400-Level Project Proposals 2015

A list of possible 400-level projects is given below. You do not have to restrict yourself to this list, and can make up your own project topic. However, to do so you will have to arrange for a supervisor. You can find more information about the research interests of the staff members on the Computer Science Department Web pages.

You should email your project selections to the project co-ordinator, Steven Mills steven@cs.otago.ac.nz by **Friday 27th February**. Before making these selections you should talk to the supervisors of the projects you are interested in. In most cases the best way is to email them to make a time to do so – supervisors' email addresses are given with each project.

In order to allocate the projects as fairly as possible we ask that you give first, second, and third choice projects, and that you choose projects from three different supervisors. If you find multiple projects that look equally appealing, then you can give lists of projects as first, second, and third choices. You will still need to include at least three different supervisors, and at least two in your first two lists.

There are two different project papers - COSC480 and COSC490. The relevant Web pages (for COSC480 and COSC490) have more details, but the short version is that if you are enrolled in the Computer Science Honours degree you take COSC490, students in the first year of a Masters take either paper, and everyone else takes COSC480. Most projects are suitable for either COSC480 or COSC490, but a few might be only suitable for COSC480, which will be indicated in the project description.

1. **Simulator for Li-Fi: Data Communications Through LED Light Bulbs**

   Yawen Chen (yawen@cs.otago.ac.nz)
   Haibo Zhang (haibo@cs.otago.ac.nz)

   LED bulbs can provide more than just illumination! Recent development of "Li-Fi" is an innovative way of wireless optical communications using the signals sent through the light bulb instead of radio waves as in Wi-Fi.The lightbulb is flicked on and off very quickly, up to billions of times per second. That flicker is so fast that the human eye cannot notice it. There are many advantages and potential applications to use Li-Fi in the future. One of the problems is how to design the layout of LED bulbs according to different application scenarios and requirements (e.g.illumination level/transmission rate/energy saving). This project aims to explore the characteristics of Li-Fi and design a simulator for Li-Fi, which is useful for the usage and development of Li-Fi. `http://thenextweb.com/insider/2014/08/21/purelifi-li-fi-vlc-led/`

2. **Simulator for Optical Network on Chip (ONoC): Core-to-Core Communication by Light**

   Yawen Chen (yawen@cs.otago.ac.nz)
   Haibo Zhang (haibo@cs.otago.ac.nz)

   With the increasing number of CPU cores on a processor chip, conventional electronic interconnects for core-to-core communication are becoming a bottleneck. Recent rapid advances have made the optical-based interconnection, Optical Network-on-Chip (ONoC), an attractive solution to breakthrough electronic interconnect limitations. This project aims to explore the characteristics of ONoC (number of cores, light communication path, topologies, and etc.), and design a simulator for Optical Network on Chip (ONoC). The simulator will be very useful for the hardware design and development. `http://insidehpc.com/2014/01/01/ibm-brings-nanophotonics-real-world-manufacturing/`.

3. **Quantifying Conceptual Density in Text**

   Anthony Robins (anthony@cs.otago.ac.nz)
   Ali Knott (alik@cs.otago.ac.nz)

   Can we identify the density of connections between individual concepts in a body of text? In other words, can we identify discrete concepts, and quantify the relationships between them? This project will involve investigating existing natural language processing tools, and perhaps developing our own.

4. **Attractor spaces in Hopfield nets (Anthony Robins)**

   Anthony Robins (anthony@cs.otago.ac.nz)

   Hopfield networks (a kind of neural network) have dynamic behaviour that can be characterised in terms of gradient descent in a multidimensional attractor space. We know a fair bit about the structure of such spaces. This project will involve implementing and exploring a Hopfield type network to further develop our understanding of its dynamic behaviour.

5. **Serial learning in neural networks (Anthony Robins)**

   Anthony Robins (anthony@cs.otago.ac.nz)

   Neural networks are very powerful learning systems. But most of them work best when all training data is available simultaneously (concurrent learning). In the real world humans and other animals learn different things on different occasions (serial learning). This project will review the current state of serial learning in neural networks, and further explore my "pseudorehearsal" solution.

6. **Extending DrJava (Anthony Robins)**

   Anthony Robins (anthony@cs.otago.ac.nz)

   DrJava is the IDE used in COMP160. It is open source. The purpose of this project is to extend the IDE to provide support for porting a Java program to an Android device. The goal is to make the process as simple as possible, so that COMP160 students can port simple graphical programs to Android. If time allows other platforms (iOS) and other extensions to DrJava could be considered.

