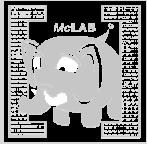


McLAB: A toolkit for static and dynamic compilers for MATLAB



Laurie Hendren
McGill University

Leverhulme Visiting Professor
Department of Computer Science
University of Oxford

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren Intro - 1

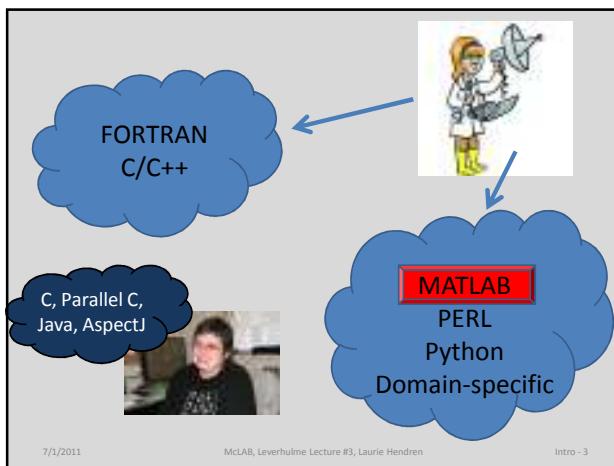
Amina Aslam
Toheed Aslam
Andrew Casey
Maxime Chevalier-Boisvert
Jesse Doherty
Anton Dubrau
Rahul Garg
Maja Frydrychowicz
Nurudeen Lameed
Jun Li
Soroush Radpour
Olivier Savary

Overview



- Why MATLAB?
- Overview of the McLAB tools
- McVM – a Virtual Machine and Just-In-Time (JIT) compiler
- McFOR – translating MATLAB to FORTRAN95

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren Intro - 2



FORTRAN
C/C++

C, Parallel C,
Java, AspectJ

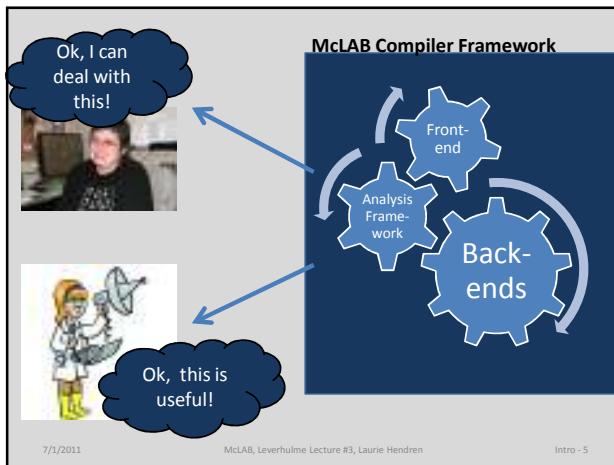
MATLAB
PERL
Python
Domain-specific

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren Intro - 3

Culture Gap

Scientists / Engineers	Programming Language / Compiler Researchers
<ul style="list-style-type: none"> • Comfortable with informal descriptions and “how to” documentation. • Don’t really care about types and scoping mechanisms, at least when developing small prototypes. • Appreciate libraries, convenient syntax, simple tool support, and interactive development tools. 	<ul style="list-style-type: none"> • Prefer more formal language specifications. • Prefer well-defined types (even if dynamic) and well-defined scoping and modularization mechanisms. • Appreciate “harder/deeper/more beautiful” programming language/compiler research problems.

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren Intro - 4



Ok, I can deal with this!

Ok, this is useful!

