The *Core* language is for *programming in the small*, by supporting the definition of types and variables denoting values of those types.

- The *Modules* language is for *programming in the large*, by grouping related Core definitions of types and variables into self-contained units, with descriptive interfaces.

The *Core* expresses details of *data structures* and *algorithms*. The *Modules* language expresses *software architecture*. Both languages are largely independent.

The Modules language

Writing a real program as an unstructured sequence of Core definitions quickly becomes unmanageable.

```
type nat = int
val zero = 0
fun succ x = x + 1
fun iter b f i =
   if i = zero then b
               else  f (iter b f (i-1))
...
(* thousands of lines later *)
fun even (n:nat) = iter true not n
```

The SML Modules language lets one split large programs into separate units with descriptive interfaces.

SML Modules

Signatures and structures

An *abstract data type* is a type equipped with a set of operations, which are the only operations applicable to that type.
Its representation can be changed without affecting the rest of the program.

- *Structures* let us *package* up declarations of related types, values, and functions.
- *Signatures* let us *specify* what components a structure must contain.
- Structures are like Core ML records `{a = 3, b = [4,5]}` but can contain definitions of types (and exceptions) in addition to definitions of values.

Signatures are to structures what types are to values.

Structures

In Modules, one can encapsulate a sequence of Core *type* and *value* definitions into a unit called a *structure*.
We enclose the definitions in between the keywords

                    struct ... end.

Example: A structure representing the natural numbers, as positive integers:

```
struct
  type nat = int
  val zero = 0
  fun succ x = x + 1
  fun iter b f i = if i = zero then b
                   else f (iter b f (i-1))
end
```

(We'll later show how to name this `IntNat`)

Concrete signatures

Signatures specify the 'types' of structures by listing the specifications of their components.

A signature consists of a *sequence* of component specifications, enclosed in between the keywords `sig ... end`.

```
sig  type nat = int
     val zero : nat
     val succ : nat -> nat
     val 'a iter : 'a -> ('a->'a) -> nat -> 'a
end
```

This signature fully describes the 'type' of `IntNat` on the previous slide.
(We'll later see signatures which *restrict* the names visible in a structure.)

The specification of type `nat` is *concrete*: it must be `int`.

Opaque signatures

On the other hand, the following signature

```
sig type nat
    val zero : nat
    val succ : nat -> nat
    val 'a iter : 'a -> ('a->'a) -> nat -> 'a
end
```

specifies structures that are free to use *any* implementation for
type nat (perhaps int, or word, or some recursive datatype).

This specification of type nat is *opaque*.

Naming structures and signatures, and dot notation

The `structure` keyword binds a structure to an identifier:

```
structure IntNat =
  struct type nat = int
         ...
         fun iter b f i = ...
  end
```

The `signature` keyword similarly names signatures:

```
signature NAT =
  sig type nat
        ...
      val 'a iter : 'a -> ('a->'a) -> nat -> 'a
  end
```

Dot notation accesses components of structures (and signatures):

```
fun even (n : IntNat.nat) = IntNat.iter true not n
```

NB: Type `IntNat.nat` is statically equal to `int`.
Value `IntNat.iter` dynamically evaluates to a closure.

Example: Polymorphic functional stacks.

```
signature STACK =
sig
  exception E
  type 'a reptype   (* <-- INTERNAL REPRESENTATION *)
  val new: 'a reptype
  val push: 'a -> 'a reptype -> 'a reptype
  val pop: 'a reptype -> 'a reptype
  val top: 'a reptype -> 'a
end ;

structure MyStack: STACK =
struct
  exception E ;
  type 'a reptype = 'a list ;
  val new = [] ;
  fun push x s = x::s ;
  fun split( h::t ) = ( h , t )
    | split _ = raise E ;
  fun pop s = #2( split s ) ;
  fun top s = #1( split s ) ;
end ;
```

```
val e = MyStack.new ;
val s0 = MyStack.push 0 e ;
val s01 = MyStack.push 1 s0 ;
val s0' = MyStack.pop s01 ;
MyStack.top s0' ;
```

Gives:

```
val e = [] : 'a MyStack.reptype
val s0 = [0] : int MyStack.reptype
val s01 = [1,0] : int MyStack.reptype
val s0' = [0] : int MyStack.reptype
val it = 0 : int
```

Signature matching

Q: When does a structure satisfy a signature?

A: The type of a structure *matches* a signature whenever it implements at least the components of the signature.

- ▶ The structure must *realise* (i.e. define) all of the opaque type components in the signature.
- ▶ The structure must *enrich* this realised signature, component-wise:
  - ⋆ every concrete type must be implemented equivalently;
  - ⋆ every value must have a more general type scheme;
  - ⋆ every structure must be enriched by a substructure.

Signature matching supports a form of *subtyping* not found in the Core language:

- ▶ A structure with more type, value, and structure components may be used where fewer components are expected.
- ▶ A value component may have a more general type scheme than expected.

Properties of signature matching

- The components of a structure can be defined in a different order than in the signature; names matter but ordering does not.

- A structure may contain more components, or components of more general types, than are specified in a matching signature.

- Signature matching is *structural*. A structure can match many signatures and there is no need to pre-declare its matching signatures (unlike "interfaces" in Java and C#).

- Although similar to record types, signatures actually play a number of different roles.

Using signatures to restrict access

We can use a *signature constraint* to provide a restricted view of a structure. E.g.

```
structure SubIntNat =
   IntNat : sig type nat
               val succ : nat->nat
               val iter : nat->(nat->nat)->nat->nat
            end
```

Here the constraint `str:sgn` prunes the structure `str` according to the signature `sgn`. So:

- ▶ `SubIntNat.zero` is faulted ("not a member")
- ▶ `SubIntNat.iter` is *less* polymorphic than `IntNat.iter`.

Transparency of "_ : _"

- Although the `_:_` operator can hide names, it does not conceal the definitions of opaque types. It is an *transparent signature constraint*.
- Thus, the fact that `SubIntNat.nat = IntNat.nat = int` remains *transparent*.
- For instance the application `SubIntNat.succ(~3)` is still well-typed, because `~3` has type `int` ... but `~3` is negative, so not a valid representation of a natural number!
- There is also an *opaque signature constraint* operator: "`_ :> _`"

Opaque signature constraints: "_ :> _"

Using the ':>' syntax, instead of the ':' syntax used earlier,
signature matching can also be used to enforce *data
abstraction* by hiding the identity of types:

```
structure AbsNat = IntNat :> sig
      type nat
      val zero: nat
      val succ: nat->nat
      val 'a iter: 'a->('a->'a)->nat->'a
    end
```

The constraint `str :> sgn` prunes `str` but also generates a
new, *abstract* type for each opaque type in `sgn`.

- The actual implementation of `AbsNat.nat` by `int` is *hidden*, so that `AbsNat.nat` $\neq$ `int`.
  `AbsNat` is just `IntNat`, but with a hidden type representation.

