

Taming MATLAB

Anton Dubrau and Laurie Hendren

McGill University

anton.dubrau@mail.mcgill.ca, hendren@cs.mcgill.ca

Abstract

MATLAB is a dynamic scientific language used by scientists, engineers and students worldwide. Although MATLAB is very suitable for rapid prototyping and development, MATLAB users often want to convert their final MATLAB programs to a static language such as FORTRAN. This paper presents an extensible object-oriented toolkit for supporting the generation of static programs from dynamic MATLAB programs. Our open source toolkit, called the MATLAB Tamer, identifies a large tame subset of MATLAB, supports the generation of a specialized Tame IR for that subset, provides a principled approach to handling the large number of builtin MATLAB functions, and supports an extensible interprocedural value analysis for estimating MATLAB types and call graphs.

Categories and Subject Descriptors D.3.4 [Processors]: Compilers

General Terms Experimentation, Languages

Keywords MATLAB, Static Compilation, Interprocedural Value Analysis, Analysis Framework

1. Introduction

MATLAB is a popular numeric programming language, used by millions of scientists, engineers as well as students worldwide[?]. MATLAB programmers appreciate the high-level matrix operators, the fact that variables and types do not need to be declared, the large number of library and builtin functions available, and the interactive style of program development available through the IDE and the interpreter-style read-eval-print loop. However, even though MATLAB programmers appreciate all of the features that enable rapid prototyping, they often have other ultimate goals. Frequently they wish to convert their MATLAB code to a more static language such as FORTRAN or C.

There are several reasons why users would like to have static code. First, users may want source code, which can be compiled into a self-contained executable with standard and freely available compilers. This allows them to easily distribute both the source and executable. Second, users may want to generate code that can be integrated into existing systems. As just one example, one of our users wanted to generate FORTRAN code that can be plugged into a weather simulation environment. Finally, users may wish to generate code that can take advantage of high-performance optimizing or parallelizing compilers (such as FORTRAN), or novel parallel systems (such as X10).¹

This paper addresses the problem of how to provide the bridge between the dynamic realities of MATLAB and the ultimate goal of wanting efficient and static programs. It is not realistic to support all the MATLAB features, but our goal is to define and provide support for a very large subset of MATLAB, which includes dynamic typing, variable numbers of input and output arguments, support for a variety of MATLAB data types including arrays, cell arrays and structs, and support for function handles and lambda expressions.

Providing this bridge presents two main challenges. The first is that MATLAB is actually quite a complex language. It evolved over many years and has non-standard type rules and function lookup semantics. The second major challenge is properly dealing with the large number of builtin and library functions, which have also been developed over time and which sometimes have unexpected or irregular behaviour.

Our solution is an open-source, extensible, objected-oriented framework, implemented in Java, as presented in Figure 1. The overall goal of the system is to take MATLAB programs as input and produce output which is suitable for static compilation, a process that we call *Taming* MATLAB. Given a .m file as input, which is the entry point, the MATLAB Tamer produces as output: (1) a Tame IR for all functions (both user and library) which are reachable from the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

OOPSLA'12, October 19–26, 2012, Tucson, Arizona, USA.
Copyright © 2012 ACM 978-1-4503-1561-6/12/10...\$10.00

¹We should note that another way to achieve performance is to use more dynamic techniques with a virtual machine and an optimizing/parallelizing JIT. For a dynamic language like MATLAB, this seems to be a good fit, and in fact our research group is developing dynamic techniques with the McVM project [3] in parallel with our work for static compilers. Somewhat to our surprise there has been much more user interest in the static compilation project, for the reasons we have given. Thus, we think that pursuing both the static and dynamic approaches are worthwhile.

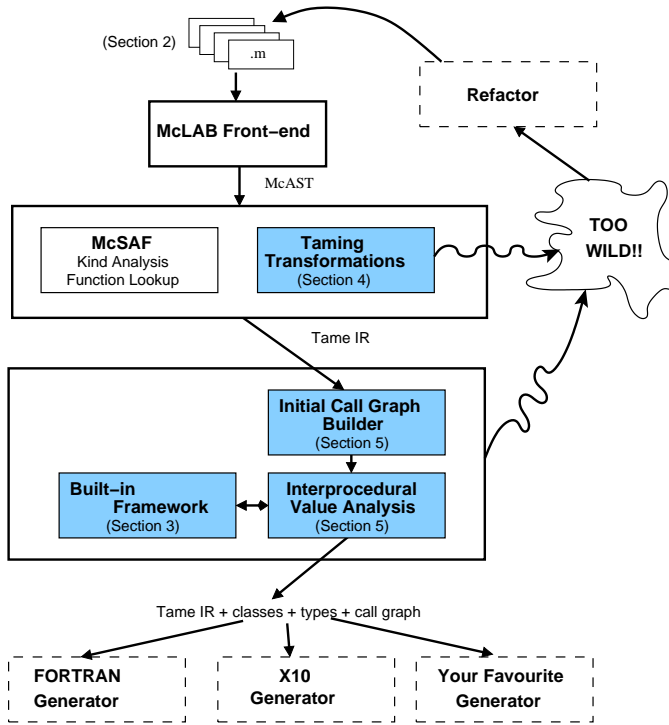


Figure 1. Overview of our MATLAB Tamer. The shaded boxes indicate the components presented in this paper. The other solid border boxes correspond to existing McLAB tools we use, and the dashed boxes correspond to ongoing projects which are using the results of this paper.

entry point, (2) a complete call graph, and (3) an estimation of classes/types for all variables.

There are some features in MATLAB that are simply too wild to handle, and so our system will reject programs using those features. The user will need to refactor their program to eliminate the use of those wild features. Thus, another important goal in our work is to define a subset of MATLAB which is as large as possible and which can be tamed without user intervention.

The main contributions of this paper are as follows.

- We present an overall design and implementation for the MATLAB Tamer, an extensible object-oriented framework which provides the bridge between the dynamic MATLAB language and static back-end compilers.
- We describe the key features of MATLAB necessary for compiler developers and for tool writers to understand MATLAB and the analyses in this paper. We hope that by carefully explaining these ideas, we can enable other researchers to also work on static tools for MATLAB. Our discussion of MATLAB features also motivates our choice of the subset of MATLAB that we aim to tame.
- We provide a principled approach to understanding, grouping, and analyzing the large number of MATLAB

builtin functions. This is a non-trivial effort that is really needed in order to be able to properly analyze MATLAB programs which typically make heavy use of library functions.

- We provide extensions to the McSAF [? ?] framework to support a lower-level and more specialized Tame IR, suitable for back-end static code generation.
- We present an interprocedural flow analysis framework that computes both abstract values and the complete call graph. This flow analysis provides an object-oriented approach which allows for extension and refinement of the abstract value representations.

The remainder of the paper is structured as follows. Sec. 2 introduces the key MATLAB features, Sec. 3 describes our approach to MATLAB builtins, Sec. 4 describes the Tame IR and transformations, Sec. 5 explains our extensible and interprocedural value analysis and call graph construction, Sec. 6 summarizes how our framework can be used, Sec. 7 provides an overview of related work and Sec. 8 concludes.

2. MATLAB - a dynamic language

In this section we describe key MATLAB semantics and features to provide necessary background for compiler writers and tool developers to understand MATLAB and its challenges, and to motivate our approach of constructing a “tame” intermediate representation and MATLAB callgraph. In each subsection we give a description followed by annotated examples using the MATLAB read-eval-print loop. In the examples, “>>” indicates a line of user input, and the following line(s) give the printed output.

2.1 Basics

MATLAB was originally designed in the 1970s to give access to features of FORTRAN (like LINPACK, EISPACK) without having to learn FORTRAN[?]. As the name MATLAB(MATRIX LABORATORY) suggests, MATLAB is centered around numerical computation. Floating point matrices are the core of the language. However, the language has evolved beyond just simple matrices and now has a type system including matrices of different types, compound types including cell arrays and structs, as well as function references.

Given its origins, MATLAB is a language that is built around matrices. Every value is a *Matrix* with some number of dimensions, so every value has an associated array shape. Even scalar values are 1×1 matrices. Vectors are either $1 \times n$ or $n \times 1$ matrices and strings are just vectors of characters. Most operations are defined for matrices, for example $a * b$ specifies matrix multiplication if both a and b are matrices. Operators are overloaded and sometimes refer to scalar operations, for example $a * b$ specifies an element-wise multiplication if a is a matrix and b is a scalar.

```
>> size(3)           % the scalar 3 is a 1x1 matrix
    1    1
>> size([1 2 3])    % a 1x3 vector
```

```

1     3
>> size([5; 6; 7; 8; 9]) % a 5x1 vector
5     1
>> size('hello world') % a string, which is a 1x11 vector
1    11
>> ['a' 'b'; 'e' 'f'] % a 2-dimensional matrix of chars
ab
ef

```

2.2 MATLAB Type System

MATLAB is dynamically typed - variables need not be declared, they will take on any value that is assigned to them. Every MATLAB value has an associated MATLAB class (henceforth we will use the name *mclass* when referring to a MATLAB class, in order to avoid confusion with the usual notion of a class). The *mclass* generally denotes the type of the elements of a value. For example, the *mclass* of an array of doubles is `double`. The default numeric *mclass* is `double`. While MATLAB also includes integer types, all numeric literals are doubles.

```

>> n = 1 % input literal and output look like an int
1
>> class(n) % however, mclass is really double, the default
double
>> class(1:100) % mclass of vector [1, 2, ..., 100] is double
double

```

MATLAB has a set of builtin *mclasses*, which can be summarized as follows:

- `double`, `single`: floating point values
- `uint8`, `uint16`, `uint32`, `uint64`, `int8`, `int16`, `int32`, `int64`: integer values
- `logical`: boolean values
- `char`: character values (strings)
- `cell`: heterogeneous arrays
- `struct`: structures
- `function handle`: references to functions

Given that by default any numerical value in MATLAB is a `double`, all values that are intended to be of a different numeric type have to be specifically converted. This also means that when combining a value of some non-`double` *mclass* with a value that is a `double`, the result will be of the non-`double` *mclass*. This leads to the surprising semantics that adding an integer and a double results in an integer, because that is the more specialized type.

```

>> x = 3; % x is implicitly a double
>> y = int8(5); % y is explicitly an integer
>> class(x) % the class of x is double
double
>> class(y) % the class of y is int8
int8
>> class(x+y) % the result of x+y is int8, not double
int8

```

2.3 MATLAB Functions and Specialization

A MATLAB function is defined in a `.m`-file which has the same name as the function.² So, for example, a function named `foo` would be defined in a file named `foo.m`, and that file needs to be placed either in the current directory, or in a directory on the MATLAB path. A `.m`-file can also define subfunctions following the main (primary) function definition in a file, but those subfunctions are only visible to the functions within the file. Functions may also be defined in a `private/` directory, in which case they are visible only to functions defined in the parent directory.