McLAB Compiler Framework

Front-end

Analysis Framework

Back-ends

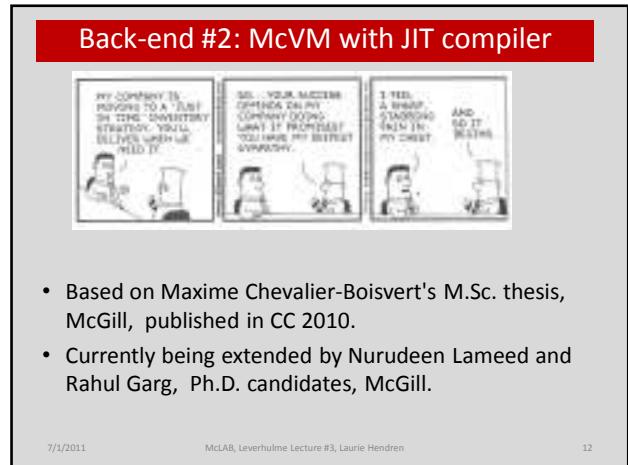
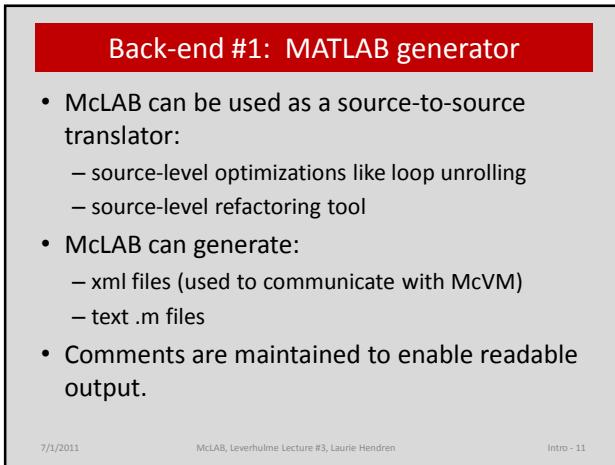
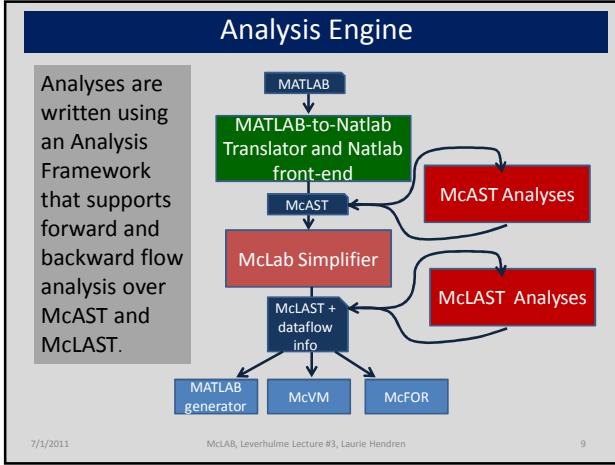
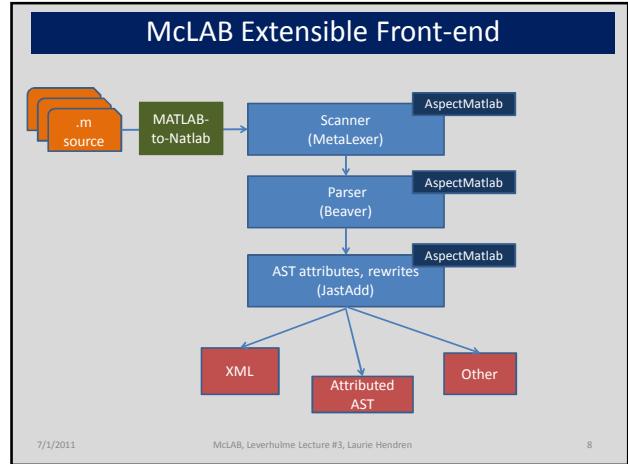
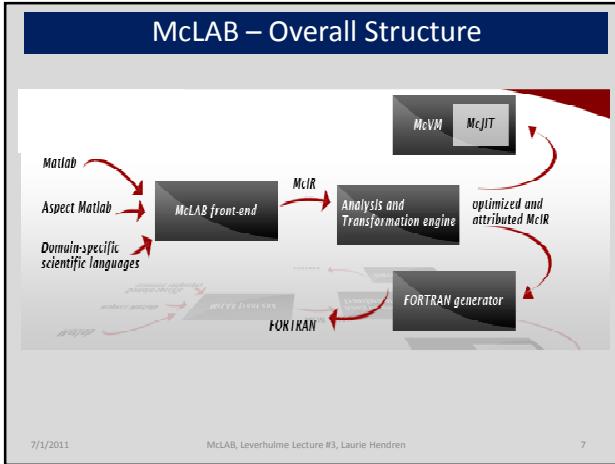
7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren Intro - 5