- `AbsNat` defines an *abstract datatype* of natural numbers: the only way to construct and use values of the abstract type `AbsNat.nat` is through the operations, `zero`, `succ`, and `iter`.
  E.g., the application `AbsNat.succ(~3)` is ill-typed: `~3` has type `int`, not `AbsNat.nat`. This is what we want, since `~3` is not a natural number in our representation.
  In general, abstractions can also prune and specialise components.

- Remark: abstractions generalise the expressive power of the Core SML `abstype` declaration (this works like `datatype`, but does not export its constructors, and so also hides the representation of a type).

Information hiding

Opaque signature constraints

```
structure MyOpaqueStack :> STACK = MyStack ;


val e = MyOpaqueStack.new ;
val s0 = MyOpaqueStack.push 0 e ;
val s01 = MyOpaqueStack.push 1 s0 ;
val s0' = MyOpaqueStack.pop s01 ;
MyOpaqueStack.top s0' ;
```

Gives:

```
val e = - : 'a MyOpaqueStack.reptype
val s0 = - : int MyOpaqueStack.reptype
val s01 = - : int MyOpaqueStack.reptype
val s0' = - : int MyOpaqueStack.reptype
val it = 0 : int
```

[Compare slide 141 which exposes `reptype` as `list`.]

Datatype and exception specifications

Signatures can also specify datatypes and exceptions:

```
structure PredNat =
  struct datatype nat = zero | succ of nat
    fun iter b f i = ...
    exception Pred
    fun pred zero = raise Pred
      | pred (succ n) = n
  end :> sig datatype nat = zero | succ of nat
            val iter: 'a->('a->'a)->(nat->'a)
            exception Pred
            val pred: nat -> nat (* raises Pred *)
        end
```

This means that clients can still pattern-match on datatype constructors, and handle exceptions.

Structures and signatures can be nested

```
structure IntNatAdd =
  struct
    structure Nat = IntNat
    fun add n m = Nat.iter m Nat.succ n
  end
  ...
signature Add =
  sig structure Nat: NAT
      val add: Nat.nat -> Nat.nat -> Nat.nat
  end
```

Then `IntNatadd: Add` holds (slide 139 defines `IntNat` and `NAT`).

To avoid nesting, it is also possible to directly `include` a signature identifier:

```
  sig  include NAT
      val add: nat -> nat -> nat
  end
```

SML Modules

- An SML *functor* is a structure that takes other structures as parameters.
  Functors are to structures what functions are to values, in SML they are first-order only (OCaml has higher-order versions).

- Functors can be *applied* to structures. The actual argument must *match* the signature of the formal parameter—so it can provide more components, of more general types but these are not available to the body of the functor.

- Functors let us write program units that can be combined in different ways. Functors can also express generic algorithms.

Functors

The functor `AddFun` below takes any implementation, `N`, of naturals and re-exports it with an addition operation.

```
functor AddFun(N:NAT) =
        struct
          structure Nat = N
          fun add n m = Nat.iter n (Nat.succ) m
        end
```

`AddFun` can be applied:

```
structure IntNatAdd = AddFun(IntNat)
structure AbsNatAdd = AddFun(AbsNat)
```

This definition of `IntNatAdd` evaluates to the same implementation as `IntNatAdd` seen earlier.

Above, `AddFun` is applied twice, but to arguments that differ in their implementation of type `nat` (`AbsNat.nat` $\neq$ `IntNat.nat`).

Example: Generic imperative stacks.

```
signature STACK =
 sig
  type itemtype
  val push: itemtype -> unit
  val pop: unit -> unit
  val top: unit -> itemtype
 end ;

exception E ;
functor Stack( T: sig type atype end ) : STACK =
 struct
  type itemtype = T.atype
  val stack = ref( []: itemtype list )
  fun push x
    = ( stack := x :: !stack )
  fun pop()
    = case !stack of [] => raise E
      | _::s => ( stack := s )
  fun top()
    = case !stack of [] => raise E
      | t::_ => t
 end ;
```

```
structure intStack
  = Stack(struct type atype = int end) ;

intStack.push(0) ;
intStack.top() ;
intStack.pop() ;
intStack.push(4) ;
```

Gives:

```
structure intStack : STACK

val it = () : unit
val it = 0 : intStack.itemtype
val it = () : unit
val it = () : unit
```

Why functors ?

Functors support:

Code reuse.

`AddFun` may be applied many times to different structures, reusing its body.

Code abstraction.

`AddFun` can be compiled before any argument is implemented.

Type abstraction.

`AddFun` can be applied to different types `N.nat`.

But there are various complications:

- Should functor application be *applicative* or *generative*?
- We need some way of specifying types as being *shared*.

Functors: generative or applicative?

The following functor is almost the identity functor, but *re-abstracts* its argument:

```
functor GenFun(N:NAT) =  N :> NAT
```

Now, suppose we apply it twice to the same argument:

```
structure X = GenFun(IntNat)
structure Y = GenFun(IntNat)
```

Question: are the types `X.nat` and `Y.nat` *compatible*?

- The *applicative* interpretation of functor application (used in OCaml) says "yes".
- The *generative* interpretation (used in SML) says "no". (Abstract types from the body of a functor are replaced by fresh types at each application. This is consistent with inlining the body of a functor at applications.)

This question is the tip of a very large iceberg (many papers).

Modules in other languages

Java also provides various structuring and hiding features:

- classes
- interfaces (with default implementations in Java 8)
- packages
- modules (from Java 9)

These offer independently developed, but overlapping, facilities—beyond this current course.

A separate issue is *versioning*. Linux packages (informal modules) come with versions, e.g. `libtest.so.1.0.1` and meta-information says which versions of a package are compatible for use within another versioned package. The similar Windows problem is 'DLL hell'. There is little linguistic support for this.

∼ *Part C* ∼

Linguistic ideas beyond Java and ML

Distributed concurrency, Scala, Monads

∼ *Topic IX* ∼

Languages for Concurrency and Parallelism

Sources of parallel computing

Five main sources:

1. Theoretical models: PRAM, BSP (complexity theory),
   CSP, CCS, π-calculus (semantic theory),
   Actors (programming model).
2. Multi-core CPUs (possibly heterogeneous—mobile
   phones).
3. Graphics cards (just unusual SIMD multi-core CPUs).
4. Supercomputers (mainly for scientific computing).
5. Cluster Computing, Cloud Computing.

NB: Items 2–5 conceptually only differ in processor-memory
communication.

Language groups

1. Theoretical models (PRAM, π-calculus, Actors, etc.).
2. C/C++ and roll-your-own using pthreads.
3. Pure functional programming ('free' distribution).
4. [Multi-core CPUs] Open/MP, Java (esp. Java 8), Open/MP, Cilk, X10.
5. [Graphics cards] CUDA (Nvidia), OpenCL (open standard).
6. [Supercomputers] MPI.
7. [Cloud Computing] MapReduce, Hadoop, Skywriting.
8. [On Chip] Verilog, Bluespec.

NB: Language features may fit multiple architectures.

