

AD-A259 417

12



Technical Report 1321

Concurrent Smalltalk on the Message-Driven Processor

Waldemar Horwat

MIT Artificial Intelligence Laboratory

DTIC
ELECTE
JAN 22 1993
S E D

DISTRIBUTION STATEMENT
Approved for public release
Distribution Unlimited

93-01047



03 1 21 035

REPORT DOCUMENTATION PAGE

Form Approved
OMB No 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 1991	3. REPORT TYPE AND DATES COVERED technical report	
4. TITLE AND SUBTITLE Concurrent Smalltalk on the Message-Driven Processor			5. FUNDING NUMBERS N00014-87-K-0825 N00014-87-K-0738 MIP-8657531	
6. AUTHOR(S) Waldemar Horwat				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Artificial Intelligence Laboratory 545 Technology Square Cambridge, Massachusetts 02139			8. PERFORMING ORGANIZATION REPORT NUMBER AI-TR 1321	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research Information Systems Arlington, Virginia 22217			10. SPONSORING MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES None				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Distribution of this document is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) <p style="text-align: center;">Million-transistor processors are being manufactured today, and soon it will be possible to put several million transistors on one integrated circuit. While memory applications of this technology are clear, it is not obvious how best to use it for computation purposes. One possibility is the architecture of the Message-Driven Processor (MDP), which consists of a 32+4-bit CPU, memory, and a network interface together on one chip. MDPs can be connected directly to each other to form a 65536-processor, message-passing, MIMD, parallel computer, the J-Machine. The MDP's architecture is unusual in that it provides a very high processing power to memory ratio.</p> <p style="text-align: right;">(continued on back)</p>				
14. SUBJECT TERMS (key words) compiler message passing parallel processing object oriented			15. NUMBER OF PAGES 240	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UNCLASSIFIED	

Block 13 continued:

Concurrent Smalltalk is the primary language used for programming the J-Machine. Concurrent Smalltalk is the the language of choice because it fits the J-Machine's fine-grain, message-passing model well. This thesis describes Concurrent Smalltalk and its implementation on the J-Machine, including the Optimist II compiler and Cosmos operating system. Optimist II can perform global optimization of programs, including inline function expansion, type inference, and global evaluation of constant expressions. Next, Cosmos and the Concurrent Smalltalk runtime environment are described. Finally, some quantitative and qualitative results are presented. The grain size (the average amount of time a method executes before suspending) was found to be about 60 instructions, and the MDP was found to execute one instruction every two or four cycles, depending on whether external DRAM is used. A number of qualitative issues are described, along with a few preliminary results for addressing difficult problems such as controlling parallelism.

Concurrent Smalltalk on the Message-Driven Processor

by
Waldemar Horwat

Submitted to the Department of Electrical Engineering and Computer Science on
May 12, 1989 in partial fulfillment of the requirements for the degree of
Master of Science in Computer Science

Updated September 26, 1991

Abstract

Million-transistor processors are being manufactured today, and soon it will be possible to put several million transistors on one integrated circuit. While memory applications of this technology are clear, it is not obvious how best to use it for computation purposes. One possibility is the architecture of the Message-Driven Processor (MDP), which consists of a 32+4-bit CPU, memory, and a network interface together on one chip. MDPs can be connected directly to each other to form a 65536-processor, message-passing, MIMD, parallel computer, the J-Machine. The MDP's architecture is unusual in that it provides a very high processing power to memory ratio.

Concurrent Smalltalk is the primary language used for programming the J-Machine. Concurrent Smalltalk is the the language of choice because it fits the J-Machine's fine-grain, message-passing model well. This thesis describes Concurrent Smalltalk and its implementation on the J-Machine, including the Optimist II compiler and Cosmos operating system. Optimist II can perform global optimization of programs, including inline function expansion, type inference, and global evaluation of constant expressions. Next, Cosmos and the Concurrent Smalltalk runtime environment are described. Finally, some quantitative and qualitative results are presented. The grain size (the average amount of time a method executes before suspending) was found to be about 60 instructions, and the MDP was found to execute one instruction every two or four cycles, depending on whether external DRAM is used. A number of qualitative issues are described, along with a few preliminary results for addressing difficult problems such as controlling parallelism.

Thesis Supervisor: William J. Dally, Ph.D.
Title: Associate Professor of Computer Science and Engineering

Acknowledgements

I would like to thank

Professor Bill Dally for his support during the difficult times as well as the easy ones and for advice without giving orders.

Andrew Chien for providing many insightful comments while this work was being done, reviewing drafts of this thesis, and sharing his ideas about the development of abstractions in concurrent programming.

Scott Wills for providing suggestions and taking his time to show me around the MIT purchasing bureaucracy.

Carl Manning for providing feedback on Concurrent Smalltalk and other ideas.

The co-authors of Concurrent Smalltalk, Andrew Chien and Scott Wills, and the original creator of the language, Bill Dally, for making a wonderful language for programming the J-Machine.

Peter Nuth for his well-developed, Canadian sense of humor and a healthy attitude towards MIT, as well as for his numerous contributions to the MDP project.

The other hardbeaners, Stuart Fiske, John Keen, Rich Lethin, Mike Noakes, and Debbie Wallach for working long hours striving to make the MDP chip a success.

Brian Totty for writing the operating system on which Cosmos is based.

Lucien Van Elsen for providing relevant ideas about micro-optimization of floating point calculations on the MDP.

Ricardo Jenez and John Wolfe for knowing everything about MIT's computers, sharing that knowledge, and for connecting UNIX computers to the Macintosh.

Ellen Spertus for getting the first dataflow compiler to work on the MDP and for rescuing this thesis by getting Word 4.0 directly from Microsoft in less than twenty hours.

Kathy Knobe for listening well and suggesting ideas for further research.

Scott Furman, Todd Dampier, and Shaun Kaneshiro for debugging the J-Machine, bringing up Cosmos on it, and adding extensions to it.

Steve Keckler and Larry Dennison for being friends.

Anant Agarwal, my advisor, for advice.

The Office of Naval Research for providing a fellowship with no strings attached and amazingly little bureaucracy.

Also, I would like to thank Bernard and Maria Horwat, my parents, for their love and support which made my education possible.

Most of all, I would like to thank God for the abilities and opportunities He has granted me.

Table of Contents

Chapter 1. Introduction.....	1
Goals	1
Second Edition	1
1.1. Hardware and Software Architecture	2
The J-Machine.....	2
The Message-Driven Processor	2
Concurrent Smalltalk	2
1.2. Overview	4
Foundations.....	4
System Overview	4
The Optimist II Compiler.....	4
Cosmos.....	5
Example	6
Implementation	6
Results	7
Caveats.....	8
Reading Guide	8
1.3. Related Work	10
Smalltalk Systems	10
Smalltalk-80.....	10
Optimized Sequential Smalltalk.....	10
ConcurrentSmalltalk	11
Actor Systems	11
Cantor.....	11
Acore.....	12
J-Machine References	13
Chapter 2. Concurrent Smalltalk.....	14
Introduction	14
Functions	14
Extracting Methods.....	15
Classes	15
Overriding Methods	16
Type Restriction.....	16
The Class Object.....	16
Local Variables	17
Types	17
Concurrency	18
Locks	19
Distributed Objects	20
Macros	21
Chapter 3. The Optimist II Compiler.....	22
Structure	23
Reading Guide	23
3.1. Data Structures	25
Utilities	25
Hcodes	25
Values	27
Types and Classes	28

Multitypes	28
Global Data Structures	29
3.2. Initial Phase	30
Reader	30
Parser	30
Macro Implementation	31
Environments	31
Concurrent Smalltalk Runtime	31
Top-Level Primitives	31
Method-Lambdas	31
Loops	32
Initial Transformations	32
Lambda-Collapsing	34
Efficiency Considerations	35
Example	35
Top-Level Evaluator	37
Interpreter	37
3.3. Optimization	38
Treewalker	38
Calling the Optimizer	38
Guide to Optimizations	38
Preparatory Transformations	40
Lambda Copier and Structural Optimizer	40
Nconcurrently Flattener	40
Continuation Expander	41
Iterative Optimizations	41
Type Specializer	41
Dataflow Optimizer	42
Constant Folder	42
Function Inlining	43
Cleanup Transformations	45
MDP-Specific Transformations	45
Global Expander	45
Addressing Mode Flattener	45
Statement Splitter	45
Built-in Optimizer	46
Instance Variable Target Transformer	46
Grab Introducer	46
Cfuture Parameter Eliminator	46
Enter/Exit Introducer	46
3.4. Code Generation	47
New Hcode Compiler Features	47
New Assembler Features	47
Global Compilation	47
Identifiers	48
IDs	48
Method Tables	48
Data Formats	48
3.5. Conclusion	49
Observations	49
Generality or Simplicity?	49
Future Plans	49
Chapter 4. The Cosmos Operating System.....	51
Design Goals	51
Functionality	51

Chapter 5. Sample Program	85
Initial Phase	86
Optimization Phase	86
Compilation Phase	91
Running Rangesum	93
 Chapter 6. Debugging	 98
Debugging Concurrent Smalltalk Code	98
Debugging MDP Code on MDPSim	98
Debugging MDP Code on a J-Machine	99
Summary	99
 Chapter 7. Performance Measurements	 101
7.1. Derived Times	102
Cosmos Estimates	102
User Program Estimates	103
Analysis	104
Standard Invocations.....	104
Tail-Forwarded Invocations	105
Latency	106
Summary	106
7.2. Measurements	107
Grain Size and Machine Load	107
Comparison with Dataflow	110
Network Load	110
Instruction Frequencies	111
7.3. Conclusion	114
Context Switching Performance	114
Summary	114
 Chapter 8. Future Evolution	 115
8.1. Features	116
Arrays	116
Overriding Primitive Selectors	116
Long Integers	117
Futures	117
Floating Point Numbers	117
True Loops	118
Inline Objects	118
8.2. Resource Management	119
Heap Compaction	119
Fanout Bottlenecks	119
Garbage Collection	120
Load Management	120
Controlling Parallelism	121
Name Spaces	123
8.3. Architectural Considerations	125
Minor Instruction Set Changes.....	125
Fast Context Saves and Restores	125
8.4. Conclusion	126

Chapter 9. Conclusion.....	127
Optimist II.....	127
Cosmos.....	127
Debugging.....	127
Performance Measurements.....	127
Future Work.....	128
Hopes.....	128
Appendix A. Concurrent Smalltalk Reference.....	129
A.1. Introduction.....	129
Goals.....	129
Format.....	130
BNF.....	130
Methods and Functions.....	130
A.2. Syntax.....	131
Tokens.....	131
Identifiers.....	131
Symbols.....	133
Constants.....	133
Comments.....	134
A.3. Programs.....	135
Constant Expressions.....	135
Global Definitions.....	135
A.4. Classes.....	137
Built-in Classes.....	137
Defining New Classes.....	137
Class Inheritance.....	137
Instance Variables.....	138
Reader and Writer Methods.....	138
Class Definition Options.....	138
Inline Classes.....	139
A.5. Methods and Functions.....	140
Introduction.....	140
Formals.....	140
Return Values.....	141
Method and Function Declarations.....	141
The Calling Process.....	141
Scoping of Local Variables.....	142
Functions.....	142
Methods.....	142
A.6. Statements.....	144
Futures and CFutures.....	144
Argument Evaluation.....	145
Application Statement.....	146
Type Assertion.....	146
Variable Bindings.....	147
Multiple Values.....	148
Syntactic Sugar.....	148
Flow of Control.....	149
Loops.....	149
Primitive Control.....	150
Returning Values.....	150
Continuations.....	150

A.7. Built-in Methods and Functions	152
Built-in Classes	152
Built-in Methods	153
Redefining Restricted Selectors	153
A.8. System and Object Operations	158
Objects	156
Copiers	156
Deallocators	156
Class Inquiries	156
A.9. Distributed Objects	158
Group and Constituents	158
Creation	158
Operations	158
A.10. Logical and Arithmetic Operations	160
Comparisons	160
Logical Operations	160
Arithmetic Operations	161
Bitwise Logical Operations	162
A.11. Locks	163
Lock Operations	163
A.12. Strings and Arrays	164
Creating Arrays	164
Operations on Entire Arrays	164
Accessing Arrays	165
A.13. Input and Output	166
Streams	166
Operations on General Streams	166
Reading	166
Writing	166
Atomicity	167
Input and Output Streams	167
Formatting	167
A.14. Macros	168
A.15. Environment	169
Errors	169
Utilities	169
Options	169
Appendix B. Using Optimist II	170
Starting the Compiler	170
Top-Level Commands	170
Utility Commands	170
Viewing Objects	170
Compiling Programs	171
Options	171
Appendix C. Using Cosmos	174
Loading Cosmos	174
Loading User Programs	174
Running Programs	174
Appendix D. MDP Architecture Summary	176
Introduction	176

Processor State	176
Data Types	178
Network Interface	179
Message Transmission	179
Fault Processing	179
Instruction Encoding	180
Instruction Set Summary	181
Appendix E. Optimist II Listing	184
Appendix F. Cosmos Listing	185
Cosmos.i	185
Cosmos.m	190
Runtime.m	218
Bibliography	221
Concurrent Smalltalk Index	224

List of Figures

Figure 1-1. Software Environment Organization.....	5
Figure 1-2. Compiling Fact.....	6
Figure 1-3. Running Fact.....	7
Figure 2-1. A simple Fibonacci program.....	14
Figure 2-2. A simple Fibonacci program as a function.....	14
Figure 2-3. The pair class.....	15
Figure 2-4. Fibonacci program with local variables.....	17
Figure 2-5. Fibonacci program with types.....	18
Figure 2-6. Lock Example.....	19
Figure 2-7. Distributed Object Example.....	20
Figure 2-8. When macro.....	21
Figure 3-1. Optimist II Organization.....	24
Figure 3-2. Exit Flattening Example.....	33
Figure 3-3. Lexical Variable State Machine.....	34
Figure 3-4. Optimizer Organization.....	39
Figure 4-1. Operating System Organization.....	52
Figure 4-2. Concurrent Smalltalk Object Representations.....	55
Figure 4-3. MDP Memory Organization.....	56
Figure 4-4. A Heap Block.....	59
Figure 4-5. XLATE Table Format.....	62
Figure 4-6. BRAT Entry Format.....	62
Figure 4-7. BRAT Table Format.....	63
Figure 4-8. Object ID Formats.....	63
Figure 4-9. Context Format.....	65
Figure 4-10. RestartContext Message.....	68
Figure 4-11. Object XLATE Table and BRAT Entries.....	69
Figure 4-12. Object Migration Messages.....	70
Figure 4-13. Object Migration Protocol.....	71
Figure 4-14. Class Object Format.....	72
Figure 4-15. Object Creation and Disposal Messages.....	73
Figure 4-16. Class/Selector Word Format.....	75
Figure 4-17. Selector Object Format.....	76
Figure 4-18. Method Manager Messages.....	76
Figure 4-19. Application Messages.....	77
Figure 4-20. Function Object Format.....	78
Figure 4-21. Reply Message Format.....	78
Figure 4-22. Closure Format.....	79
Figure 4-23. Distributed Object Group ID.....	80
Figure 4-24. Looking up a Constituent in a Sparse Distributed Object.....	81
Figure 4-25. Looking up a Constituent in a Dense Distributed Object.....	82
Figure 5-1. The Rangesum Program.....	85
Figure 5-2. Rangesum Interactive Session.....	86
Figure 5-3. Rangesum Macroexpansion.....	87
Figure 5-4. Initial Rangesum Hcode.....	87
Figure 5-5. Hcode after Initial Transformations.....	88
Figure 5-6. Locally Optimized Hcode.....	88
Figure 5-7. Hcode after Global Constant Propagation.....	89
Figure 5-8. Optimized Average Hcode.....	89
Figure 5-9. Rangesum with Average Inlined.....	90
Figure 5-10. Rangesum after General Optimizations.....	90
Figure 5-11. Final Hcode.....	91
Figure 5-12. MDPSim Output File.....	93
Figure 5-13. Rangesum Call File.....	94
Figure 5-14. MDPSim Transcript.....	97

Figure 7-1. Function Invocation Latency	105
Figure 7-2. Factorial Program	107
Figure 7-3. Rangesum2 Program	107
Figure 7-4. Sort Program	108
Figure 7-5. Factorial Program used in Dataflow	110
Figure 8-1. A Doubly-Recursive Fibonacci Program	121
Figure 8-2. Progress of a Sequential Computation	122
Figure 8-3. Progress of a Parallel Computation	122
Figure 8-4. Modified Fib Assembly Language Function	124
Figure A-1. Scopes of Identifiers	132
Figure A-2. Hierarchy of built-in classes	153
Figure C-1. An Injected Application Message	174
Figure D-1. The MDP Register Set.	177
Figure D-2. The MDP Data Types.....	178
Figure D-3. The MDP Normal Addressing Modes.....	181
Figure D-4. The MDP Register Oriented Addressing Modes.....	182

List of Tables

Table 3-1. Hcodes	26
Table 3-2. Conditions	27
Table 3-3. Rvalues	27
Table 3-4. Lvalues	28
Table 3-5. Types	28
Table 3-6. Lambda-Collapser Example Results	36
Table 3-7. Permanent Root Objects	38
Table 3-8. Identifier Prefixes	48
Table 4-1. Criticalities	58
Table 4-2. XLATE and BRAT Associations	61
Table 7-1. Selected Cosmos Routine Instruction Counts.....	102
Table 7-2. Selected User Action Instruction Counts	104
Table 7-3. Performance Measurements	109
Table 7-4. Static Instruction Frequencies.....	111
Table 7-5. Dynamic Instruction Frequencies.....	112
Table 7-6. Memory Access Frequencies	113
Table A-1. Predefined Constants.....	133
Table A-2. Built-in Classes	152
Table A-3. Restricted Selectors	154
Table A-4. Identities among Primitive Methods	155
Table B-1. Options.....	172
Table D-1. MDP Faults	180
Table D-2. MDP Instructions	182

Chapter 1. Introduction

Goals

This thesis describes the Concurrent Smalltalk language and its implementation on the Message-Driven Processor. Concurrent Smalltalk, also known as CST, is a concurrent version of the object-oriented programming language Smalltalk [20]. The implementation consists of a global, optimizing compiler and a streamlined operating system for the J-Machine.