Goals of the McLab Project

- Improve the understanding and documentation of the semantics of MATLAB.
- Provide front-end compiler tools suitable for MATLAB and language extensions of MATLAB.
- Provide a flow-analysis framework and a suite of analyses suitable for a wide range of compiler/soft. eng. applications.
- Provide back-ends that enable experimentation with JIT and ahead-of-time compilation.

Enable PL, Compiler and SE Researchers to work on MATLAB

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren Intro - 6



McVM-McJIT

- The dynamic nature of MATLAB makes it very suitable for a VM/JIT.
- MathWorks' implementation does have a JIT, although technical details are not known.
- McVM/McJIT is an open implementation aimed at supporting research into dynamic optimization techniques for MATLAB.

7/1/2011

McLAB, Leverhulme Lecture #3, Laurie Hendren

Backends- 13

Design Choices for JITs

- Interpreter + JIT compiler with various levels of optimizations.
- Fast JIT for naïve code generation + optimizing JIT with various levels of optimizations.
- McVM uses the 1st option because it simplifies adding new features, if a feature is not yet supported by the JIT it can back-up to the interpreter implementation, which is easy to provide.

7/1/2011

McLAB, Leverhulme Lecture #3, Laurie Hendren

14

McVM Design

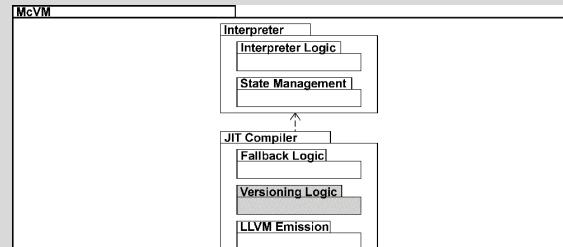
- A basic, but fast, interpreter for the MATLAB language.
- A garbage-collected JIT Compiler as an extension to the interpreter.
- Easy to add new data types and statements by modifying only the interpreter.
- Supported by the LLVM compiler framework and some numerical computing libraries.
- Written entirely in C++; interface with the McLAB front-end via a network port.

7/1/2011

McLAB, Leverhulme Lecture #3, Laurie Hendren

Backends- 15

McVM Organization

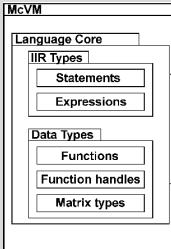


7/1/2011

McLAB, Leverhulme Lecture #3, Laurie Hendren

16

McVM Organization

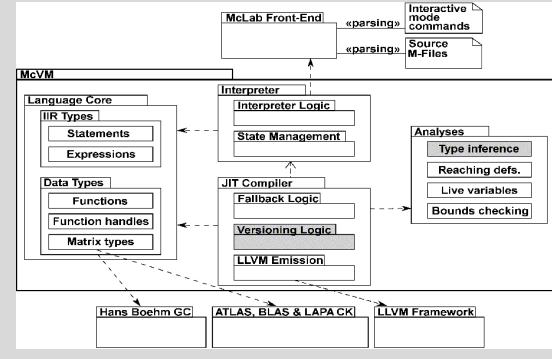


7/1/2011

McLAB, Leverhulme Lecture #3, Laurie Hendren

17

McVM Organization



7/1/2011

McLAB, Leverhulme Lecture #3, Laurie Hendren

18

MATLAB Optimization Challenges

```
float sumvals(float start, float step, float stop)
{
    float i = start;
    float s = i;

    while (i < stop)
    {
        i = i + step;
        s = s + i;
    }

    return s;
}
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

19

MATLAB Optimization Challenges