Painful facts of parallel life

1. Single-core clock speeds have stagnated at around 3GHz for the last ten years. Moore's law continues to give more transistors (hence multi-core, many-core, giga-core).
2. Inter-processor *communication* is far far far more expensive than *computation* (executing an instruction).
3. Can't the compiler just take my old C/Java/Fortran (or ML/Haskell) program and, *you know*, parallelise it? Just another compiler optimisation? NO! (Compiler researchers' pipe-dream/elephants' graveyard.)

Takeaway: optimising performance requires exploiting parallelism, you'll have to program this yourself, and getting it wrong gives slow-downs and bugs due to races.
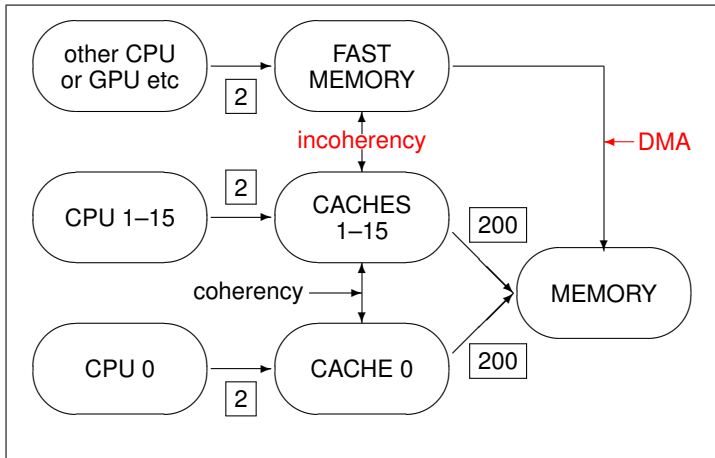
CPU → MEMORY
1 cycle
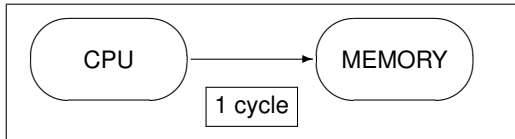
(This model was pretty accurate in 1985.)
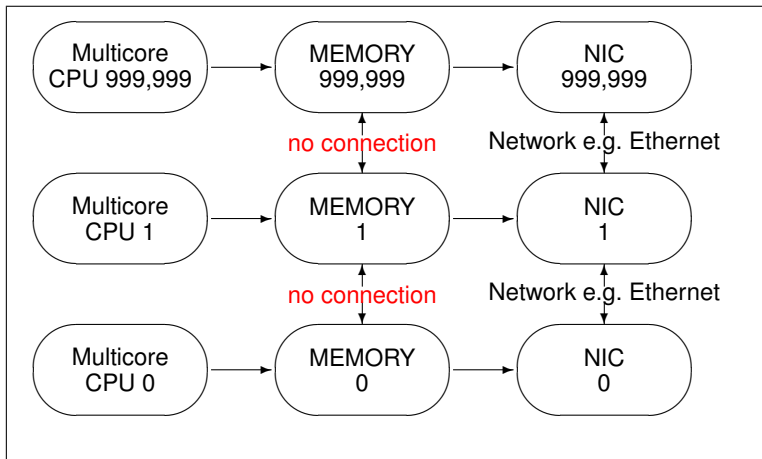
CPU → L1 cache → L2 cache → MEMORY
2   10   200
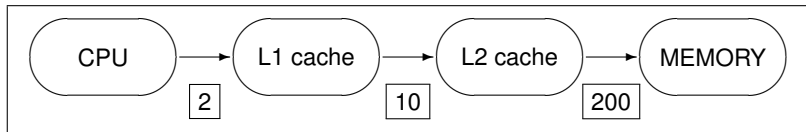
Multi-core-chip memory models

Today's model (cache simplified to one level):

A Compute Cluster or Cloud-Computing Server

Multicore CPU 999,999	→ MEMORY 999,999	→ NIC 999,999
	no connection	Network e.g. Ethernet
Multicore CPU 1	→ MEMORY 1	→ NIC 1
	no connection	Network e.g. Ethernet
Multicore CPU 0	→ MEMORY 0	→ NIC 0

(The sort of thing which Google uses.)

Lecture topic: what *programming abstractions*?

- ► We've got a large (and increasing) number of processors available for use within each 'device'
- ► This holds at multiple levels of scale (from on-chip to on-cloud). "Fractal" or "self-similar".
- ► Memory is local to processor units (at each scale)
- ► Communication (message passing) between units is much slower than computation.

Question: what are good programming abstractions for a system containing lots of processors?
Answer: no complete answer, rest of this lecture gives fragmentary answers.

What hardware architecture tells us

- Communication latency is far higher than instruction execution time (2–6 orders of magnitude)
- So, realistically a task needs to have need at least 10^4 instructions for it to be worth moving to another CPU.
- *Long-running independent computations* fit the hardware best.
- *"Shared memory" is an illusion.* At the lowest level it is emulated by message passing in the cache-coherency protocol.
- Often best to think of multi-core processors as distributed systems.

Communication abstractions for programming

- ▶ "Head in sand": What communication – I'm just using a multi-core CPU?
- ▶ "Principled head in sand": the restrictions in my programming language means I can leave this to someone else (or even the compiler).
- ▶ Just use TCP/IP.
- ▶ Shared memory, message passing, RMI/RPC?
- ▶ Communication is expensive, expose it to programmer (no lies about 'shared memory').

Reflection: does a language implicitly create a programmer mental model of communication?

Concurrent, Parallel, Distributed

These words are often used informally as near synonyms.

- ▶ Distributed systems have separate processors connected by a network, perhaps on-chip (multi-core)?

- ▶ 'Parallel' suggests multiple CPUs or even SIMD, but "parallel computation" isn't clearly different from "concurrency".

- ▶ Concurrent behaviour *can* happen on a single-core CPU (e.g. Operating System and threads), Theorists often separate 'true concurrency' (meaning parallel behaviour) from 'interleaving concurrency'.

SIMD, MIMD

- Most parallel systems nowadays are MIMD.
- GPUs (graphical processor units) are a bit of an exception; several cores execute the same instructions, perhaps conditionally based on a previous test which sets per-processor condition codes.
- Programming Languages for GPUs (OpenCL, CUDA) emphasise the idea of a single program which is executed by many tasks. A program can enquire to find out the numerical value of its task identifier, originally its (x, y) co-ordinate, to behave differently at different places (in addition to having separate per-task pixel data).

Theoretical model – process algebra

- ▶ CCS, CSP, Pi-Calculus (calculus = "simple programming language"). E.g.
- ▶ Atomic actions α, $\overline{\alpha}$, can communicate with each other or the world (non-deterministically if multiple partners offered). Internal communication gives special internal action τ.
- ▶ Behaviour $p ::= 0 \mid \alpha.p \mid p + p \mid p|p \mid X \mid rec\ X.p$ (Deadlock, prefixing, non-determinism, parallelism, recursive definitions, also (not shown) parameterisation/hiding and value-passing.)
- ▶ Typical questions: "is $\alpha.0|\beta.0$ the same as $\alpha.\beta.0 + \beta.\alpha.0$" and
  "what does it mean for two behaviours to be equal"

Part II course: 'Topics in Concurrency'.