This thesis covers quite a broad scope of the implementation of Concurrent Smalltalk, including subjects ranging from issues in parallel programming in general and the design of Concurrent Smalltalk itself to some of the fine points of the design and optimization of the MDP architecture. The goal of the thesis is to demonstrate a working implementation of Concurrent Smalltalk on the Message-Driven processor. Although the implementation is not yet complete, it does provide hooks for all of the advertised functionality of Concurrent Smalltalk and is based on solid ground. Versions of the implementation are running on recently manufactured MDP chips, and I hope that the programs described herein will survive and evolve for the next five years.

Another goal of this thesis was to discover and, whenever possible, fix design flaws in the MDP architecture and language specification so as to make an implementation of Concurrent Smalltalk practical. Several errors in the MDP architecture and Concurrent Smalltalk were found, as well as numerous bugs in the simulation tools used to verify the hardware.

The next section gives a brief overview of the J-Machine hardware and the Concurrent Smalltalk language. It is followed by an outline of the software bridging the gap between Concurrent Smalltalk and the MDP hardware—the Optimist II compiler and the Cosmos operating system. The relationship of this work to others' in fine grain concurrent computation is then described.

Second Edition

This work was originally a Master's thesis completed in May 1989. It has been updated for the state of Optimist II compiler, Cosmos operating system, and MDPSim 7.0 simulator as of the end of May 1991. The Optimist II compiler now produces better code, and several Cosmos routines, especially the CFUT fault handler, have been sped up. Furthermore, Cosmos has been updated for a few minor architectural revisions.

The compiler and operating system have been evolving rapidly in the past few months due to the recent availability of MDP chips. This document does not include these newest changes, which include support for hardware I/O, debugging aids, and workarounds for first-silicon chip bugs, as they have little effect on the ideas in this work. Other members of the Concurrent VLSI Architecture group, including Scott Furman, Rich Lethin, Todd Dampier, Shaun Kaneshiro, John Keen, and Mike Noakes, are now working on CST applications and Cosmos enhancements such as floating-point arithmetic, queue overflow handling, and garbage collection. These will be published in separate documents as they are completed.

1.1. Hardware and Software Architecture

The J-Machine

Million-transistor processors being manufactured today, and soon it will be possible to put several million transistors on one integrated circuit. While memory applications of this technology are clear, it is not obvious how best to use it for computation purposes. One possibility is the architecture of the Message-Driven Processor (MDP), which consists of a 32+4-bit¹ CPU, memory, and a network interface together on one chip. MDPs can be connected directly to each other to form a 65536-processor, message-passing, MIMD, parallel computer, the J-Machine [14]. The network is a three-dimensional mesh fast enough to provide communication between the farthest pair of processors on a 65536-processor J-Machine in a few microseconds—on an unloaded network an 8-word message can be transmitted from one corner of the J-Machine to the other in just 4 microseconds. The processors are optimized for sending and receiving messages; a processor can be working on a message even before the entire message has arrived. The MDP's architecture is unusual in that it provides a very high processing power to memory ratio.

The Message-Driven Processor

The MDP has a register-based architecture and operates on 32-bit data words with 4-bit tags. Tags are essential in efficiently supporting late binding for object-oriented languages such as Concurrent Smalltalk. In addition, tags are necessary for garbage collection and valuable for debugging programs.

The MDP is message-based. In its normal mode of operation, the MDP listens on the network for messages. When it receives a message from the network, it stores the message in a FIFO input message queue and dispatches on the address given in the first word of the message. Messages are used for all communication tasks, including function and method calls, replies, object transfers, and other synchronization facilities.

A detailed but slightly obsolete description of the MDP architecture is in [16]; a updated summary is presented in Appendix D. MDPSim [24] [25] is an instruction level simulator, assembler, and debugger used to run MDP assembly language programs and test the operating system.

Concurrent Smalltalk

Concurrent Smalltalk is the primary language used to program the J-Machine. One of the main goals of designing Concurrent Smalltalk was to take advantage of the J-Machine's unique features. A new software architecture was needed that would efficiently support fine-grain, message-passing computation. Whereas some existing parallel computers have message routing times measured in milliseconds, the routing time for a message sent from one end of even a large J-Machine to another is on the order of several microseconds. Operating system overhead on processing and dispatching that message of more than a few microseconds is not acceptable.

Concurrent Smalltalk introduces concurrency to standard Smalltalk by evaluating arguments to method calls in parallel as well as allowing the computation of the value of a variable to proceed in parallel with the other computations of a method until the variable's value is actually needed. Furthermore, Concurrent Smalltalk adds *distributed objects* to Smalltalk. A distributed object is an object that can process many methods at the same time without any serialization bottlenecks other than those required by the algorithm in use. Although

¹Each word consists of 32 bits of data and a 4-bit tag.

standard objects can also process several methods simultaneously, they can only dispatch on one method at a time¹.

Concurrent Smalltalk is an ideal language for programming the J-Machine because it is easy to parallelize and yields small, fine-grain methods as well as a considerable amount of flexibility in the system software implementation. The methods dealing with a particular class can travel to the data object as opposed to the data traveling to the code. Concurrent Smalltalk also provides excellent facilities for creating data abstractions; the Optimist II compiler amplifies this power by providing global optimizations so performance does not suffer because abstractions are used.

Another advantage of Concurrent Smalltalk is that it is low-level enough to be useful in implementing parts of the J-Machine runtime system, while being at a level high enough that the programmer does not have to worry about the infamous problems of parallel process synchronization and deadlocks. In fact, once the data structures are defined properly, programming in Concurrent Smalltalk feels much like programming in a standard sequential language.

¹This restriction is relaxed for immutable standard objects because they may be copied at the operating system's discretion. Nevertheless, a distributed object can be mutable and still have no synchronization bottlenecks.

1.2. Overview

Foundations

Some of the pieces comprising the Concurrent Smalltalk environment were available before this thesis was done. A primitive compiler was available [21], as were a description of the operating system kernel [38], several descriptions of the language [13] [21] [17], and an MDP assembly language simulator (MDPSim 5.2) [24]. Unfortunately, none of the pieces really fit together—the various versions of the language were inconsistent, the output of the compiler was incompatible with the untested operating system kernel, which itself was written for an obsolete version of the MDP architecture [23].

It became clear that it would be easier to design the language, the compiler, and the operating system from scratch than to try to fit the existing pieces together. Nevertheless, the existing code and ideas were useful as guides to which approaches would likely yield good results and which techniques should be abandoned. I took advantage of this opportunity to extend Concurrent Smalltalk to support several programming styles and add functions, closures, continuations, arrays, nested local variables, and inline classes to produce a language with a compact implementation yet powerful libraries. The new features did not complicate implementation; in fact, by providing a small set of fundamental primitives, the new features often simplified the implementation of existing functionality, a phenomenon noticed in the design of the Scheme language [31] [1].

The contributions of this thesis include:

- A redesign of the Concurrent Smalltalk language.
- *Optimist II*, a new Concurrent Smalltalk compiler and interpreter.
- *Cosmos* (*Concurrent Smalltalk Operating System*), an operating system that supports Concurrent Smalltalk on the MDP.
- Runtime libraries for Concurrent Smalltalk.
- Modifications to MDPSim, the MDP assembler/simulator, to facilitate downloading programs, simplify debugging, and collect performance measurements.
- Modifications to the MDP architecture that make it more suitable for Concurrent Smalltalk.

I am indebted to Scott Wills and Andrew Chien for helping with the redesign of the Concurrent Smalltalk language, and Richard Lethin, John Keen, and Stuart Fiske for helping with the MDP architecture changes. Professor William Dally supervised the project.

System Overview

The Optimist II Compiler

The Optimist II compiler continues in the tradition of the Optimist compiler by compiling Concurrent Smalltalk to assembly code that is as small as possible without sacrificing speed. In addition, Optimist II contains an interactive Concurrent Smalltalk interpreter that is useful for prototyping and debugging Concurrent Smalltalk programs at the source level. Optimist II is also a platform for experimenting with compiler optimizations. Global optimizations such as function inlining and the reduction of method calls to function calls were added and found to be highly successful.

The compiler itself is divided into several phases, which are described in more detail in Chapter 3. It produces an MDPSim command file which can be downloaded into MDPSim and run on a simulated J-Machine.

Cosmos

Cosmos is the operating system used on the Message-Driven Processor to support code output by Optimist II. Many of the ideas in Cosmos are borrowed from JOSS [38] written by Brian Totty—JOSS introduced the concept of a Birth/Residence Address Table (BRAT) and the protocol for migrating object between processors. Nevertheless, Cosmos's code bears little resemblance to JOSS.

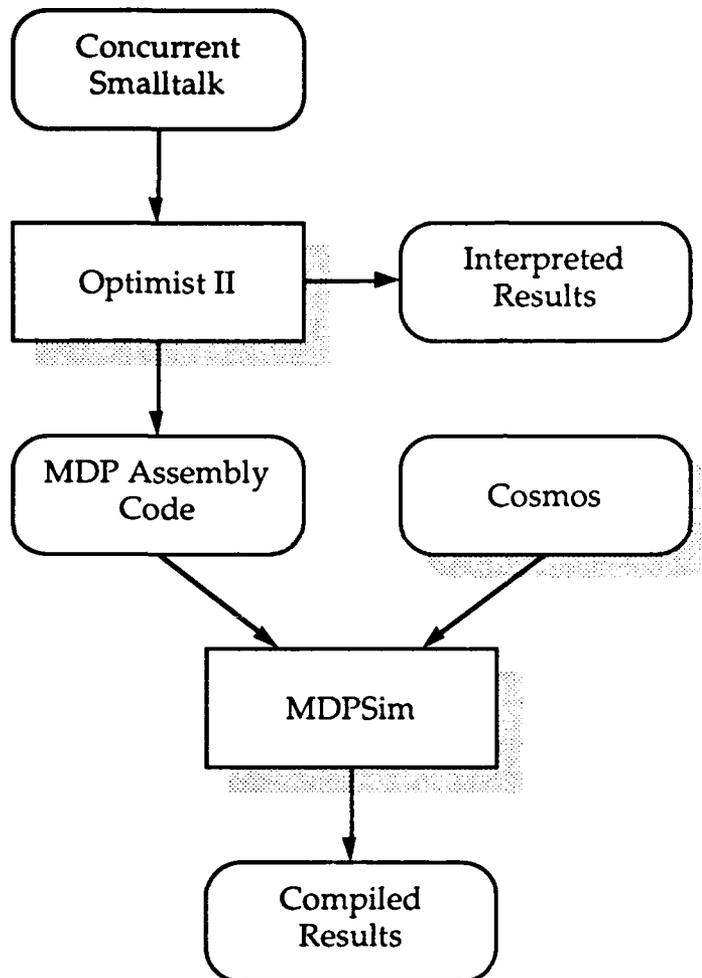


Figure 1-1. Software Environment Organization

A Concurrent Smalltalk program can be either compiled or interpreted by the Optimist II compiler. Interpretation is useful to debug Concurrent Smalltalk programs and interactively experiment with language features. When a Concurrent Smalltalk program is compiled, it is loaded into MDPSim, a J-Machine simulator, together with the Cosmos operating system. MDPSim will then run the program to obtain its results as well as program performance statistics.

The main goals of Cosmos were to make a working operating system, make it as efficient as possible, and make it simple, all subject to the time constraints of a Master's thesis. Those three goals have been achieved to a large extent, in that the operating system does work, and simple programs have been run on it. Unfortunately, controlling a large parallel computer is a difficult task, and Cosmos still falls short in many ways which are described in Chapter 8. In particular, higher-level resource management and load balancing issues are yet to be ade-

quately addressed. Nevertheless, Cosmos is a good start and a platform for experimenting with the more difficult problems.

Example

A very simple example of the use of the system to compile and run a factorial program is listed below. Please refer to chapter 5 for a more detailed example of the transformations in the compiler and Appendices B and C for information about using the compiler and the operating system.

```
CST:(defun fact (n)
      (if (<= n 1)
          1
          (* n (factorial (- n 1)))))
#<Cst-Lambda 5090060 FACT>
CST:(fact 3)
```

```
When interpreting: (FACT 3)
Error: Unbound global FACTORIAL
> Break:
> Type Command-/ to continue, Command-. to abort.
1 > Continuing...Fatal error: Can't apply #<Nil>
> Break:
> Type Command-/ to continue, Command-. to abort.
1 > Continuing...
CST:(defun fact (n)
      (if (<= n 1)
          1
          (* n (fact (- n 1)))))
#<Cst-Lambda 4920924 FACT>
CST:(fact 4)
#<Integer 24>
CST:(compile fact "NewFact.mdp")
```

Figure 1-2. Compiling Fact

The user entered a factorial function, corrected an error in it, tested it on a sample input, and then compiled it into MDP assembly code in the *NewFact.mdp* file. The user's input is shown in bold.

First the user starts the compiler and enters the compiler's interactive mode (see Appendix B) as shown in Figure 1-2. He enters the fact function and runs it only to find an error—fact's recursive call should be to fact, not factorial. The user corrects the error and then uses the compiler's interpreter to successfully compute the factorial of 4.

Afterwards the user compiles fact to MDP assembly code, quits Optimist II, and launches MDPSim, where he loads the object file, and calls fact on 4 to get the correct answer—24 (Figure 1-3). The stats command can then be used to determine some running statistics, such as the frequencies of instructions executed, the amount of parallelism used, and the total time taken to run the program. Starting from a cold start, fact takes 725 steps on a 2x2x1 J-Machine to compute its answer.

Implementation

The Optimist II compiler is written in CLOS [27], the Common Lisp Object System. Except for the use of the LOOP iteration macro [7], Optimist II adheres to standard Common Lisp as specified in [35] and amended in [6] and in the amendments specified by the Common Lisp Cleanup Committee that were available at the time of this writing. The LOOP macro is itself written in standard Common Lisp, so Optimist should run on any machine with a faithful implementation of Common Lisp. A slightly modified version of the 12/7/88 version of Xerox's PCL was used to implement a subset of CLOS before Apple Common Lisp 2.0 became available.

```

MDPSim -x 2 -y 2 -msize 0x1000 ::Cosmos:Cosmos.m NewFact.mdp

Message-Driven Processor Simulator
Version 7.0 Rev B
Accompanies MDP Architecture Document 11B
Written by Waldemar Horwat
Architecture Updates by Brian Totty and Jerry Larivee
UROPs for Bill Dally

4 MDPs present.

@0..3)MESSAGE fact4
Message)MSG:msgApplyFunction|5
Message){fFact}
Message)4
Message)IONODE
Message)0
Message)END
@0..3)inject fact4@3
@0..3)resetstats
@0..3)run
Tick 724 Received priority 0 message:
    OBJ:$801BF004 u=1 f=0 offset=$006EC=Reply length=$0004
    INT:$0000FC00 = 64512
    INT:$00000000 = 0
    INT:$00000018 = 24

@0..3)stats
725 ticks executed.
... More statistics ...

