

Lecture 17 – Dataflow Analysis (and more!)

THEORY OF COMPILATION

Eran Yahav

www.cs.technion.ac.il/~yahave/tocs2011/compilers-lec17.pptx

Reference: Dragon 9, 12

1

Last week: dataflow

- General recipe
 - design the domain
 - transfer functions
 - determine join operation (may/must?)
 - direction: forward/backward?
- note initial values

2

Analyses Summary

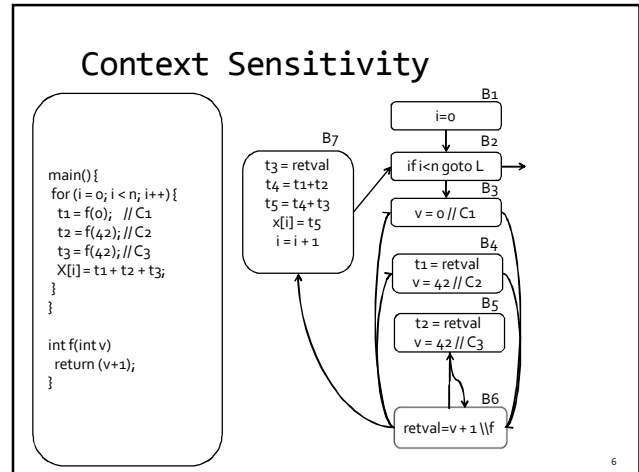
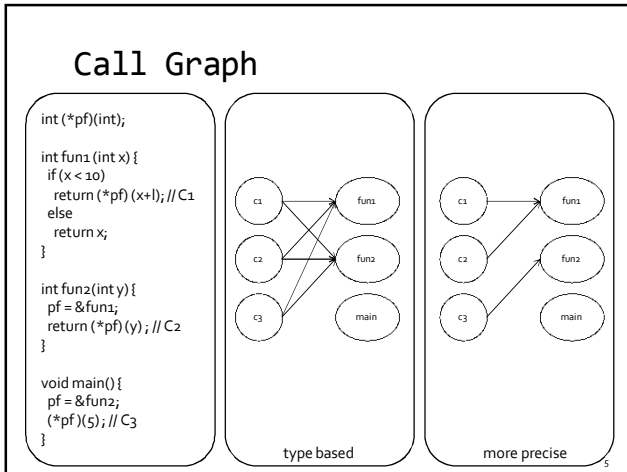
	Reaching Definitions	Available Expressions	Live Variables
L	$\wp(\text{Var} \times \text{Lab})$	$\wp(\text{AExp})$	$\wp(\text{Var})$
\sqsubseteq	\subseteq	\supseteq	\subseteq
\sqcup	\cup	\cap	\cup
\perp	\emptyset	AExp	\emptyset
Initial	$\{(x,?) \mid x \in \text{Var}\}$	\emptyset	\emptyset
Entry labels	{ init }	{ init }	final
Direction	Forward	Forward	Backward
F	$\{ f: L \rightarrow L \mid \exists k, g: f(\text{val}) = (\text{val} \setminus k) \cup g \}$		
f_{lab}	$f_{lab}(\text{val}) = (\text{val} \setminus \text{kill}) \cup \text{gen}$		

3

Interprocedural Analysis

4

- The effect of calling a procedure is the effect of executing its body

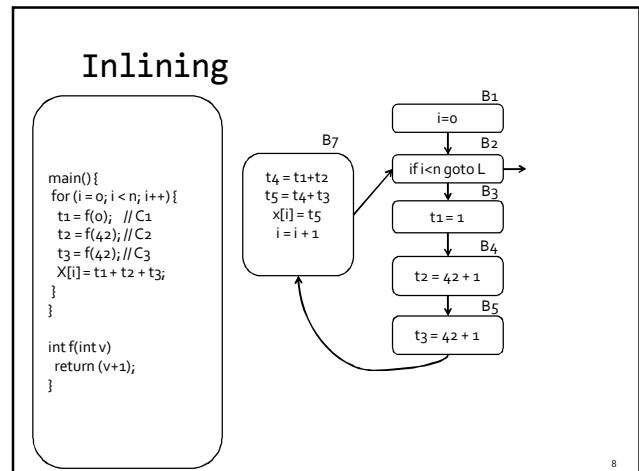


Solution Attempt #1

- Inline callees into callers
 - End up with one big procedure
 - CFGs of individual procedures = duplicated many times
- Good: it is precise
 - distinguishes different calls to the same function
- Bad
 - exponential blow-up, not efficient
 - doesn't work with recursion

```

main() { f(); f(); }
f() { g(); g(); }
g() { h(); h(); }
h() { ... }
    
```