7. **Developing Android App for Wearable EEG Devices**

   Zhiyi Huang (zhuang@cs.otago.ac.nz)

   Electroencephalography (EEG) is a recording of electrical activity from human brain. It contains a lot of interesting and important bio-information from human body which are yet to be discovered. Recently there are quite a few off-the-shelf, wearable and portable EEG devices available, such as Brainlink, MUSE, and Emotiv. Some interesting mobile Apps were developed for gaming and meditation. However, more interesting Apps could be designed and developed. In this project, we will make our first attempt to develop a simple Andriod App for recording and processing EEG on Android phones. It involves programming with Android Studio and BluetoothSocket to set up Bluetooth connection with EEG devices. If time allows, advanced Apps based on EEG will be designed and developed. This project has external support from Psychology.

8. **Building a Scalable Distributed Entropy Infrastructure**

   David Eyers (dme@cs.otago.ac.nz)
   Advisors: Jim Cheetham and Paul Campbell as appropriate

   Software used in computer security often requires a source of randomness, e.g. for the generation of cryptographic keys. While modern CPUs include components that sample their physical environment to derive random numbers, the recent disclosures about US NSA and other similar organisations' interference have thrown doubt over the trustworthiness of these components.

   Jim Cheetham and Paul Campbell developed the OneRNG `http://www.onerng.info` as a USB-connected source of entropy that can be verified, both in terms of hardware and software.

   This project is to design and develop a network-based, scalable random number generator. This will be useful for cases such as cloud-hosted servers, into which you cannot plug your own hardware random number generators, but to which you are able to deliver a network stream.

   When higher rates of random data are required than can be safely provided, just plug in more OneRNGs, and the system should automatically reconfigure itself to use them.

9. **Monitoring the Quality of a Random Number Stream**

   David Eyers (dme@cs.otago.ac.nz)
   Advisors: Jim Cheetham and Paul Campbell as appropriate

   The OneRNG `http://www.onerng.info` has been developed by Paul Campbell and Jim Cheetham as a USB-connected source of entropy, the implementation of which can be verified both in terms of hardware and software. The OneRNG uses two independent noise sources, so that it is harder for an attacker to manipulate the environment of the device, in an attempt to maliciously influence its noise sources.

   This project will use and extend existing analysis packages, such as "dieharder", that assess the randomness of a random number stream. The goal is to develop a framework for long-term randomness analysis that is tuned specifically for the OneRNG, e.g. to detect shifts in behaviour over time of one or the other of its noise sources.

10. **QR Codes for Security**

    David Eyers (dme@cs.otago.ac.nz)

    QR codes are a class of 2D barcodes of various sizes and error correction capabilities that have a public, extensive and extensible schema for data that can be encoded in them. Open source software is available to generate and detect QR codes.

This project focuses on the development of software that uses QR codes for computer security functions. Forming a line-of-sight coupling between the QR code display (possibly even just on paper) and a camera has the advantage that so called "man in the middle" attacks are difficult compared to communication over a data network. This increases the convenience of operations such as delegation of security privileges.

A previous project student successfully implemented the one-way transport of Kerberos tickets through QR codes. A promising alternative target is the OAuth2 protocol for decentralised authorisation, and/or two-way QR code interactions.

11. **Web Pipes**

    David Eyers (dme@cs.otago.ac.nz)

    Pipes are frequently used within shell scripting on Unix-like operating systems. They are a simple, yet highly effective inter-process communication paradigm. Shells create anonymous pipes between process invocations sequenced using the '|' character. Named pipes allow pipe stages to be connected together at a time independent from the time that they are created: the pipe endpoints sit in the filesystem. Pipes can span multiple computers, by including invocations of tools such as the Secure Shell (SSH), netcat, or wget within pipelines.

    This project integrates named pipes into web browsers. This goes beyond simply using HTTP(S) for the data transport of a pipe: wget and friends can already do this. The browser integration is to allow pipe creation, configuration, and local and remote pipe linking to be done with user involvement, and with the involvement of the browser's JavaScript engine

12. **An Open Representation for Histories of Application Actions**

    David Eyers (dme@cs.otago.ac.nz)
    Advisor: John Gillanders

    Unlimited undo/redo has vastly improved user experiences of software. However the representation of history is often only intended for transient use, i.e. saving a document, or quitting / crashing an application loses the history. This project aims to create an open format for persistent journals of application actions on a document. An extension would be to develop independent applications that can cooperatively utilise the journal of a single document. Many technologies can be drawn on that achieve at least some of the aims of the project, such as Apple's Core Data framework, database write-ahead logs, source code revision systems, filesystem journals and complex event processing systems.