```
function s = sumvals(start, step, stop)
    i = start;
    s = i;

    while i < stop
        i = i + step;
        s = s + i;
    end

function caller()
    a = sumvals(1, 1, 10^6);
    b = sumvals([1 2], [1.5 3], [20^5 20^5]);
    c = [a b];

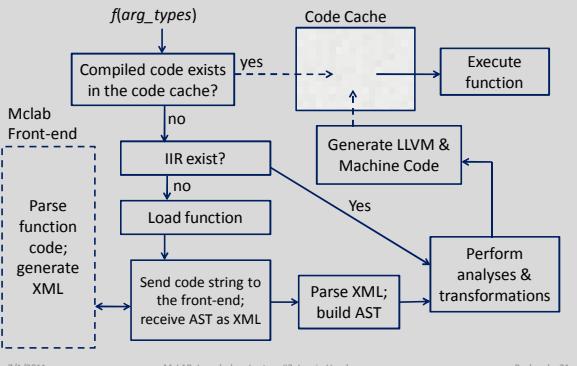
    disp(c);
end
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

20

McJIT: Executing a Function



7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

Backends- 21

Just-In-Time Specialization (1)

```
>> a = sumvals(1, 1, 10^6);
>> b = sumvals([1 2], [1.5 3], [20^5 20^5]);
>> a = sumvals(1, 1, 500);
>> c = [a b];
>> disp(c);
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

22

Just-In-Time Specialization (2)

```
>> a = sumvals(1, 1, 10^6);
    ↓
Interpreter.runCommand("a = sumvals(1, 1, 10^6);")
    ↓
Interpreter.callFunction(sumvals, [1, 1, 10^6]);
    ↓
JIT.callFunction(sumvals, [1, 1, 10^6]);
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

23

Just-In-Time Specialization (3)

```
JITCompiler.callFunction(sumvals, [1, 1, 10^6]);
    ↓
sumvalsJIT = JITCompiler.compileFunction(sumvals, [<scalar int>, <scalar int>, <scalar int>]);
    ↓
function s = sumvals(start <scalar int>, step <scalar int>, stop <scalar int>)
    i = start;
    s = i;

    while i < stop
        i = i + step;
        s = s + i;
    end
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

24

Just-In-Time Specialization (4)

```
JITCompiler.callFunction(sumvals, [1, 1, 10^6]);
sumvalsJIT = JITCompiler.compileFunction(sumvals, [scalar int], [scalar int], [scalar int]);
function s <scalar int> = sumvals(start <scalar int>, step <scalar int>, stop <scalar int>
i <scalar int> = start;
s <scalar int> = i;
while i < stop
    i <scalar int> = i + step;
    s <scalar int> = s + i;
end
end
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

25

Just-In-Time Specialization - 2nd example

```
>> a = sumvals(1, 1, 10^6);
>> b = sumvals([1 2], [1.5 3], [20^5 20^5]);
>> a = sumvals(1, 1, 500);
>> c = [a b];
>> disp(c);
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

26

JIT – second specialization (1)

```
JITCompiler.callFunction(sumvals, [[1 2], [1.5 3], [20^5 20^5]]);
sumvalsJIT2 = JITCompiler.compileFunction(sumvals, [1x2 int], [1x2 real], [1x2 int]);
function s = sumvals(start <1x2 int>, step <1x2 real>, stop <1x2 int>
i <1x2 int> = start;
s <1x2 int> = i;
while (i <1x2 int>) < (stop <1x2 int>
    i <1x2 real> = i + step;
    s <1x2 real> = s + i;
end
end
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

27

JIT – second specialization (2)

```
JITCompiler.callFunction(sumvals, [[1 2], [1.5 3], [20^5 20^5]]);
sumvalsJIT2 = JITCompiler.compileFunction(sumvals, [1x2 int], [1x2 real], [1x2 int]);
function s <1x2 real> = sumvals(start <1x2 int>, step <1x2 real>, stop <1x2 int>
i <1x2 int> = start;
s <1x2 int> = i;
while (i <1x2 real>) < (stop <1x2 int>
    i <1x2 real> = i + step;
    s <1x2 real> = s + i;
end
end
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

28

JIT – third specialization same as first