Solution Attempt #2

- Build a "supergraph" = inter-procedural CFG
- Replace each call from P to Q with
 - An edge from point before the call (call point) to Q's entry point
 - An edge from Q's exit point to the point after the call (return pt)
 - Add assignments of actuals to formals, and assignment of return value
- Good: efficient
 - Graph of each function included exactly once in the supergraph
 - Works for recursive functions (although local variables need additional treatment)
- Bad: imprecise, "context-insensitive"
 - The "unrealizable paths problem": dataflow facts can propagate along infeasible control paths

9

Unrealizable Paths



10

Interprocedural Analysis


```

begin
proc p() is1
  [x := a + 1]2
end3
[a = 7]4
[call p()]56
[print x]7
[a = 9]8
[call p()]910
[print a]11
end
    
```

- Extend language with begin/end and with [call p()]^{clab}_{rlab}
- Call label clab, and return label rlab


11

IVP: Interprocedural Valid Paths



- IVP: all paths with matching calls and returns
- And prefixes

Valid Paths



13

Interprocedural Valid Paths

- IVP set of paths
 - Start at program entry
- Only considers matching calls and returns
 - aka, valid
- Can be defined via context free grammar
 - $\text{matched} ::= \text{matched } (, \text{matched }) | \epsilon$
 - $\text{valid} ::= \text{valid } (, \text{matched } | \text{matched}$
 - *paths* can be defined by a regular expression

The Join-Over-Valid-Paths (JVP)

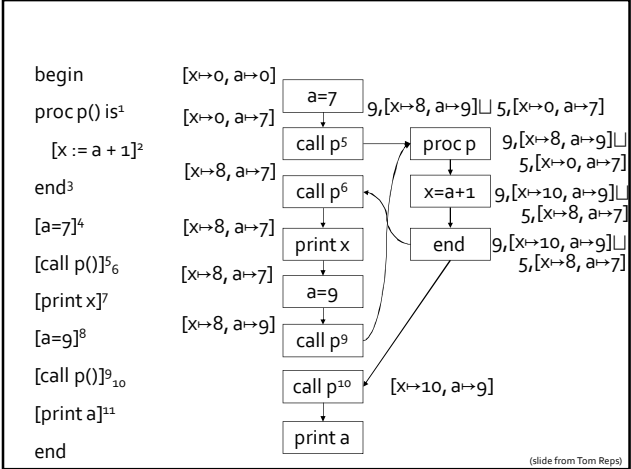
- $\text{vpaths}(n)$ all valid paths from program start to n
- $\text{JVP}[n] = \sqcup \{ \llbracket e_{i_1} e_{i_2} \dots, e \rrbracket \mid (e_{i_1} e_{i_2} \dots, e) \in \text{vpaths}(n) \}$ (initial)
- $\text{JVP} \sqsubseteq \text{JFP}$
 - In some cases the JVP can be computed
 - (Distributive problem)

Sharir and Pnueli '82

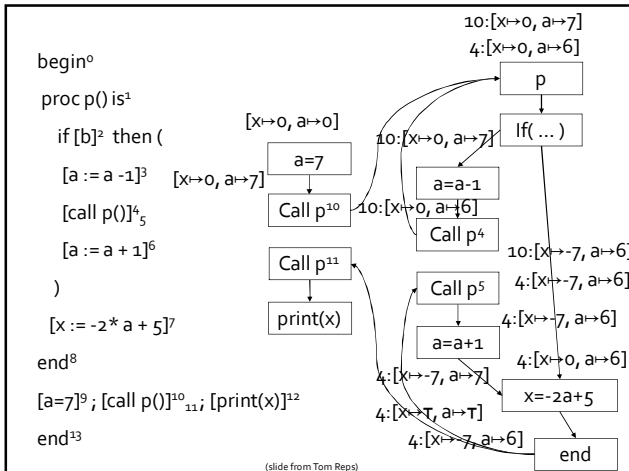
- Call String approach
 - Blend interprocedural flow with intra procedural flow
 - Tag every dataflow fact with call history
- Functional approach
 - Determine the effect of a procedure
 - E.g., in/out map
 - Treat procedure invocations as "super ops"