13. **Giguesaur**

Geoff Wyvill (geoff@cs.otago.ac.nz)
David Eyers (dme@cs.otago.ac.nz)

Our objective is to create a collaborative puzzle game played on a virtual board. Each player uses an iPad. The iPad screens show what the camera sees, but also the pieces of a giant jigsaw puzzle. The players build the jigsaw by picking up the virtual pieces and dropping them in the right places.

Building Giguesaur needs three student projects to be run in cooperation:

(a) **Localisation of the iPad** Design a pattern of markers on the ground by which each iPad can determine where it is. The software must then superimpose the images of the jigsaw pieces in the right place. Part of this project will investigate the pros and cons of different marker types.

(b) **Communication between the iPads** Build a distributed communication system so all the iPads 'know' where all the pieces are and get updated when anything changes. Players should be able to join and leave the game at any time while it is being played. This might be possible without a central server or might need a master iPad or another computer to create the network.

(c) **User Interface and Game Control** If the game is to be easy to learn and engaging to play, it will need a good user interface. This project will develop a touch and gesture interface that makes the game easy to play.

Clearly the three developers will have to work closely together and define rules by which the different parts of the project interact. If successful, this game will be exhibited at the International Science Festival in 2016.

More information can be found at
`http://www.cs.otago.ac.nz/graphics/Giguesaur.html`.

14. **Testing a Model of Attention to Properties**

Ali Knott (alik@cs.otago.ac.nz)
Lech Szymanski (lechszym@cs.otago.ac.nz)
Advisor: Jamin Halberstadt, Dept of Psychology

When we look at an object, we can attend to particular properties it has: for instance, we can notice that a cup is shiny, or dirty, or we can notice that a dog is brown, or furry. No-one knows exactly how the brain implements this attentional operation - but one idea is that after we categorise an object (identifying an object category like 'cup' or 'dog'), we can *inhibit* the properties that are most closely associated with the chosen category, emphasising properties of the object that are somewhat unusual for objects of the identified type. For instance, after identifying an object as a cup, if we inhibit the properties most strongly associated with cups, we can draw attention to properties of the cup that are not found in prototypical cups, like 'dirty', or 'shiny'. The aim of this project is to test this idea about inhibition in an experiment. The project will involve two phases: first a review of literature about how objects and types are represented in the brain, and then the design and running of a behavioural experiment to test the idea about inhibition.

15. **Structuring the Cognitive Map**

Ali Knott (alik@cs.otago.ac.nz)
Advisor: David Bilkey, Dept of Psychology

An area of the brain called the hippocampus contains a set of neurons called 'place cells', that provide a representation of the agent's local environment. Each place cell is most active when the agent is at a particular place in the environment, regardless of the agent's orientation; collectively, the set of place cells represent the set of places in the environment. Of course, different environments have different spatial structures: a corridor is long and narrow, an office room may be rectangular, a garden may be relatively open. The set of place cells must be organised to capture the spatial structure of the current environment, and this organisation must be changed every time the agent moves from one known environment to another. The aim of this project is to develop a model of the function that structures the map of place cells. The project will involve two phases: first a review of research about how place cells reorganise when the agent moves from one environment to another, and second, the design and implementation of a neural network model of a function that imposes a spatial structure on a set of place cells.

16. **Size and Shape in an Object Classification Model**

    Ali Knott (alik@cs.otago.ac.nz)

    Size adjectives (e.g. 'big', 'small') and shape adjectives (e.g. 'tall', 'thin', 'fat') denote relative concepts, not absolute ones: they are defined in relation to the normal size and shape of objects of a given type. For instance, when we say an ant is big, we mean it's big *for an ant*; when we say an elephant is small, we mean it's small *for an elephant*. (In absolute terms a big ant is still much smaller than a small elephant.) These properties are particularly interesting for an object classifier, because object classification must recognise an object's type regardless of its size and shape - yet we must still have expectations about the 'normal' size and shape of objects of different types, in order to be able to use relative adjectives. One idea is that our template for a given object type can be transformed (enlarged, shrunk, stretched) to fit the particular token object being classified, and that relative adjectives describe the nature of the transformations that are applied. The aim of this project is to develop a model of the mechanism that applies this transformation. The project will involve two stages: first, a review of research about spatial transformation operations during object classification, and then implementation of a selected model of these operations.