MATLAB allows overriding or specializing operations and functions to operate on specific *mclasses*. This is accomplished by defining the function in a file inside a specially named directory which starts with the character `@` followed by the name of the *mclass*. For example, one could create a specialized function `firstWord` defined for Strings, by creating a file `@char/firstWord.m` somewhere on the MATLAB path. Functions that are specialized in such a way have precedence over non-specialized functions, but they do not have precedence over inner functions, subfunctions (defined in the same file) or private functions (defined in the `/private` directory). So, in our example, if there existed two definitions of `firstWord.m`, one general implementation somewhere on the MATLAB path, and one specialized implementation in a directory `@char` on the MATLAB path, then a call to `firstWord` with a `char` argument will result in a call to `@char/firstWord.m`, whereas a call with an argument with any other *mclass*, will result in a call to the general `firstWord.m` definition.

When calling a function that has *mclass*-specialized versions with multiple arguments of different *mclasses*, MATLAB has to resolve which version of the function to call. There doesn't exist a standard inheritance relationship between the builtin *mclasses*. Rather, MATLAB has the notion of a *superior* or *inferior* class. We were unable to find a succinct summary of these relationships, so we generated a MATLAB program which exercised all cases and which produced a `.dot` file describing all relationships, with all transitive relationships removed. Figure 2 shows the relationships between different builtin *mclasses*, showing superior classes above inferior classes. Note that some *mclasses* have no defined relationship, for example, there are no defined inferior/superior relationships between the different integer *mclasses*. Further, note that `double`, being the default *mclass*, is inferior to integer *mclasses*. Also, the compound *mclasses* (`struct` and `cell`), are superior to all matrix *mclasses*.

When resolving a call with multiple arguments, MATLAB finds the most superior argument, and uses its *mclass* to resolve the call. If multiple arguments have no defined superior/inferior relationships, MATLAB uses the leftmost superior argument. For example, if a function is called with three ar-

²In the case where the name of the file and the function do not match, the name of the file takes precedence.

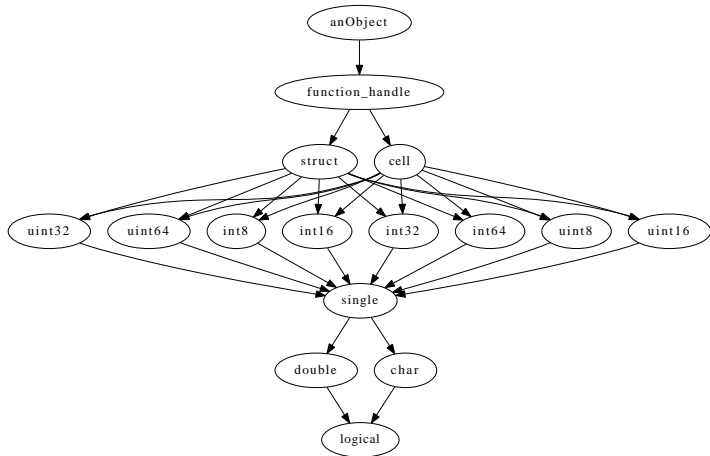


Figure 2. Superior/inferior class relationships for MATLAB

arguments with the mclasses (`double`, `int8`, `uint32`), in that order, MATLAB attempts to find a specialized version for mclass `int8`. If none is found, MATLAB attempts to find a non-specialized version.

The class specialization semantics for MATLAB means that if one intends to build a complete callgraph, i.e. resolve all possible call edges, one has to find all possible MATLAB classes for all arguments, and one must safely approximate the lookup semantics of functions, including the correct lookup of specialized functions using the `mclass` and the superior/inferior mclass relationships from Figure 2.

2.4 MATLAB Classes

It is important to note that the `mclass` of a value does not completely define its type. For example, numeric MATLAB values may be real or complex, and matrices have an array shape. Both of these properties are defined orthogonally to the notion of its `mclass`. Although a computation can ask whether a value is complex or real, and can ask for the shape of an array, the lookup semantics solely depend on the `mclass`, which is effectively just a name. Within the MATLAB language, there is no dedicated class of values to represent mclasses. Usually, strings (`char` vectors) are used to denote mclasses. For example, `ones(3,2, 'single')`, will call the builtin function `'ones'` and create a 3×2 array of unit values of mclass `single`.

2.5 Function Handles

MATLAB values with `mclass` `function_handle` store a reference to a function. This allows passing functions as arguments to other functions. Function handles can either be created to refer to an existing function, or can be a lambda expression. Lambda expressions may also encapsulate state from the current workspace via free variables in the lambda expression.

```
>> f = @sin % a function handle to a named function
f = @sin
>> g = @(x) exp(a*x) % a lambda with a free variable "a"
```

```
g = @(x)exp(a*x)
```

Function handles, and especially lambdas, are useful in numerical computing, for example when calling numerical solvers, as illustrated below.

```
f = @(t,y) D*t + c; % set up derivative function
span = [0 1]; % set interval
y0 = [0:0.1:10]'; % set initial value
result = ode23s(f,span,y0);
% use matlab library function to solve
```

When building a callgraph of a program that includes function handles, one needs to propagate function handles through the program interprocedurally in order to find out which variables may refer to function handles, and to find associated call edges.

2.6 Function Parameters and Arguments

MATLAB uses call-by-value semantics, so that each parameter denotes a fresh copy of a variable.³ This simplifies interprocedural analyses for static compilation as calling a function cannot directly modify local variables in the caller.

In MATLAB, function arguments are optional. That is, when calling a function one may provide fewer arguments than the function is declared with. However, MATLAB does not have a declarative way of specifying default values, nor does it automatically provide default values. That is, a parameter corresponding to an argument that was not provided will simply be unassigned and a runtime error will be thrown if an unassigned variable is read.

MATLAB does provide the function `nargin` to query how many arguments have been provided to the currently executing function. This allows the programmer to use the value of `nargin` to explicitly assign values to the missing parameters, as illustrated below.

```
function [result1, result2] = myFunction(arg1,arg2)
    if (nargin < 1)
        arg1 = 0;
    end
    if (nargin < 2)
        arg2 = 1;
    end;
    ...
end
```

As shown above, MATLAB also supports assigning multiple return variables. A function call may request any number of return values simply by assigning the call into a vector of lvalues. Just like the function arguments, the return values don't all need to be assigned, and a runtime error is thrown if a requested return value is not assigned. MATLAB provides the `nargout` function to query how many results need to be returned.

Clearly a static compiler for MATLAB must deal with optional arguments in a sound fashion.

³Actual MATLAB implementations only make copies where actually necessary, using either lazy copying when writing to an array with reference count greater than 1, or by using static analyses to determine where to insert copies[?].

2.7 Wild Dynamic Features

Whereas features like dynamic typing, function handles, and variable numbers of input arguments are both widely used and possible to tame, there are other truly wild dynamic features in MATLAB that are not as heavily used, are sometimes abused, and are not amenable for static compilation.

These features include the use of scripts (instead of functions), arbitrary dynamic evaluation (`eval`), dynamic calls to functions using `feval`, deletion of workspace variables (`clear`), assigning variables at runtime in the caller scope of a function (`assignin`), changing the function lookup directories during runtime (`cd`) and certain introspective features. Some of these can destroy all available static information, even information associated with other function scopes than where these features occur.

Our approach to these features is to detect them and helping programmers to remove them via refactorings. Some refactorings can be automated. For example, McLAB already supports refactorings to convert scripts to functions and some calls to `feval` to direct function calls[?]. Other refactorings may need to be done by the programmer. For example, the programmer may use `cd` to change directory to access some data file, not being aware that this also changes the function lookup order. The solution in this case is to use a path to access the data file, and not to perform a dynamic call to `cd`. We have also observed many cases where dynamic `eval` or `feval` calls are used because the programmer was not aware of the correct direct syntax or programming feature to use.⁴ For example, `feval` is often used to evaluate a function name passed as a String, where a more correct programming idiom would be to use a function handle.

2.8 Illustrative Example

The example in Figure 3 exemplifies some of the ideas presented in the previous section. It is a complete, valid MATLAB program, performing some numerical computation⁵, using common features and functionality. It is a program using multiple user-defined functions, as well as MATLAB builtin functions, and one MATLAB library function (the call to `toeplitz` in function `Dxx`).

Note the use of `single` values. These single values are created only in the first function (`main`). But they are superior to `double` values used in the other functions, and all the values in the actual computation are a result of operating on the incoming `single` values, so the overall computation is done using `single` values. Essentially functions performing the actual computations are agnostic to the mclasses of

⁴This is at least partly due to the fact that older versions of MATLAB did not support all of the modern features.

⁵While the presented example is a complete, valid MATLAB program, and uses features and the kind of functions one would commonly see used in MATLAB, it does not represent a meaningful way to solve the contained numerical problem, because it is highly unstable. We created this example for the purposes of illustrating different MATLAB language ides, and for brevity.

```
function result = main(N)
    % set up sytem conditions/variables
    rod.L = single(10);
    rod.Ta = single(3);
    rod.Tb = single(4);
    a = .2;
    A = 0.05;
    Tspan = [0 20];
    % solve system
    result = solveSystem(rod,a,A,N,Tspan);
end

% solves the heat equation for some rod of length rod.L,
% initially at temperature rod.Ta, rod.Tb, using N subdivisions.
% uses default values for Tspan, Tsteps if they are undefined
function out=solveSystem(rod,a,A,N,Tspan,Tsteps)
    switch nargin
        case 4
            Tspan = [0, 10]; Tsteps = 100;
        case 5
            Tsteps = 100;
    end
    h = rod.L/(N+1); % step size in x
    X = [h:h:rod.L-h].'; % create x axis with subdivisions
    U0 = X; % allocate u0
    mid = round(length(X)/2);
    U0(1:mid) = rod.Ta; % set left to Ta
    U0(mid:end) = rod.Tb; % set right to Tb
    [D,c] = Dxx(N,rod.Ta,rod.Tb,h); % set up derivative matrix
    f = @(t,u) a*(D*u + c) + A; % right hand side of ODE
    [T,out] = RungeKutta2(f,Tspan,U0,Tsteps); % solve ODE
end

% creates a centered-in-space finite difference differentiation
% matrix for the second spatial derivative.
function [D,c]=Dxx(N,a,b,h)
    D= toeplitz([-2;1;zeros(N-2,1)],...
                [-2,1,zeros(1,N-2)]./h^2);
    c = [a/h^2; zeros(N-2,1); b/h^2];
end

% Runge-Kutta 2 implementation solving the system
% y'=f(t,y), y(a)=y0
% over tspan=[a,b], using function f, and N integration steps
function [X,Y]=RungeKutta2(f,tspan,y0,N)
    a=tspan(1); % initial time
    b=tspan(2); % final time
    h=(b-a)/N; % step size
    Y = [y0 zeros(length(y0),N)]; % result matrix, with y0
    X = (a:h:b)'; % set times
    for j=1:N % integration steps
        k1 = h*f(X(j), Y(:,j));
        k2 = h*f(X(j)+ h, Y(:,j) + k1);
        Y(:,j+1)=Y(:,j) + 1/2*(k1 + k2);
    end
end
```

Figure 3. MATLAB numerical example

the arguments (including the MATLAB function `toeplitz` itself).

Note also the use of the following MATLAB features presented in the previous discussion:

- use of a structure (in `main`)

- use of a lambda expression encapsulating local variables to construct a function reference (in `solveSystem`)
- use of optional input arguments (in `solveSystem`), where optional parameters are assigned using a switch statement on `nargin`.
- use of multiple return values (in `Dxx`).