The Call String Approach

- Record at every node a pair (l, c) where l ∈ L is the dataflow information and c is a suffix of unmatched calls
- Use Chaotic iterations
- To guarantee termination limit the size of c (typically 1 or 2)
- Emulates inline (but no code growth)
- Exponential in size of c



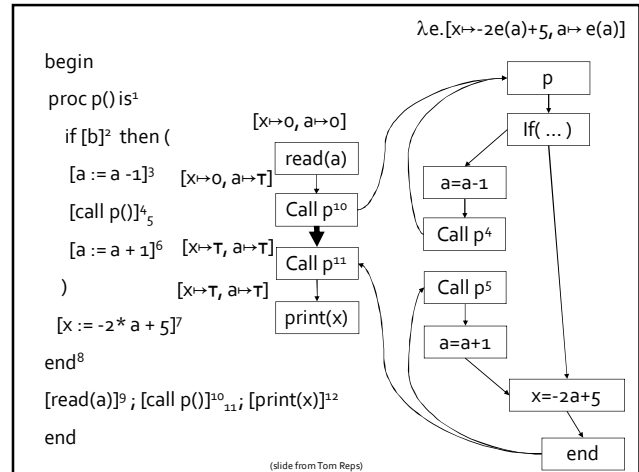
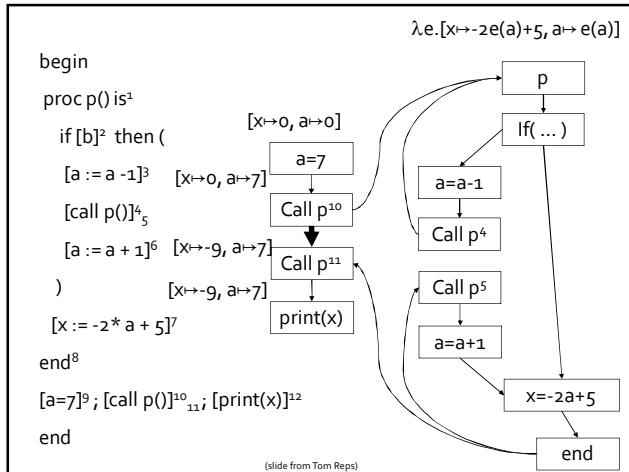
(slide from Tom Reps)



(slide from Tom Reps)

The Functional Approach

- The meaning of a procedure is mapping from states into states
- The abstract meaning of a procedure is function from an abstract state to abstract states

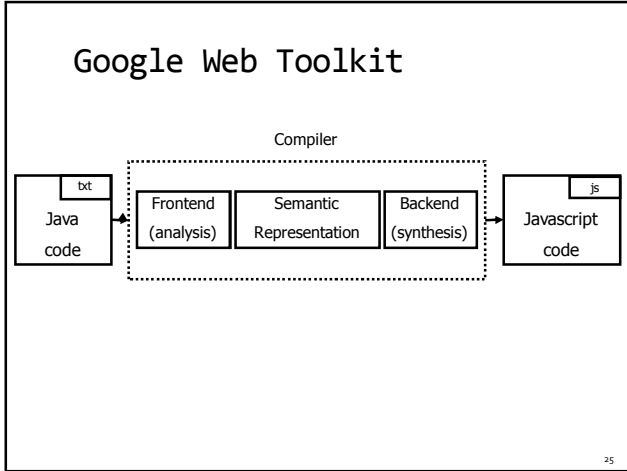


Functional Approach: Main Idea

- Iterate on the abstract domain of functions from L to L
- Two phase algorithm
 - Compute the dataflow solution at the exit of a procedure as a function of the initial values at the procedure entry (functional values)
 - Compute the dataflow values at every point using the functional values
- Computes JVP for distributive problems

Meanwhile, IRL

- new compilers for new languages
- new compilers for old languages
 - e.g., Java->JavaScript
- new uses of compiler technology
- ...