17. **Hand Gesture Recognition for a Model of Infant Action Perception**

    Ali Knott (alik@cs.otago.ac.nz)
    Advisor: Mark Sagar (University of Auckland)

    This project involves a collaboration with Mark Sagar, formerly of Weta Digital, now at the Laboratory for Animate Technologies in Auckland. Marks group has developed a baby CGI agent: http://www.abi.auckland.ac.nz/en/about/our-research/animate-technologies.html My lab is collaborating with Marks group to simulate the babys perceptual, motor and language skills. The baby agent and its human interlocutor can jointly act on 2D objects on a touch screen. The goal of the current project is to build a system that can monitor the trajectory and shape of a human users hand as it interacts with objects on the screen. Its a computer vision project. There will be two cameras, viewing the user from the front and from the side. The first goal is to track the users hand in real time, and compute its trajectory as the user executes a hand action. The second goal is to classify the shape of the users hand, assuming a simple repertoire of possible hand shapes: an open palm, a closed fist, a pointed finger, and a hand-grasp preshape.

18. **We can Know Everything that Happens in our House in Real Time, Remotely**

    Haibo Zhang (haibo@cs.otago.ac.nz)

    Have you every worried about your house when you were away, because you were uncertain whether you have turned off the oven and have locked the door? Have you ever worried about your grandparents who live alone? Through the combination of wireless sensor networking and smartphone technologies, you can easily monitor you house remotely via your smartphone, e.g. detect fire and falls of the elderly. In last year, we have developed a prototype for such a system. In this year we are going to develop a smart client program for smartphones and a protocol for transferring photos from sensor to smartphone. This project can host two students, with each working on one function. The first one requires experience on Android programming.

19. **Tracking the Motion of the Human Body using 3D Inertial Measurement Units**

    Haibo Zhang (haibo@cs.otago.ac.nz)
    with some subset of:
    David Eyers (dme@cs.otago.ac.nz)
    Zhiyi Huang (zhuang@cs.otago.ac.nz)
    Lech Szymanski (lechszym@cs.otago.ac.nz)
    Steven Mills (steven@cs.otago.ac.nz)

    Human motion tracking has a wide range of applications such as in sports, gaming, and medical rehabilitation. In this project we are going to develop a motion capture system using wireless sensors and 3D inertial measurement units. Please watch the following YouTube videos to get a feeling of such systems.

    - https://www.youtube.com/watch?v=nRzS39UtTF4
    - https://www.youtube.com/watch?v=b9TQcv0Asys
    - https://www.youtube.com/watch?v=tyiuocPvUes

    We have already designed the hardware, implemented the communication protocol, and carried out some preliminary work on animating the captured body motions using Blender. In this project we will add more functions and produce a demo for real-time body motion tracking. Background and experience with computer graphics and Python programming would be helpful to work on this project.

20. **Having fun with programmable hardware.**

    Richard O'Keefe (ok@cs.otago.ac.nz)
    David Eyers (dme@cs.otago.ac.nz)

    A Field Programmable Gate Array is reprogrammable hardware. You can write a hardware design in a high level language (the two main ones are VHDL and Verilog) and compile it, down load it into an FPGA, and run it. VHDL and Verilog can also be compiled to silicon, to make application specific integrated circuits, so FPGAs are often used to prototype things. One can design and implement ones own CPU, but it is more useful to implement an algorithm directly without any instruction interpretation overhead.

    We have a couple of FPGA boards and some textbooks and free software.

    It would be particularly interesting to see if these things can be used in information retrieval. For example, the inventor of the widely used Porter stemming algorithm has devised a special purpose programming language for writing stemmers; is it possible to translate stemmers into VHDL so they can be compiled into hardware? Even more basically, Unicode is enormously complicated, even a simple thing like asking "are these two strings equal" is hard. Can we implement that efficiently on an FPGA?

    If you have something of your own you'd like to try, that would be welcome too.