```

Figure 1-3. Running Fact

The user loaded the fact object code and typed a few magic incantations that invoked the fact function on the input 4 (the third word in the injected message). The result 24 (the fourth word in the ejected message) was returned after 725 steps on a 4-node J-Machine. Most of the time was spent distributing the fact code throughout the J-Machine; the second time it only takes 498 steps to compute the answer (some code is still being distributed), the third time takes 289 steps, and afterwards the execution time is about 265 steps.

Optimist II was developed on a Macintosh using Apple Common Lisp 1.2.2 and 2.0 written by Coral Software Corp (now merged with Apple Computer, Inc.). It runs on a 5-megabyte Macintosh II, although 8 megabytes are recommended and at least 16 are needed to run Optimist II and MDPSim simultaneously.

Cosmos is written in MDP assembly language [16]. MDPSim [24] [25] was used as an assembler and simulator for a small J-Machine.

All of the software needed to compile and run Concurrent Smalltalk programs exists on both a Macintosh II platform and on Sun workstations.

Results

The primary result of this work is a demonstration of a working implementation of Concurrent Smalltalk on a J-Machine. In addition, a number of secondary results were obtained. These include the qualitative and quantitative benefits of optimizations in the Optimist II compiler, data on the expected grain size (the number of instructions executed in response to a message), and a number of qualitative observations about the shortcomings of the current system. The results did not always come out as expected. For example, the finding that the grain size is about 60 instructions was surprising; it was expected to be much lower. Code statistics indicate that the MDP will take about 1.9 cycles per instruction, although most instructions execute in 1 cycle; if slow external DRAM is used to hold user programs and data, the MDP could take as many as 3.5 cycles per instruction. Network loading calculations indicate that network congestion will become a concern when the size of the J-Machine exceeds

343 nodes; either a faster network or some means of exploiting locality will be needed for larger J-Machines.

The quantitative results are listed in Chapter 7, while the qualitative ones are in Chapter 8. Chapter 8 may seem a little pessimistic, but many of the current shortcomings listed there would not have been found had this work not been done; furthermore, the current implementation of Cosmos provides a great, highly accurate platform for research into the issues presented there.

Caveats

Due to the availability of only a finite amount of time for writing this thesis, which could potentially involve an infinite amount of work, some features could not be included in the current implementation of Concurrent Smalltalk. The biggest omission is the lack of garbage collection—if enough storage isn't reclaimed, the machine will fail. Garbage collection, although interesting, was omitted to keep this project to a reasonable size—a good garbage collector and load manager would require more effort than is desirable for a Master's thesis.

Full futures were also not implemented. They were omitted from the interpreter in the compiler because simulating them is difficult on a sequential machine in a sequential language (Common Lisp). Futures were omitted from the run-time system because of the considerable amount of work needed to implement all the fault handlers and special cases involved. Nevertheless, almost all Concurrent Smalltalk programs still attain reasonable parallelism through the use of cfutures¹, which are fully operational.

Other features that were not implemented are I/O facilities at both the Optimist II and Cosmos levels and runtime support for local (non-distributed) arrays and floating point numbers. I/O facilities, while useful, do not contribute much to the project and are easy to add later. Local arrays and floating point numbers are supported by the Optimist II compiler but not the runtime system; supporting them at the runtime level will require writing MDP assembly language; no major surprises are expected there.

Some of the optional features of Concurrent Smalltalk were not included due to a lack of time. All class inline declarations are currently ignored; I anticipate that it will be possible to inline objects inside other objects sometime in the future, but that is not a high priority at this time. The omission of class inlining does not change the semantics of Concurrent Smalltalk programs. Function inlining is more useful, and it does work now.

Reading Guide

The remainder of this chapter describes related work in fine-grain concurrent computation. The succeeding chapters delve into various aspects of the system, starting from the top—Chapter 2, **Concurrent Smalltalk**, provides an introduction to the Concurrent Smalltalk language in general. Chapter 3, **The Optimist II Compiler**, describes the Concurrent Smalltalk compiler and interpreter. Chapter 4, **The Cosmos Operating System**, describes the operating system. To avoid overlap, the compiler features documented in [21] are not documented here; thus, it might be helpful to consult [21] when reading Chapter 3.

Chapter 5, **Sample Program**, traces the progress of a sample program from the Concurrent Smalltalk source level down to object code. Chapter 6, **Debugging**, provides some debugging techniques for Concurrent Smalltalk and MDP programs. Chapters 7, **Performance Measurements**, and 8, **Future Evolution**, present the results of this work. Chapter 7 contains quantitative measurements of the performance of Cosmos and the compiled code, while Chapter 8 describes some of the less tangible, qualitative shortcomings of the current system and ideas for correcting them. Chapter 9, **Conclusion**, concludes the main body of the thesis.

¹A cfuture, also called a context future, is a local future which cannot be passed outside the function without being touched (i.e. replaced by its value).

The appendices parallel the main chapters with more detailed information. Appendix A, **Concurrent Smalltalk Reference**, is the most important, for it contains the specification of Concurrent Smalltalk. Appendix B, **Using Optimist II**, provides a detailed description of the Optimist II features not listed in Appendix A. Similarly, Appendix C, **Using Cosmos**, is a guide to running Cosmos on MDPSim; the latest MDPSim reference manual [25] should also be consulted when running Cosmos. Appendix D, **MDP Architecture Summary**, summarizes the current version of the MDP architecture. Finally, Appendix F, **Cosmos Listing**, contains a listing of the entire operating system.

Since this thesis also serves as a reference manual for Concurrent Smalltalk, Chapter 2 and Appendices A and B have been indexed. The index appears at the end of the thesis.

1.3. Related Work

The ideas of optimizing Smalltalk and running object-oriented software on concurrent, fine-grain systems are not new, but they have not been integrated previously to the extent found on the J-Machine. While most of the efforts concentrated on either optimizing Smalltalk for conventional computers or developing radically new programming methodologies, Concurrent Smalltalk presents a somewhat conventional Smalltalk environment to the programmer (with a few new features such as futures and distributed objects), which is at the same time efficiently implemented on a fine-grain parallel computer.

A major contribution of this work is the actual optimized implementation of Concurrent Smalltalk on an assembly language architecture. While theoretical studies and simulations in higher-level languages can yield asymptotic and qualitative results, an implementation yields the constant factors determining a system's performance. These performance measurements are an important part of this work, as they indicate the relative costs of the primitive operations and can be used to gauge the true performance of a concurrent computer.

Smalltalk Systems

Smalltalk-80

Early Smalltalk-80 optimization efforts such as [18] concentrated on optimizing Smalltalk within the constraints of the byte code interpreter. In addition, the work was limited by the Smalltalk-80 constraints of making contexts and methods program-visible data structures, which required some effort to convert between the optimized and standardized versions of the structures. Several context optimizations are also presented in [18], including determining which contexts which can be referred to as first-class data objects and which contexts can be pointed to by blocks. Most contexts do not fall into either category, and they can be placed on the stack. Such optimizations are now also commonly done in Lisp compilers [36].

Whereas early Smalltalk-80 implementations were constrained to compatibility with byte codes and were run on stack machines, Concurrent Smalltalk is bound by neither constraint. The formats of contexts and method code are not defined in the language, and there are no portable means to store a pointer to a context in a programmer-visible variable. Thus, Optimist II and Cosmos can use the most efficient format for a context or even several different formats if they so desire. Furthermore, the MDP is not a stack-based machine, so there are no clear advantages to determining which contexts will be live for a long time. Also, contexts are fully self-contained, so a closure cannot refer to a context. Finally, several techniques are used to optimize closures. As will be seen in Chapter 3, when a closure is created, either the lexical variables are copied into the closure, or a common object is made to which both the context and the closure refer.

Optimized Sequential Smalltalk

A few years later it became clear that global analysis and optimization were necessary to optimize Smalltalk programs further. Optimizing Smalltalk well required an ability to convert method dispatches into more efficient function calls, which led rise to several type systems for Smalltalk [5] [26]. When a type system could be applied to a Smalltalk program, the compiler could optimize it by a factor of 5 to 10 over interpreted Smalltalk. The main compiler optimizations of TS [26] are similar to those of Optimist II: Both TS and Optimist II can convert a message send into a case statement of procedure calls, substitute functions inline, and optimize tail recursion. In addition, TS can beta-reduce blocks, which Optimist II currently cannot do. On the other hand, Optimist II contains a number of other powerful dataflow optimizations (see Chapter 3 and [21]) commonly found in C compilers, which make its assembly language output close to optimal. Moreover, Optimist II can evaluate large constant expressions at compile time, and it can infer types of variables, allowing it to produce

good code even though type declarations in Concurrent Smalltalk are completely optional. TS, on the other hand, has difficulties combining typed code with untyped code.

The MDP hardware also plays an important role in making Optimist II efficient. By providing tags and checking them on primitive operations, the MDP architecture frees Optimist II from the difficult and often unrewarding process of analyzing programs trying to determine information such as whether an integer variable could contain a large-integer (an integer which does not fit into a single 32-bit word) or whether the arguments to + are known to be numbers. Although this information is generally difficult to determine, in most cases integers are small and the arguments to arithmetic primitives are usually numbers, so hardware tag-checking is the right approach to this problem. Thanks to the MDP hardware, even if Optimist II cannot determine the type of some expression, performance does not suffer too much.

CONCURRENTSMALLTALK

A recent language close to Concurrent Smalltalk and having an almost identical name is CONCURRENTSMALLTALK [39] [40] independently developed by Yasuhiko Yokote and Mario Tokoro. CONCURRENTSMALLTALK shares with Concurrent Smalltalk the cfuture facility (called a CBox in CONCURRENTSMALLTALK) and the ability to process messages asynchronously. In addition, CONCURRENTSMALLTALK defines atomic objects, which Concurrent Smalltalk does not have but can easily emulate using locks. On the other hand, Concurrent Smalltalk includes distributed objects, which CONCURRENTSMALLTALK does not provide. Furthermore, the implementation of Concurrent Smalltalk is more optimized. Whereas CONCURRENTSMALLTALK is implemented as a byte code interpreter, Concurrent Smalltalk compiles to assembly language.

The two languages have somewhat different flavors. CONCURRENTSMALLTALK is very close to Smalltalk-80, and most of the concurrent features are add-ons that have to be explicitly requested by the programmer. Concurrent Smalltalk makes concurrency the default, and the programmer has to explicitly request sequential processing if he wants it. At the same time, the MDP hardware assists Concurrent Smalltalk by making the use of concurrency very cheap. For example, a hardware tag is provided that implements cfutures in Concurrent Smalltalk using much less overhead than cboxes in CONCURRENTSMALLTALK.

In [40] several changes to the original CONCURRENTSMALLTALK are discussed. Blocks are treated differently depending on whether they were created by atomic objects' contexts or not. Concurrent Smalltalk's model of only having one kind of object and using locks where necessary to make atomic transactions does not lead to these difficulties. Finally, secretary objects were introduced to CONCURRENTSMALLTALK to keep track of which threads are waiting for a resource. An equivalent facility is used internally in locks in Concurrent Smalltalk.

Actor Systems

Another recent development in object oriented programming was the rise in actor systems [2]. An actor system is a programming paradigm in which simple self-contained entities called actors communicate with each other to run a program. Much of the program's content is held in the interconnections among the actors. From the implementation standpoint, Concurrent Smalltalk shares many of the ideas with actor systems, but the language itself is not designed exclusively as an actor language. Instead, Concurrent Smalltalk is as a language closer to Smalltalk and Lisp, but it is possible to write actor-like programs in Concurrent Smalltalk without too much trouble.

Cantor

Cantor [4] is both a programming language and a formalism for reasoning about the problems that arise in fine-grain, message-passing parallel computers. In Cantor each object (the Cantor equivalent of a Concurrent Smalltalk context) can only perform a bounded amount of computation on receiving a message, and that computation is atomic. Also, messages sent

from one object to another are guaranteed to arrive in the original order. Concurrent Smalltalk is similar to Cantor at the implementation level—when a message is sent to a context, it performs a bounded amount of computation¹, perhaps sends a few more messages, and then either suspends or waits for the next message. The state of a computation is composed mostly of idle objects and messages traveling between objects, with only a few objects executing. Hence, at a superficial level, a Concurrent Smalltalk object code program is a Cantor program. Nevertheless, the Concurrent Smalltalk object code program is more complicated because it might fault while performing the computation of the next state. One can view this possibility as either computation being non-atomic or treating faults as if they were message sends and suspends, preserving the Cantor model. Another distinction is that Concurrent Smalltalk does not guarantee that messages between a pair of objects will arrive in the order in which they were sent.

Probably the best relationship between Concurrent Smalltalk and Cantor is that Concurrent Smalltalk is a high-level language that compiles to Cantor-like object code. At the source level, Concurrent Smalltalk frees the programmer from the myriad of error-prone synchronization details found in Cantor. Concurrent Smalltalk encapsulates the Cantor concept of future flow into a few easy-to-use primitives such as `touch` and `nconcurrently`. At the same time, Concurrent Smalltalk presents the appearance of global and nested data structures (such as lexical scoping of local variables) which are compiled into interacting objects.

Nevertheless, Cantor is a good theoretical model for computation on the J-Machine. For example, the load balancing and management results in [4] are expected to also apply to the J-Machine. However, the J-Machine can also suffer from problems not discussed in [4], such as having too much parallelism. Some of the load balancing issues are presented in Chapter 8.

Acore

Acore [30], an “actor core language,” is another recent actor language. Like Cantor, it provides an environment in which a computation is done by interacting actors with limited abilities; however, actors in Acore can compute arbitrary functions to determine state, and Acore has a notion of a transaction (a message send and a reply), which greatly simplifies programming.

Acore and Concurrent Smalltalk are similar in many ways. Both languages implement message sends, replies, concurrent evaluation of subexpressions, local variables, static scoping, and instance objects (called actors in Acore). However, there are also a few differences. Due to its Smalltalk-80 heritage, Concurrent Smalltalk permits local variables to be altered, while Acore does not; both languages allow mutation of instance variables. In addition, Acore implements a sponsorship mechanism for higher-order control of the course of a computation and a complaint mechanism for handling exceptions. It remains to be seen whether these mechanisms will be necessary in Concurrent Smalltalk².

Acore is compiled into Pract, which is a form of an actor assembly language, whereas Concurrent Smalltalk is compiled into MDP assembly language. As a result of this difference, some actions which are cheap in one language are expensive in the other, which affects the language design. Actor creation is very cheap in Acore, while instance object creation, moderately expensive in Concurrent Smalltalk, is avoided whenever possible. On the other hand, futures are fairly expensive in Acore, while they are very cheap in Concurrent Smalltalk; thus, Concurrent Smalltalk creates a future (or a cheaper cfuture) as a result of every non-primitive function call, achieving maximum concurrency within a method in most cases. Acore, on the other hand, often has to do a relatively expensive join operation. For the same

¹As will be discussed in Chapters 5 and 10, the amount of computation done by a Concurrent Smalltalk process on receiving a message truly is bounded, but it is done for a more prosaic reason than keeping a clean model—user Concurrent Smalltalk methods are not allowed to loop without a message send somewhere to break the loop to prevent the incoming message queues on an MDP from overflowing if the loop lasts for a long time. Also, long, indivisible loops would degrade latency for other messages that are waiting in an MDP's incoming message queue.

reason, futures are transparent in Concurrent Smalltalk, while they are programmer-visible in Acore¹.

The two languages use the same mechanism for calling messages. When a Concurrent Smalltalk process or an Acore actor makes a function or method call, it passes a continuation to which results should be sent. The continuation includes both a process and a slot within that process in which the result should be stored.

J-Machine References

[13] and [14] are good descriptions of the philosophy of the J-Machine project and the early Concurrent Smalltalk language; [15] is a recent status report on the MDP from the hardware perspective. [22] describes some of the experiences gained from designing the previous version of Concurrent Smalltalk and implementing the first-generation Optimist compiler. [10] contains a nontrivial program written in an older dialect of Concurrent Smalltalk. [8] and [9] describe *Concurrent Aggregates*, a higher-level language than Concurrent Smalltalk for programming the J-Machine. [33] and [34] describe a parallel project to implement dataflow on the J-Machine. Finally, [41] and [42] analyze the desirability of supporting the more common existing parallel programming paradigms on the J-Machine.