In the following sections we will illustrate the relevant points using this example. The complete result of running the Tamer on the example in Figure 3 is shown in Appendix A.

2.9 Summary

In this section we have outlined key MATLAB features and semantics, especially concentrating on the definition of `mclass` and function lookup. Our approach is to tame as much of MATLAB as possible, including support for function pointers and lambda definitions. Capturing as much as possible of the evolved language is not just useful to allow access to a wider set of MATLAB features for user code. Also, a significant portion of MATLAB's extensive libraries are written in MATLAB itself, and make extensive use of some of the features discussed above. Since we implement the MATLAB lookup semantics, and allow the inclusion of the MATLAB path, our callgraph will automatically include available MATLAB library functions. Thus, implementing more features will also benefit users who do not make direct use of advanced features.

3. Framework for Builtins

One of the strengths of MATLAB is in its large library, which doesn't only provide access to a large number of matrix computation functions, but packages for other scientific fields. Even relatively simple programs tend to use a fair number of library functions. Many library functions are actually implemented in MATLAB code. To provide their functionality, the callgraph construction needs to include any MATLAB function on the MATLAB path, if it is available. Thus we can provide access to a large number of library functions as long as we can support the language features they use. However, hundreds of MATLAB functions are actually implemented in native code. We call these functions builtins or builtin functions. Every MATLAB operator (such as `+`, `*`) is actually a builtin function; the operations are merely syntactic sugar for calling the functions that represent the operations (like `'plus'`, `'mtimes'`). Thus, for an accurate static analysis of MATLAB programs, one requires an accurate model of the builtins, and a way to deal with them. Consider that the example presented in Figure 3 at first sight only seems to use some basic matrix arithmetic and some comparison operators; however, the full program, including the embedded MATLAB library function call, refers to a total of twenty-six

builtins.⁶ In this section we describe how we have modeled the builtins and how we integrate the analysis into the static interprocedural analysis framework.

3.1 Learning about Builtins

As a first step to build a framework of builtin functions, we need to identify builtins, and need to find out about their behaviour, especially with respect to `mclasses`.

3.1.1 Identifying Builtins:

To make the task of building a framework for builtins manageable, we wanted to identify the most commonly used builtin functions and organize those into a framework. Other builtins can be added incrementally, but this initial set was useful to find a good structure.

To identify commonly used builtins we used the `McBENCH` framework[?] to find all references to functions that occur in a large corpus of over 3000 MATLAB programs.⁷ We recorded the frequency of use for every function and then using the MATLAB function `exist`, which returns whether a name is a variable, user-defined function or builtin, we identified which one of these functions is a builtin function. This provided us with a list of builtin functions used in real MATLAB programs, with their associated frequency of use. We selected approximately three hundred of the most frequent functions, excluding very dynamic functions like `eval` as our initial set of builtin functions.⁸

3.1.2 Finding Builtin Behaviours:

In order to build a call graph it is very important to be able to approximate the behaviour of builtins. More precisely, given the `mclass` of the input arguments, one needs to know a safe approximation of the `mclass` of the output arguments. This behaviour is actually quite complex, and since the implementation of MATLAB 7 is the defacto specification of the behaviour we decided to take a programmatic approach to exploring and determining the behaviours.⁹

We developed a set of scripts that generated random MATLAB values of all combinations of builtin `mclasses`, and called selected builtins using these arguments. If different random values of the same `mclass` result in consistent resulting `mclasses` over many trials, the scripts record the associ-

⁶`mtimes`, `vertcat`, `single`, `transpose`, `conj`, `round`, `colon`, `lt`, `isequal`, `mrdivide`, `length`, `horzcat`, `mpower`, `end`, `plus`, `minus`, `nargin`, `class`, `ones`, `message`, `uminus`, `zeros`, `ne`, `rdivide`, `ctranspose`, `warning`

⁷This is the same set of projects that are used in [?]. The benchmarks come from a wide variety of application areas including Computational Physics, Statistics, Computational Biology, Geometry, Linear Algebra, Signal Processing and Image Processing.

⁸The complete list can be found at www.sable.mcgill.ca/mclab/tamer.html

⁹We also used the MATLAB 7 documentation, which unfortunately includes fairly little type information, and ad-hoc trials to determine the behaviour of builtins.

ated mclass propagation for builtins in a table, and collect functions with the same mclass propagation tables together. Examples of three such tables are given in Figure 4.¹⁰

As compared with type rules in other languages, these results may seem a bit strange. For example, the “-” entry for `plus(int16,int32)` in Figure 4(a) shows that it is an error to add an `int16` to and `int32`. However adding an `int64` to a double is allowed and it results in an `int64`. Also, note that although the three tables in Figure 4 are similar, they are not identical. For example, in Figure 4(a), multiplying a logical with a logical results in a double, but using the power operator with two logicals is an error. Finally, note that the tables are not always symmetrical. In particular, the `f64` column and row in Figure 4(b) are not the same.

The reader may have noticed how the superior/inferior mclass relationships as shown in figure Figure 2 seem to resemble the implicit type conversion rules for MATLAB builtin functions. For example, when adding an integer and a double, the result will be double. However, it is not sufficient to model the implicit MATLAB class conversion semantics by just using class-specialized functions and their relationships. Many MATLAB builtins perform explicit checks on the actual runtime types and shapes of the arguments and perform different computations or raise errors based on those checks.

Through the collection of a large number of tables we found that many builtins have similar high-level behaviour. We found that some functions work on any matrix, some work on numeric data, some only work on floats, and some work on arbitrary builtin values, including cell arrays or function handles.

3.2 Specifying Builtins

To capture the regularities in the builtin behaviour we arranged all of the builtins in a hierarchy - a part of the hierarchy is given in Figure 5. Leaves of the hierarchy correspond to actual builtins and upper levels correspond to abstract groups which share some sort of similar behaviour. The motivation is that some flow analyses need only specify the abstract behaviour of a group, and the flow analysis framework will automatically apply the correct (most specialized) behaviour for a specific builtin.

To specify builtins and their relationships, we developed a simple domain-specific language. One just needs to specify the name of a builtin. If the builtin is abstract (i.e. it refers to a group of builtins), the parent group has to be specified. If no parent is specified, the specified name is an actual builtin, belonging to the group of the most recently specified abstract builtin. This leads to a very compact representation, allowing builtins to be specified on one line each, as illustrated by the following snippet of the builtin specification:

```
floatFunction; matrixFunction
properFloatFunction; floatFunction
unaryFloatFunction; properFloatFunction
elementalUnaryFloatFunction; unaryFloatFunction
sqrt
realsqrt
erf
...
improperFloatFunction; floatFunction
...
```

The builtin framework takes a builtin specification and generates a set of Java classes whose inheritance relationship reflects the specified tree. It also generates a visitor class, which allows annotating methods to Builtins using the visitor pattern - a pattern that is already extensively used in the McSAF framework[??] upon which the Tamer is built.

We categorize the MATLAB builtin functions according to many properties, such as mclass, arity, shape, semantics. This means that different analyses or attributes can be specify at exactly the required category. It also means that when adding builtins that do fit in already existing categories, one does not need to add all required attributes or flow equations.

3.3 Specifying Builtin attributes

It is not sufficient to just specify the existence of builtins; their behaviour needs to be specified as well. In particular, we need flow equations for the propagation of mclasses. Thus the builtin specification language allows the addition of properties. A property is just a name, with a set of arguments that follow it. A specific property can be defined for any builtin, and it will trigger the addition of more methods in the generated Java code as well as the inclusion of interfaces. In this way, any property defined for an abstract builtin group is defined for any builtin inside that group as well, unless it gets overridden.

3.3.1 The “Class” attribute

The first property we defined was the property `Class`. When specified for a builtin, it forces the inclusion of the Java interface `ClassPropagationDefined` in the generated Java code, and will add a method that returns an mclass flow equation. The mclass flow equation itself is specified as an argument to the `Class` attribute using a small domain specific language that allows matching argument mclasses, and returns result mclasses based on matches. An example snippet is given below which shows the specification of mclass flow equations for unary functions taking numeric arguments. Functions in that group accept any numeric argument and return a result of the same mclass (`numeric->0`), a char or logical argument will result in a double.

```
unaryNumericFunction; properNumericFunction;
-Class(numeric->0, \
    char|logical->double)

elementalUnaryNumericFunction; unaryNumericFunction;
-abstract
```

¹⁰To save space we have not included the whole table, we have left out the columns and rows for unsigned integer mclasses and for handles. All result tables can be found at www.sable.mcgill.ca/mclab/tamer.html

	i8	i16	i32	i64	f32	f64	c	b
i8	i8	-	-	-	-	i8	i8	-
i16	-	i16	-	-	-	i16	i16	-
i32	-	-	i32	-	-	i32	i32	-
i64	-	-	-	i64	-	i64	i64	-
f32	-	-	-	-	f32	f32	f32	f32
f64	i8	i16	i32	i64	f32	f64	f64	f64
c	i8	i16	i32	i64	f32	f64	f64	f64
b	-	-	-	-	f32	f64	f64	f64

(a) plus, minus, mtimes, times, kron

	i8	i16	i32	i64	f32	f64	c	b
i8	i8	-	-	-	-	-	i8	-
i16	-	i16	-	-	-	-	i16	-
i32	-	-	i32	-	-	-	i32	-
i64	-	-	-	i64	-	-	i64	-
f32	-	-	-	-	f32	-	f32	f32
f64	i8	i16	i32	i64	f32	f64	f64	f64
c	i8	i16	i32	i64	f32	f64	f64	f64
b	-	-	-	-	f32	f64	f64	-

(b) mpower, power

	i8	i16	i32	i64	f32	f64	c	b
i8	i8	-	-	-	-	i8	i8	-
i16	-	i16	-	-	-	i16	i16	-
i32	-	-	i32	-	-	i32	i32	-
i64	-	-	-	i64	-	i64	i64	-
f32	-	-	-	-	f32	f32	f32	f32
f64	i8	i16	i32	i64	f32	f64	f64	f64
c	i8	i16	i32	i64	f32	f64	f64	f64
b	-	-	-	-	f32	f64	f64	-

(c) mldivide, mrdivide, ldivide, rdivide, mod, rem, mod

Figure 4. Example mclass results for groups of Built-in binary operators. Rows correspond to the mclass of the left operand, columns correspond to the mclass of the right operand, and the table entries give the mclass of the result. The labels i8 to i64 represent the classes int8 through int64, f32 is single, f64 is double, c is char, and b is logical. Entries of the form “-” indicate that this combination is not allowed and will result in a runtime error.

```

real
imag
abs
conj;; MatlabClass(logical->error, natlab)
sign;; MatlabClass(logical->error, natlab)

```

We have noticed some irregularities in the pure MATLAB semantics, and our specification sometimes removes those. In order to keep a record of the differences we use the `MatlabClass` specification which allows us to specify the exact MATLAB semantics - and thus provides an exact definition and documentation of MATLAB class semantics. In the example above, we specify that the functions `conj` and `sign` have different MATLAB semantics: they disallow logical arguments, which will result in an error.