21. **Latent Methods**

    Richard O'Keefe (ok@cs.otago.ac.nz)

    I've recently come up with an idea called "latent methods". These are methods that you define in some class, but they do not become available until a subclass implements the interface(s) they are declared to need. This project would involve

    (a) critiquing the idea and comparing it with other approaches like multiple inheritance,
    (b) implementing latent methods in some OO programming language, possibly by modifying a compiler or writing a preprocessor, and
    (c) writing test cases for your implementation that can be used as examples to show the usefulness (or otherwise!) of the idea.

22. **Semantic Relationships in Source Code**

    Richard O'Keefe (ok@cs.otago.ac.nz)

    Some programming languages, notably C#, Java, and Smalltalk, have a standard way to attach "annotations" to code. This is usually used to provide the compiler and run-time system with more information like “this reference should not be null”, “this is a test case”, or “this is part of an Object-Relational Mapping”. The nearest we tend to get to something for the programmer’s edification is “don’t use this any more”.

But chunks of source code are related to each other and to external documents in ways that the compiler and run time system do not need to know about, while maintenance programmers would love to know. The idea is to take an existing body of code that you are not familiar with, read it, and note what kinds of links are either informally specified or entirely missing but important to you, so that we can build up this catalogue. You should also review what has been done in Java and C# and Smalltalk. For more work, you might construct an actual set of annotation definitions that Java (or C# or Smalltalk) programmers might use and provide some tools (perhaps by extending an existing IDE) to exploit them.

23. **Speeding up LaTeX**

Richard O’Keefe (ok@cs.otago.ac.nz)

LaTeX is an open source typesetting package used by many (particularly academic researchers) due to the high quality of its output – particularly mathematical notation, its clean separation of document content from formatting, its vast extension library, and its plain-text source format. Much of LaTeX’s typesetting quality is due to it *not* being a WYSIWYG system: when generating output, it can try multiple different ways to place and hyphenate words in a paragraph to optimise the number of words per line for the best readability. This project aims to apply analysis of how LaTeX runs, to speed up its operation for particular use cases, for example the generation of PowerPoint-like presentations. The pipe-dream is incremental LaTeX recompilation.

24. **Processing patents**

Richard O’Keefe (ok@cs.otago.ac.nz)

We have a copy of a collection of 100,000 US patents. Perhaps the most striking thing about patents is how weird the language is. There are some ideas about how we could measure it: few occurrences of ‘I’ and ‘you’, long complex sentences, more nominalisations. There are some ideas that haven’t been tried yet, like sentiment analysis: do they use terms of approval or disapproval more or less than ordinary text?

We have about 10,000 documents in English from the European parliament, several years’ worth of IEEE articles in English, and a snapshot of the Wikipedia to compare the patents with.

25. **Bringing Command Shells and Terminals Closer Together**

    Richard O'Keefe (ok@cs.otago.ac.nz),
    David Eyers (dme@cs.otago.ac.nz)

    Apart from being able to edit the current command, most functions of the command shells that run within today's terminal emulators assume little more functionality than what is provided by a Teletype. That is, output from commands scrolls off the top of the terminal emulator (and possibly into scroll-back region, if active).

    Enriching the connection between shells and terminal emulators can unlock so much more functionality. For example, why not track the difference between command entry, program input and program output? This would allow for folding and unfolding of the output of commands. Why not provide back-scroll regions for each command independently showing, say, an updating 4-line tail of multiple commands that are all running simultaneously.

    By independently recording the output of each previous command, future text processing could be done directly from this recorded output, navigating to it using the enriched shell, and avoiding the need to actually re-run the command that produced the output.

26. **Image Analysis for Core Samples**

    Steven Mills (steven@cs.otago.ac.nz)
    Brendan McCane (mccane@cs.otago.ac.nz)

    Geologists often take core samples of rocks for analysis, and recently colleagues in the Geology department have been using CT scans to analyse core samples taken from NZ's Alpine Fault. Currently this is a manual process, identifying and classifying fractures in the rock. These fractures appear as dark regions in the 3D image data, and the purpose of this project is to explore techniques for automatically finding and describing these fractures.

27. **Making 3D Line Drawings from Images of Buildings**

    Steven Mills (steven@cs.otago.ac.nz)

    There are established techniques for creating 3D models from collections of images. These are usually based on point features in the images, which leads to a representation in terms of 3D point clouds. Scenes of and inside buildings, however, are often rich in line features. This project will look at existing techniques for using line features in 3D reconstruction, implement some of them, and work towards making 3D line sketches from images of buildings.