²A complaint mechanism could be built on top of Concurrent Smalltalk by using the multiple-value return feature—one of the values could denote a continuation to which exceptions should be routed. Acore uses a similar implementation to handle exceptions.

¹Nevertheless, a language that hides futures could be built on top of Acore.

Chapter 2. Concurrent Smalltalk

Introduction

A Concurrent Smalltalk program is a sequence of top-level definitions. Figure 2-1 shows a sample program that calculates Fibonacci numbers using double recursion.

```
(Defmethod fib Integer ()
  (if (<= self 2)
    1
    (+ (fib (- self 1)) (fib (- self 2))))))
```

Figure 2-1. A simple Fibonacci program

This program calculates Fibonacci numbers using double recursion. Although it does not use the most efficient algorithm to calculate Fibonacci numbers, it does illustrate Concurrent Smalltalk's implicit concurrency.

The program is a single *method* associated with the *selector* `fib` and *class* `integer`. The fact that the method takes no arguments other than the integer receiver is indicated by the empty list, `()`, on the first line. The following three lines contain the body of the method. `Self` represents the *receiver object*, which is the number to which `fib` was applied. The `if` statement checks whether that number is less than or equal to 2. If so, `fib` returns 1. Otherwise, `fib` returns the sum of `(fib (- self 1))` and `(fib (- self 2))`, which are computed concurrently. This concurrent evaluation of arguments is one of the important differences between Concurrent Smalltalk and sequential Smalltalk.

`Fib` can be invoked by calling it on an integer (the receiver object):

```
(fib 30)
```

`Fib` would then calculate and return the answer 832040. If `fib` had any more arguments, they would be included after the receiver object, as in:

```
(fib 30 x y z)
```

Functions

The Fibonacci program was defined as a method. It is also possible to define it as a function, as in Figure 2-2. A function is a method not associated with any class or selector. Although in this example methods and functions are equivalent, in other cases, such as in iterators, functions may be more useful than methods.

```
(Defun ffib (n)
  (if (<= n 2)
    1
    (+ (ffib (- n 1)) (ffib (- n 2)))))
```

Figure 2-2. A simple Fibonacci program as a function

Functions have no receiver object, so the parameter `n` has to be specified explicitly.

The syntax for a method and a function call is the same, so `ffib` would also be called by:

```
(ffib 30)
```

The meaning of applying `ffib` to arguments (30 in this case) depends on whether `ffib` is a selector or a function. If `ffib` were a selector, a method lookup would be done to determine the class of the first argument and then call the method corresponding to the selector and that class, while if `ffib` is a function, it is called directly.

Extracting Methods

A manual method lookup can be done using the `method` primitive. `Method` takes two parameters, a selector and a class, and returns a function which performs the same action as the method. For example, the method shown in Figure 2-1 can be extracted using

```
(method fib integer)
```

The result behaves just like the `ffib` function in Figure 2-2. It can be called using

```
((method fib integer) 30)
```

A method extracted in this way does not have to be a direct method of the class; it can be an inherited method.

Classes

A Concurrent Smalltalk class is a type; the two words are used interchangeably in the language definition¹. A few built-in classes are predefined; these include symbols, booleans, integers, floating point numbers, characters, functions, and other classes. A complete list is given in table A-2. All classes are subclasses of the class `object`.

The `defclass` primitive can be used to add user-defined classes. A class definition consists of a list of superclasses and zero or more new instance variables. Each instance object of that class contains those instance variables. The user may also define a number of methods for that class. A simple class that implements Lisp-like lists is shown in Figure 2-3.

```
(Defclass pair (object) car cdr)

; (Defmethod car pair () car)
; (Defmethod cdr pair () cdr)
; (Defmethod get-car pair () car)
; (Defmethod get-cdr pair () cdr)
; (Defmethod put-car pair (value):pair (set car value) self)
; (Defmethod put-cdr pair (value):pair (set cdr value) self)

(Defun cons (first second):pair
  (put-car-cdr (new pair) first second))

(Defmethod put-car-cdr pair (first second):pair
  (cset car first)
  (cset cdr second)
  self)
```

Figure 2-3. The `pair` class

The six methods that are commented out by semicolons are defined automatically by `defclass` (in addition to a few others described in Section A.4). `car` and `get-car` do the same thing; both are defined because `car` is more convenient, but it cannot be used in the body of a method of class `pair` because static scoping shadows the method `car` by the instance variable `car`.

The `:pair` constructs define the result types of the methods. They are unnecessary, but they do improve efficiency and allow rudimentary type checking.

The class `pair` is defined on the first line of Figure 2-3. The `defclass` primitive specifies the class name (`pair`), the superclasses (`(object)`), and the *instance variables* (`car` and `cdr`).

Whenever a class `c` is defined, a *class predicate* and *reader* and *writer* methods are defined automatically, as well other, less-used methods described in Section A.4. The *class predicate* is a function named `c?` that accepts one argument `a` and returns `true` if `a` is a member of class `c` (or one of its subclasses) and `false` otherwise. Also, for each instance variable `x` of `c`,

¹Nonetheless, the words *type* and *class* have slightly different meanings in the discussion of the compiler in Chapter 3.

the methods `x`, `get-x`, and `put-x` are defined. The first two methods take an instance object `o` as an argument and return the value of `x` in `o`, while `put-x` takes two arguments, an instance object `o` and a new value `v` of `x`, and assigns `v` to `x` in `o`. The methods `x` and `get-x` are known as *reader methods*, while `put-x` is called a *writer method*. The writer methods return `o`, the object to which the value is written.

After a class is defined, additional methods may be defined for it. In the above example, a method `put-car-cdr` is defined for the class `pair`. `put-car-cdr` sets the value of a pair's `car` and `cdr` variables and returns the pair. Inside a method, the receiver's instance variables can be accessed by their names.

Overriding Methods

Consider a class `c2` which is a subclass of `c1`. When a class `c2` defines a method `m2` with the same selector `s` as a method `m1` of `c1`, the class `c2` is said to be *overriding* the method `m1`. When selector `s` is applied to an object of class `c2` or one of its descendants, method `m2` will be used instead of `m1`.

Nevertheless, sometimes it is desirable to call `m1` on an object of class `c2`. For example, method `m2` might want to call the method it is overriding. An overridden method `m1` can be called by performing a manual method lookup using the form `(method s c1)`. The resulting method can be called normally.

Type Restriction

The type of an overriding method must be a subtype of the type of the overridden method. For instance, in the above example the type of `m2` must be a subtype of the type of `m1`. This means that both methods must have the same number of arguments, the types of the arguments of the overriding method must be supertypes (superclasses) of the types of the arguments of the overridden method, and the result type of the overriding method must be a subtype (subclass) of the result type of the overridden method. If any argument of the overridden method is declared inline or using any other declaration, either explicitly or by default, the corresponding argument of the overriding method must have the same type and declarations. The results of violating the above rules are undefined. The compiler may issue errors if the above rule is violated, but it is not guaranteed to do so.

The above restrictions apply only to methods being overridden. There are no restrictions on methods with the same name declared for disjoint classes (i.e. classes which are not subclasses of each other).

The Class Object

Methods of class `object` are very similar to functions. There are two main differences between functions and methods of class `object`:

- A method of class `object` can be overridden by a method of a more specific class. For example, if `cons` in Figure 2-3 is defined as a function, no other function or method may be called `cons`. On the other hand, if it is defined as a method of class `object`, it may be overridden by a method `cons` defined for integers. However, a method may not be overridden by a function.
- A function that takes no parameters can be defined, while a method must always take at least one parameter—the instance object.

In the interest of code maintenance and readability, it is recommended that functions be used in cases when overriding makes no sense; parameter functions to iterators fall into this category. On the other hand, if overriding a function might be desirable, that function should be defined as a method of `object`. It is not clear whether overriding `cons` (Figure 2-3)

would be useful, so it might be defined either as a function or a method, depending on one's taste.

Local Variables

A method or a function can declare local variables using the `clet` or `let` statements or their derivatives. For example, the function `fib` from Figure 2-1 could be rewritten using two local variables as in Figure 2-4.

```
(Defmethod Integer lfib ()
  (if (<= self 2)
    1
    (clet
      ((a (lfib (- self 1)))
       (b (lfib (- self 2))))
      (+ a b)))
```

Figure 2-4. Fibonacci program with local variables

The above program is equivalent to the one in Figure 2-1 and actually compiles into the same code.

Local variables declared with a `clet` or a `let` statement have a scope which is the body of the `clet` or `let` statement (except for the bindings themselves). `Clet` and `let` statements can be nested. Local variables can be altered using a `cset` or a `set` statement; the difference between the two will be explained in the **Concurrency** section below.

Types

The types (i.e. classes) of various values can be declared explicitly. Such declarations serve three purposes:

- Types allow the compiler to generate faster code by allowing it to perform operations such as method lookup at compile time.
- The compiler can perform type checking to find simple errors such as passing a value of one type to a function that is expecting a value of a different type.
- Declaring types of function parameters and results serves to document the code.

For the purposes of type inclusion, a type is its own supertype and subtype.

Due to the common use of generic types, the compiler's type checking is necessarily limited. In particular, when an expression of type `t1` is assigned to a variable of type `t2` or passed as a parameter to a function that expects type `t2`, the compiler usually will give an error or a warning if `t1` is not `t2`, `t1` is not a superclass of `t2`, and `t2` is not a superclass of `t1`. This does not mean, however, that the semantics of function parameter and return type declarations are any different from their standard interpretations—when a function parameter is declared type `t`, every value passed as that parameter must be a member of type `t`, and when a function result is declared type `t`, the function must return a value that is a member of type `t` as that result—the only difficulty is that the compiler is not able to do full type checking, so it usually follows the rules outlined above.

For example, `integer` and `boolean` are both subclasses of the `object` and `magnitude` classes (see Figure A-2), but they are otherwise unrelated to each other. Thus an `integer` can be passed to a function that expects an `object`, an `object` can be passed to a function that expects an `integer`, but a `boolean` cannot be passed to a function that expects an `integer`. The second possibility, passing a more general type to a function that expects a less general one, is included to handle the common case of extracting values from general storage class. One could, for example, keep a pair of integers and desire to add the pair's `car` and `cdr` together. Since a `pair` is a generic data structure, it can contain values of type `object`;

a compiler has no simple way of knowing at compile time that the pair will contain integers, so the best it can deduce is that the pair's car and cdr are objects.

Types can be declared as follows:

- To specify the type of a local or an instance variable, follow the variable name with a colon and its type. Several locals can be declared using the same type by separating their names with commas.
- To specify the type of a function or method formal, follow the formal name with a colon and its type. Several formals can be declared using the same type by separating their names with commas.
- To specify the result type of a function or method, follow the list of formals with a colon and the result type¹.
- A type of an intermediate result can be specified using a type-assertion statement².

The three kinds of declarations are illustrated in Figure 2-5, yet another copy of the Fibonacci program. All untyped variables, parameters, and functions and methods are typed object by default.

```
(Defun tfib (n:integer):integer
  (if (<= n 2)
      1
      (clet
        ((a:integer (tfib (- self 1)))
         (b:integer (tfib (- self 2))))
        (+ a b)))
```

Figure 2-5. Fibonacci program with types

There are three type declarations here. In order, they are a declaration of the parameter type of *n*, a declaration of *tfib*'s result type, and declarations of the types of the local variables *a* and *b*.

Concurrency

Concurrency is expressed in Concurrent Smalltalk in several ways:

- **Concurrent argument evaluation.** In


```
(+ (big-computation 3) (time-sink 738))
```

 the expressions *big-computation* and *time-sink* can be evaluated in parallel.
- Expressions in *concurrently* statements may be evaluated concurrently. The expressions in *parallel* statements are always evaluated concurrently.
- The variable bindings in *clet* and *let* statements can also be evaluated concurrently. For example, the expressions *big-computation* and *time-sink* can be evaluated concurrently in


```
(cset a (big-computation 3))
(cset b (time-sink 738))
(+ a b)
```

 as well as in


```
(let ((a (big-computation 3))
       (b (time-sink 738)))
      (+ a b))
```
- The computations in assignments using *cset* and in function calls whose result values are unused can be done concurrently with neighboring statements.

¹See also **return values** in section A.5 for a description of specifying types of multiple results.

- The computations done for *futures* are always evaluated in parallel.

The action of a *cset* can be thought of as storing a promise (known as a *cfuture*) to calculate the value of a variable. For example, after

```
(cset a (big-computation 3))
```

is executed, *a* will contain either the value of *(big-computation 3)* or a *cfuture* promising to deliver that value when it is needed. If *a* contains a *cfuture*, *(big-computation 3)* is evaluated in parallel by a different task. At the same time, execution of the method can proceed and the method can perform another time-consuming task. It will not have to wait for *(big-computation 3)* to complete until the value of *a* is needed.

Sometimes it is desirable to explicitly wait until the value of an expression is available before continuing. This is called either *touching* or *forcing* the expression. Touching or forcing an expression that evaluates to a normal value does nothing. Touching or forcing an expression that evaluates to a *cfuture* causes evaluation to wait until the value of the *cfuture* is available. Finally, touching an expression that evaluates to a future does nothing, while forcing it causes evaluation to wait until the value of the future is available. The resulting value is then touched or forced again until the touch or force operation does not change it.

An expression can be touched using the *touch* statement and forced using the *force* statement. Since built-in methods and functions usually touch or force their arguments, touching and forcing are rarely done explicitly.

The reference manual in Appendix A defines more precise semantics for what expressions may or may not be evaluated in parallel.

Locks

```
(defclass resource (object)
  1:lock
  ... other fields)

(defmethod init resource ()
  (cset 1 (new-simple-lock)) ;Creates an initially available lock
  ... other initialization code)

(defun new-resource ()
  (init (new resource)))

(defmethod access resource (parameters
  (acquire 1)
  ... code to perform the access using parameters ...
  (release 1))

(defmethod access2 resource (parameters
  (with-locks (1)
  ... code to perform the access using parameters ...))
```

Figure 2-6. Lock Example

This example defines a class *resource* that contains a lock. Every call to *access* acquires the lock when it starts and releases it when done, so the code in the middle of the *access* method cannot be interrupted by another *access* method. The *with-locks* macro is a convenient shorthand for acquiring and releasing locks; the *access* method could have been rewritten as *access2*.

Locks are used to synchronize computation by Concurrent Smalltalk programs. Locks are especially useful around critical sections of code where only one process may access a resource; a process that wants the resource acquires a lock before accessing the resource and releases it when it is done. Two variants of locks are provided. *Simple-locks* are fast locks which, however, perform poorly when many processes are waiting for a resource; *simple-*

²See section A.6.

locks should be used in situations in which the probability of contention for a resource is small. Queueing-locks are slower locks designed to handle a large amount of contention.

As an example of the use of locks, suppose one wants to restrict the use of a resource so that only one process can access it at a time. To accomplish this exclusion, a lock can be associated with the resource, in which case every process should acquire the lock before using the resource and release it when done. Figure 2-6 shows sample code used to access the resource.

Distributed Objects

```
(defclass distarray (distobj)
  value)

(defun new-distarray (size:integer)
  (new distarray size))

(defmethod get distarray (index:integer)
  (get-value (co group index)))

(defmethod put distarray (index:integer new-value)
  (set (get-value (co group index)) new-value))

(defmethod size distarray ()
  (logical-limit self))
```

Figure 2-7. Distributed Object Example

This example defines a class `distarray` used for distributed arrays. The `get` method returns the element at position `index` in the array; since each constituent contains only one element of the array, the `get` method returns the value in the constituent specified by the given index. Similarly, the `put` method routes the message to the constituent specified by `index`, where it stores `new-value`. The `size` method simply returns the array's size.

Whereas standard objects serialize messages sent to them¹, distributed objects can accept and process many messages at a time. A distributed object is comprised of an array of constituent objects and a common, group name. When a message is sent to the group name, the operating system routes it to a constituent of its choosing. The constituent can then process the message or send it to another constituent; constituents know how to address each other. The `co` primitive is used to find a particular constituent of a distributed object, while the `group` instance variable can be read to determine the group name of a distributed object given one of its constituents.

For example, a large array might be implemented as a distributed object. When a `get` message is sent to the array to read a value of a particular element, the message is routed to one of the constituents. That constituent examines the given index and forwards the message to the constituent containing the element, which reads and returns the value.