3.3.2 Other Attributes

It is possible to add new kinds of attributes to the builtin specification language. One merely has to add a function to the builtin generator with a specific function interface and the same name as the new attribute. This function has to return Java code that will be inserted in the generated Builtin class. The function may also force the generated Builtin class to implement a certain interface (i.e. to signify that certain methods were added to it). We will use this facility to add more information about builtins or flow equations for the operations they represent. It will conversely allow us to quickly provide information about builtins that we add to the framework.

3.4 Summary

We have performed an extensive analysis of the behaviour of MATLAB builtin functions. Based on that we developed a framework that allows to specify MATLAB builtin functions, their relationships and properties such as flow equations in a compact way. This framework is extensible both by allowing the quick addition of more builtin functions; and by allowing to specify information and behaviour for builtin functions. This can be done either adding new properties to the framework itself; or by implementing visitor classes.¹¹

4. Tame IR

As indicated in Figure 1, we build upon the McSAF [?] framework by adding taming transformations and by producing a more specialized Tame IR.

The McSAF framework provides both a high-level AST and a lower-level AST called McLAST in which many simplifications have been performed, including simplifying the left-hand and right-hand sides of assignment statements, transforming the implicit control-flow of short-circuit operators to explicit control flow, resolving the precise meaning of the `end` operator, and simplifying the `if` and `for` control constructs.

To produce an easily analyzable Tame IR we have made three important additions: (1) generating more specialized AST nodes, (2) translating `switch` statements to equivalent conditional statements, and (3) transforming lambda ex-

¹¹The complete specification of builtins, documentation of the specification and diagrams of all builtins is available at www.sable.mcgill.ca/mclab/tamer.html

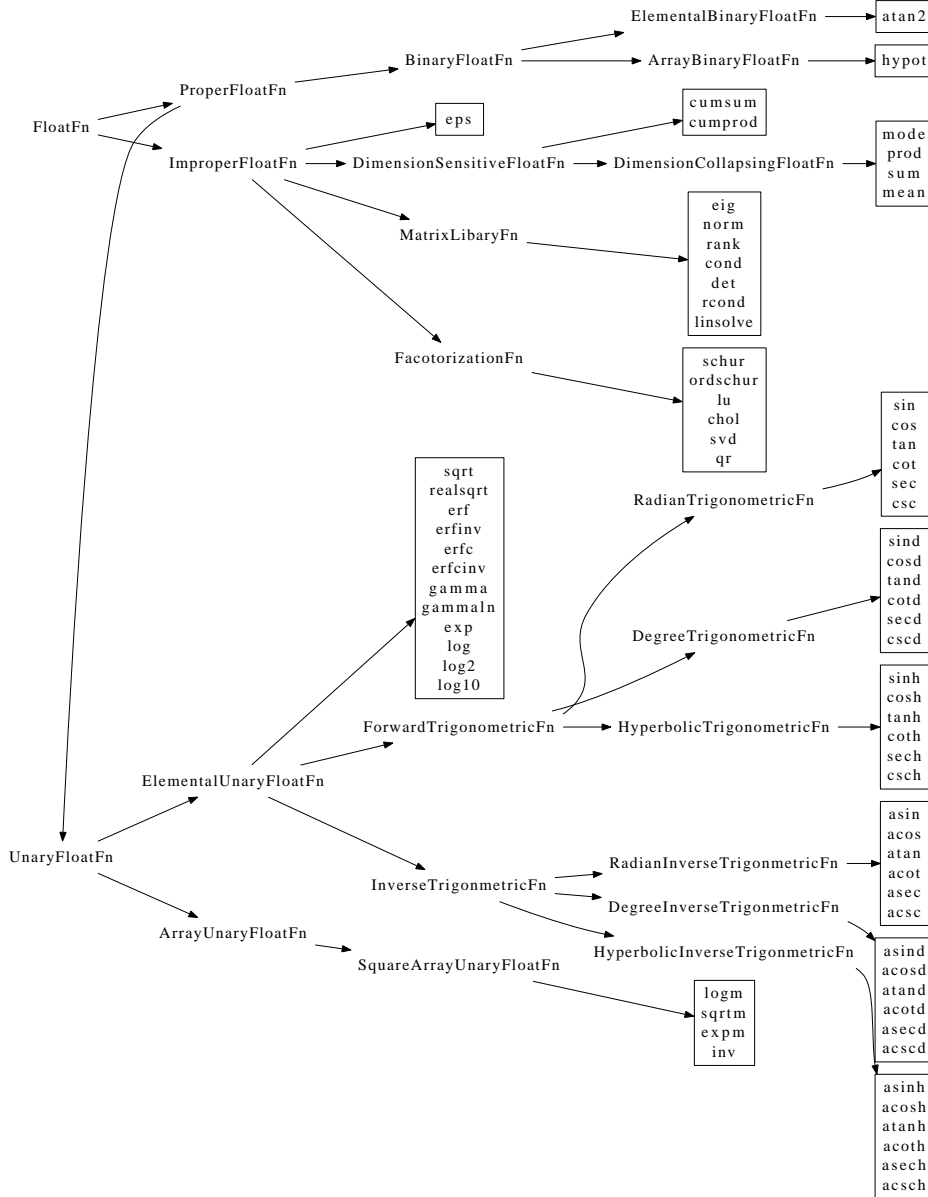


Figure 5. Subtree of builtin tree, showing all defined floating point builtins of MATLAB

pressions to analyzable equivalents. One might wonder why these transformations are not already part of McSAF, the framework upon which the Tamer is built. The important point is that McSAF must handle all of MATLAB, whereas for our Tame IR we can make restrictions that are reasonable for the purposes of static compilation. This allows us to make the Tame IR more specialized and enables more simplifying transformations.

4.1 Specialized AST nodes

One goal for our Taming framework was to produce an IR that is very simple to analyze, and has operations that are low-level enough to map fairly naturally to static languages

like FORTRAN. As one example, in McSAF there is only one kind of assignment statement, assigning from an expression to an lvalue expression. The Tame IR has many more specialized cases as illustrated in Figure 6. We also extended McSAF’s analysis framework to recognize these new IR nodes, so flow equations can be specified for them. Note how the Tame IR has a different statement for a function call or an array indexing operation. In MATLAB these use the same syntax, a parameterized expression. We use the kind analysis[?] to resolve names to being a function or variable, but there are rare cases when this is not possible. Tame MATLAB will reject these cases.

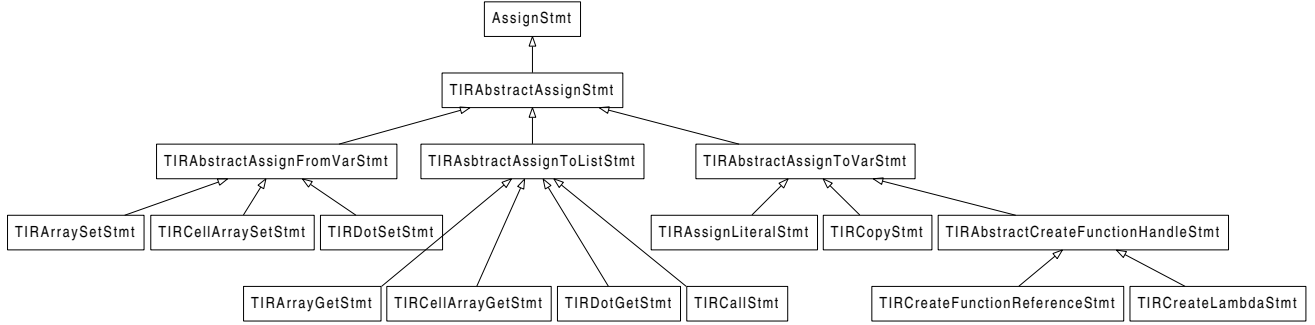


Figure 6. Specializations of the assignment statement

4.2 Lambda Simplification

MATLAB supports lambda expressions. In order to be compatible with the Tame IR, their bodies need to be converted to a three address form in some way. MATLAB lambda expressions are just a single expression (rather than, say, statement lists), so we extract the body of the lambda expression into an external function. The lambda expression still remains, but will encapsulate only a single call, all whose arguments are variables. For example, the lambda simplification will transform the expression in Figure 7(a) to the code in Figure 7(b).

<pre>function outer ... f = @(t,y) D*t + c ... end</pre>	<pre>function outer ... f = @(t,y) lambda1(D,c,t,y) ... end function r = lambda1(D,c,t,y) r = D*t + c end</pre>
--	--

(a) lambda (b) transformed lambda

Figure 7. Transforming lambda expressions

The new lambda expression encapsulates a call to the new function `lambda1`. Note that the first two arguments are variables from the workspace, the remaining ones are the parameters of the lambda expression. In the analyses, we can thus model the lambda expression using partial evaluation of the function `lambda1`. To make this transformation work, the generated function must return exactly one value, and thus Tame MATLAB makes the restriction that lambda expressions return a single value (of course that value may be an array, struct or cell array).

4.3 Switch simplification

As illustrated in Figure 8(a), MATLAB has support for very flexible switch statements. Unlike in other languages, all case blocks have implicit breaks at the end. In order to specify multiple case comparisons for the same case block, MAT-

LAB allows using cell arrays of case expressions, for example `{2, 3}` in Figure 8(a). Indeed, MATLAB allows arbitrary case expressions, such as `c` in the example. If `c` refers to a cell array, then the case will match if any element of the cell array matches. Without knowing the static type and size of the case expressions, a simplification transformation is not possible. Thus, to enable the static simplification shown in Figure 8(b) we add the constraint for the Tame MATLAB that case-expressions are only allowed to be syntactic cell arrays.

