

Advances in Programming Languages

Coursework Assignment Topics

Ian Stark

School of Informatics
The University of Edinburgh

Thursday 27 September 2018
Semester 1 Week 2



Outline

- 1 Exams, coursework, and homework
- 2 Assignment topics
- 3 Writing bibliographic references
- 4 Assignment timing and format
- 5 Academic integrity

Final grades are based on an exam (80%) and a written coursework assignment (20%).

The exam is in April/May, and has a standard “choose two questions out of three” format. Past papers are available through the course web page.

The examinable material for this course is the **content of the lectures** and their accompanying **homework exercises**. The following will not be assessed in the examination:

- Any guest lectures;
- The written coursework;
- Further references in the course blog.

The written coursework requires investigation of a topic in programming languages and writing an 8–10 page report with example code.

The aims of the homework exercises set in lectures include:

- To prepare for forthcoming lectures
- To review lecture material
- To give some context for understanding the lectures
- To provide other sources and views on the lecture topic
- To help you learn by exploring the subject

Crucially, these effects arise from **doing** the exercises. They are not assessed, nor would this necessarily be useful: they are intended to be largely self-validating (i.e. you can tell when you have succeeded).

Although your coursework reports will be assessed, most of this also applies there: the purpose is for you to find out and learn new things.

Homework (1/2)

Read This



Luca Cardelli

https://is.gd/cardelli_types

Type Systems: Section 1 “Introduction”.

Chapter 97 of *The Computer Science and Engineering Handbook*, 2nd Edition

Watch This

https://is.gd/wadler_pat_video (Video, 43m)

Propositions as Types (Recorded at Strange Loop, September 2015)

Phil Wadler, Professor of Theoretical Computer Science, Edinburgh University



Code This

... see next slide.

Homework (2/2)

Java has *subtyping*: a value of one type may be used at any more general type. So `String` \leq `Object`, and every `String` is an `Object`. This isn't always straightforward.

```
String[] a = { "Hello", "world" };           // Array of strings
Object[] b = a;                             // Array of objects (every string is an object)
b[0] = Boolean.FALSE;                       // Assign object to array of objects
String s = a[0];                            // Fetch string from array of strings
System.out.println(s.toUpperCase());        // Convert string to upper-case
```

- 1 Build a Java program around this.
- 2 Compile it. Run it.
- 3 What happens, and when? Can you explain why?
- 4 How might you change the Java language to prevent this?
- 5 Pick another object-oriented language: what happens when you try this there?

Outline

- 1 Exams, coursework, and homework
- 2 Assignment topics**
- 3 Writing bibliographic references
- 4 Assignment timing and format
- 5 Academic integrity

Parallel Performance Portability with Lift

The *Lift* language is designed to support high-level data-parallel programming in a way that can be efficiently executed on a range of different architectures. The aim is that a programmer can write a single high-level functional program that describes a desired computation on massive datasets and then compile it to run efficiently on a variety of specialist parallel hardware — whether CPU, GPU, or other accelerator.

Existing approaches are often closely tied to a specific target device and cannot easily take advantages of new GPU generations or other architectural changes. Lift provides a system of *rewrite rules* so that the same language primitives can be given hardware-specific optimisations that will benefit all applications written in the language.

Dynamic Information Flow Policies in Jeeves

The *Jeeves* language provides a way for programmers to describe and enforce security policies for code. These policies manage *information flow* in a program: identifying what values depend on what others, and in particular where confidential information may travel during execution. In *Jeeves* these policies are written separately from program code, and the language implementation ensures that policies are enforced when the program is executed.

Jeeves uses *faceted* expressions to track information flow at runtime, which means it can efficiently identify *implicit flows*. It is implemented as an *embedded domain-specific language* for Python and Scala.

Programming Quantum Computation with Quipper

The field of *quantum computation* considers devices that exploit differences between classical and quantum physics. There are several cases where *quantum algorithms* for particular problems are known to give complexity improvements over any classical algorithm. These are often described in terms of *quantum gates* and *quantum circuits*: quantum analogues of the same low-level components in classical computers.

Quipper is a high-level language for describing quantum circuits. It is implemented as an embedded domain-specific language in Haskell, and provides support for generating circuits and simulating their behaviour on a classical computer.

Query Expressions for Language-Integrated Database Access in F#

Microsoft's F# language provides several facilities for building and high-level manipulation of computations and metacomputations. In particular *computation expressions* allow libraries to define domain-specific sublanguages for particular kinds of computation.

One example of this is *query expressions*, a way to write code describing access to the tables and records of classic relational databases. These query expressions can combine F# code with familiar operators from database query languages like SQL, which can then be compiled to run efficiently on different kinds of data source.

Probabilistic Programming for Statistical Inference in Stan

The *Stan* platform provides a language and libraries for setting up complex statistical models and computing inferences for these. Programs written in the Stan language describe probabilistic systems. It's possible with Stan to execute direct simulations of these systems: but also to carry out Bayesian inference and other analysis to understand their behaviour beyond a single execution run.

There is also the recent *SlicStan* language which uses information flow analysis to overcome some restrictive aspects of Stan programming. This gives a more expressive way to write probabilistic models while still using the efficient Stan inference machinery.

I learned about the language from key *SlicStan* researcher *Maria Gorinova*, now a PhD student at Edinburgh and APL participant in 2016/17.