Theoretical model – PRAM model

- PRAM: parallel random-access machine.
- N shared memory locations and P processors (both unbounded); each processor can access any location in one cycle.
- Execute instructions in lock-step (often SIMD, but MIMD within the model): fetch data, do operation, write result.
- Typical question: "given n items can we sort them in $O(n)$ time, or find the maximum in $O(1)$ time"
- BSP (bulk-synchronous parallel) model refines PRAM by adding costs for communication and synchronisation.

Part II course: 'Advanced Algorithms'.

Oldest idea: Threads

Java threads – either extend Thread or implement Runnable:

```
class PrimeRun implements Runnable {
    long minPrime;
    PrimeRun(long m) { minPrime = m; }
    public void run() {
        // compute primes larger than minPrime
    }
}
...
p = new PrimeRun(143);  // create a thread
new Thread(p).start();  // run it
```

Posix's pthreads are similar.

Threads, and what's wrong with them

▶ Need explicit synchronisation. Error prone.

▶ Because they're implemented as library calls, the compiler (and often users) cannot work out where they start and end.

▶ pthreads as OS-level threads. Need context switch. Heavyweight.

▶ Various lightweight-thread systems. Often non-preemptive. Blocking operations can block all lightweight operations sharing the same OS thread.

▶ Number of threads pretty hard-coded into your program.

▶ Nice IBM article:
https://www.ibm.com/developerworks/library/
j-nothreads/

Language support: Cilk

Cilk [example from Wikipedia]

```
cilk int fib (int n)
{   int x,y;
    if (n < 2) return n;
    x = spawn fib (n-1);
    y = spawn fib (n-2);
    sync;
    return x+y;
}
```

Compiler/run-time library can manage threads. Neat
implementation by "work stealing". Can adapt to hardware.
X10 (IBM) adds support for partitioned memory.

Language support: OpenMP

OpenMP [example from Wikipedia]

```c
int main(int argc, char *argv[]) {
    const int N = 100000;
    int i, a[N];
    #pragma omp parallel for
    for (i = 0; i < N; i++)
        a[i] = 2 * i;
    return 0;
}
```

The directive "omp parallel for" tells the compiler "it is safe to do the iterations in parallel".
Fortran "FORALL INDEPENDENT".

Clusters/Cloud Computing

- Memory support for threads, Cilk, OpenMP centres around a shared address space. (Even if secretly multi-core machines behave like distributed machines.)
- What about clusters? Cloud Computing?
- More emphasis on message-passing . . .

Software support for message passing: MPI

▶ MPI = Message Passing Interface [nothing to do with OpenMP]

▶ "de-facto standard for communication among processes that model a parallel program running on a distributed memory system." [no shared memory].

▶ Standardised API calls for transferring data and synchronising iterations. Message passing is generally synchronous, suitable for repeated sweeps over scientific data.

▶ Emphasis on message passing (visible and expensive-looking to user) means that MPI programs can work surprisingly well on multi-core, because they encourage within-core locality.

Software support for message passing: Erlang

- ▶ Shared-nothing language based on the actor model (asynchronous message passing).
- ▶ Dynamically typed, functional-style (no assignment).
- ▶ Means tasks can just commit suicide if they feel there's a problem and someone else fixes things, including restarting them
- ▶ Relatively easy to support hot-swapping of code.

Cloud Computing (1)

Can mean either "doing one computer's worth of work on a server instead of locally". Google Docs. Or . . .

Cloud Computing (2)

. . . massively parallel combinations of computing, e.g.
MapReduce invoked by a search engine.

- ► MapReduce can match a search term against many
  computers (Map) each holding part of Google index of
  words, and then combine these result (Reduce).
- ► Reduce here means parallel reduce (tree-like, logarithmic
  cost), not *foldl* or *foldr* from ML.
- ► Functional style (idempotency) useful for error resilience
  (errors happen often in big computations). Try to ensure
  computation units are larger than cost of transmitting
  arguments and results. (also: Skywriting project in
  Cambridge)

Embarrassingly Parallel

Program having many separate sub-units of work (typically more than the number of processors) which

- do not interact (no communication between them, not even via shared memory)
- are large

Example: the map part of MapReduce.

Functional Programming

In pure functional programming every tuple (perhaps an argument list to an application) can be evaluated in parallel. So functional programming is embarrassingly parallel? Not in general (i.e. not enough for compilers to be able to choose the parallelism for you). Need to find sub-executions with:

- ▶ little data to transfer at spawn time (because it needs copying, even if memory claims to be shared);
- ▶ a large enough unit of work to be done before return

Probably only certain stylised code.

Garbage Collection

While we're talking about functional programming, and as
garbage collection has previously been mentioned . . .
Just how do we do garbage collection across multiple cores?

- ▶ Manage data so that data structures do not move from one
  processor to another?
- ▶ "Stop the world" GC with one big lock doesn't look like it
  will work.
- ▶ *Parallel GC:* use multiple cores for GC.
  Concurrent GC: do GC while the mutator (user's program)
  is running. Hard?
- ▶ Incremental? Track imported/exported pointers?

How have languages adapted to parallel hardware?

- ▶ Hardware change (multi-core parallelism in 2005) is one of the major changes ever, comparable with adding *index registers* and *interrupts* from the 1960s.
- ▶ It took many years between the early languages we discussed to get to the globals-stack-heap model from the 1970s to today (for sequential processors).
- ▶ Programming language support for parallelism today seems rather fragmentary in comparison. Several incomparable solutions.
  Is 2005 too recent for languages to have got the full answer? Or . . . ?

Partial language evolution: functional programming ideas

- ► Concurrency and mutable data structures seem too hard for programmers to correctly use together in large programs.
- ► Pure functional languages can feel too 'hair shirt' for commercial use.
- ► Observation: languages have evolved to use ideas from functional programming such as higher-order functions and reduced exposure of mutable data structures in APIs used concurrently. Let's look at Streams in Java 8.

Java 8: Internal vs. External iteration

Can't trust users to iterate over data. Consider traditional *external iteration*. It's easy to start by writing

```
for (i : collection)
{   // whatever
}
```

and then get lazy. Do we really want to write this?

```
for (k = 0; k<NUMPROCESSORS; k++}
{   spawn for (i : subpart(collection,k))
    {   // whatever
    }
}
sync;
// combine results from sub-parts here
```

And what about the maintenance load?

Internal vs External iteration (2)

▶ Previous slide was *external iteration*. It's hard to parallelise (especially in Java where iterators have shared mutable state).

▶ The Java 8 Streams library encourages *internal iteration* – keep the iterator in the library, and use ML-like stream operation to encode the body of the loop

```
maxeven = collection.stream().parallel()
                    .filter(x -> x%2 == 0)
                    .max();
```

▶ The library can optimise the iteration based on the number of threads available (and do a better job than users make!).

