

# COMP 621 Suggested Projects - Fall 2015

Due Dates:

- Project descriptions available and signups, Monday Oct 19 to Monday Oct 26
- Project proposal due, Monday Nov 2 (pdf by email)
- Project update #1 due, Friday Nov 13 (pdf by email)
- Project update #2 due, Friday Nov 20 (pdf by e-mail)
- Project presentations start: Wednesday November 25
- Project presentations end: Monday December 7
- Project reports due, Monday December 7th to Wednesday December 16 (pdf by e-mail, zipped source)

## 1 Introduction

The purpose of the final project is to give you a chance to focus on a specific problem, and to work out the solution to this problem in detail. Your project should give you the opportunity to review the current literature in your chosen area, to implement either the ideas you read about or to implement your own improvements on those ideas, and to evaluate the effectiveness of the ideas. You are encouraged to try to develop improvements over the current state-of-the-art. Even if your ideas do not turn out better, it is still worth trying (after all, that is what research is all about). It is worth noting that publishing novel work done in this course is not out of the question.

I have listed the various possible project areas below. Students should work individually. Projects must be done specifically for this course and should not contain material used for other course projects.

The projects will be given out on a first-come, first-served basis. You may sign-up for your selected project in the spreadsheet announced in the e-mail. Two students may choose the same project, but they must do individual projects using complementary solutions.

You are also free to design your own project, or to change a suggested project to more suit your interests. If you wish to modify or design your own project, please contact the instructor to discuss the project.

Your project proposal should be done using the document outline in:

<http://www.sable.mcgill.ca/~hendren/621/Docs>

Latex documents are preferred, but not absolutely necessary. In any case, you should follow the overall organization given in <http://www.sable.mcgill.ca/~hendren/621/Docs/prop621.tex>.

Your final report should follow the outline defined in <http://www.sable.mcgill.ca/~hendren/621/Docs/rep621.tex>.

Please submit .pdf files of the form `prop621_yourname.pdf` and `rep621_yourname.pdf` to `hendren@cs.mcgill.ca`.

You are also asked to submit two project updates. These should be short (approx 1 page) .pdf documents which summarize the current state of your project, discusses the plan for the next 10 days to two weeks, and discusses any

changes to your project plan. These documents should be called `update1_yourname.pdf` and `update2_yourname.pdf` and should also be e-mailed to `hendren@cs.mcgill.ca`.

Please ensure that all e-mails have a subject title that starts with "COMP621".

## 2 Suggested Projects

### 2.1 Sparse Matrix optimization for MATLAB

Most of our work to date has concentrated on optimizing dense matrices. However, MATLAB also supports sparse matrices which can give good performance improvements (when used correctly). See [http://www.mathworks.com/help/pdf\\_doc/otherdocs/simax.pdf](http://www.mathworks.com/help/pdf_doc/otherdocs/simax.pdf).

This project would be to:

- understand and identify under what circumstances sparse matrices should be used
- understand how the MATLAB source code should be written so as to maximize the benefit of using sparse matrices
- identify a collection of at least 4-5 benchmarks (computation kernels) that use sparse matrices. You may find the aspect described in the AOSD 10 paper to be useful here. See <http://www.sable.mcgill.ca/mclab/aspectmatlab/>
- identify analyses and transformations that are useful for transforming MATLAB programs so as to benefit from sparse matrices
- as much as possible, implement the analyses and transformations
- study the performance benefit on the benchmarks

### 2.2 Rewrite loops to calls to library routines or vector statements

Problem Statement: MATLAB is an interactive programming environment for scientific computing. It is designed for vector and matrix computation. Naively written MATLAB code can suffer a performance slowdown. One way to address the performance problem of high-level domain-specific languages, such as Octave or MATLAB is to use optimized library routines instead performing operations on vectors and matrices iteratively (using loops).