```
>> a = sumvals(1, 1, 10^6);
>> b = sumvals([1 2], [1.5 3], [20^5 20^5]);
>> a = sumvals(1, 1, 500);
>> c = [a b];
>> disp(c);
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

29

Type and Shape Inference

- In MATLAB, work with incomplete information
 - Dynamic loading: working with incomplete program
 - Dynamic typing : variables can change type
- Know argument types, what can we infer?
 - Propagate type info to deduce locals type and return type
- Forward dataflow analysis
 - Based on abstract interpretation
 - Structure-based fixed point
 - Annotates AST with type info

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

30

Flow Analysis Summary

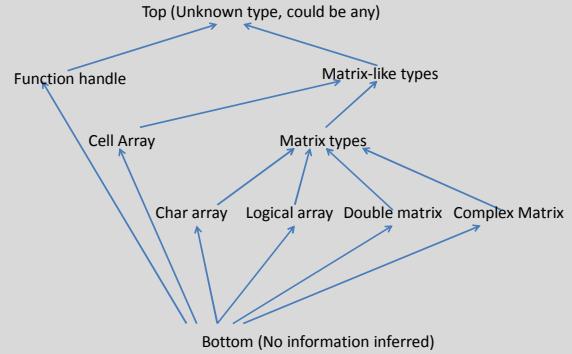
- Start from known argument types
- Propagate type information forward
- Use a transfer function for each expression
 - Transfer functions provided for all primitive operators
 - Library functions provide their own transfer functions
 - Function calls resolved, recursively inferred
- Assignment statements can change var. types
- Merge operator
 - Union + filtering

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

31

Lattice of McVM types



7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

Backends-32

Type Abstraction Properties

- Collection of simple abstractions
 - Specific features computed in parallel
- Represent variable types with 8-tuples:

```
<overallType, is2D, isScalar, isInteger,
 sizeKnown, size, handle, cellTypes>
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

33

Type Abstraction Properties

```
a = [1 2];
```

```
type(a) =
<overallType = double, is2D = T, isScalar = F,
isInteger = T, sizeKnown = T, size = (1,2),
handle = null, cellTypes = {}>
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

34

Variable Types: Type Sets

```
if (c1)
  a = [1 2];
elseif (c2)
  a = 1.5;
else
  a = 'a';
end

type(a) =
<overallType = double, is2D = T, isScalar = F,
isInteger = T, sizeKnown = T, size = (1,2),
handle = null, cellTypes = {}>,
<overallType = double, is2D = T, isScalar = T,
isInteger = F, sizeKnown = T, size = (1,1),
handle = null, cellTypes = {}>,
<overallType = char, is2D = T, isScalar = T,
isInteger = T, sizeKnown = T, size = (1,1),
handle = null, cellTypes = {}>
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

35

Type Set Filtering

```
type(a) =
<overallType = double, is2D = T, isScalar = F,
isInteger = T, sizeKnown = T, size = (1,2),
handle = null, cellTypes = {}>,
<overallType = double, is2D = T, isScalar = T,
isInteger = F, sizeKnown = T, size = (1,1),
handle = null, cellTypes = {}>
```

Becomes:

```
type(a) =
<overallType = double, is2D = T, isScalar = F,
isInteger = F, sizeKnown = F, size = (),
handle = null, cellTypes = {}>
```

7/1/2011

McLAb, Leverhulme Lecture #3, Laurie Hendren

36

Transfer Functions