GWT Compiler Optimization

```

public class ShapeExample implements EntryPoint {
    private static final double SIDE_LEN_SMALL = 2;
    private final Shape shape = new SmallSquare();
    public static abstract class Shape {
        public abstract double getArea();
    }
    public static abstract class Square extends Shape {
        public double getArea() { return getSideLength() * getSideLength(); }
        public abstract double getSideLength();
    }
    public static class SmallSquare extends Square {
        public double getSideLength() { return SIDE_LEN_SMALL; }
    }
    public void onModuleLoad() {
        Shape shape = getShape();
        Window.alert("Area is " + shape.getArea());
    }
    private Shape getShape() { return shape; }
}
    
```

(source: http://dl.google.com/dl/200g/pres/Th_3043_TheStoryofourCompile-ReadingtheTeaLeavesoftheGWTCompilerforanOptimizedFuture.pdf) 26

GWT Compiler Optimization

```

public class ShapeExample implements EntryPoint {
    public void onModuleLoad() {
        Window.alert("Area is 4.0");
    }
}
    
```

(source: http://dl.google.com/dl/200g/pres/Th_3043_TheStoryofourCompile-ReadingtheTeaLeavesoftheGWTCompilerforanOptimizedFuture.pdf) 27

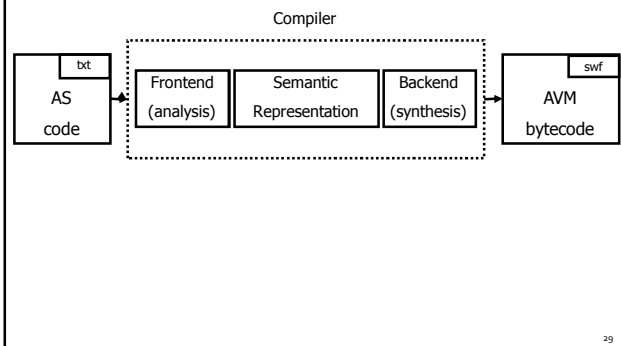
Adobe ActionScript

/*AS*/ First introduced in Flash Player 9, **ActionScript 3** is an object-oriented programming (OOP) language based on ECMAScript—the same standard that is the basis for JavaScript—and provides incredible gains in runtime performance and developer productivity. **ActionScript 2**, the version of ActionScript used in Flash Player 8 and earlier, continues to be supported in Flash Player 9 and Flash Player 10.

ActionScript 3.0 introduces a new highly optimized ActionScript Virtual Machine, AVM2, which dramatically exceeds the performance possible with AVM1. As a result, ActionScript 3.0 code executes up to 10 times faster than legacy ActionScript code.

(source: http://www.adobe.com/devnet/actionscript/articles/actionscript3_overview.html) 28

Adobe ActionScript



29

Adobe ActionScript: Tamarin

The goal of the "Tamarin" project is to implement a high-performance, open source implementation of the ActionScript™ 3 language, which is based upon and extends ECMAScript 3rd edition (ES3). ActionScript provides many extensions to the ECMAScript language, including packages, namespaces, classes, and optional strict typing of variables. "Tamarin" implements both a high-performance just-in-time compiler and interpreter.

The Tamarin virtual machine is used within the Adobe® Flash® Player and is also being adopted for use by projects outside Adobe. The Tamarin just-in-time compiler (the "NanoJIT") is a collaboratively developed component used by both Tamarin and Mozilla TraceMonkey. The ActionScript compiler is available as a component from the open source [Flex SDK](#).

30

Mozilla SpiderMonkey

- SpiderMonkey is a fast interpreter
- runs an untyped bytecode and operates on values of type `jsval`—type-tagged double-sized values that represent the full range of JavaScript values.
- SpiderMonkey contains two JavaScript Just-In-Time (JIT) compilers, a garbage collector, code implementing the basic behavior of JavaScript values...
- SpiderMonkey's interpreter is mainly a single, tremendously long function that steps through the bytecode one instruction at a time, using a switch statement (or faster alternative, depending on the compiler) to jump to the appropriate chunk of code for the current instruction.