While it is tempting for those with experience in traditional programming languages to use explicit loop constructs for mathematical operations, this temptation should be resisted strenuously. For example, instead of:

```
port_val = 0;
for j = 1:n
    port_val = port_val + (holdings(j) * prices(j));
end
```

Instead:

```
port_val = holdings*prices;
or
port_val = mtimes(holdings,prices);
```

The latter is more succinct, far clearer, and will run much faster if the appropriate library calls are made by the MATLAB system. The goal of this project is to develop one or more analyses that detect loops that correspond to known and efficient library routines, or to equivalent vectorized operations and to automatically rewrite the loops to higher-level code that will be compiled to an efficient library routine.

See the paper at <http://portal.acm.org/citation.cfm?id=1267941> for one relevant reference.

## 2.3 Design and implement an analysis to identify for loops in MATLAB that can be safely transformed to parfor loops

In the MiX10 compiler (and other parallel backends) we can take advantage of target language's superior parallel performance compared to MATLAB, for the parfor loop construct. However, for various reasons, many MATLAB programmers use a for loop where they can actually take advantage of a parfor loop. It would be nice to identify such for loops and automatically transform them to parfor loops. You can read details about parfor at <http://www.mathworks.com/help/distcomp/getting-started-with-parfor.html#brdq6p-1>.

This analysis can be made cleverer by not just looking what's inside the loop but also analysing code outside the for loop or maybe by making some safe transformations to the loop; for example, The code below cannot be transformed to a parfor loop because the value of 'd' after the loop depends on order of execution of iterations. A parfor loop, instead of throwing an error, will force 'd' to retain its value before the loop. If we know that the definition of 'd' inside the loop is never used outside the loop, we can transform it to parfor.

```
clear A
d = 0; i = 0;
parfor i = 1:4
    d = i*2;
    A(i) = d;
end
A
d
```

You can read more about reduction variables at <http://www.mathworks.com/help/distcomp/advanced-topics.html>.

One might also consider developing an aspect which could observe loops and find those that behave in a way in which is suitable for parfor loops, and then use that information to guide the programmer to convert the loops to parfor loops.

## 2.4 Improving AspectMatlab Performance or Functionality

AspectMatlab can be improved either by improving its performance or by adding new functionality.

In Assignment 1 we have seen that weaving aspects can introduce significant runtime overheads. Thus, one potential project would be to design new code generation strategies or optimizations, in order to reduce the overheads. You would start with the most recent version Previously written aspects, plus aspects used in Assignment #1, can be used for benchmarks.

The second option would be to improve AspectMatlab by adding new functionality. This could include new patterns or pattern operators, or it could include matching traces (for example, a simple version of tracematches as implemented for AspectJ).

## 2.5 Domain-specific language for specifying dataflow analyses

In assignment #2, you are exposed to McSAF, McLab's static analysis framework. While McSAF makes the job of writing a dataflow analysis simpler, there is still a fair amount of boilerplate code involved. Also, McSAF is only used by McLab's Java-based tools; the McVM project, implemented in C++, does not use it. Actually the McVM project doesn't have an analysis framework (although one is planned); all of its analyses are implemented from scratch, and contain lots of duplicate code. Even if it did, we would still have to specify analyses twice; once for McSAF, once for McVM. Even now we have multiple implementations of e.g. reaching definitions and live variables analysis.

The purpose of this project is to address these issues with an extra level of abstraction, by designing and implementing a small declarative language for specifying dataflow analyses. An analysis could be specified at a relatively high-level as seen in class, by describing the data, the direction, the merging operation, and providing dataflow equations for different kinds of AST nodes. Coming up with a good natural syntax for this is part the project. A compiler would take such a description as input and generate code for the analysis (assuming the existence of an analysis framework to code against). A McSAF backend should be provided, but the compiler should be extensible enough to allow us to add (for example) a backend for McVM's analysis framework once it's ready.

This is probably a hard problem in general. We are restricting ourselves to analyses operating on the Matlab (or Natlab) AST, not any arbitrary tree. Also, we are assuming that a suitable (i.e. McSAF-like) analysis framework exists for the targets we generate code for.