Figure 2-7 shows a simple example of the use of distributed objects to create a distributed array. Each constituent contains only one element of the array to keep this example short; a better implementation would use a `simple-array` at each constituent to reduce the number of constituents needed.

The advantages of using a `distarray` class like the one in Figure 2-7 is that many accesses can be made to the array simultaneously; they do not have to pass through a common bottleneck to access the array. In addition, as will be clarified in Section 3.3, the `get` and `put` methods do not access any instance variables of `distarray` themselves, so they could be inlined wherever they are called²; thus, reading or writing the `distarray` in Figure 2-7 could

¹Except for a few special cases such as immutable objects and messages which do not need to access an object's data to execute, only one message may be processing on a standard object at a time.

²The compiler's handling of `group` would have to change a little to permit this optimization; the compiler currently treats `group` solely as an instance variable, but there is no intrinsic reason why the compiler could not provide a by-pass path that checks whether a method was called on a group ID (as opposed to a constituent ID) and just uses the

involve only two message sends, which is no less efficient than reading or writing a simple-array.

Macros

Concurrent Smalltalk provides a macro facility which can be used to extend the language. A macro consists of a pattern and a replacement. The pattern can contain variables or keywords. If it matches with an expression, that expression is replaced by the replacement, which can be either another pattern or a Common Lisp function¹. Much of the language itself has been implemented in terms of macros. Figure 2-8 contains a sample macro which defines a when form that is the equivalent of a Common Lisp when.

```
(defmacro (when ?test . ?body)
  (if ?test
      (begin . ?body)))
```

Figure 2-8. When macro

The when form defined by this macro takes a test and a number of statements comprising the body. If the test is true, the statements are executed one after another, as in begin. If the test is false, when returns nil. This macro takes advantage of the fact that if returns nil if there is no else-clause and the condition is false. The Lisp dot notation is used to indicate that the body forms the rest of the given list.

group ID if it was provided instead of always using the group instance variable. When this optimization is implemented, distributed arrays such as the one above will be as efficient as simple arrays.

¹Concurrent Smalltalk functions may be added as replacements later, when the entire compiler and development system is rewritten in Concurrent Smalltalk.

Chapter 3. The Optimist II Compiler

Optimist II is an optimizing compiler for the Concurrent Smalltalk language described in Appendix A. The compiler generates assembly language code for the Message-Driven Processor.

Optimist II is based on the Optimist compiler described in [21]. Optimist included many standard optimizations such as register variable assignment, dataflow analysis, copy propagation, and dead code elimination [3] [43] that are used in compilers for conventional processors. In addition, Optimist included fork and join mergers that try to merge similar (not necessarily identical) statements on both sides of conditionals, a powerful move eliminator, and numerous code generator optimizations to accommodate various idiosyncrasies of the MDP.

Optimist II is a substantial improvement over the Optimist compiler. While Optimist supported only a small subset of an early Concurrent Smalltalk language, Optimist II implements almost the entire new Concurrent Smalltalk language. Some language features supported by Optimist II that were not present in the original Optimist include:

- Method lookup (Optimist could compile method code but could not associate a method with a selector)
- Global variables
- Class and variable declarations
- Macros
- Lambdas and closures
- Multiple inheritance of classes
- Distributed objects
- Multiple return values
- Nonlocal exits
- Functions
- Methods referencing more than one object at a time
- Synchronization primitives
- Arrays
- Methods overriding primitive selectors such as +
- Compile-time evaluation of expressions

Furthermore, Optimist II contains an interactive language environment, including a Concurrent Smalltalk interpreter and facilities to view code in various stages of compilation. Optimist II gives helpful warnings and errors when it encounters questionable language constructs. It also includes entire new categories of optimization, including type inference and global program optimizations. Finally, Optimist II's code generator has been updated to conform to and optimize for MDP Architecture version 11B [16]¹ instead of Optimist's Architecture 10 [23].

¹This reference is to MDP Architecture version 11. Version 11B has not been published yet.

The only language features listed in Appendix A missing from Optimist II are full futures and I/O facilities. It is expected that they will be added later, when the operating system is updated to support them. In addition, some optional features of the language such as inline objects and first-class continuations have not been implemented, although facilities have been provided that will simplify their implementation in the future.

Structure

Figure 3-1 shows the overall structure of the compiler. Concurrent Smalltalk code is read and parsed by the reader and parser, transformed by the preoptimizer, and saved in the global environment. It can be either interpreted using the global environment or optimized further by the optimizer and then compiled into MDP assembly code by the compiler and assembler. The treewalker controls the compilation process and prevents unused modules and objects from being compiled and assembled.

Reading Guide

The **Data Structures** section introduces the common data structures used in the Optimist II compiler. A few data structures such as digraphs and hcode appear throughout the compiler, and familiarity with them is assumed in the later sections.

The next three sections discuss the three main components of the compiler environment: The **Initial Phase** includes facilities to read Concurrent Smalltalk expressions and compile them into hcode (an intermediate code format), interpret that hcode, and maintain the global Concurrent Smalltalk environment. This phase executes until the user requests a compilation of the program to MDP assembly code, at which time the other two phases are invoked. Most of the optimizations in Optimist II are done in the **Optimization phase**, although a few appropriate optimizations are scattered in the other phases. The **Code Generation phase** compiles the optimized hcode into MDP assembly language and outputs that assembly language, together with immediate objects, class descriptors, and method tables, after performing a few final optimizations. The output of the Code Generation phase can be read directly into MDPSim. The code generator and MDPSim share the task of linking programs. Finally, the **Summary** section summarizes the important ideas in the compiler.

Chapter 5, **Sample Program**, shows the progress of a sample program through various phases of the compiler, and it may be helpful to illustrate some of the optimizations.

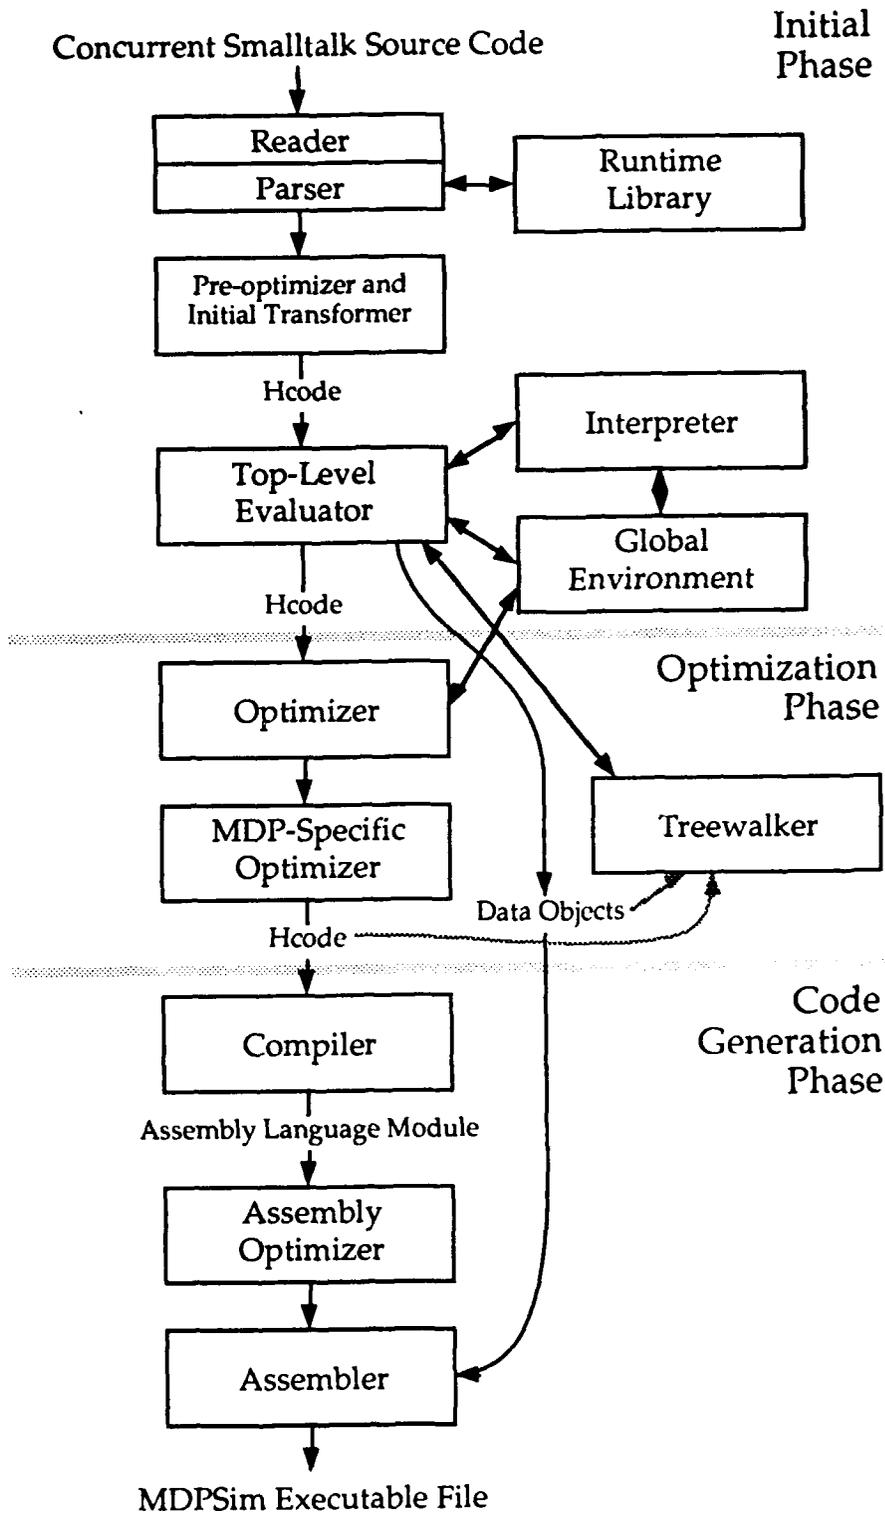


Figure 3-1. Optimist II Organization

Table 3-4. Lvalues

Lvalue	Specializers	Notes
Local	Name, scope, etc.	A local variable.
Global	Name	A global variable.
Continuation-ref	Continuation or Context and Displacement	A reference to a continuation specified either as a continuation rvalue or as a pair of context and displacement rvalues (See Section 3.3).

All rvalues are instances of the rvalue CLOS class, all lvalues are instances of the lvalue CLOS class, and all values are instances of the value and rvalue CLOS classes. CLOS's multiple inheritance is used to define objects that are both rvalues and lvalues or other combinations of the above.

Types and Classes

Table 3-5. Types

Type	Specializers	Notes
Class		Any class is also a type.
Continuation-type	Continuation-type	A type based on the continuation class that represents a continuation that will return a value of the continuation-type type.
Displacement-type	Continuation-type	A type based on the displacement class that represents a displacement field of a continuation that will return a value of the continuation-type type.

A Concurrent Smalltalk class is a Concurrent Smalltalk value that is an instance of the class class. Classes are implemented in Optimist II as instances of the cst-class CLOS class. In addition to itself being a value, a class also represents a set of values. For example, the class `integer` represents the set of all integers, which includes, among others, the values 4 and -17. The class `null` represents the singleton set `{nil}`. The class `class` represents the set of all Concurrent Smalltalk classes, including itself.

In addition to classes, Optimist II includes types which provide finer discrimination than classes for describing sets of values. Types are listed in Table 3-5. Currently a type is either a class or a continuation that returns an object of some type. A type can be always projected to a class; the base-class Lisp generic function performs this conversion. A type that is also a class projects to itself, while a continuation type projects to the class `continuation`. Although a class is always a value, a type is not necessarily a value.

Multitypes

When describing the possible contents of variables, Optimist II uses the concept of a *multitype*. A multitype is a list of zero or more types; a value is a member of a multitype (satisfies that multitype) if it is a member of one of its types. No value satisfies a null multitype, while every value satisfies a multitype that has `object` as one of its types. Routines are provided to calculate unions (least upper bound) and intersections (greatest lower bound) of multitypes and simplify representations of multitypes. Since multitypes are not necessarily closed under those operations, the `lub` and `glb` routines may conservatively enlarge their multitype results.

Global Data Structures

Two atomic environments, the global environment and the class environment, contain most of the state of the Concurrent Smalltalk interpreter. The global environment contains all Concurrent Smalltalk globals, parameters, and constants, while the class environment contains all known Concurrent Smalltalk classes. The global environment is linked to the class environment, so the latter is searched if an identifier is not found in the global environment.

The classes are themselves heavily linked together. Each class object has lists of its immediate superclasses and subclasses and all of its superclasses and subclasses, as well as a metaclass, a description of its instance variables, and sundry options such as whether the class is immutable. To allow typed recursive data structures, an “undefined” class structure is created when a class name is encountered in a program without being defined. An “undefined” class can turn into a normal class when the class is defined; CLOS’s change-class construct is very valuable here. A substantial number of classes have to be updated whenever a new Concurrent Smalltalk class is defined, but compilation speed does not seem to suffer because of this. The heavy linking of classes made defining a bootstrapping subset of Concurrent Smalltalk classes challenging; some CLOS objects had to be created with the wrong classes and then transformed to the right classes. Once the bootstrapping subset of Concurrent Classes was defined, defining the remaining classes on top of it was easy.

A method is associated with both a class and a selector. There is no single method table in Optimist II; instead, whenever a method is added, it is added to the selector’s list of methods hashed by class and the class’s list of methods hashed by selector. Thus, a selector knows all of the methods defined for it, as does a class. Methods are not replicated in these hash tables unless a method is added more than once; instead, the lookup-method function, which returns a method associated with a class and a selector, searches the superclasses when a method is not defined for a selector and a class; an ambiguous selector error is signalled if there is more than one superclass and they are associated with differing methods.

Current settings of the options are also kept in a global data structure. Each option is declared as a dynamic Lisp variable, and a list of all options and their default values is kept in an object. The `#&name` reader macro expands into a reference of the option named `name`.

Concurrent Smalltalk symbols are not accumulated in any data structure; however, when a Lisp symbol is used as a Concurrent Smalltalk symbol, its `cst-symbol` property is set to the Concurrent Smalltalk symbol object to ensure that that object is reused if the symbol is referenced again; otherwise, `(eq 'sym 'sym)` would be false according to the interpreter. Number objects are not reused, so `(eq 13 13)` is false according to the interpreter¹, but `(clet ((x 13)) (eq x x))` is true.

¹Nevertheless, compiled code will currently return true if `eq` is used to compare two equal integers. The action of `eq` on numbers is purposely not defined in Concurrent Smalltalk to allow an implementation of a bignum package.

3.2. Initial Phase

The initial phase of the compiler reads the Concurrent Smalltalk input and converts it into a rough hcode form. Several early transformations have to be done on the resulting hcode before it becomes suitable for optimizations.

The most complicated early transformations create statically scoped functions. The initial phase determines parameter interfaces for lexical variable displays [3] used by closures, and it does a considerable amount of work to pick those interfaces well. Delaying this decision would have made manipulation of functions in that stage very difficult; the advantages of splitting nested functions into components early are that every function is self-contained and completely owns its local variables—no other function can alter or examine the local variables.

Reader

A customized Common Lisp reader is used to read the Concurrent Smalltalk programs. The customizations consist of using a special readtable and reading all Concurrent Smalltalk names into the CST package. The readtable is used to implement the special characters in the Concurrent Smalltalk syntax. Most special characters expand into lists; for example, `! a` expands into `(! a)`. Some character tokens such as `:`, `::`, and `,` (comma) expand into symbols with the same names.

The CST package is used to prevent conflicts between Concurrent Smalltalk symbols and any symbols the compiler or the Common Lisp environment might be using. For instance, `nil` is just the name of a constant (which happens to have the value `'nil`) in Concurrent Smalltalk; `nil` is not confused with the Lisp `nil`, which also represents an empty Lisp list. Since the colon has a special readtable meaning in Concurrent Smalltalk mode, Concurrent Smalltalk symbols are restricted to the CST package.

Read macros have been inserted into both the Common Lisp readtable and the Concurrent Smalltalk one to facilitate easy switching between the two tables. The `#$` macro in standard Lisp input reads the next token in Concurrent Smalltalk mode, while `#^` can be used inside a `#$`-expression to switch back to Lisp mode. In addition, the `#L` macro in Concurrent Smalltalk mode reads a list expression and returns a two-element list with the symbol `lisp` as its first element and the expression read as the second.