```

a = 5 * 10e11;
b = 1.0e12 * [0.8533 1.7067];
c = [a b];

type(a) =
<overallType = double, is2D = T, isScalar = T, isInteger = T,
sizeKnown = T, size = (1,1), handle = null, cellTypes = {}>

type(b) =
<overallType = double, is2D = T, isScalar = F, isInteger = F,
sizeKnown = T, size = (1,2), handle = null, cellTypes = {}>

type([a b]) =
<overallType = double, is2D = T, isScalar = F, isInteger = F,
sizeKnown = T, size = (1,3), handle = null, cellTypes = {}>

```

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 37

Experimental Results

- 20 benchmark programs
- FALCON, OTTER, etc. Some made by McLAB
- Measured
 - Dynamic availability of type info.
 - Number of versions compiled
 - Compilation time
 - 0.55s per benchmark, on average

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 38

Results of Type Analysis

Type Inference Efficiency

Category	Percentage (%)
Top sets	~15
Singletton sets	~85
Scalars known	~92
Mat. size known	~65

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 39

How many versions...

Functions and Versions Compiled

Benchmark	# functions	# versions
adapt	2	2
clos	2	2
dich	2	2
diff	2	2
edit	2	2
fild	2	2
fft	2	2
fiff	2	2
crni	3	3
mbtr	3	3
nb1d	3	3
nb3d	3	3
nnet	4	4
capr	5	5
nfc	5	5
schr	8	9
play	6	10
sdku	10	11
svd	11	15
beul	9	16

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 40

Performance Results

- Experimental setup
 - Core 2 Quad Q6600, 4GB RAM
 - Ubuntu 9.10, kernel 2.6.31, 32-bit
 - All timings averaged over 10 runs
- Comparing
 - McVM interpreter, McVM JIT w/ spec.
 - MATLAB R2009a
 - GNU Octave 3.0.5
 - McFor / GNU Fortran 4.4.1

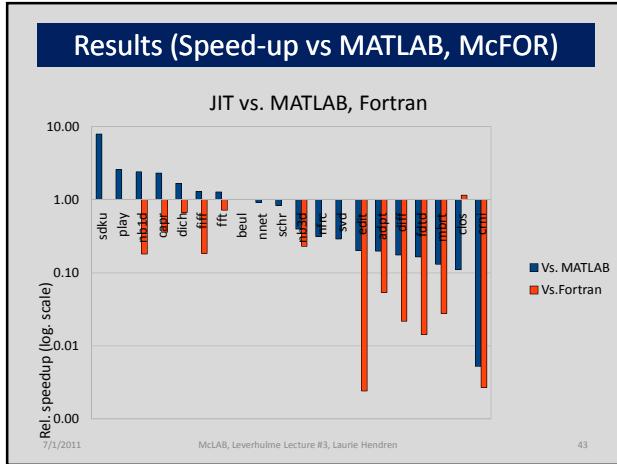
7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 41

Results (Speed-up vs Interpreters)

JIT Speedup

Benchmark	Vs. Interp. (Rel. speedup)	Vs. Octave (Rel. speedup)
capr	~1000	~1200
fiff	~500	~800
sdku	~10	~100
clos	~1	~5
edit	~1	~5
crni	~1	~5
schr	~8	~10
nb1d	~1	~5
nb3d	~1	~5
nnet	~4	~5
capr	~1	~5
nfc	~5	~5
play	~6	~10
sdku	~10	~15
svd	~1	~5
beul	~9	~15

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 42



- ### About Slow Benchmarks
- **crni** benchmark takes 1321s to execute in McVM, 6.95s in MATLAB
 - ~4x faster than Octave, but still slow
 - Why?
 - Scalars known 68.7%, one of the lowest ratios
 - Unknown types propagated through entire benchmark
 - Weakness of type inference system to be fixed in future work
- 7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 44

Back-end #3: McFOR FORTRAN95 generator

I COULD RESTRUCTURE THE PROGRAM'S FLOW OR USE ONE LITTLE 'GOTO' INSTEAD.

EH, SCREW GOOD PRACTICE. HOW BAD CAN IT BE?

goto main_sub5;
COMPILE

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 45

- ### Goals of McFOR
- Handle as large a sub-set of MATLAB as possible, while staying in the "static" setting.
 - Generate code that can be effectively compiled by modern FORTRAN compilers.
 - Make the generated code readable by programmers.
 - Allow longer compile times and whole program analysis.
 - Limit the need for type annotations.
- 7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 46