Java Collections vs Streams

Collections

- ▶ 'Trusty old Java looping'; `for`-loops, internally using mutable state (Iterators).
- ▶ Note that Java 8 did not add `map` and `filter` to the Collections classes. Why?

Streams

- ▶ Almost ML-style understanding. Items need not be in memory all at the same time (lazy).
- ▶ But additional laziness: a stream pipeline, e.g. `.filter().max()` above, traverses the stream only once.
- ▶ `.parallel()` authorises parallel/arbitrary order of calls to the stream pipeline. Good for parallelism, beware if you use side-effecting functions on stream elements!

It's quite practical to use `.stream` on a collection and then process it as a stream and then convert back.

∼ *Topic X* ∼

Functional-style programming meets
object-orientation

References:

- ▶ *Scala By Example* by M. Odersky. Programming
  Methods Laboratory, EPFL, 2008.

- ▶ *An overview of the Scala programming language* by

  M. Odersky *et al*. Technical Report
  LAMP-REPORT-2006-001, Second Edition, 2006.

Big picture

Historically, object-oriented and functional languages were seen as disjoint programming paradigms.

> `en.wikipedia.org/wiki/Programming_paradigm`

- ▶ Recently, mainstream OO languages have acquired functional features (Java, C++, C#).
- ▶ Why?
- ▶ Influences of JavaScript and, more directly, Scala.
- ▶ Multi-core parallelism: we just saw how internal iteration can move iteration code into a library, which receives the 'body of the loop' as a lambda expression.
- ▶ Other programming paradigms like *futures*, *reactive programming* can be added to a language as library features once lambdas are present.

Scala

- ▶ Invented at EPFL by a group led by Martin Odersky; first release 2004.
  Implementation shared the JVM: interoperability with Java.
- ▶ Minimal expression-oriented language, control-flow operators defined using call-by-name.
- ▶ Smalltalk-style "everything is an object".
  Neat unification of OO and functional language features.
  Function closures are objects.
- ▶ Rich polymorphic static type system, including types without NULL. (Tony Hoare described his (1965) invention of the NULL value as his "billion-dollar mistake".)
- ▶ Lots of other goodies: `trait`s for tamed multiple-inheritance, pattern-matching on objects, actor model for concurrency, singleton objects instead of *static* class members, . . .
- ▶ Rich source of ideas for other languages (e.g. Java 8).

A procedural language!

```
def qsort( xs: Array[Int] ) {
  def swap(i: Int, j:Int) {
    val t = xs(i); xs(i) = xs(j); xs(j) = t          }
  def sortaux(l: Int, r: Int) {
    val pivot = xs( (l+r)/2 ); var i = l; var j = r
    while (i <= j) {
      while ( lt( xs(i), pivot ) ) i += 1
      while ( lt( xs(j), pivot ) ) j -= 1
      if ( i<=j ) { swap(i,j); i += 1; j -= 1 }    }
    if (l<j) sortaux(l,j)
    if (j<r) sortaux(i,r)                                    }
  sortaux(0,xs.length-1)
}
```

Remarks:
- ▶ Syntax: reserved word for definitions, colon for types.
- ▶ Array selections are written in functional notation.
  (Achieved by clever use of implicit coercions but gives the
  appearance that arrays inherit from functions.)
- ▶ Block structure.

A declarative language!

```
def qsort[T] (xs: Array[T]) (lt: (T,T)=>Boolean): Array[T]
= if ( xs.length <= 1 ) xs
  else {
    val pivot = xs( xs.length/2 )
    Array.concat(qsort(xs filter (x => lt(x,pivot))) lt,
                       xs filter (x => x == pivot)       ,
                 qsort(xs filter (x => lt(pivot,x))) lt)
  }
```

Remarks:

- ► Polymorphism.
- ► Type declarations can often be omitted because the compiler can infer it from the context.
- ► Higher-order functions, lambdas (spelt '=>').
- ► The binary operation $e \star e'$ is always interpreted a the method call $e. \star (e')$.

Parameter passing; control structures

Scala uses call-by-value by default, but it switches to call-by-name evaluation if the parameter type is preceded by `=>` (a function arrow with no parameters; expressions of this type are simply called whenever they are used).

Imperative control structures

A functional implementation of `while` loops (this mirrors the ML definition of `while` loops, but is expressible within Scala itself):

```scala
def whileLoop( cond: => Boolean )( comm: => Unit )
  { if (cond) { comm ; whileLoop( cond )( comm ) } }
```

How lambdas are implemented: Java

Running example:

```
import java.util.function.Function;
public class ClosureDemo {

    static int k = 10;    // note free variables copied, not aliased
    static String s = "14 char string";
    static int test(Function<Integer,Integer> f) {
      return f.apply(1);
    }
    public static void main(String[] args) {
        System.out.println(test(x -> k+x));
        k = 20;
        System.out.println(test(x -> k+x));
        System.out.println(test(x -> s.length()-x));
    }
}
```

- ▶ Lambdas are internally (almost) syntactic sugar for anonymous classes (but much easier for programmers).
- ▶ Lambdas have (limited) type inference – see above example.

Implementation by anonymous classes

```
@FunctionalInterface
interface Function<T,R> {
   R apply(T arg);
}
...
    public static void main(String[] args) {
        System.out.println(test(new Function<Integer, Integer>() {
            public Integer apply(Integer x) {
                return k + x;
            }
        }));
        k = 20;
        System.out.println(test(new Function<Integer, Integer>() {
            public Integer apply(Integer x) {
                return k + x;
            }
        }));
        System.out.println(test(new Function<Integer, Integer>() {
            public Integer apply(Integer x) {
                return s.length() - x;
            }
        }));
```

Implementation by named classes

```
class FunctionI implements Function<Integer,Integer> {
  int myk;
  FunctionI(int k) { myk = k; }
  public Integer apply(Integer x) { return myk + x; }
}

class FunctionS implements Function<Integer,Integer> {
  String mys;
  FunctionS(String s) { mys = s; }
  public Integer apply(Integer x) { return mys.length() - x; }
}

public class ClosureDemoClass {
    static int k = 10;
    static int test(Function<Integer,Integer> f) {
      return f.apply(1);
    }
    public static void main(String[] args) {
        System.out.println(test(new FunctionI(k)));
        k = 20;
        System.out.println(test(new FunctionI(k)));
        System.out.println(test(new FunctionS("14 char string")));
    }
}
```

Function types

In Scala it's easy: ML `t1*...*t`*k* `-> t0` encodes as

   Function*k*`[t1,...,t`*k*`,t0]`

for example: `function2[Int, Array[Int], Bool]`.

In Java it's similar, except that `Function1` is spelt `Function`
and `Function2` is spelt `BiFunction`.

Worse still in Java, there's the boxed/unboxed issue for
generics:you can write `Function<Integer,Boolean>`, but not
`Function<Integer,boolean>`. So for efficiency the latter has
a special 'function' type

   `Predicate<Integer>`

It's even more ad hoc as you don't `.apply` a `Predicate`,
you `.test` it.