(source: <https://developer.mozilla.org/En/SpiderMonkeyInternals>)

31

Mozilla SpiderMonkey: Compiler

- consumes JavaScript source code
- produces a *script* which contains bytecode, source annotations, and a pool of string, number, and identifier literals. The script also contains objects, including any functions defined in the source code, each of which has its own, nested script.
- The compiler consists of
 - a random-logic rather than table-driven lexical scanner
 - a recursive-descent parser that produces an AST
 - a tree-walking code generator
- The emitter does some constant folding and a few codegen optimizations

(source: <https://developer.mozilla.org/En/SpiderMonkeyInternals>)

32

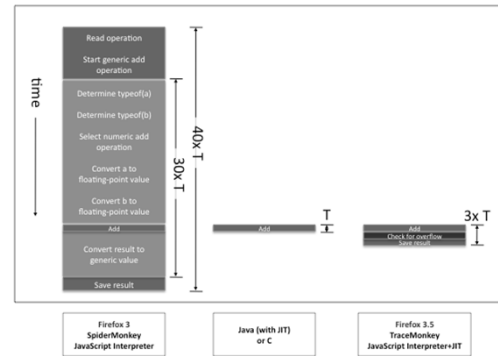
Mozilla SpiderMonkey

- SpiderMonkey contains a just-in-time trace compiler that converts bytecode to machine code for faster execution.
- The JIT works by detecting commonly executed loops, tracing the executed bytecodes in those loops as they run in the interpreter, and then compiling the trace to machine code.
- See the page about the [Tracing JIT](#) for more details.
- The SpiderMonkey GC is a mark-and-sweep, non-conservative (exact) collector.

(source: <https://developer.mozilla.org/En/SpiderMonkey/Internals>)

33

Mozilla TraceMonkey



(source: <http://hacks.mozilla.org/2009/07/tracemonkey-overview/>)

34

Mozilla TraceMonkey

- Goal: generate type-specialized code
- Challenges
 - no type declarations
 - statically trying to determine types mostly hopeless
- Idea
 - run the program for a while and observe types
 - use observed types to generate type-specialized code
 - compile traces
- Sounds suspicious?

35

Mozilla TraceMonkey

- Problem 1: "observing types" + compiling possibly more expensive than running the code in the interpreter
- Solution: only compile code that executes many times
 - "hot code" = loops
 - initially everything runs in the interpreter
 - start tracing a loop once it becomes "hot"

36

Mozilla TraceMonkey

- Problem 2: past types do not guarantee future types... what happens if types change?
- Solution: insert type-checks into the compiled code
 - if type-check fails, need to recompile for new types
 - code with frequent type changes will suffer some slowdown

37

Mozilla TraceMonkey

```
function addTo(a, n){
  for (var i = 0; i < n; ++i)
    a = a + i;
  return a;
}
```

```
var to = new Date();
var n = addTo(0,
10000000);
print(n);
print(new Date() - to);
```

```
a = a + i; // a is an integer number (0 before, 1 after)
++i; // i is an integer number (1 before, 2 after)
if (!(i < n)) // n is an integer number (10000000)
  break;
```

```
trace_1_start:
++i; // i is an integer number (0 before, 1 after)
temp = a + i; // a is an integer number (1 before, 2 after)
if (lastOperationOverflowed())
  exit_trace(OVERFLOWED);
a = temp;
if (!(i < n)) // n is an integer number (10000000)
  exit_trace(BRANCHED);
goto trace_1_start;
```

38

Mozilla TraceMonkey

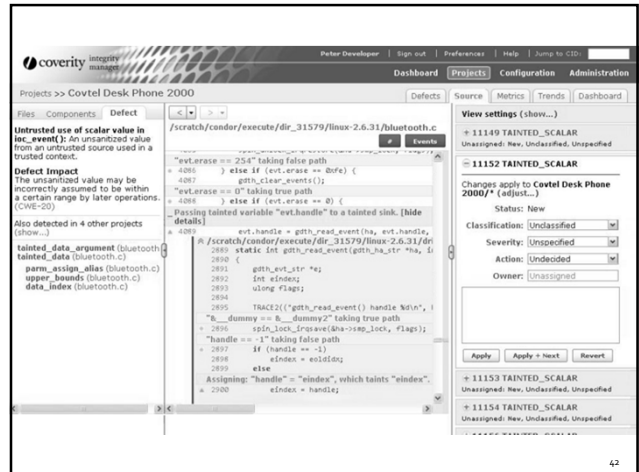
System	Run Time (ms)
SpiderMonkey (FF3)	990
TraceMonkey (FF3.5)	45
Java (using int)	25
Java (using double)	74
C (using int)	5
C (using double)	15

39

Static Analysis Tools

- Coverity
- SLAM
- ASTREE
- ...

40



Lots and lots of research

- Program Analysis
- Program Synthesis
- ...
- Next semester
 - Project – contact me if you’re interested
 - Advanced course in program analysis

THE END