<pre>switch n case 1 ... case {2, 3} ... case c ... otherwise ... end</pre>	<pre>t = n if (isequal(t,1)) ... elseif (isequal(t,2) isequal(t,3)) ... elseif (isequal(t,c)) ... else ... end</pre>
---	---

(a) switch (b) transformed switch

Figure 8. Transforming switch statements

4.4 Example

The `solveSystem` function from the example in Figure 3 gets transformed into the code shown in Figure 9. The code has been transformed into a three-address version, with many temporaries, and all expressions have been turned into calls to the equivalent builtin functions. Note in particular how the switch on `nargin()` gets transformed into a series of if-statements (lines 5-20), and how the lambda expression is turned into a function. The original lambda expression is represented as a single call (on line 60) to a newly introduced function (lines 65-73) that holds the body of the lambda expression. Comments in the original code are preserved in empty statements that hold the comments (`TIRCommentStmt`). This means that no other statement contains any comments, and every statement can

```

1 function [out] = solveSystem(rod,a,A,N,Tspan,Tsteps)
2 % solves the heat equation for some rod of length rod.L,
3 % initially at temperature rod.Ta, rod.Tb, using N subdivisions.
4 % uses default values for Tspan, Tsteps if they are undefined
5 [mc_t10] = nargin();
6 mc_t34 = 4;
7 [mc_t12] = isequal(mc_t10, mc_t34);
8 if mc_t12
9     mc_t36 = 0;
10    mc_t37 = 10;
11    [Tspan] = horzcat(mc_t36, mc_t37);
12    Tsteps = 100;
13 else
14    mc_t38 = 5;
15    [mc_t11] = isequal(mc_t10, mc_t38);
16    if mc_t11
17        Tsteps = 100;
18    else
19        end
20    end
21    [mc_t17] = rod.L;
22    mc_t39 = 1;
23    [mc_t18] = plus(N, mc_t39);
24    [h] = mrdivide(mc_t17, mc_t18);
25    % step size in x
26    mc_t21 = h;
27    mc_t22 = h;
28    [mc_t24] = rod.L;
29    mc_t25 = h;
30    [mc_t23] = minus(mc_t24, mc_t25);
31    [mc_t20] = colon(mc_t21, mc_t22, mc_t23);
32    [mc_t19] = vertcat(mc_t20);
33    [X] = ctranspose(mc_t19);
34    % create x axis with subdivisions
35    U0 = X;
36    % allocate u0
37    [mc_t27] = length(X);
38    mc_t40 = 2;
39    [mc_t26] = mrdivide(mc_t27, mc_t40);
40    [mid] = round(mc_t26);
41    [mc_t13] = rod.Ta;
42    mc_t41 = 1;
43    [mc_t15] = colon(mc_t41, mid);
44    U0(mc_t15) = mc_t13;
45    % set left to Ta
46    [mc_t14] = rod.Tb;
47    mc_t28 = mid;
48    mc_t42 = 1;
49    mc_t43 = 1;
50    [mc_t29] = end(U0, mc_t42, mc_t43);
51    [mc_t16] = colon(mc_t28, mc_t29);
52    U0(mc_t16) = mc_t14;
53    % set right to Tb
54    mc_t30 = N;
55    mc_t31 = rod.Ta;
56    mc_t32 = rod.Tb;
57    mc_t33 = h;
58    [D, c] = Dxx(mc_t30, mc_t31, mc_t32, mc_t33);
59    % set up derivative matrix
60    f = (@(t, u) lambda_1(a, D, c, A, t, u));
61    % right hand side of ODE
62    [T, out] = RungeKutta2(f, Tspan, U0, Tsteps);
63 end
64
65 function [mc_t0] = lambda_1(a, D, c, A, t, u)
66 mc_t148 = a;
67 [mc_t150] = mtimes(D, u);
68 mc_t151 = c;
69 [mc_t149] = plus(mc_t150, mc_t151);
70 [mc_t146] = mtimes(mc_t148, mc_t149);
71 mc_t147 = A;
72 [mc_t0] = plus(mc_t146, mc_t147);
73 end

```

Figure 9. Tame IR version of the function `solveSystem` from the example in Figure 3

get transformed without having to worry about preserving comments.

5. Interprocedural Value Analysis and Call Graph Construction

The core of the MATLAB Tamer is the *value analysis*. It's an extensible monolithic context-sensitive inter-procedural forward propagation of abstract MATLAB values. For every program point, it estimates what possible values every variable can take on. Most notably it finds the possible set of mclasses. It also propagates function handle values. This allows resolution of all possible call edges, and the construction of a complete call graph of a tame MATLAB program.

The value analysis is part of an extensible interprocedural analysis framework. It contains a set of modules, one building on top of the other. All of them can be used by users of the framework to build analyses.

- The **interprocedural analysis framework** (section 5.1) builds on top of the Tame IR and the McSAF intraprocedural analysis framework. It allows the construction of

interprocedural analyses by extending an intraprocedural analysis built using the McSAF framework. This framework works together with a callgraph object implementing the correct MATLAB look up semantics. An analysis can be run on an existing callgraph object, or it can be used to build new callgraph objects, discovering new functions as the analysis runs.

- The **abstract value analysis** (section 5.2), built using the interprocedural analysis framework, is a generic analysis of abstract MATLAB values. The implementation is agnostic to the actual representation of abstract values, but is aware of MATLAB mclasses. It can thus build a call graph using the correct function lookup semantics including function specialization.
- We provide an implementation of **composite values** like cell arrays, structures and function handles, which is generic in the implementation of abstract matrix values (section 5.4). This makes composite values completely transparent, allowing users to implement very fine-grained abstract value analyses by only providing

an abstraction for MATLAB values which are matrices, thus simplifying developing new analyses.

- Building on top of all the above modules and putting everything together, we provide an abstraction for all MATLAB values, which we call simple values (section 5.5). Since it includes the function handle abstractions, this can be used by users to build a complete tame MATLAB callgraph. This is the **concrete value analysis**, whose results are presented in section 5.7.

5.1 The Interprocedural Analysis Framework

The interprocedural Analysis framework is an extension of the intraprocedural flow analyses provided by the McSAF framework. It is context-sensitive to aid code generation targeting static languages like FORTRAN. FORTRAN's polymorphism features are quite limited; every generated variable needs to have one specific type. The backend may thus require that every MATLAB variable has a specific known mclass at every program point. Functions may need to be specialized for different kinds of arguments, which a context-sensitive analysis provides at the analysis level.

In the Tamer framework an interprocedural analysis is a collection of interprocedural analysis nodes, which represent a specific intraprocedural analysis for some function and some context. The context is usually a flow representation of the passed arguments. Every such interprocedural analysis node produces a result set using the contained intraprocedural analysis.

Every interprocedural analysis has an associated callgraph object, which may initially contain only one function acting as the entry point for the program. The interprocedural analysis requires a context or argument set for the entry point function.

The analysis starts by creating an interprocedural analysis node for the entry point function and the associated context, which triggers the associated intraprocedural flow analysis. As the intraprocedural flow analysis encounters calls to other functions, it has to create context objects for those calls, and ask the interprocedural analysis to analyze the called functions using the given context. The call also gets added to the set of call edges associated with the interprocedural analysis node.

As the interprocedural analysis has to analyze newly encountered calls, the associated functions are resolved, and loaded into the callgraph if necessary. The result is a complete callgraph, and an interprocedural analysis.

The interprocedural analysis framework supports simple and mutual recursion by performing a fixed point iteration within the first recursive interprocedural analysis node.

5.2 Introducing the Value Analysis

The abstract value analysis is a forward propagation of generic abstract MATLAB values. The mclass of any abstract value is always known.

A specific instance of a value analysis may use different representations for values of different mclasses. For example, function handle values may be represented in a different way than numeric values. This in turn means that values of different Matlab classes can not be merged (joined).

5.2.1 Mclasses, Values and Value Sets:

To define the value analysis independently of a specific representation of values, We first define the set of all mclasses:

$$C = \{\text{double}, \text{single}, \text{logical}, \text{cell}, \dots\}$$

For each mclass, we need some lattice of values that represent estimations of MATLAB values of that class:

$$V_{mclass} = \{v : v \text{ approximates a value with mclass } mclass\}, \\ mclass \in C$$

We require that merge operations are defined, so

$$\forall v_1, v_2 \in V_{mclass}, v_1 \wedge v_2 \in V_{mclass}.$$

We can not join values of different mclasses, because their actual representation may be incompatible. So V_C is not a lattice.

In order to allow union values for variables, i.e. to allow variables to have more than one possible mclass, we estimate the value of a MATLAB variable as a set of pairs of abstract values and their mclasses, where the mclasses are disjoint. We call this a value set. More formally, we define a value set as:

$$ValueSet = \{(mclass_1, v_1), \dots, (mclass_n, v_n) : \\ class_i \neq class_j, class_i \in C, v_i \in V_{class_i}\}$$

Or the set of all possible value sets given a set V of lattices for every mclass.

$$S_V = \{\{(mclass_k, v_k) : \\ mclass_i \neq mclass_j, v_i \in V_{mclass_i}, k \in 0..n\} : \\ 0 \leq n \leq |C|\}$$

This is a lattice, with the join operation which is the simple set union of all the pairs, but for any two pairs with matching mclasses, their values get joined, resulting in only one pair in the result set.

While the notion of a value set allows the analysis to deal with ambiguous variables, still building a complete callgraph and giving a valid estimation of types, having ambiguous variables is not conducive to code generation for a language like FORTRAN. So

```

if (..)
    t = 4;
else
    t = 'hi';
end

```

results in t having the abstract value

$$\{(\text{double}, 4), (\text{char}, 'hi')\}.$$

This example is not tame MATLAB.

5.2.2 Flow Sets:

We define a flow set as a set of pairs of variables and value sets, i.e.

$$flow = \{(var_1, s_1), (var_2, s_2), \dots, (var_n, s_n) : s_i \in S_V, var_i \neq var_j\}$$

and we define an associated look-up operation

$$flow(var) = s \text{ if } (var, s) \in flow$$

This is a lattice whose merge operation resembles that of the value sets.

Flow sets may be *nonviable*, representing non-reachable code (for statements after errors, or non-viable branches). Joining any non-*bottom* flow set with the *nonviable* set results in the viable flow set. joining *bottom* and *nonviable* results in *nonviable*.

5.2.3 Argument and Return sets:

The context or argument set for the interprocedural analysis is a vector of values representing argument values. Arguments are not value sets, but simple values $v \in V_C$ with a single known mclass c . When encountering a call, the analysis has to construct all combinations of possible argument sets, construct a context from that and analyze the call for all such contexts. For example, if we reach a call

`r = foo(a, b)`

with a flow set

$$\{(a, \{(double, v_1), (char, v_2)\}), (b, \{(logical, v_3)\})\},$$

the value analysis constructs two contexts, from (v_1, v_3) and (v_2, v_3) , and analyzes function `foo` with each context. Note how the dominant argument for the first context is `double`, whereas it is `char` for the second. If there exist mclass specialized versions for `foo`, then this results in call edges to, and analysis of, two different functions.

More formally, for a call $func(a_1, a_2, \dots, a_n)$ at program point p , with the input flow set f_p , we have the set of all possible contexts

$$allargs = f_p(a_1) \times f_p(a_2) \times \dots \times f_p(a_n) = \prod_{1 \leq i \leq n} f_p(a_i).$$

The interprocedural analysis needs to analyze $func$ with all these contexts and merge the result,

$$R = \bigwedge_{arg \in allargs} analyze(func, arg).$$

To construct a context, the value analysis may simplify (push up) values to a more general representation. For example, if the value abstraction includes constants, the push up operation may turn constants into *top*. Otherwise, the number of contexts for any given function may grow unnecessarily large.

The result of analyzing a function with an argument set is a vector of value sets, where every component represents a returned variable. They are joined by component-wise joining of the value sets. In the value analysis we require that for a particular call, the number of returned variables is the same for all possible contexts.

5.2.4 Builtin Propagators:

Every implementation of the value abstractions needs to provide a builtin propagator, which provides flow equations for builtins. If B is the set of all defined builtin functions $\{plus, minus, sin, \dots\}$, then the builtin propagator P_V for some representation of values V_C is a function mapping a builtin and argument set to a result set.

$$P_V : B \times \bigcup_{n \in \mathbb{N}} (V_C)^n \rightarrow \bigcup_{n \in \mathbb{N}} (S_V)^n$$

The builtin framework provides tools to help implement builtin propagators by providing builtin visitor classes. The framework also provides attributes for builtin functions, for example the class propagation information attributes.