## 2.6 Array Growth Optimization in Matlab

Matlab is a popular programming language for scientific computing. Several features of the language contribute to its appeal. Two of these features are dynamic typing and dynamic array growth. The two features contribute to the flexibility of the Matlab language and can aid rapid program development. However, dynamic array growth presents performance challenges. Consider the following Matlab function.

```
function grow0()
    s = 0;
    for i=1:10
        a(i) = sin(i);
        s = s + a(i);
    end
    disp('Sum of ');
    disp(a);
    disp(['is: ', num2str(s)]);
end
```

At each iteration of the loop, array `a`, which is undefined before the loop, will be expanded by 1 and the `sin` of the current `i` is assigned to the location in `a` indexed by `i`. This can cause frequent re-allocation of the array and can be a source of significant runtime overhead. After the execution of the loop, array `a` will contain 10 elements. Array concatenation operation can also grow an array dynamically. For instance,

```
a=0;
for i=1:10
    a = [a, i];
end
disp(a);
```

Array `a` grows with the loop's iterations. At each iteration, the current value of `i` is appended to the array so that the array grows in length by 1. After the execution of the loop, array `a` has 11 elements: 0 to 10. Other forms of

dynamic array growth in Matlab exist. Multidimensional arrays can also grow dynamically.

### **2.6.1 Problem Statement**

The problem to be addressed is the development an array growth optimization. This project will involve:

- an investigation of different forms of array growth in Matlab;
- the development of analyses and transformations that can detect dynamic array growths;
- a technique for estimating the optimal size of a growing array; and
- the transformation that pre-allocates arrays using the results of the analyses.

### **2.6.2 Potential Optimization Benefits**

The ability to estimate the optimal size of an array statically can enable the pre- allocation of the array before a loop, which can lead to a significant performance improvement. It can also simplify the elimination of redundant array bound checks generated to ensure that array accesses are within the bounds of the array.

## **2.7 Lean and Mean Framework for McLab**

In assignmnt #2 you used the current McLab framework to implement an analysis. This project is intended to design and implement a new analysis framework that is either: (1) simpler to use, or (2) provides more efficient solvers.

If you choose to design a simpler-to-use framework, then your challenge would be to demonstrate that analyses currently implemented in the McSAF framework could be more easily specified in your system. If you choose to implement a faster solver, then you need to demonstrate that your solver is faster on a number of analyses/benchmarks.

## **2.8 Improvements in McLab Back-ends**

There are two backends for McLab, Mc2For ([www.sable.mcgill.ca/mclab/projects/mc2for](http://www.sable.mcgill.ca/mclab/projects/mc2for)) which generates Fortran code, and MiX10 which generates X10 code ([www.sable.mcgill.ca/mclab/projects/mix10/](http://www.sable.mcgill.ca/mclab/projects/mix10/)).

The challenge in this project would be to examine the output of one of these back-ends to identify how the generated code code be improved, and then to implement the necessary analyses/transformations to implement that improvement. Your experimental results would compare the size and/or speed of the original code with the improved code provided by your system.

## **2.9 A new Tamer analysis**

In assignment #2 you used the McSAF framework. There is an even lower level framework based on the TameIR, which allows for interprocedural analysis. Currently there are analyses to approximate types, shapes, whether or not a result is complex, whether or not a variable can be stored in an integer, and some initial range analysis (what range of values can a variable take on).

Your challenge in this project it to identify a new interprocedural analysis problem that could be used to further improve the back-ends that build on TameIR. For example, a good range analysis would be useful for better determining if array accesses might be out of bounds.

### **3 Propose a Project**

You may also propose a new project that is not listed here. If you choose something else, it should have some element of "program" analysis in it, although the "program" being analyzed could be something from a different domain such as models or games.

If you are proposing your own project, please e-mail the instructor ([hendren@cs.mcgill.ca](mailto:hendren@cs.mcgill.ca)) with a brief project proposal similar to the ones above, and add your project title to the list below.