Outline

- 1 Exams, coursework, and homework
- 2 Assignment topics
- 3 Writing bibliographic references**
- 4 Assignment timing and format
- 5 Academic integrity

The GHC compiler for Haskell does not support dynamically linked libraries on Windows [1].

1. GHC Users Guide.

http://www.haskell.org/ghc/docs/latest/html/users_guide/index.html

(Not true since GHC 7.0) [Citation Needed]

Facebook signed up 200 million users in its first five years [NYT].

[NYT] “Is Facebook Growing Up Too Fast?” New York Times

The *Links* programming language compiles client code into Javascript (Cooper et al. 2007).

Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop.

Links: Web Programming Without Tiers.

In *Formal Methods for Components and Objects: Proceedings of the 5th International Symposium FMCO 2006*. Lecture Notes in Computer Science 4709, pages 266–296.

Springer-Verlag, 2007. DOI: [10.1007/978-3-540-74792-5_12](https://doi.org/10.1007/978-3-540-74792-5_12)

Moore's law says that the number of components in an integrated circuit doubles every year¹.

1. Wikipedia

Using *events* can be an effective technique for concurrent programming[†].

[†] John Ousterhout. Why Threads are a Bad Idea (for most purposes). In *Proceedings of the USENIX 1996 Technical Conference*. January 1996.

Outline

- 1 Exams, coursework, and homework
- 2 Assignment topics
- 3 Writing bibliographic references
- 4 Assignment timing and format**
- 5 Academic integrity

Dates and submission

Week 2 Thursday 27 September	Topics announced
Week 4 Friday 12 October	Outline draft due by 4pm
Week 5 Friday 19 October	Feedback on outline draft
Week 8 Friday 9 November	Final report due by 4pm
Week 10 Friday 23 November	Feedback on final report

Each report must be submitted electronically as a PDF document. The recommended method for creating these is `pdflatex` with the `article` document class.

In addition, `LibreOffice` is freely available for Windows, Linux and Mac, installed on Informatics machines, and can write PDF. Mac OS X natively creates PDF. Microsoft Office 2007 and later can save as PDF.

Outline Draft

Your submitted draft should contain the following.

- Your student matriculation number;
- The topic you have chosen;
- Notes on background and what you plan to write;
- A screenshot by you of the selected system in action.
- Three suitable references;
- For each reference, a paragraph summarizing what it says.

One reference must be to a published paper; the other two may be also, but could also be white papers, web tutorials, manuals, or similar. In all cases provide enough information for someone else to obtain the document.

To create the screenshot, you will need to have your chosen system downloaded, installed, and running on a suitable machine.

Suggested outline for final report

Heading Title, date, student number

Abstract This report describes ...

Introduction Content summary, overview of report structure

Context The problem domain

⟨Main topic⟩ What it is, how it works, advantages and limitations

Example Annotated code, explanation, screenshot

Salt: The example must in some way concern social media (e.g. chat messages, image posts, social networks, ...)

Resources For each notable resources used (article, tutorial, manual), give a summary in your own words of what it contains

Related work Other approaches to the problem

Conclusion What ⟨topic⟩ does, good and bad points

Bibliography Full references for all resources used

Total around 8–10 A4 pages. See the course web pages for further details.

Outline

- 1 Exams, coursework, and homework
- 2 Assignment topics
- 3 Writing bibliographic references
- 4 Assignment timing and format
- 5 Academic integrity

University of Edinburgh Undergraduate Assessment Regulations

Regulation 29

All work submitted for assessment by students is accepted on the understanding that it is the student's own effort without falsification of any kind.

University of Edinburgh Undergraduate Assessment Regulations

Regulation 30

30.1 Marks or grades can only be given for original work by students at the University. Plagiarism is the act of copying or including in one's own work, without adequate acknowledgement, intentionally or unintentionally, the work of another or one's own previously assessed original work. It is academically fraudulent and an offence against University discipline. . . .

University of Edinburgh Undergraduate Assessment Regulations

Regulation 30

30.2 It is academically fraudulent and an offence against University discipline for a student to invent or falsify data, evidence, references, experimental results or other material contributing to any student's assessed work or for a student knowingly to make use of such material. It is also an offence against University discipline for students to collude in the submission of work that is intended for the assessment of individual academic performance or for a student to allow their work to be used by another student for fraudulent purposes.

University of Edinburgh Undergraduate Assessment Regulations

Regulation 29

All work submitted for assessment by students is accepted on the understanding that it is the student's own effort without falsification of any kind.

<http://www.inf.ed.ac.uk/teaching/tar>

See also:

- University guidance

<http://www.ed.ac.uk/academic-services/students/conduct/academic-misconduct>

- Informatics statement

<http://web.inf.ed.ac.uk/infweb/admin/policies/academic-misconduct>

Working practices

- Start with a blank document; all the words must be yours.
- Do not cut and paste from other documents.
 - Except for direct quotations, which must have source declared.
- Do not let others read your text; nor read theirs (except as directed).

Aims of this assignment

- To learn about the chosen topic
- To improve researching and learning skills
- To demonstrate said knowledge and skills

The tangible outcome is a document, composed and written by you, demonstrating what you have learnt.

Homework

Before the next lecture find an online tutorial for each of the assignment topics:

- Parallel performance portability with [Lift](#)
- Dynamic information flow policies in [Jeeves](#)
- Programming quantum computation with [Quipper](#)
- Query expressions for language-integrated database access in [F#](#)
- Probabilistic programming for statistical inference in [Stan](#)

Send me your list of five links and I will summarize them for the class.

Use them to help choose your topic.