Parser

The parser parses the input expressions into a prototypical hcode form. The parser is a recursive descent macro evaluator. Each primitive in Concurrent Smalltalk is implemented as a macro. There are three main kinds of macros: normal macros substitute Concurrent Smalltalk text with other literal Concurrent Smalltalk text as described in Section A.14, non-terminal macros substitute Concurrent Smalltalk text with Concurrent Smalltalk text produced by a Lisp function, and terminal macros read Concurrent Smalltalk text and perform an action such as emitting hcodes. Furthermore, macros can be restricted to evaluate at the top level only.

The parser, when asked to parse an expression, compares it against macros in its macro list in reverse chronological order until it finds a match; when a match occurs, the macro is expanded as above. If the macro was not a terminal one, the resulting text is expanded again until either no macro matches the text or a terminal macro is expanded. If no macro applies, the text must be a symbol, which is looked up in the current lexical environment. If the symbol is not found in the current environment, it is assumed to be an undefined global unless it

happens to be one of the Concurrent Smalltalk primitive names or the `warn-free-references`¹ declaration is in effect, in which case an error or a warning is given.

Macro Implementation

Since the parser is an intensive user of macros, a fast implementation of macros is used to make the parser in the compiler fast. Macros are stored in linked lists hashed by the first non-variable symbol in the macro pattern; macros with no such symbols are stored in a separate list. Thus, relatively few macros have to be examined for a given piece of Concurrent Smalltalk text. Furthermore, the macros themselves are compiled Lisp functions that check that their patterns are satisfied and, if so, compute the text replacement or perform their terminal actions. Compiling macros avoids the costly interpreted unification step during pattern matching. The `make-macro-text` function in the Environment file compiles a macro into a Lisp function.

If a macro contains an `@` directive in its pattern, the macro expander calls itself recursively on the text matching the `@` directive. In this case it does not allow terminal macro expansion on that text.

Environments

While the parser is generating code, it frequently needs to determine the meanings of identifiers. It uses linked environments to keep track of statically scoped identifiers such as the names of local variables and continuations. The last local environment is linked to the global environment to cause a search of the global and class environments when an identifier is not defined locally. Optimist II distinguishes local variables according to whether they are `eq` to each other or not. Thus, no alpha-renaming is necessary anywhere in the parser. Also, a lambda may reference local variables it captured from an enclosing lambda. Since most of the optimizations cannot handle externally visible local variables, such local variables are “unshared” before the optimization pass is invoked.

Concurrent Smalltalk Runtime

Most of the Concurrent Smalltalk directives described in Appendix A are macros which expand into either other Concurrent Smalltalk primitives or hidden primitives. The Runtime file contains a listing of all macros used by Concurrent Smalltalk.

Top-Level Primitives

Most Concurrent Smalltalk top-level primitives listed in Appendix A expand into the directive `hcodes` and are evaluated at expression interpretation time. Directive `hcodes` may be interpreted but not compiled; to ensure that no directive will be compiled, directives are prohibited inside `lambdas` (and, of course, any constructs which expand into `lambdas`). A few directives such as `include`, `top-level set`, and `defclass`² are evaluated by the reader; those directives must be placed at the top level—they may not be nested in any expression except a top-level `begin`, which evaluates its arguments sequentially at the top level.

Method-Lambdas

A method-lambda of a class `c` expands into a lambda with a formal `self` of type `c` prepended to the method-lambda's formals and a `(_with-object (self:c) ...)` form surrounding the body of the lambda. The `_with-object` form establishes bindings in the parser's environment that associ-

¹See Appendix B.

²`Defclass` isn't really evaluated by the reader; nevertheless, it must be a top-level form because it expands into a top-level `begin` containing the internal class definition followed by definitions of accessor and predicate methods. The internal class definition has to have been *interpreted* before the accessor method definitions are *read*; otherwise, the reader will complain about an undefined class. Grouped forms not at the top level and not in a top-level `begin` are read as a group and then interpreted as a group.

ate names of C's instance variables to ivar-refs of the corresponding instance variables pointed by the self object. The action of `_with-object` is analogous to that of the symbol-macrolet construct in CLOS [6].

Optimist II does not restrict a lambda to referencing only one instance object; in fact, through inlining of method-lambdas or accessor methods, a lambda can reference many objects at the same time. Objects may also be referenced through the use of `_with-object` directly in Concurrent Smalltalk code, but this practice is discouraged, as it uses a nonstandard feature of the language and gains no real functionality.

Loops

Although Optimist II can optimize and output code with loops in it, loops are currently not implemented this way. The problem is that a Concurrent Smalltalk function with a loop in it might execute for a long time and not allow any other messages to be processed at its node. To prevent this problem, loops are implemented as closures which pass themselves as arguments—`while (< i 10) (set i (+ i 1))` expands into:

```
(clet ((_loop
      (lambda ((_loop-arg: function &no-leak))::_while)
      (if (< i 10)
          (set i (+ i 1))
          (return _while 'nil))
      (_loop-arg)))
      (_loop _loop))
```

The `_loop` function is called and passed itself as an argument. If `i` is less than 10, `_loop` increments `i` and calls its argument tail-recursively; otherwise, it returns `nil` to the caller. The tail-recursive call breaks the long invocation of the function.

The compiler is not yet sophisticated enough to detect that the value of the `_loop` variable never changes, so the `_loop-arg` argument to the internal function can be eliminated and the function could call itself recursively directly.

Initial Transformations

Immediately after the hcode is created by the parser, a transformation and an optimization are done on it. The first transformation flattens all exit hcodes out of every newly created lambda. Exit hcodes are generated by the `exit` Concurrent Smalltalk primitive, which may also be a result of the expansion of a `return` statement. Each exit hcode in the lambda is removed and the preceding statement linked to the digraph's root dinode to indicate that the execution of the lambda should terminate at that point. Sometimes exit hcodes can be found nested inside `nconcurrently` hcodes; if that is the case, the exit flattener moves as many of the `nconcurrently`'s threads outside as it needs to remove all exit hcodes from the `nconcurrently`. Then it flattens the exits as usual. An example is shown in Figure 3-2.

Simple structural optimizations are done immediately after the exits are flattened. These optimizations do not depend on dataflow analysis and can, therefore, be done before lexical variables are untangled. The optimizations consist of the following transformations:

- If statements with identical consequents and alternatives are deleted.
- If statements conditioned on constants are deleted, and resulting dead code, if any, eliminated.
- Move statements with identical sources and destinations are deleted.
- Assert-type statements on constants are checked and deleted. The compiler generates an error if an assertion fails.

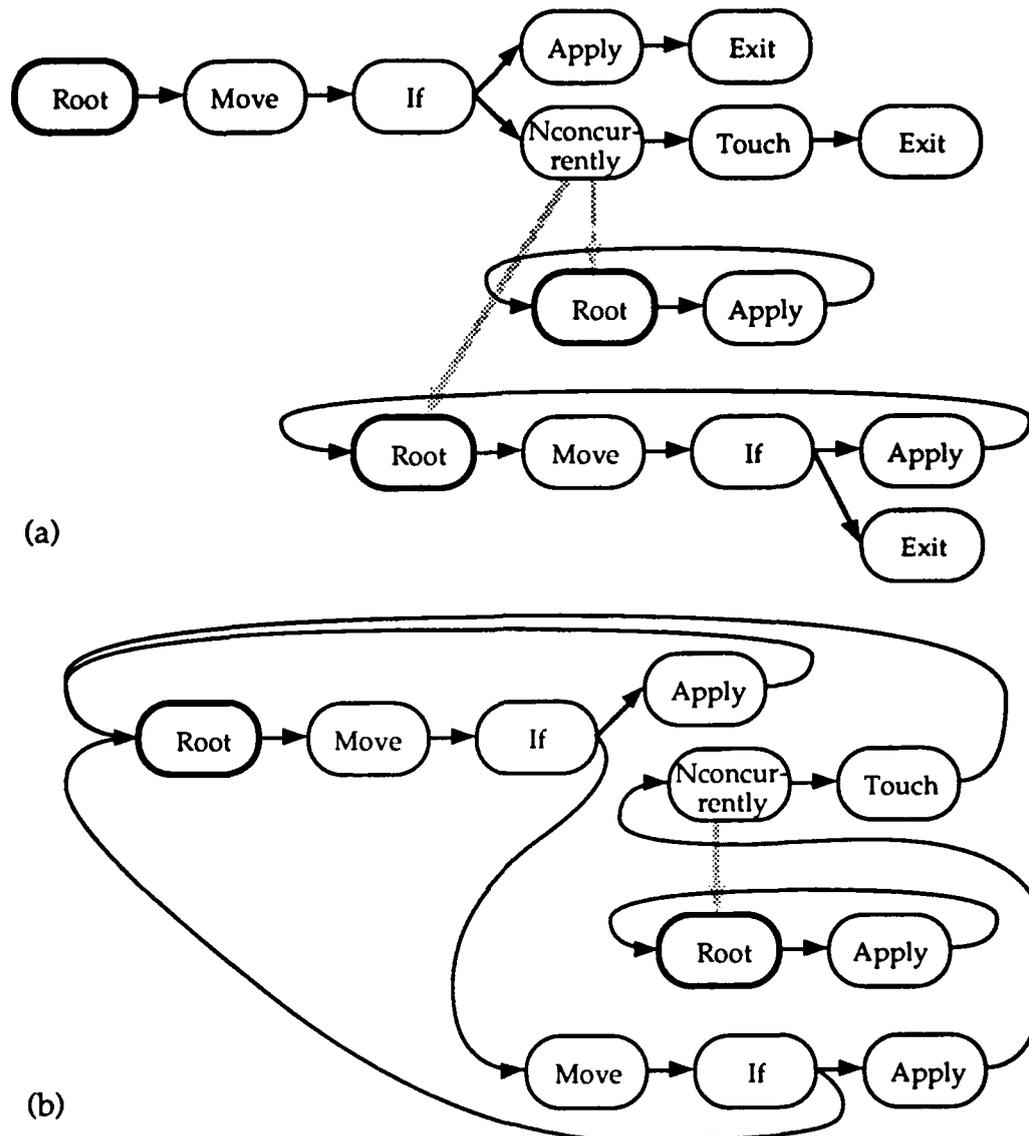


Figure 3-2. Exit Flattening Example

Exit statements are inserted by the parser in all places in which the execution of a lambda should terminate. As the first transformation, those exit statements are removed and replaced with links back to the root of a digraph. For example, part (a) shows the main body of a lambda with two sub-digraphs that are the threads of a nconcurrently. After exit removal (b), all exit paths are linked back to the root of the main body of the lambda, which also required the inlining of one of the nconcurrently's threads.

- Touch and force statements on constants are deleted.
- Empty nconcurrently statements are deleted.
- One-thread nconcurrently statements are replaced by their threads.

The structural optimizations are done for two reasons: First, structural optimizations shorten the hcode, using less memory in the later compiler stages and making them run faster. Second, structural optimizations may remove some variable references, improving the quality of the code produced by lambda-collapsing and the nconcurrently flattener in the optimization phase.

Lambda-Collapsing

Lambda-collapsing is the process of unnesting nested lambdas. After lambda-collapsing, each lambda has exclusive access to its local variables. Lambda-collapsing becomes difficult when the inner lambdas reference the outer lambdas' local variables and continuations. Since continuations are restricted local variables, they will not be discussed here further. Lambda-collapsing occupies most of the Preoptimizer file. Since lambda-collapsing is a complex process, an illustrative example is provided at the end of this section.

The lambda-collapser (the assign-lexicals Lisp function) examines each outermost lambda in the hcode produced by the initial transformations. For each outermost lambda L it looks at the lambdas N_1, N_2, \dots, N_k nested in L and their free variables. Each nested lambda N_i is considered to also include any lambdas nested in it. Thus, if, say, N_2 contains a lambda $N_{2,1}$ that references a variable x that is not defined in $N_{2,1}$ or N_2 , then x is a free variable of both $N_{2,1}$ and N_2 . If a nested lambda N_i does not reference any free variables, it is a self-contained lambda and a first-class data object and does not present any difficulties here. Otherwise, N_i is the code portion of a closure.

The lambda-collapser first calculates the sets of free variables read and written by N_i . Next, the lambda-collapser considers each local variable x_j of L. A local x_j is called a *mutable lexical* if it is either (1) written by any N_i or (2) read by any closure N_i and written by L after the closure N_i has been created by L and before the closure was called for the last time. Mutable lexicals of the first kind are easy to determine by scanning every N_i and checking which free variables are written in any hcode in it. To determine mutable lexicals of the second kind, the lambda-collapser solves a few dataflow problems on L. In effect, to each variable x_j in L, it assigns a state machine S_j (Figure 3-3) and uses the dataflow problem solver to run S_j through all possible control paths in L. If S_j ever enters state 4, x_j is a mutable lexical of the second kind. The state machine assumes that any local variable x_j that is modified after the creation of a live¹ closure which reads x_j is a mutable lexical. Since the compiler cannot currently determine when a lambda finishes executing, it cannot optimize local variables that are modified by L only after the closures have completed execution.

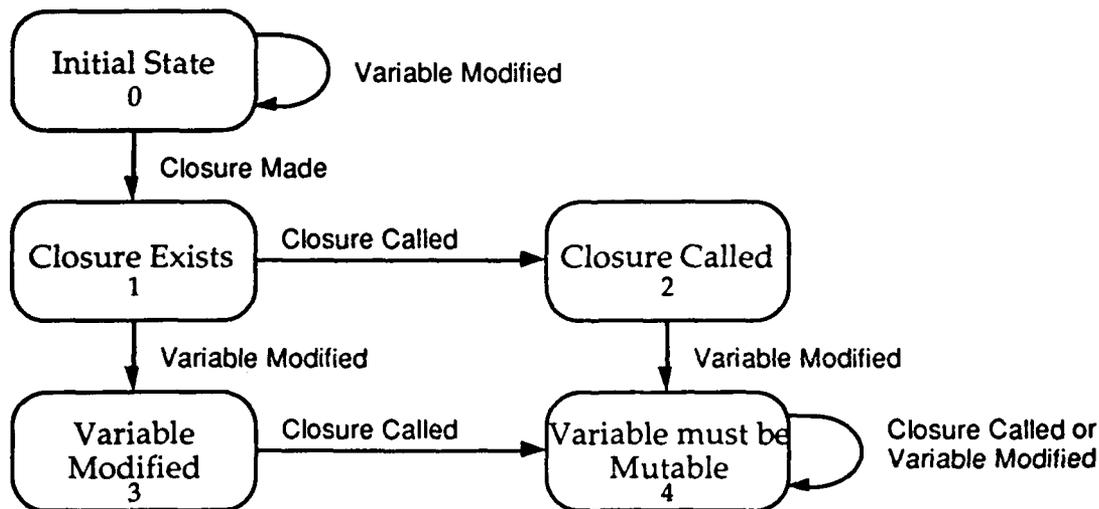


Figure 3-3. Lexical Variable State Machine

Each local variable starts in state 0 at the beginning of the lambda. For each local variable every possible path of control flow is traversed and a state updated as above. If the variable ever enters state 4, it must be a mutable lexical of the second kind—the variable's value cannot be saved with the closure when the closure is made.

¹If the closure is not called, it is not a live closure, and the variable is not necessarily a mutable lexical.

Any variable that is free in one of the lambdas N_i and is not a mutable lexical is an *immutable lexical*. Once all of the free variables in the sublambdas of L have been classified, the lambdas are separated.

Each sublambda N_i of L that has free variables is assigned a number of *display* parameters in addition to the normal parameters it has. The values of the display parameters are determined at the time a closure of N_i is created. Immutable lexicals are stored directly in the display, while mutable lexicals are stored in an object whose pointer is passed in the display. More than one such object may be present if N_i uses mutable lexicals from several levels of enclosing lambdas.

Once the display parameters are assigned to the sublambdas, the code of L is modified to store the display parameters into a closure whenever one is created, and the N_i 's are modified to use the display parameters instead of referencing L 's locals directly. If L has any mutable lexicals, it creates an object containing them upon entry and treats mutable lexicals as if they were instance variables of that object; any mutable lexicals that are also parameters of L are copied into that object as soon as it is created. The object containing mutable lexicals is itself mutable, so only one copy of it per invocation of L can be present on the J-Machine. The object is not disposed because Optimist II cannot determine the temporal lifetime of a closure; the object and the closures have to be garbage-collected.