Data structures in functional and object-oriented style

In ML-like languages we use `datatype` to define (say) trees:

```
datatype Tree = Leaf of t₁
              | Branch of Tree * Tree * t₂;
```

In OO languages we define an `abstract class` for `Tree` and subclass it as `TreeLeaf` and `TreeBranch`.

These define the same data structure at machine-code level – even the discriminator tags (recall the Pascal lecture) are present. ML implements these as numeric case tags, Java as pointers to the run-time-type information (virtual method table).

How does the language paradigm affect programming?

Case study (I)

```scala
abstract class Expr {    // Scala code - but the Java should be clear
  def isNumber: Boolean
  def isSum: Boolean
  def numValue: Int
  def leftOp: Expr
  def rightOp: Expr
}
class Number( n: Int ) extends Expr {
  def isNumber: Boolean = true
  def isSum: Boolean = false
  def numValue: Int = n
  def leftOp: Expr = error("Number.leftOp")
  def rightOp: Expr = error("Number.rightOp")
}
class Sum( e1: Expr; e2: Expr ) extends Expr {
  def isNumber: Boolean = false
  def isSum: Boolean = true
  def numValue: Int = error("Sum.numValue")
  def leftOp: Expr = e1
  def rightOp: Expr = e2
}
def eval( e: Expr ): Int = {
  if (e.isNumber) e.NumValue
  else if (e.isSum) eval(e.leftOp) + eval(e.rightOp)
  else error("bad expression")
}
```

Case study (I): critique

- It's clumsy; we could use `instanceof`.
- It's hard to maintain.
- It's hard to extend:
  - adding a new *form of data* (e.g. `Prod`) requires every existing class and function to be changed.
- It's easy to extend:
  - adding a new *operation on the data* (e.g. `print`) requires just adding one function.
  - ML-style pattern matching would shorten the code.

But one advantage is that it highlights the $n \times m$ complexity of n functions operating on m forms of data.

Case study (II)

We could rephrase this more 'object-orientedly' [Scala code, but the Java should be clear]:

```scala
abstract class Expr {
  def eval: Int
}
class Number( n: Int ) extends Expr {
  def eval: Int = n
}
class Sum( e1: Expr; e2: Expr ) extends Expr {
  def eval: Int = e1.eval + e2.eval
}
```

This implementation is easily extensible with *new forms of data*:

```scala
class Prod( e1: Expr; e2: Expr ) extends Expr {
  def eval: Int = e1.eval * e2.eval
}
```

But, a new operation on the existing data *still* requires every existing class to be changed.

The expression problem

Constrast this with ML-style datatypes and functions over them:

- ▶ a new operation on the existing data (datatypes) is easy;
  but
- ▶ a new form of data requires every `case` match on that
  datatype to be changed.

The language-design problem of allowing a data-type definition
where one can add new cases to the datatype and new
functions over the datatype (without requiring ubiquitous
changes) is known as the 'expression problem':

<div align="center">

en.wikipedia.org/wiki/Expression_problem

</div>

Case study (III): Scala case classes

```scala
abstract class Expr
case class Number( n: Int ) extends Expr
case class Sum( e1: Expr, e2: Expr ) extends Expr
case class Prod( e1: Expr, e2: Expr ) extends Expr
```

▶ Case classes implicitly come with a constructor function, with the same name as the class. So can write (e.g.)

```scala
Sum( Sum(Number(1),Number(2)), Number(3) )
```

The `match` keyword takes as argument a number of `case`s:

```scala
def eval( e: Expr ): Int
  = e match
    { case Number(x) => x
      case Sum(l,r) => eval(l) + eval(r)
      case Prod(l,r) => eval(l) * eval(r)
    }
```

▶ If none of the patterns matches, the pattern matching expression is aborted with a `MatchError` exception.

- ▶ Case classes and case objects implicitly come with implementations of methods `toString`, `equals`, and `hashCode`.
- ▶ Case classes implicitly come with nullary accessor methods which retrieve the constructor arguments.
- ▶ Case classes allow the constructions of *patterns* which refer to the case class constructor (see previous slide).

Scala case classes are essentially ML data types and allow both functional and object-oriented styles of programming to be expressed and (relatively) easily refactored.

Generics and polymorphism in Scala

▶ Scala has a fairly powerful type inferencer which allows one to omit type parameters to polymorphic functions and constructors.

▶ Classes in Scala can have type parameters (as in Java generics), but as in Java, mixing generics with subtyping poses the question:

> *if* `T` *is a subtype of a type* `S`*, should* `Array[T]` *be a subtype of the type* `Array[S]`*?*

In Scala, generic types like the following one also have *non-variant* subtyping by default.

```
class Array[A] {
  def apply( index: Int ): A
  def update( index: Int, elem: A )
    ...
}
```

One can enforce *covariant* subtyping by prefixing a formal type parameter with a '+'. The prefix '−' indicates *contravariant* subtyping.

This is *declaration-site* variance rather than Java's *use-site* variance declarations.

Scala uses a conservative approximation to verify soundness of variance annotations: a covariant type parameter of a class may only appear in covariant position[7] inside the class. Hence, the following class definition is rejected:

```
class Array[+A] {
  def apply( index: Int ): A         // OK
  def update( index:Int , elem: A )  // bad
    ...
}
```

So, how do we define a type like this?

---

[7]Mathematically this is "left of an even number of function arrow types".

Generic types
Type bounds

The following attempt to define a covariant functional stack fails
(as it would be unsound just like the Java array example
earlier):

```
abstract class Stack[+A]  // covariant declaration
{ def push( x:  A )  // A in contravariant position
                     // hence rejected
     : Stack[A]
   = new NonEmptyStack(x,this)
  def top: A
  def pop: Stack[A] }
```

Scala achieves this by using type bounds (`B >: A` below)
echoing Java's use of `extends` and `super` in its use-site
variance declarations:

```scala
abstract class Stack[+A] {
  def push[B >: A]( x: B ): Stack[B]
    = new NonEmptyStack(x,this)
  def top: A
  def pop: Stack[A]
}
class NonEmptyStack[+A](elem: A, rest: Stack[A])
      extends Stack[A] {
  def top = elem
  def pop = rest
}
```

Traits

Mixins and traits are not on the syllabus in 2018/19.

- A `trait` is a special form of an abstract class that does not have any value (as opposed to type) parameters for its constructor and is meant to be combined with other classes.

```
trait IntSet {
  def incl( x:Int ): IntSet
  def contains( x:Int ): Boolean
}
```

Traits may be used to collect signatures of some functionality provided by different classes.