5.3 Flow Equations

In the following subsection we will show a sample of flow equations to illustrate the flow analysis. We assume a statement to be at program point p , with incoming flow set f_p . The flow equation for program point p results in the new flow set f'_p

- $var_t = var_s$:

$$f'_p = f_p \setminus \{(var_t, f_p(var_t))\} \cup \{(var, f_p(var_s))\}$$

- $var = l$, where l is a literal with mclass c_l and value representation v_l :

$$f'_p = f_p \setminus \{(var, f_p(var))\} \cup \{(var, \{(c_l, v_l)\})\}$$

- $[t_1, t_2, \dots, t_m] = func(a_1, a_2, \dots, a_n)$, a function call to some function $func$:
with

$$call_{func, args} = \begin{cases} P_V(b, args) & \text{- if } func \text{ with } args \text{ refers to a builtin } b \\ analyze(f, args) & \text{- if } func(args) \text{ refers to a function } f \end{cases}$$

we set

$$R = \bigwedge_{args \in f_p(a_1) \times f_p(a_2) \times \dots \times f_p(a_n)} call_{func, args}$$

then

$$f'_p = f_p \setminus \bigcup_{i=1}^m \{(t_i, f_p(t_i))\} \cup \bigcup_{i=1}^m \{(t_i, R_i)\}$$

Note that when analyzing a call to a function in an m-file, the argument values will be pushed up. For calls to builtins, the actual argument values will be used, effectively in-lining the behaviour of builtin functions.

5.4 Structures, Cell Arrays and Function Handles

We implemented a value abstraction for structs, cell arrays and function handles (which we call `AggrValue` internally). This abstraction is again modular, this one with respect to the representation of matrix values (i.e. values with mclass `double`, `single`, `char`, `logical` or `integer`). Structures, cell arrays and function handles act as containers for other values, making them effectively transparent. A user may provide a fine-grained abstraction for just matrix values and combine it with abstraction of composite values to implement a concrete value analysis.

5.4.1 `struct`, `cell`:

For structures and cell arrays, there are two possible abstractions:

- *tuple*: The exact index set of the `struct`/`cell` is known and every indexing operation can be completely resolved statically. Then the value is represented as a set of pairs $\{(i_1, s_1), (i_2, s_2), \dots, (i_n, s_n) : i_k \in I, s_n \in S_V\}$, where I is an index set - integer vectors for cell arrays, and fieldnames for structs.
- *collection*: Not all indexing operations can be statically resolved, or the set of indices is unknown. In this case, all value sets contained in the `struct` or `cell` are merged together, and the representation is a single value set $s \in S_V$.

The usual representation for a structure is a tuple, because usually all accesses (dot-expressions) are explicit in the code and known. Cell arrays are usually a collection, because the index expressions are usually not constant. But cell arrays tend to have homogeneous mclass values, so there is some expectation that any access of a `struct` or `cell` results in some unambiguous mclass and thus allows static compilation.

5.4.2 `function_handle`:

As explained in section 2.5, function handles can be created either by referring to an existing function, or by using a lambda expression to generate an anonymous function using a lambda expression. The lambda simplification (presented in section 4.2) reduces lambda expressions to single calls.

We model all function handles as sets of function handle pairs. A function handle pair consists of a reference to a function and a vector of partial argument value sets. A function handle value may thus refer to multiple possible function/partial argument pairs.

Given some flow set f_p defined at the program point p ,

- $g = @\sin$
results in

$$f'_p = f_p \setminus (g, f_p(g)) \cup \{(g, \{(function_handle, \{(sin, ())\})\})\}$$

- $g = @(t, y) \lambda\text{d}a1(D, c, t, y)$

results in

$$f'_p = f_p \setminus (g, f_p(g)) \cup \{(g, \{(function_handle, \{(\lambda\text{d}a1, (f_p(D), f_p(c)))\})\})\}$$

Note that function handles get invoked at array get statements, rather than calls. That is because the tame IR is constructed without mclass information, and will correctly interpret a function handle as a variable. When the target of an array get statement is a function handle, the analysis inserts one or more call edges at that program point, referring to the functions contained in the function handle.

5.5 The Simple Matrix Abstraction

Using the value abstraction for structures, cell arrays and function, we implemented a concrete value abstraction by adding an abstraction for matrix values, which we call simple matrix values. On top of the required mclass, this abstraction merely adds constant propagation for scalar doubles, strings (char vectors), and scalar logicals.

This allows the analysis of MATLAB code utilizing optional function arguments using the builtin function `nargin` , and some limited dynamic features utilizing strings. For example, a call like `ones(n, m, 'int8')` can be considered tame.

This implementation represents the concrete value analysis that is used to construct complete callgraphs.

5.6 Example

The concrete value analysis using simple matrix values is able to build a complete callgraph for our example introduced in Sec. 2.8, and find a unique mclass for every variable at every program point. It includes and analyzes the MATLAB library function `toeplitz`. This includes resolution of structures as well as function handles created by a lambda expression.

The analysis finds the following function calls, with their associated arguments and return values:

```
main:
  double -> single

solveSystem:
  struct{Ta=single, Tb=single, L=single},
  double, double, double, double -> single

Dxx:
  double, single, single, single
  -> single, single

toeplitz:
  double, double -> double

RungeKutta2:
  (handle, @lambda_1(double, single, single,
                    double,...),
  double, single, double -> double, single
```

```
lambda_1:
    double, single, single, double, double,
    single -> single
```

Note the structure and the function handle, and the mix of `single` and `double` values.

We have included the complete program, which is the result of the tamer, in appendix A. It is shown in the Tame IR form, with the result of the value analysis annotated to every assignment statement.

The complete resolution of every value is only possible because we include constant propagation. For example, it allows the value analysis to know which case for the switch statement in `solveSystem` is valid. Also, the included function `toeplitz` includes the lines

```
[xclass] = class(x);           % xclass=(char,double)
[mc_t90] = zeros(xclass);     % mc_t90=(double)
```

Without knowing that the variable `xclass` holds the string “double”, the following call to `zeros` wouldn’t be known to return a `double`, and the value analysis would have to assume the result could be any numerical value (`double`, `single`, or any integer).

5.7 Applying the Value Analysis

In order to exercise the framework, we applied it to the set of benchmarks we have previously used for evaluating McVM/McJIT[?], a dynamic system. The benchmarks and results are given in Table 1. About half of the benchmarks come from the FALCON project[?] and are purely array-based computations. The other half of the benchmarks were collected by the McLAB team and cover a broader set of applications and use more language features such as lambda expressions, cell arrays and recursion. The columns labeled #Fn correspond to the number of user functions, and the column labeled #BFn corresponds to the number of builtin functions used by the benchmark. Note the high number of builtins. The column labeled “Wild” indicates if our system rejected the program as too wild. Only the `sdku` benchmark was rejected because it used the `load` library function which loads arbitrary variables from a stored file. It is likely that we should provide a tamer version of `load`. The column labeled “Mclass” indicates “unique” if the interprocedural value propagation found a unique `mclass` for every variable in the program. Only three benchmarks had one or more variables with multiple different `mclasses`. We verified that it was really the case that a variable had two different possible classes in those three cases.

Although the main point of this experiment was just to exercise the framework, we were very encouraged by the number of benchmarks that were not wild and the overall accuracy of the basic interprocedural value analysis. We expect many other analyses to be built using the framework, with different abstractions. By implementing them all in a common framework we will be able to compare the different approaches.

6. Building on the Tamer

The Tamer framework, as presented in this paper, is intended both to support further analyses and to provide the building blocks for a variety of back-ends.

6.1 Using the Tamer for further analyses

There are several ways in which the Tamer supports extensions or further analyses.

6.1.1 Implementing new Value Abstractions

As discussed in Sec. 5.5, we have built a concrete value analysis that is used to build the callgraph, which is built on top of the composite value abstraction for structures, cell arrays and function handles. The most common case to extend the value analysis is to provide a different implementation for matrix values, while reusing the abstraction for the composite values.

Many flow analyses that attempt to estimate some aspect of MATLAB values are only concerned with MATLAB values that are matrices. They are basically independent of the behaviour of composite values, except that these may contain matrices. This can be exactly addressed by providing a new matrix abstraction: the user only has to provide the abstraction for matrix values, a factory class that can produce these matrix values, and a class that defines the flow equations for builtins for that matrix abstraction. The whole machinery of the value analysis can be reused, as well as composite values - they simply contain whatever matrix value abstraction is being used, as needed.

For example, an analysis writer may want to implement an analysis that tracks whether matrices are sparse or not (i.e. contain mostly zeros). Sparse matrices can be represented with special data types that can leverage performance sparseness. This property is only relevant to matrices, so the analysis can reuse the implemented composite values and only provide an abstraction for matrices. The user can combine their information about matrices with other components that we have already implemented (like constants) to get a richer or more detailed abstraction.

Users could also implement a complete new abstraction for values, but still reuse the analysis itself.

6.1.2 Implementing other Interprocedural Analyses

The interprocedural analysis framework provides a convenient framework for developing new flow analyses, forward or backward, that must be interprocedural, which are not necessarily based on flowing abstract representations of MATLAB values forward. For example, a user may implement an interprocedural side-effect analysis, computed on a callgraph that was produced by the value analysis.

All interprocedural analysis use McSAF-based intraprocedural analysis operating on the Tamer IR.

Name	Description	Source	#Fn	#BFn	Features	Wild	Mclass
adpt	<i>Adaptive quadrature</i>	Numerical Methods	1	17		no	unique
beul	<i>Backward Euler</i>	McLAB	11	30	lambda	no	unique
capr	<i>Capacitance</i>	Chalmers EEK 170	4	12		no	unique
clos	<i>Transitive Closure</i>	Otter	1	10		no	unique
crni	<i>Tridiagonal Solver</i>	Numerical Methods	2	14		no	unique
dich	<i>Dirichlet Solver</i>	Numerical Methods	1	14		no	unique
diff	<i>Light Diffraction</i>	Appelbaum (MUC)	1	13		no	unique
edit	<i>Edit Distance</i>	Castro (MUC)	1	6		no	unique
fdtd	<i>Finite Distance Time Domain</i>	Chalmers EEK 170	1	8		no	unique
fft	<i>Fast Fourier Transform</i>	Numerical Recipes	1	13		no	multi
fiff	<i>Finite Difference</i>	Numerical Methods	1	8		no	unique
mbrt	<i>Mandelbrot Set</i>	McLAB	2	12		no	unique
mils	<i>Mixed Integer Least Squares</i>	Chang and Zhou	6	35		no	unique
nb1d	<i>1-D Nbody</i>	Otter	2	9		no	unique
nb3d	<i>3-D Nbody</i>	Otter	2	12		no	unique
nfrc	<i>Newton Fractal</i>	McLAB	4	16		no	unique
nne	<i>Neural Net</i>	McLAB	3	16	cell	no	unique
play	<i>Minimax Search</i>	McLAB	5	26	recursive, cell	no	multi
rayt	<i>Raytracer</i>	Aalborg (Jensen)	2	28		no	unique
sch2	<i>Sparse Schroed. Eqn Solver</i>	McLAB	8	32	cell, lambda	no	unique
schr	<i>Schroedinger Eqn Solver</i>	McLAB	8	31	cell, lambda	no	unique
sdku	<i>Sudoku Puzzle Solver</i>	McLAB	8		load	yes	
sga	<i>Vectorized Genetic Algorithm</i>	Burjorjee	4	30		no	multi
svd	<i>SVD Factorization</i>	McLAB	11	26		no	unique

Table 1. Results of Running Value Analysis

6.2 Using the Tamer for back-ends

The tamer provides a convenient IR, a complete call graph, and the mclass information for each variable. This provides a good starting point for developing back-ends. In the McLAB group we are developing two such back-ends, one for generating FORTRAN and another for generating X10 code. Other research groups are also starting to use the Tamer framework, including a project for a back-end for Modelica.