After the above transformation, L has exclusive access to its locals. Since some of the N_i 's could themselves have locals used by their sublambdas, the lambda-collapser calls itself recursively on every lambda and closure contained in L , even if that lambda did not have any external free variables.

Efficiency Considerations

There are several advantages for using immutable lexicals instead of mutable lexicals:

- Immutable lexicals are stored directly in a closure's display, so the closure has immediate access to their values.
- Closures are immutable objects. If many closures are executing simultaneously, many copies of the closures and their immutable lexicals can be made. On the other hand, if many copies of a closure with a mutable lexical are executing, the copies will be contending for the single object containing that lexical's current value.
- The outer lambda can store immutable lexicals in its context or in registers, while it has to allocate an object for mutable lexicals and keep their values there.

In order to ensure that lexically scoped variables are immutable lexicals, the programmer should check that their values are not altered after any closures which might reference them are created.

Example

Consider the following code:

```
(defun outer (x)
  (clet ((y 3)
        (z 4)
        (t 1))
    (clet ((inner1
           (lambda ()
             ((lambda () (cset x z)) ;inner1
              (write x y)))
          (inner2
           (lambda ()
             (write y)))
          (inner3
           (lambda (a)
             (write a))))
      (if (zero? x)
          (inner1)
          (cset x 5))
      (cset z 3)
      (inner 2)
      (write x y z t))))
```

The lambda-collapser first determines that the outer lambda has no free variables, so it is made into a normal function instead of a closure. Next it examines the three sublambdas within outer: inner1, inner2, and inner3. Inner1 will become a closure because it has three free variables, x, y and z. It writes to x, so x becomes a mutable lexical; although inner1 does not write to y and z, another lambda might, so y's and z's statuses are unknown. Inner2 will also become a closure because it has one free variable, y, whose status is still unknown. Since inner3 has no free variables, it becomes a normal function.

Next the lambda-collapser runs the state machines on the x, y, and z locals in outer; outer also has other locals such as t, inner1, inner2, and inner3, but those are not referenced by any inner lambdas. X is already known to be a mutable lexical of the first kind. Y is not written anywhere after inner1 and inner2 are created, so it is an immutable lexical. Z is written after the inner1 closure is created, and the compiler makes it a mutable lexical of the second kind. Unfortunately, the compiler does not realize that z is altered only after inner1 finishes executing; if it were smarter, it could have made z an immutable lexical. Finally, the lambda-collapser creates the displays and alters the code of the lambdas to produce a parameter-passing pattern shown in Table 3-6.

Table 3-6. Lambda-Collapser Example Results

Name	Outer	Inner1	Inner11	Inner2	Inner3
Parameters	x (copied into lexical-object)				a
Returns	continuation-0	continuation-1	continuation-2	continuation-3	continuation-4
Display		lexical-object y	lexical-object	y	
Locals	y t inner1 inner2 inner3 lexical-object				a

	lexical-object
Instance Variables	x z

Top-Level Evaluator

Lambda-collapsing was the last preliminary hcode transformation. At this point the hcode is in a format understood by the interpreter. If it found no syntax errors, Optimist II now evaluates the Concurrent Smalltalk expression it just read by running the expression's hcodes through the hcode interpreter. If the expression contained any directives, the interpreter executes them at this time.

Interpreter

The interpreter is a simple hcode interpreter for executing Concurrent Smalltalk programs. The interpreter is completely sequential. Except for full futures and some unimplemented input/output facilities, the interpreter is a valid Concurrent Smalltalk implementation—the Concurrent Smalltalk definition allows cfutures to be touched at the implementation's discretion, so a completely sequential Concurrent Smalltalk interpreter trivially “touches” each cfuture as soon as it is created. While the interpreter never achieves any parallelism, it couldn't use parallelism if it had any because it is running on a sequential computer.

The interpreter in Optimist II was provided for three reasons:

- It is a powerful constant expression evaluator for expressions encountered while compiling Concurrent Smalltalk programs.
- It is the most interactive Concurrent Smalltalk environment, allowing methods and functions to be changed almost instantly.
- It permits debugging of Concurrent Smalltalk programs before they are compiled into MDP assembly language.
- It maintains the Concurrent Smalltalk global environment and permits interactive examination of that environment.

Currently the interpreter can only interpret unoptimized hcode; however, a bypass hcode path could be added to transfer optimized hcode back to the interpreter. This bypass is not quite as simple as it sounds because the format of continuations changes during optimization.

3.3. Optimization

As long as no MDP code output is desired, Optimist II does not leave its first phase. Only when a compile command is issued does Optimist II enter its second phase, its first goal being to determine just what it should compile. Every compile command requires a *root set* of objects that should be compiled. The compiler uses the treewalker to automatically determine the minimum amount of code that has to be compiled and loaded in order to permit running the functions in the root set on the J-Machine.

Treewalker

The root set specified in the compile command is passed to the treewalker, which appends it to its own permanent root set of objects which must always be compiled (Table 3-7). The treewalker then calls the optimizer on each code object in its set and scans the optimized hcode (if the object is not code, the treewalker scans it directly). If, while scanning, it encounters an object not in its current set of objects, it adds that object to its set, optimizes it if necessary, and scans it. The process continues until every object referenced by any object in the treewalker's set is also in that set. At that point the second phase of the compiler has completed and the treewalker calls the compiler's third stage to compile and assemble each object in the set and print the resulting MDPSim code into a text file.

Table 3-7. Permanent Root Objects

<code>_closure</code>	<code>boolean</code>	<code>character</code>	<code>#:class</code>	<code>context</code>
<code>#:continuation</code>	<code>displacement</code>	<code>distobj</code>	<code>distributed-class</code>	<code>#:false</code>
<code>float</code>	<code>funct</code>	<code>function</code>	<code>global</code>	<code>integer</code>
<code>magnitude</code>	<code>null</code>	<code>number</code>	<code>object</code>	<code>primitive-class</code>
<code>real</code>	<code>selector</code>	<code>standard-class</code>	<code>symbol</code>	<code>#:true</code>

These objects are emitted in the output assembly file regardless of which objects were compiled. `_closure`, `context`, `displacement`, `#:continuation`, and `global` are internal Optimist II classes.

Calling the Optimizer

The optimizer is called simply by requesting the value of the hcode or mdp-hcode CLOS slot in a Concurrent Smalltalk lambda (`cst-lambda`). If the lambda has already been optimized, these slots contain the optimized hcode and hcode optimized for the MDP, respectively. If not, those slots are unbound, and CLOS calls the optimizer to calculate their values. Thus, a lambda's optimized hcode can be requested repeatedly by the treewalker or the optimizer without a performance penalty. To prevent infinite loops, a semaphore keeps a function optimizing a lambda from requesting that lambda's optimized hcode. One of the consequences of this rule is that a function may not be inlined inside itself.

Guide to Optimizations

The transformations done by the optimizer are summarized in Figure 3-4. The transformations can be divided roughly into two classes: general hcode optimizations and MDP-specific optimizations and transformations. The general optimizations occupying the first half of the optimizer produce optimized hcode. If MDP assembly code output is desired, the second half of the optimizer is invoked to convert a number of hcode constructs into simpler, MDP-specific ones. For example, the second half of the optimizer converts globals into references to global objects, CAS built-ins into code that explicitly compares and sets values, and three-argument sums into two two-argument sums. The order of optimization is critical; expansion of CASes into compare-and-set code could not have been done in the first half of the optimizer because there was no way to assure its atomicity.

have a longer range, Optimist II could be more liberal with the use of MDP register R0 to hold values between statements¹. A smarter register allocator could assign a variable to a register for part of its lifetime. The peephole optimizer could replace branches to SUSPEND instructions with SUSPEND instructions themselves. The implementation of closures could be made faster. The compiler could automatically detect side-effect-free and no-leak functions; this information might permit it to explicitly deallocate some objects such as closures if it could prove that they could not be referenced again. Overall, though, it seems that, except for loops which are deliberately broken to avoid hogging processors, no more than a few per cent more performance can be squeezed out of the code generated by Optimist II; however, since the operating system overhead time overwhelms the execution time in Concurrent Smalltalk methods, there might be room for improvement through coordinated compiler and operating system changes.

¹In Architecture 10, all but the shortest branches required the value of R0 to be altered, rendering that register practically useless for holding values between statements.

Chapter 4. The Cosmos Operating System

Design Goals

The Cosmos operating system was designed primarily as a support kernel for running Concurrent Smalltalk programs on the J-Machine. Nevertheless, Cosmos is not specialized to Concurrent Smalltalk, and many of the operating system's components could be used to support a general message-passing environment.

The goals in designing the operating system were, in order:

1. To make a working operating system.
2. To make the operating system as efficient as possible.
3. To make the operating system as simple and flexible as possible.

The design of the operating system also had to be small enough to allow both it and most of the Optimist II compiler to be written in one semester; for this reason garbage collection and load management facilities were not included in the operating system. Several steps were taken to achieve goal (1), including the criticality system and the debugging techniques described later. The criticality system is an organized accounting method used to ensure that no re-entrancy problems occur when operating system routines call each other. Features were added to MDPSim to detect and signal race conditions known as hazards. To achieve goal (2), the entire operating system kernel was written in hand-optimized assembly language. Poor J-Machine performance can no longer be blamed solely on the operating system. Goal (3) was achieved by providing general data structures that are reused in many components of the system.

Functionality

The operating system assists Concurrent Smalltalk programs by providing the following services:

- Initialization and setup of the J-Machine.
- Providing fault handlers for faults needed to keep the J-Machine running.
- Global function calls and returns.
- Looking up methods corresponding to class-selector or object-selector pairs.
- Context allocation and deallocation facilities and conventions.
- Local and global object allocation, deallocation, lookup, and migration facilities. Mutable objects exist on only one node at a time, while immutable objects can exist on many nodes at a time; all but the primary copy can be purged when extra memory is needed.
- Support for distributed objects as defined in Concurrent Smalltalk.
- Support for Concurrent Smalltalk primitives such as determining the type of an object.
- Calls assisting in the creation and evaluation of closures.
- An integer division routine.
- Debugging and consistency-checking facilities.

ting the locked flag prevents an object from migrating or being deleted during critical protocol sections.

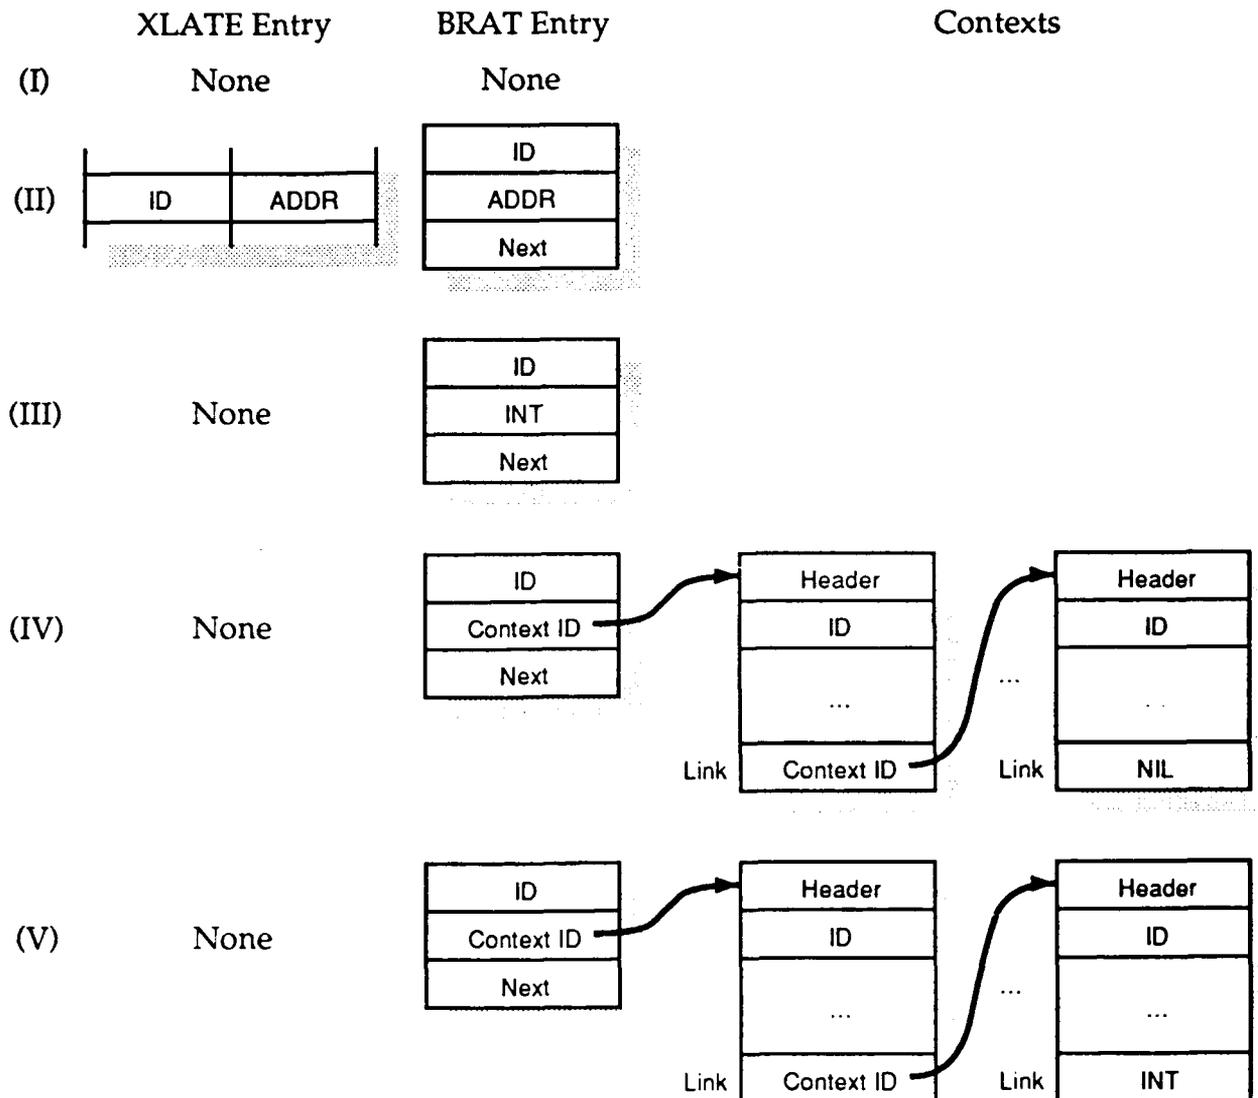


Figure 4-11. Object XLATE Table and BRAT Entries

There are five possible BRAT table states for a particular object. Each object must have a BRAT entry on its home node. The XLATE table entry, where specified, is optional. The states are as follows:

- I. The object does not exist on this node, and its whereabouts are unknown.
- II. The object exists on this node. Its physical address is given.
- III. The object does not exist on this node, but it is believed to reside on the node specified by the integer.
- IV. The object does not exist on this node, but the contexts linked to its BRAT entry are waiting for its arrival.
- V. The object does not exist on this node, but the contexts linked to its BRAT entry are waiting for its arrival, and the object is believed to reside on the node specified by the integer.

Only states II, III, and V are allowed on an object's home node, while only states I, II, and IV are allowed on the other nodes.

If the data in the CFUT-tagged word was zero or negative, the control manager halts the computer because an uninitialized variable was accessed. On startup, all memory in the MDP's heap is cleared to CFUT:-1.

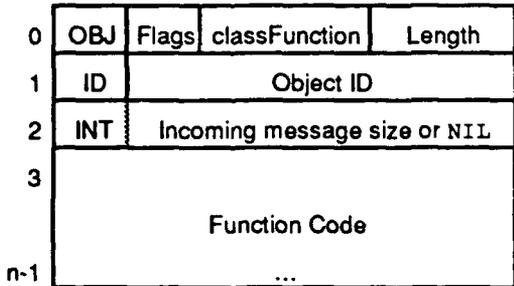


Figure 4-20. Function Object Format

The function object contains the code for a function. Registers A0 and ID0 point to the function while it is executing. The third word contains the size of the message expected by the function or NIL if the size is not known or the function expects a variable number of arguments. The compiler initializes that word, but the operating system does not check it against the size of the message that invoked the function; that check would add at least five instructions to the function dispatch time.