Mixin-class composition

Every class or object in Scala can inherit from several traits in addition to a normal class. E.g.

```scala
trait AbsIterator[T] {
  def hasNext: Boolean
  def next: T
}
trait RichIterator[T] extends AbsIterator[T] {
  def foreach( f: T => Unit ): Unit =
    while (hasNext) f(next)
}
class StringIterator( s: String )
  extends AbsIterator[Char] {
    private var i = 0
    def hasNext = i < s.length
    def next = { val x = s charAt i; i = i+1; x }
  }
```

Traits can be used in all contexts where other abstract classes appear; however only traits can be used as *mixins* (introduced using the keyword `with` below):

```
object Test {
  def main( args: Array[String] ): Unit = {
    class Iter extends StringIterator(args(0))
                with    RichIterator[Char]
    val iter = new Iter
    iter.foreach(System.out.println)
  }
}
```

The class `Iter` is constructed from a *mixin composition* of the parents `StringIterator` (called the *superclass*) and `RichIterator` (called a *mixin*) so as to combine their functionality.

The class `Iter` inherits members from both `StringIterator` and `RichIterator`.

NB: Mixin-class composition is a form of multiple inheritance, but avoids the 'diamond' problem of C++ and similar languages (where a class containing a field appears at multiple places in the inheritance hierarchy).

```
class A  public: int f;
class B: public A  ...
class C: public A  ...
class D: public B,C  ... f ...
// Is this B::f or C::f?
// And do B and C both have an f of type A,
//   or share a single one?
```

[End of section on mixins and traits.]

Value types

A *value type* is a type representing an a pure value – such values can be copied, compared and subfields extracted but subfields cannot be mutated. Complex numbers and immutable arrays are examples, but mutable arrays are not. Strings of course should be values, but C and C++ historically have had issues as to whether they are `char[]` or `const char[]`.

▶ In Java, every value is either of primitive type (spelt with lower case) or is a reference to an object. Assignment '`=`' copies bit patterns for primitives, but only copies references for object types. Equality '`==`' differs similarly.

▶ Worse, every reference-type object has *identity*, so that `new Complex(1.0,2.0)` allocates a new *distinguishable* object.

Value types and languages

When talking about value types, we want a type which only admits functional update (copy but changing a field), not mutability and no identity.

Why are value types useful? Increasing use of functional-style APIs (partly due to the need to exploit parallelisms).

Think about defining a class `Complex` and note that every arithmetic operation is likely to do an allocation. What we want is the equivalent of `struct { double re,im; }` in C. C# (Microsoft's Java-style language) provides such values:

> *Variables [of value type] directly contain values. Assigning one value-type variable to another copies the contained value. This differs from the assignment of reference-type variables, which copies a reference to the object but not the object itself."* [C# reference manual]

Modelling value types in Java

We can model a value type in Java in two ways

- ▶ use `final` to stop modification; or
- ▶ use *deep copy* to snapshot values on operations.

Both of these are fragile during program enhancement.

For example a routine may copy an object to another object, but program enhancement may change a value type into a reference type, resulting in code elsewhere getting a value which is *part copied from, and part aliased with,* a source object.

- ▶ An object-copy routine which simply copies the fields of an object is known as *shallow copy*
- ▶ An object-copy routine which recursively copies the fields of an object is known as *deep copy*

Value types and inefficient representation

Think how we might represent an array of size 100 of complex values in Java.

- ▶ We'll have an array of 100 pointers (around 800 bytes) and 100 complex values (perhaps 24 bytes each).
- ▶ But at machine level we really just want to store 200 `double` values (only 1600 bytes).

This costs Java programs in memory use, cache effects and instructions executed.

- ▶ Java would benefit from C#-like value types.

Value types form part of the 'Valhalla project' originally planned for Java 10:

    http://en.wikipedia.org/wiki/Project_Valhalla_
                    (Java_language)

Incidentally, the Wikipedia entry on value types needs improving (April 2018):

        https://en.wikipedia.org/wiki/Value_type

Aliasing and mutability

Mixing aliasing and mutability – as Java does – can enable subtle bugs. Adding concurrency enables more subtle bugs (and more-subtle bugs!).

There appear to be two 'easy-to-use' sweet spots to use current languages:

- ▶ use mutability freely, but use deep copy so that mutations don't accidentally update some other data structure.
- ▶ use aliasing freely, but make data structures immutable.

Ideally, future languages will allow these two techniques to be mixed and perhaps also allow aliasing and mutability to be used together locally without impacting whole-program maintainability. But for now, all I can do is to highlight the issues.[8]

---

[8]And perhaps mention 'separation logic' for the interested reader.

Topic XI

Miscellaneous (entertaining) concepts

Overview

- Monadic I/O
- Generalised Algebraic Data Types (GADTs)
- Continuation-passing style (CPS) and call/cc
  Wikipedia:Call-with-current-continuation
- Dependent types (Coq, Agda, Idris)

I/O in functional languages

The ML approach: "evaluation is left-right, just let side-effecting I/O happen as in C or Java".

- ▶ Breaks *referential transparency* ('purity'), e.g. that $e + e$ and **let** $x = e$ **in** $x + x$ should be equal.
- ▶ Not an appropriate solution for lazy languages (order of side effects in arguments in a function call would depend on the detail of the called function).
- ▶ Haskell is a lazy function language ("laziness keeps us pure").

Everything I say about I/O applies to other side-effecting operations, e.g. mutable variables, exceptions, backtracking . . .

Giving different types to pure and impure functions

Suppose we have two types $A \to B$ for functions with no side-effects and $A \rightsquigarrow B$ for impure functions.

- ▶ Then everything with a side effect would be visible in its type (contagious).

- ▶ (Later we might have ways of hiding "locally impure" functions within a pure function, but this doesn't apply to I/O.)

- ▶ Instead of writing $A \rightsquigarrow B$ we write $A \to B\,M$ where M is a unary type constructor called a *monad*.
  Factors the idea of 'call a function' and 'do the resulting *computation*'.

Syntax: ML type constructors are postfix so we write *t list* or *B M* whereas Haskell writes *M B* and perhaps *List t*.

So, how does I/O work?

In ML we might read from and write to stdin/stdout with
(writing $\rightsquigarrow$ for emphasis):

```
MLrdint: unit ⤳ int
MLwrint: int ⤳ unit
```

Using monads (either in Haskell or ML) these instead have type

```
rdint: int IO
wrint: int -> unit IO
```

- ▶ The unary type constructor `IO` is predefined, just like the binary type constructor '`->`'.
- ▶ Shouldn't we have `rdint: unit -> int IO`?
  We could, but this would be a bit pointless—writing takes an argument and gives a computation, but reading from stdin is just a computation.

Composing I/O functions
or Sequencing I/O effects

- In ML we could just write e; e' to sequence the side effects of e and those in e'. But now all 'effects' like I/O are part of monadic values, and no longer 'side effects'—so this doesn't work

- Instead every monad `M` (including `IO` as a special case) has two operators: one to sequence computations and one to create an empty computation.

```
(>>=): 'a M -> ('a -> 'b M) -> 'b M
       (* infix, pronounced 'bind' *)
return: 'a -> 'a M
```

Digression – Monad Laws

The idea of monad originates in mathematics, so these operators have axioms relating $>>=$ and `return` in the *mathematics*; these are seem as laws which all well-behaved *programming* monads satisfy.

```
[left unit]
  m >>= return     =  m
[right unit]
  return x >>= f   =  f x
[associativity]
  (m >>= f) >>= g  =  m >>= λx.(f x >>= g)
```

(The bind operator '$>>=$' syntactically groups to the left so the brackets in the final line are redundant.)