We would like to encourage other compiler writers interested in MATLAB to consider using the Tamer, so that they can focus on issues in their backend. For example, a recent MATLAB compiler system called MEGHA (“MATLAB Execution on GPU based Heterogeneous Architectures”)[?] focuses on translating programs for parallel execution on GPU and CPU. They use a relatively small subset of MATLAB that does not even include user functions; the whole program is assumed to be one script. Using the Tamer they could support a wider set of features, even if the Tamer were just used to pre-process programs into a single inlined function (which is one possible output).

7. Related Work

There are several categories of related work. First, we have the immediate work upon which we are building. The McLAB project already provided the front-end and the McSAF [?] analysis framework, which provided an important basis for the Tamer. We also learned a lot from McLAB’s

previous McFOR project[5] which was a first prototype MATLAB to FORTRAN95 compiler. McFOR supported a smaller subset of the language, did not have a comprehensive approach to the builtin functions, and had a much more ad hoc approach to the analyses. However, it really showed that conversion of MATLAB to FORTRAN95 was possible, and that FORTRAN95 is an excellent target language. In this paper we have gone back to the basics and defined a much larger subset of MATLAB, taken a more structured and extensible approach to building a general toolkit, tackled the problem of a principled approach to the builtins, and defined the interprocedural analyses in a more rigorous and extensible fashion. The next generation of McFOR, as well as other backends, can now be built upon these new foundations.

Although we were not able to find publicly available versions, there have been several excellent previous research projects on static compilation of MATLAB which focused particularly on the array-based subset of MATLAB and developed advanced static analyses for determining shapes and sizes of arrays. For example, FALCON [?] is a MATLAB to FORTRAN90 translator with sophisticated type inference algorithms. Our Tamer is targeting a larger and more modern set of MATLAB that includes other types of data structures such as cell arrays and structs, function handles and lambda expressions, and which obeys the modern semantics of MATLAB 7. We should note that FALCON handled interprocedural issues by fully inlining all of the the code. MaJIC[?], a MATLAB Just-In-Time compiler, is patterned after FAL-

CON. It uses similar type inference techniques to FALCON, but they are simpler to fit the JIT context. MAGICA [?] is a type inference engine developed by Joisha and Banerjee of Northwestern University, which is written in Mathematica and is designed as an add-on module used by MAT2C compiler [?]. We hope to learn from the advanced type inference approaches in these projects and to implement similar approximations using our interprocedural value analysis.

The previously mentioned MEGHA project[?] provides an interesting approach to map MATLAB array operations to CPUs and GPUs, but only supports a very small subset of MATLAB.

There are also commercial compilers, which are not publicly available, and for which there are no research articles. One such product is the *MATLAB Coder* recently released by MathWorks[?]. This product produces C code for a subset of MATLAB. According to our preliminary tests, this product does not appear to support cell arrays except in very specific circumstances, nor does it support a general form of lambda expressions, and was therefore unable to handle quite a few of our benchmarks. However, the key differences with our work is that we are designing and providing an extensible and open source toolkit for compiler and tool researchers. This is clearly not the main goal of proprietary compilers.

There are other projects providing open source implementations of MATLAB-like languages, such as Octave[1] and Scilab[?]. Although these add valuable contributions to the open source community, their focus is on providing interpreters and open library support and they have not tackled the problems of static compilation. Thus, we believe that our contributions are complementary.

Other dynamic languages have had very successful efforts in defining compilable subsets and statically analyzing whole programs. For example RPython[2] uses a similar approach to ours, defining a reduced set of python that can be statically compiled, requiring that all variables have a unique possible type, while providing an analysis and compiler to compile that subset. RPython does not produce multiple versions of a function if it is called with different arguments; so every function can only be used with one set of argument types.

DiamondbackRuby (DRuby) is a static type inference toolkit for Ruby [4], mostly with the goal to gain the advantage of static languages to report potential errors ahead of time. Ruby, like MATLAB, is a dynamic, interpreted language, with many library functions in native code - which may also have different behaviours depending on the incoming argument types. Thus DRuby has to provide type information for builtin functions. In order to that, DRuby includes a type annotation language, which can also be used to specify types for functions with difficult behaviour. DRuby constrains the set of supported language features to enable the static analysis, but allows some of them by inserting runtime checks to still be able to support them. These are included in

such a way as to help users identify where exactly the error occurred.

8. Conclusions

This paper has introduced the MATLAB Tamer, an extensible object-oriented framework for supporting the translation from dynamic MATLAB programs to a Tame IR, call graph and class/type information suitable for generating static code. We provided an introduction to the features of MATLAB in a form that we believe helps expose the semantics of mclasses and function lookup for compiler and tool writers. We tackled the somewhat daunting problem of handling the large number of builtin functions in MATLAB by defining an extensible hierarchy of builtins and a small domain-specific language to define their behaviour. We defined a Tame IR and added functionality to McSAF to produce the IR and to extend the analysis framework to handle the new IR nodes introduced. Finally, we developed an extensible interprocedural analysis framework and an extensible value analysis that can be used to build a complete callgraph, which estimates the mclass of every variable. We provide these frameworks, for users to implement further interprocedural analyses, or as an entry point for static MATLAB compilers.

Our initial experiments with the framework are very encouraging and we are now working on using the framework to implement back-ends, and we hope that others will also use the framework for a variety of static MATLAB tools.¹² We also plan to continue developing the value analysis to add richer abstractions for shape and other data structure properties. Finally, as a part of a larger project on benchmarking MATLAB, we hope to expand our set of benchmarks and to further examine which features might be tamed, and to extend our set of automated refactorings.

References

- [1] GNU Octave. <http://www.gnu.org/software/octave/index.html>.
- [2] D. Ancona, M. Ancona, A. Cuni, and N. D. Matsakis. Rpython: a step towards reconciling dynamically and statically typed oo languages. In *DLS '07: Proceedings of the 2007 symposium on Dynamic languages*, pages 53–64, New York, NY, USA, 2007. ACM.
- [3] M. Chevalier-Boisvert, L. Hendren, and C. Verbrugge. Optimizing MATLAB through Just-In-Time Specialization. In *International Conference on Compiler Construction*, pages 46–65, March 2010.
- [4] M. Furr, J.-h. D. An, J. S. Foster, and M. Hicks. Static type inference for Ruby. In *Proceedings of the 2009 ACM symposium on Applied Computing, SAC '09*, pages 1859–1866, New York, NY, USA, 2009. ACM.

¹²The Tamer and McLAB release is available at http://www.sable.mcgill.ca/mclab/software/software_releases_mclab.html.

[5] J. Li. McFor: A MATLAB to FORTRAN 95 Compiler. Master's thesis, McGill University, August 2009.