Challenges

- Determining which identifiers are variables and which are functions.
- Finding static types which match those of FORTRAN.
- Mapping high-level MATLAB array operations to the FORTRAN95 equivalents.
- Handling reshaping implicit in MATLAB operations, including concatenation.

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 47

Handling Incompatible Types

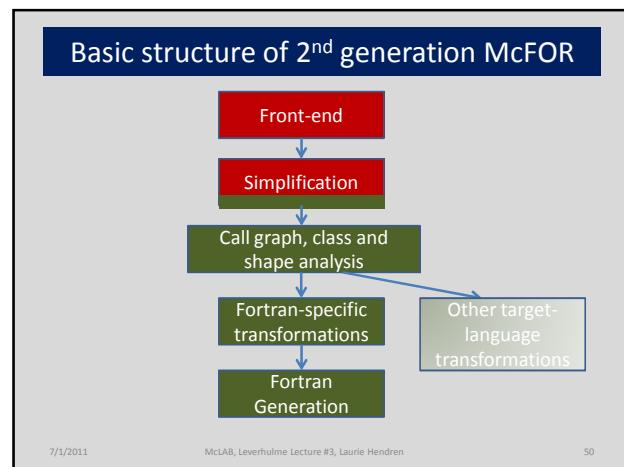
Alpha Nodes	Eliminating Alpha Nodes
<pre>x = 0; if (i>0) S1: x = foo(i); else S2: x = bar(i); end x = alpha(S1,S2); y = x;</pre>	<pre>x = 0; if (i>0) x = foo(i); else x1 = bar(i) y = x1;</pre>

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 48

Value Range Propagation

MATLAB	FORTRAN
n = floor(11.5); //n=[n,n]	n = foor(10.5);
for i = 1:n // i=[1,n]	DO i = 1,n
x = 1+2*i; // x=[3,1+2*n]	x = (1+(2*i));
	IF((.NOT.ALLOCATED(A))) THEN
	ALLOCATE(A((1+(2*n))));
	END IF
A(x) = i; // Size(A)=	A(x) = i;
end	END DO
n = fix(n/2); // n = [n,n]	n = fix(n/2);
	ARRAYBOUNDSCHECKING(A,[n+1]);
A(n+1) = n;	A(n+1) = n;

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren Intro - 49



Related Work

- Procedure cloning: Cooper et al. (1992)
- MATLAB type inference: Joisha & Banerjee (2001)
 - Suggested for error detection
- MATLAB Partial Evaluator: Elphick et al. (2003)
 - Source-to-source transformation
- MaJIC: JIT compilation and offline code cache (2002)
 - Speculative compilation MATLAB to C/Fortran
- Psyco: Python VM with specialization by need (2004)
- TraceMonkey: JIT optimization of code traces (2009)

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 51

Ongoing Work

- McVM:**
 - profile-guided optimization and re-optimization with on-stack replacement
 - target GPU/multi-core
- McFOR:**
 - "decompile" to more programmer-friendly FORTRAN95
 - refactoring toolkit to help restructure "dynamic" features to "static" features

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 52

Conclusions

- McLAB is a toolkit to enable PL, Compiler and SE research for MATLAB
- front-end for language extensions
- analysis framework
- three back-ends including McVM and McFOR

<http://www.sable.mcgill.ca/mclab>

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 53

Intermediate Representation

EXTRA SLIDES

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren 54

Internal Intermediate Representation

- A simplified form of the Abstract Syntax Tree (AST) of the original source program
- It is machine independent
- All IIR nodes are garbage collected

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren Backends- 55

IIR: A Simple MATLAB Program

.m file

```
function a = test(n)
    a = zeros(1,n);
    for i = 1:n
        a(i) = i*i;
    end
end
```

IIR form

```
function [a] = test(n)
    a = zeros(1, n);
    $t1 = 1; $t0 = 1;
    $t2 = $t1; $t3 = n;
    while True
        $t4 = ($t0 <= $t3);
        if ~$t4
            break;
        end
        i = $t0;
        a(i) = (i * i);
        $t0 = ($t0 + $t2);
    end
end
```

7/1/2011 McLAB, Leverhulme Lecture #3, Laurie Hendren Backends- 56