Using the IO monad

In a system using monadic I/O, for example Haskell, the read-eval-print loop not only deals with pure values (integers, lists and the like), but has a special treatment for values in the `IO` monad. Given a value of type `t IO`, it:

- performs the side effects in the monad; then
- prints the resulting value of type `t`.

So the question is: how do we make a computation which (say) reads an integer, adds one to it, and then prints the result?

Using the IO monad – examples

Read an integer, adds one to it, and print the result (using ML syntax):

```
rdint >>= (fn x => wrint(x+1));
```

Do this twice:

```
let val doit = rdint >>= (fn x => wrint(x+1))
in doit >>= (fn () => doit);
```

Note that `doit` has type `unit IO`, so we use »= to use `doit` twice.

NB: *computations* are not *called* as they are not functions; they are *combined* using `>>=` as in the above examples.

Using the IO monad – examples (2)

Read pairs of numbers, multiply them, printing the sum of products so far, until the product is zero.

```
fun foo s = rdint >>= fn x =>
            rdint >>= fn y =>
            if (x*y = 0) then return ()
            else wrint(s + x*y) >>= fn () =>
              foo(s + x*y);
foo 0;
```

Note the type of `foo` is `int -> unit IO`.
Note also the use of `return` to give a 'do nothing' computation.

Practical use

The Mirage OS (most recent Computer Lab spin-out acquired by Docker) is written in OCaml in monadic style. But most monadic-style programs are written in Haskell. GHCi is just like the ML read-eval-print loop, but remember:

- Haskell syntax uses upper case for constants (types, constructors) and lower case for variables.
- Haskell swaps ':' and '::' relative to ML
- Type constructors are prefix
- `fn x=>e` is written `\x->e`

Haskell also provides `do`-notation to allow programmers to write imperative-looking code which de-sugars to uses of `return` and `>>=`.

Other monads

- ▶ Many other unary type constructors have natural monadic structure (at least as important as `IO`).
- ▶ For example, using Haskell syntax, `List t` and `Maybe t`. Another important one is `State s t` of computations returning a value of type `t`, but which can mutate a state of type `s`. (Subtlety: the monad is legally the unary type `State s` for some given s.)
- ▶ Haskell overloads `>>=` and `return` to work on all such types (Haskell's *type class* construct facilitates this).
- ▶ The common idea is 'threading' some idea of state implicitly through a calculation.

Haskell example using `List` in GHCi:

```
[1,2,3] >>= \x->if x==2 then return 5 else [x,x]
[1,1,5,3,3]
```

Generalised Algebraic Data Types (GADTs)

OCaml data type (just like datatype in ML):

```
type 'a mylist = MyNil | MyCons of 'a * 'a mylist;;
```

Can also be written

```
type 'a mylist = MyNil : 'a mylist
              | MyCons : 'a * 'a mylist -> 'a mylist;;
```

Why bother (it's longer and duplicates info)?

How about this:

```
type _ exp = Val : 'a -> 'a exp
           | Eq : 'a exp * 'a exp -> bool exp
           | Add : int exp * int exp -> int exp
```

[This uses OCaml syntax, but Haskell also has GADTs]
Allows `bool exp` values to be checked that `Add`, `Eq` etc. are
used appropriately. E.g.

```
Val 3: int exp √
Val true: bool exp √
Add(Val 3, Val 4): int exp √
Add(Val 3, Val true) ✗
Eq(Val true, Val false): bool exp √
Eq(Val 3, Val 4): bool exp √
Eq(Val 3, Val true) ✗
```

Can't do this in SML (special case of *dependent types* later this
lecture).

Can even write `eval` where the *type* of the result depends on the *value* of its type:

```
fun eval(Val(x)) = x
 |  eval(Eq(x,y)) = eval(x) = eval(y)
 |  eval(Add(x,y)) = eval(x) + eval(y);

eval: 'a exp -> 'a
```

(Some type-checking dust being swept under the carpet here.)

Reified continuations

Make calling continuation appear to be a value in the language (reifying it).
Reminder on continuation-passing style (CPS), perhaps mentioned in Compiler Construction. Can see a function of type $t_1 \rightarrow t_2$ as a function of type

$$(t_2 \rightarrow \texttt{unit}) \rightarrow (t_1 \rightarrow \texttt{unit})$$

Or uncurrying

$$(t_2 \rightarrow \texttt{unit}) \times t_1 \ \rightarrow \ \texttt{unit}$$

(One parameter of type t_1 and the other saying what to do with the result t_2 – like *argument* and *return address*!)

Instead of

```
fun f(x) = ... return e ...
print f(42)
```

we write

```
fun f'(k, x) = ... return k e' ...
f'(print, 42)
```

In CPS all functions return `unit` and all calls are now tail-calls (so the above isn't just a matter of adjusting a return statement). Sussman and Steele papers from the 1970's ("Lambda the ultimate XXX").

Reified continuations (2)

- A function with *two* continuation parameters rather than one can act as normal return vs. exception return. (Or Prolog success return vs failure return.)
- But we don't want to write all our code in CPS style. So: *call/cc* "call with current continuation". Lots of neat programming tricks in a near-functional language.
- Reified? The continuation used at the meta-level (semantics) to explain how a language operates is exposed as an object-level (run-time value).

Reified continuations (3)

Core idea (originally in Scheme, a form of Lisp):

`fun f(k) = let x = k(2) in 3;`

In ML this function 'always returns 3'. E.g.

`> f(fn x=>x);`

But `callcc(f)` returns 2!

- ▶ The return address/continuation used in the call to `f` is reified into a side-effecting function value `k` which represents the "rest of the computation after the call to `f`".
- ▶ Some similarity with `f(fn x => raise Foo);`

Dependent types

Suppose *f* is a curried function of *n* boolean arguments which returns a boolean. How do we determine if *f* is a tautology (always returns true)?

```
fun taut(0,f) = f
  | taut(n,f) = taut(n-1, f true) and also
                taut(n-1, f false);
```

Works nicely in dynamically typed languages. Fails to type-check in ML. Why?

- ▶ The *type* of the second argument depends on the *value* of the first.
- ▶ Dependent type systems can capture this (languages like Coq, Agda and Idris).
- ▶ But type checking can require theorem proving.

Dependent typing of `taut` in Idris

```
-- BF n is the type Bool -> ... -> Bool -> Bool
--                       _____/
--                               n
BF : Nat -> Type
BF Z = Bool
BF (S n) = Bool -> BF n

taut : (n : Nat) -> BF n -> Bool
taut Z f = f
taut (S n) f = taut n (f False) &&
                taut n (f True)
```

- Dependent types allow compile-time checking of things which we would expect to cause exceptions, e.g.
  ```
  hd : (n:Nat) -> List (S n) alpha -> alpha
  ```
- More Idris code on the course website.
- Something to think about: are ML exceptions just dynamically checked types?