A. Example Result

The following code is what the Tamer produces for the example introduced in section Sec. 2.8. Since MATLAB is a proprietary system, we opted to omit the body of the library function that was included in the callgraph.

```
% args: {N=(double)}
function [result] = main(N)
    mc_t4 = 10; % mc.t4=(double,10.0)
    [mc_t1] = single(mc_t4); % mc.t1=(single)
    rod.L = mc_t1; % rod=struct{L=(single)}
    mc_t5 = 3; % mc.t5=(double,3.0)
    [mc_t2] = single(mc_t5); % mc.t2=(single)
    rod.Ta = mc_t2; % rod=struct{Ta=(single),L=(single)}
    mc_t6 = 4; % mc.t6=(double,4.0)
    [mc_t3] = single(mc_t6); % mc.t3=(single)
    rod.Tb = mc_t3; % rod=struct{Ta=(single),Tb=(single),L=(single)}
    a = .2; % a=(double,0.2)
    A = 0.05; % A=(double,0.05)
    mc_t8 = 0; % mc.t8=(double,0.0)
    mc_t9 = 20; % mc.t9=(double,20.0)
    [Tspan] = horzcat(mc_t8, mc_t9); % Tspan=(double)
    [result] = solveSystem(rod, a, A, N, Tspan);
    % result=(single)

end
% results: [(single)]

% args: {rod=struct{Ta=(single), Tb=(single), L=(single)},
% a=(double), A=(double), N=(double), Tspan=(double)}
function [out]=solveSystem(rod,a,A,N,Tspan,Tsteps)
    [mc_t10] = nargin(); % mc.t10=(double,5.0)
    mc_t34 = 4; % mc.t34=(double,4.0)
    [mc_t16] = isequal(mc_t10,mc_t34); % mc.t16=(logical,false)
    if mc_t16
        mc_t36 = 0; % non-viable
        mc_t37 = 10; % non-viable
        [Tspan] = horzcat(mc_t36, mc_t37); % non-viable
        Tsteps = 100; % non-viable
    else
        mc_t38 = 5; % mc.t38=(double,5.0)
        [mc_t15] = isequal(mc_t10,mc_t38); % mc.t15=(logical,true)
        if mc_t15
            Tsteps = 100; % Tsteps=(double,100.0)
        else
            end
        end
    end
    [mc_t17] = rod.L; % mc.t17=(single)
    mc_t39 = 1; % mc.t39=(double,1.0)
    [mc_t18] = plus(N, mc_t39); % mc.t18=(double)
    [h] = mrdivide(mc_t17, mc_t18); % h=(single)
    mc_t21 = h; % mc.t21=(single)
    mc_t22 = h; % mc.t22=(single)
    [mc_t24] = rod.L; % mc.t24=(single)
    mc_t25 = h; % mc.t25=(single)
    [mc_t23] = minus(mc_t24, mc_t25); % mc.t23=(single)
    [mc_t20] = colon(mc_t21, mc_t22, mc_t23);
    % mc.t20=(single)
    [mc_t19] = vertcat(mc_t20); % mc.t19=(single)
    [X] = ctranspose(mc_t19); % X=(single)
    U0 = X; % U0=(single)
    [mc_t27] = length(X); % mc.t27=(double)
    mc_t40 = 2; % mc.t40=(double,2.0)
    [mc_t26] = mrdivide(mc_t27, mc_t40); % mc.t26=(double)
    [mid] = round(mc_t26); % mid=(double)
```

```
[mc_t11] = rod.Ta; % mc.t11=(single)
mc_t41 = 1; % mc.t41=(double,1.0)
[mc_t13] = colon(mc_t41, mid); % mc.t13=(double)
U0(mc_t13) = mc_t11; % U0=(single)
[mc_t12] = rod.Tb; % mc.t12=(single)
mc_t28 = mid; % mc.t28=(double)
mc_t42 = 1; % mc.t42=(double,1.0)
mc_t43 = 1; % mc.t43=(double,1.0)
[mc_t29] = end(U0, mc_t42, mc_t43); % mc.t29=(double)
[mc_t14] = colon(mc_t28, mc_t29); % mc.t14=(double)
U0(mc_t14) = mc_t12; % U0=(single)
mc_t30 = N; % mc.t30=(double)
mc_t31 = rod.Ta; % mc.t31=(single)
mc_t32 = rod.Tb; % mc.t32=(single)
mc_t33 = h; % mc.t33=(single)
[D, c] = Dxx(mc_t30, mc_t31, mc_t32, mc_t33);
% D=(single), c=(single)

f = (@(t, u) lambda_1(a, D, c, A, t, u));
% f=(handle,@lambda.1((double),(single),(single),(double),...))
[T, out] = RungeKutta2(f, Tspan, U0, Tsteps);
% T=(double), out=(single)

% args: {N=(double), a=(single), b=(single), h=(single)}
function [D, c] = Dxx(N, a, b, h)
    mc_t62 = 2; % mc.t62=(double,2.0)
    [mc_t51] = uminus(mc_t62); % mc.t51=(double)
    mc_t63 = 2; % mc.t63=(double,2.0)
    [mc_t53] = minus(N, mc_t63); % mc.t53=(double)
    mc_t64 = 1; % mc.t64=(double,1.0)
    [mc_t52] = zeros(mc_t53, mc_t64); % mc.t52=(double)
    mc_t65 = 1; % mc.t65=(double,1.0)
    [mc_t46] = vertcat(mc_t51, mc_t65, mc_t52);
    % mc.t46=(double)
    mc_t66 = 2; % mc.t66=(double,2.0)
    [mc_t48] = uminus(mc_t66); % mc.t48=(double)
    mc_t67 = 2; % mc.t67=(double,2.0)
    [mc_t50] = minus(N, mc_t67); % mc.t50=(double)
    mc_t68 = 1; % mc.t68=(double,1.0)
    [mc_t49] = zeros(mc_t68, mc_t50); % mc.t49=(double)
    mc_t70 = 1; % mc.t70=(double,1.0)
    [mc_t47] = horzcat(mc_t48, mc_t70, mc_t49);
    % mc.t47=(double)
    [mc_t44] = toeplitz(mc_t46, mc_t47); % mc.t44=(double)
    mc_t71 = 2; % mc.t71=(double,2.0)
    [mc_t45] = mpower(h, mc_t71); % mc.t45=(single)
    [D] = rdivide(mc_t44, mc_t45); % D=(single)
    mc_t60 = a; % mc.t60=(single)
    mc_t72 = 2; % mc.t72=(double,2.0)
    [mc_t61] = mpower(h, mc_t72); % mc.t61=(single)
    [mc_t54] = mrdivide(mc_t60, mc_t61); % mc.t54=(single)
    mc_t73 = 2; % mc.t73=(double,2.0)
    [mc_t59] = minus(N, mc_t73); % mc.t59=(double)
    mc_t74 = 1; % mc.t74=(double,1.0)
    [mc_t55] = zeros(mc_t59, mc_t74); % mc.t55=(double)
    mc_t57 = b; % mc.t57=(single)
    mc_t75 = 2; % mc.t75=(double,2.0)
    [mc_t58] = mpower(h, mc_t75); % mc.t58=(single)
    [mc_t56] = mrdivide(mc_t57, mc_t58); % mc.t56=(single)
    [c] = vertcat(mc_t54, mc_t55, mc_t56); % c=(single)

end
% results: [(single), (single)]

% args: {c=(double), r=(double)}
function [t] = toeplitz(c, r)
    [mc_t80] = nargin(); % mc.t80=(double,2.0)
    ... Code Omitted
```

```

t(:) = mc_t77; % t=(double)
end
% results: [(double)]

% args: {f=(handle, @lambda.1((double),(single),(single),(double),...)),
%        tspan=(double), y0=(single),
N=(double)}
function [X, Y] = RungeKutta2(f, tspan, y0, N)
mc_t139 = 1; % mc.t139=(double,1.0)
[a] = tspan(mc_t139); % a=(double)
mc_t140 = 2; % mc.t140=(double,2.0)
[b] = tspan(mc_t140); % b=(double)
[mc_t116] = minus(b, a); % mc.t116=(double)
mc_t117 = N; % mc.t117=(double)
[h] = mrdivide(mc_t116, mc_t117); % h=(double)
mc_t118 = y0; % mc.t118=(single)
[mc_t120] = length(y0); % mc.t120=(double)
mc_t121 = N; % mc.t121=(double)
[mc_t119] = zeros(mc_t120, mc_t121); % mc.t119=(double)
[Y] = horzcat(mc_t118, mc_t119); % Y=(single)
[mc_t122] = colon(a, h, b); % mc.t122=(double)
[X] = transpose(mc_t122); % X=(double)
mc_t145 = 1; % mc.t145=(double,1.0)
for j = (mc_t145 : N);
mc_t123 = h; % mc.t123=(double)
[mc_t125] = X(j); % mc.t125=(double)
[mc_t126] = Y(:, j); % mc.t126=(single)
[mc_t124] = f(mc_t125, mc_t126); % mc.t124=(single)
[k1] = mtimes(mc_t123, mc_t124); % k1=(single)
mc_t127 = h; % mc.t127=(double)
[mc_t133] = X(j); % mc.t133=(double)
mc_t134 = h; % mc.t134=(double)
[mc_t129] = plus(mc_t133, mc_t134); % mc.t129=(double)
[mc_t131] = Y(:, j); % mc.t131=(single)
mc_t132 = k1; % mc.t132=(single)
[mc_t130] = plus(mc_t131, mc_t132); % mc.t130=(single)
[mc_t128] = f(mc_t129, mc_t130); % mc.t128=(single)
[k2] = mtimes(mc_t127, mc_t128); % k2=(single)
[mc_t135] = Y(:, j); % mc.t135=(single)
mc_t142 = 1; % mc.t142=(double,1.0)
mc_t143 = 2; % mc.t143=(double,2.0)
[mc_t137] = mrdivide(mc_t142, mc_t143); % mc.t137=(double)
[mc_t138] = plus(k1, k2); % mc.t138=(single)
[mc_t136] = mtimes(mc_t137, mc_t138); % mc.t136=(single)
[mc_t114] = plus(mc_t135, mc_t136); % mc.t114=(single)
mc_t144 = 1; % mc.t144=(double,1.0)
[mc_t115] = plus(j, mc_t144); % mc.t115=(double)
Y(:, mc_t115) = mc_t114; % Y=(single)
end
end
% results: [(double), (single)]

% args: {a=(double), D=(single), c=(single), A=(double),
%        t=(double), u=(single)}
function [mc_t0] = lambda_1(a, D, c, A, t, u)
mc_t148 = a; % mc.t148=(double)
[mc_t150] = mtimes(D, u); % mc.t150=(single)
mc_t151 = c; % mc.t151=(single)
[mc_t149] = plus(mc_t150, mc_t151); % mc.t149=(single)
[mc_t146] = mtimes(mc_t148, mc_t149); % mc.t146=(single)
mc_t147 = A; % mc.t147=(double)
[mc_t0] = plus(mc_t146, mc_t147); % mc.t0=(single)
end
% results: [(single)]

```

B. TIR Grammar

In this appendix we present the abstract syntax tree structure corresponding to the grammar for the tame IR.

We have listed all tame IR nodes, together with the parent class and the nodes they contain. All tame IR nodes either extend McSAF's AST nodes, or other Tame IR Nodes. This means that any Tame IR node is also a valid McSAF AST node. Tame IR nodes may contain other AST nodes, because they are effectively little AST subtrees. Users of the Tame IR should not modify IR Nodes, except the `TIRStatementList`. They should also only use the accessor methods provided by the Tame IR interfaces. By following these conventions users will always create correct Tame IR because the constructors of the Tame IR nodes enforce the constraints of the Tame IR. Figure 10 shows all IR statement nodes that are derived from assignments (Figure 6 in Sec. 4.1 shows their hierarchy as a tree), Figure 11 shows all remaining statement nodes, Figure 12 shows all Tame IR nodes that may contain statements, Figure 13 shows all remaining tame IR nodes.

node	extends	contains
TIRReturnStmt	ReturnStmt	-
TIRBreakStmt	BreakStmt	-
TIRContinueStmt	ContinueStmt	-
TIRGlobalStmt	GlobalStmt	List<Name> names
TIRPersistentStmt	PersistentStmt	List<Name> names
TIRCommentStmt	EmptyStmt	-

Figure 11. Non-assignment statements

node	extends	contains
TIRFunction	Function	List<Name> outputParams, String name, List<Name> inputParams, List<HelpComment> helpComments, TIRStmtList stmts, List<TIRFunction> nestedFunctions
TIRStmtList	List<Stmt>	List<TIRStmt> statements
TIRIfStmt	IfStmt	Name ConditionVar, TIRStmtList IfStmts, TIRStmtList ElseStmts
TIRWhileStmt	WhileStmt	Name condition, TIRStmtList body
TIRForStmt	ForStmt	Name var, Name lower, (Name inc), Name upper, TIRStmtList stmts

Figure 12. Compound Structures

node	extends	contains
TIRCommaSeparatedList	List<Expr>	List<Expr> elements

Figure 13. Other Tame IR Nodes

node	extends	contains
TIRAbstractAssignStmt	AssignStmt	–
TIRAbstractAssignFromVarStmt	TIRAbstractAssignStmt	Name rhs
TIRArraySetStmt	TIRAbstractAssignFromVarStmt	Name arrayVar, TIRCommaSeparatedList indices, Name rhs
TIRCellArraySetStmt	TIRAbstractAssignFromVarStmt	Name arrayVar, TIRCommaSeparatedList indices, Name rhs
TIRDotSetStmt	TIRAbstractAssignFromVarStmt	Name dotVar, Name field, Name rhs
TIRAbstractAssignToListStmt	TIRAbstractAssignStmt	IRCommaSeparatedList targets
TIRArrayGetStmt	TIRAbstractAssignToListStmt	Name lhs, Name rhs, TIRCommaSeparatedList indices
TIRCellArrayGetStmt	TIRAbstractAssignToListStmt	Name cellVar, TIRCommaSeparatedList targets, TIRCommaSeparatedList indices
TIRDotGetStmt	TIRAbstractAssignToListStmt	TIRCommaSeparatedList lhs, Name dotVar, Name field
TIRCallStmt	TIRAbstractAssignToListStmt	Name function, TIRCommaSeparatedList targets, TIRCommaSeparatedList args
TIRAbstractAssignToVarStmt	TIRAbstractAssignStmt	Name lhs
TIRAssignLiteralStmt	TIRAbstractAssignToVarStmt	Name lhs, LiteralExpr rhs
TIRCopyStmt	TIRAbstractAssignToVarStmt	Name lhs, Name rhs
TIRAbstractCreateFunctionHandleStmt	TIRAbstractAssignToVarStmt	Name lhs, Name function
TIRCreateFunctionReferenceStmt	TIRAbstractCreateFunctionHandleStmt	Name lhs, Name function
TIRCreateLambdaStmt	TIRAbstractCreateFunctionHandleStmt	Name lhs, Name function List<Name> lambdaParameters, List<Name> enclosedVariables

Figure 10. Assignment Statments