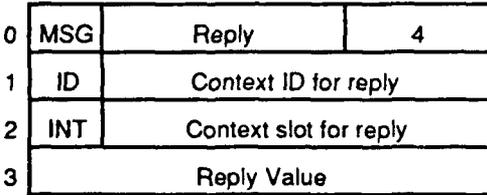


Figure 4-21. Reply Message Format

The Reply message carries the reply value to the specified slot in the specified context. The context ID and reply slot may not be NIL—if they were NIL in the Apply message, no Reply message is sent.

Functions return results to their callers via Reply messages (Figure 4-21). If a function returns multiple values, it sends one Reply message for each value returned. The Reply handler on the caller's node performs the following processing when it receives the message:

1. The value from the message is stored over the cfuture in the caller's context. However, if the slot indicated in the Reply message did not originally contain a cfuture, the Reply handler halts because some function replied twice to the same slot or the compiler generated incorrect code.
2. The CFUT-tagged link field in the caller's context is checked against the slot number of the newly updated slot. If the numbers match, the context is resumed; otherwise, the Reply handler exits because the context is waiting for some other event.

Actually, for reasons of efficiency the check in (1) is done only if the slot number in (2) doesn't match.

Utilities

The operating system kernel currently contains three utilities: a divide routine, a closure maker, and a closure evaluator. The Divide system call divides one integer by another and returns the quotient and remainder using the sign conventions described in Appendix A. The divide routine includes considerable overhead to evaluate all signed 32-bit results correctly, including special cases such as dividing -\$80000000 by 1 or -1 because a large-integer

Tree message is sent to the zeroth constituent of the distributed object (which does not exist yet, but the Co function can calculate its ID anyway). When that message returns, the DID is returned to the caller.

```
(NewDistobjTree class:Class size:integer ID:distobj start,logDelta:integer):null  
Function
```

NewDistobjTree creates constituents numbered start through $(start+2^{\logDelta}-1)$ of the distributed object with the DID ID and then returns. Each constituent has group, index, and logical size instance variables, which are initialized to the appropriate values; size is the logical size. NewDistobjTree works by creating the constituent start if logDelta is zero or by recursing itself on the two halves of its range if logDelta is positive.

The current implementation will have to be extended on a larger system so as not to bottleneck node 0, but it is adequate for small and medium-range systems.

4.4. Summary

The Cosmos operating system provides the software extension to the MDP architecture needed to run Concurrent Smalltalk programs. The operating system is comprised of a kernel resident on each MDP and a set of Concurrent Smalltalk functions written in either MDP assembly language or Concurrent Smalltalk.

The operating system is built in layers which include the heap manager, BRAT manager, object manager, context manager, global object manager, method manager, control manager, utilities, and MDP and CST runtime systems. Efficiency and re-entrancy problems were recurring issues in the design of the operating system kernel. The criticality system was developed to deal with the re-entrancy and double faulting problems. In addition, many routines are inlined in other routines to make the efficiency reasonable and avoid double faults and re-entrancy problems (in some cases a system call cannot call another system call but can use it inlined because there are no more free data registers on the MDP; global variables cannot be used as temporaries in routines running at criticality less than 2).

The operating system facilities were streamlined and simplified compared with those proposed in [38]. The emphasis was on making resource allocation decisions as late as possible. Thus, the size of the BRAT is varied dynamically at run time instead of being fixed at operating system compile time as in [38]. The object migration protocol has been streamlined compared with the one in [38]. The resource wait table in [38] has been eliminated entirely; the BRAT manager is a general-purpose mechanism that can perform the same task better.

Finally, a scheme for quickly addressing constituents of distributed objects was designed. The scheme is very fast and requires only knowledge of a group ID to find either some nearby constituent or any given constituent. Disadvantages of the scheme include the necessity of rounding the size of a distributed object up to the nearest power of two and a resulting decreased load on the higher-numbered MDPs in the J-Machine. Means of circumventing these disadvantages were explored.

Chapter 5. Sample Program

This chapter presents the progress of a simple program through the various stages of compilation. Unfortunately, it is difficult to write a simple sample program that exercises all of the features of a compiler. Instead of trying to write a contrived sample program that exercised as many features as possible, I decided that a simpler program that exercised the major optimizations would make a better example. If an illustration of a more esoteric optimization is desired, one can write an appropriate Concurrent Smalltalk program, compile it with Optimist II, and watch the intermediate output.

The source program, listed in Figure 5-1, returns the sum of the integers from 0 to n. Figure 5-2 shows a transcript of the interactive Optimist II session in which the program was entered, tested on a few inputs, and then compiled.

```
(defmethod average integer (b:integer)
  (// (+ self b) 2))

(defmethod average boolean (b:boolean)
  false)

(defmethod rangesum integer (high)
  (if (= self high)
    self
    (let ((middle (average self high)))
      (+ (rangesum self middle)
         (rangesum (+ middle 1) high))))))

(defun sum (n)
  (rangesum 0 n))
```

Figure 5-1. The Rangesum Program

The sum function adds the integers from 0 to n, inclusive. The rangesum method adds the integers from self to high, inclusive. The average method returns the average of two integers; the definition of average for booleans was included just to confuse the compiler a bit.

```
CST:(+ 2 2)
#<Integer 4>
CST:(include)
#<Cst-Lambda 5024988 SUM>
CST:(sum 0)
#<Integer 0>
CST:(sum 1)
#<Integer 1>
CST:(sum 2)
#<Integer 3>
CST:(sum 10)
#<Integer 55>
CST:(average 3 5)
#<Integer 4>
CST:(average true false)
#<False>
CST:(sum 100)
#<Integer 5050>
CST:(rangesum 10 13)
#<Integer 46>
CST:(compile sum "::fact:Rangesum.mdp")

Optimizing #<Cst-Lambda 4713968 CST::SUM>
Expanded continuations
Folded constants
Forwarded replies

Optimizing #<Cst-Lambda 4711636 CST::RANGESUM>
Collapsed nconcurrentlys
Expanded continuations
Specialized local type:
Deleted moves
Deleted touches
Folded constants

Optimizing #<Cst-Lambda 4709940 CST::AVERAGE>
Expanded continuations
```

Concurrent Smalltalk on the Message-Driven Processor

```
Specialized local types
Deleted locals

Back to #<Cst-Lambda 4711636 CST::RANGESUM>
Substituted inlines
Specialized local types
Deleted moves
Deleted touches
Propagated values
Deleted dead definitions
Deleted locals

Back to #<Cst-Lambda 4713968 CST::SUM>
Deleted locals
Inserted ENTER and EXIT
Split statements
Optimized built-ins
Inserted ENTER and EXIT

Generating code

Assembling
Initialized vlocs

Printing
Assigned labels

Generating code

Assembling
Inserted branches
Initialized vlocs
Compacted SENDs

Printing
Assigned labels
#<Cst-Lambda 4713968 SUM>
```

Figure 5-2. Rangesum Interactive Session

The Rangesum file was read in the (include) directive, at which time the user interactively chose the file name using a Macintosh dialog. A few functions were then tested, after which point the file was compiled.

The following sections will illustrate the actions of some of the compiler's optimizations on the program in Figure 5-1. Please refer to Chapter 3 and [21] for explanations of the transformations.

Initial Phase

The initial phase of the compiler first performs a few macro expansions on the input program, compiles the program into hcode, and then performs some transformations on that hcode to get it into a form that the rest of the compiler can use. Figure 5-3 shows the macroexpansions which are done by the Optimist II parser, and Figure 5-4 shows the hcode produced by the parser. To save space, only the transformations on the rangesum method will be shown from this point on.

Optimization Phase

The Optimist II optimization phase performs local and global optimizations on the program. The order of the optimizations can be seen in the transcript in Figure 5-2; the compiler often interrupts the optimization of one function to optimize another because it wants to inline the second function in the first.

The first transformation done by the optimization phase is the collapsing of nconcurrently and the expansion of continuations to the two-variable format, yielding the hcode in Figure 5-5. The threads of the nconcurrently are inlined in the function's main body, and the nconcurrently statement is removed. Then, since an MDP continuation is actually two words (a context ID and an offset within that context where the return value should be stored), each continuation variable is replaced by two variables.

Concurrent Smalltalk on the Message-Driven Processor

```

REF REV selPLUS=TAGO:subSEL<<subtagN|(0&mX)<<sX|(0&mY)<<sY|(0&mZ)<<sZ|(0&m3)<<s3|(0&m4)<<s4|(0&m5)<<s5
REF REV selEQUAL=TAGO:subSEL<<subtagN|(1&mX)<<sX|(1&mY)<<sY|(1&mZ)<<sZ|(1&m3)<<s3|(1&m4)<<s4|(1&m5)<<s5
REF REV selAsh=TAGO:subSEL<<subtagN|(2&mX)<<sX|(2&mY)<<sY|(2&mZ)<<sZ|(2&m3)<<s3|(2&m4)<<s4|(2&m5)<<s5
REF REV fRangesum=ID:(-1&mX)<<sX|(-1&mY)<<sY|(-1&mZ)<<sZ|(-1&mS)<<sS
REF REV fSum=ID:(-2&mX)<<sX|(-2&mY)<<sY|(-2&mZ)<<sZ|(-2&mS)<<sS

```

```

MODULE cObject
DC MSG:hdrCopyable|cStandard_Class<<offsetN|5
DC TAGO:subCLASS<<subtagN|cObject
DC MSG:cObject<<offsetN|2
DC 1
DC TAGO:subCLASS<<subtagN|cObject
END

```

```

MODULE cClass
DC MSG:hdrCopyable|cPrimitive_Class<<offsetN|6
DC TAGO:subCLASS<<subtagN|cClass
DC NIL
DC 2
DC TAGO:subCLASS<<subtagN|cClass
DC TAGO:subCLASS<<subtagN|cObject
END

```

```

MODULE cStandard_Class
DC MSG:hdrCopyable|cPrimitive_Class<<offsetN|7
DC TAGO:subCLASS<<subtagN|cStandard_Class
DC NIL
DC 3
DC TAGO:subCLASS<<subtagN|cStandard_Class
DC TAGO:subCLASS<<subtagN|cClass
DC TAGO:subCLASS<<subtagN|cObject
END

```

... MODULEs for the rest of the classes deleted ...

```

MODULE selPLUS
DC MSG:hdrCopyable|cSelector<<offsetN|3
DC {selPLUS}
DC 0
END

```

```

MODULE selEQUAL
DC MSG:hdrCopyable|cSelector<<offsetN|3
DC {selEQUAL}
DC 0
END

```

```

MODULE selAsh
DC MSG:hdrCopyable|cSelector<<offsetN|3
DC {selAsh}
DC 0
END

```

```

MODULE fRangesum
DC MSG:hdrCopyable|cFunction<<offsetN|28
DC {fRangesum}
DC 6
MOVE {2,A3},R0 ; 3
MOVE {2,A3},R3 ; 3.5
EQUAL R3,{3,A3},R1 ; 4
BT R1,^L001 ; 4.5
ADD R3,{3,A3},R1 ; 5
ASH R1,-1,R3 ; 5.5
ADD R3,1,R2 ; 6
MOVE R2,R0 ; 6.5
CALL objectNode ; 7
DC MSG:msgApplyFunction:6 ; 8
SEND20 R1,R0 ; 9
DC {fRangesum} ; 10
SEND20 R0,R2 ; 11
SEND0 {3,A3} ; 11.5
MOVE 6,R0 ; 12
SEND2E0 {1,A1},RC ; 12.5
WTAG RC,6,R0 ; 13
MOVE RC,{6,A1} ; 13.5
MOVE {2,A3},R0 ; 14
CALL objectNode ; 14.5
DC MSG:msgApplyFunction:6 ; 15
SEND20 R1,R0 ; 16
DC {fRangesum} ; 17
SEND0 R0 ; 18
SEND20 {2,A3},R3 ; 18.5
MOVE 7,R0 ; 19
SEND2E0 {1,A1},RC ; 19.5
WTAG RC,6,R0 ; 20
MOVE RC,{7,A1} ; 20.5
MOVE {7,A1},R2 ; 21
ADD R2,{6,A1},R1 ; 21.5
MOVE R1,{2,A3} ; 22

```

```

L001:  MOVE    [4,A3],R2          ; 23
      BN:IL  R2,~L002          ; 23.5
      DC    MSG:msgReply|4     ; 24
      SEND20 R2,R0             ; 25
      SEND0  R2                ; 25.5
      SEND0  [5,A3]           ; 26
      SENDEC [2,A3]           ; 26.5
L002:  SUSPEND                ; 27
      END

      MODULE fSum
      DC    MSG:hdrCopyable|cFunction<<offsetN|10
      DC    (fSum)
      DC    5
      MOVE  0,R0               ; 3
      CALL  objectNode         ; 3.5
      DC    MSG:msgApplyFunction|6 ; 4
      SEND20 R1,R0            ; 5
      DC    (fRangesum)       ; 6
      SEND0  R0                ; 7
      SEND0  0                 ; 7.5
      SEND0  [2,A3]           ; 8
      SEND0  [3,A3]           ; 8.5
      SENDEC [4,A3]           ; 9
      SUSPEND                ; 9.5
      END

DOWNLOAD cObject
DOWNLOAD cClass
DOWNLOAD cStandard_Class
DOWNLOAD cPrimitive_Class
DOWNLOAD cDistributed_Class
DOWNLOAD cSymbol
DOWNLOAD cNull
DOWNLOAD cFuncnt
DOWNLOAD cSelector
DOWNLOAD cMagnitude
DOWNLOAD cCharacter
DOWNLOAD cNumber
DOWNLOAD cReal
DOWNLOAD cInteger
DOWNLOAD cBoolean
DOWNLOAD cFalse
DOWNLOAD cTrue
DOWNLOAD cFloat
DOWNLOAD cFunction
DOWNLOAD c_Closure
DOWNLOAD cContext
DOWNLOAD cDisplacement
DOWNLOAD cContinuation
DOWNLOAD cGlobal
DOWNLOAD cDistobj
DOWNLOAD selPLUS
DOWNLOAD selEQUAL
DOWNLOAD selAsh
DOWNLOAD fRangesum
DOWNLOAD fSum

RUN

```

Figure 5-12 MDPSim Output File

Except for Cosmos, this file contains all code and data necessary to run sum on a J-Machine. The file starts with class number definitions, which are followed by definitions of the classes themselves, including the class hierarchy. The selectors are defined next, followed by code and MDPSim statements that download all of the code, selector, and class modules to the simulated J-Machine. The RUN command runs the J-Machine until all modules have been loaded.

Only the functions and selectors necessary to run the program have been compiled. For example, neither average method has been included because, after optimization, neither is necessary to run sum. Similarly, all method dispatches have been optimized out, so there is no need to include the definition of the rangesum selector.

Running Rangesum

Before rangesum can be run on MDPSim, a file holding the calls that will be done needs to be defined; the file that was used is shown in Figure 5-13. Each MESSAGE directive defines an ApplyFunction message that can be used to call the sum function. The argument is the third word of the message, while the fourth and fifth words contain a magic continuation that cause the Reply message to be printed by MDPSim in the listener window. The MESSAGE definitions can also be entered into MDPSim manually.

Concurrent Smalltalk on the Message-Driven Processor

Once the calls file is written, MDPSim can be started and used to run sum on a sample input. An example session is shown in Figure 5-14, in which the input 10 is tried on sum, and the statistics observed. The results will be discussed in more detail in Chapter 7.

```
MESSAGE sum1
MSG:msgApplyFunction|5
{fSum}
1
IONODE
0
END

MESSAGE sum10
MSG:msgApplyFunction|5
{fSum}
10
IONODE
0
END

MESSAGE sum50
MSG:msgApplyFunction|5
{fSum}
50
IONODE
0
END
```

Figure 5-13. Rangesum Call File

Three messages have been defined for calling the sum function with the arguments 1, 10, and 50. IONODE is an integer constant predefined by MDPSim and denotes the address of the MDP serving as the I/O node between the J-Machine and the outside world. In MDPSim, the I/O node simply prints every message it receives.

```
MDPSim -x 2 -y 2 -msize 0x1000 ::Cosmos:Cosmos.m RangeSum.mdp RangeSum.calls
```

```
Message-Driven Processor Simulator
Version 7.0 Rev B
Accompanies MDP Architecture Document 11B
Written by Waldemar Horwat
Architecture Updates by Brian Totty and Jerry Larivee
UROPs for Bill Dally
```

```
4 MDPs present.
```

```
@0..3|watch fault all
@0..3|resetstats
@0..3|inject sum10@1
@0..3|run
Fault: @ 1: (faultXlate0)
Fault: @ 1: (BBBW) S008B = DC fItXLATE ;XLATE
Fault: @ 1: (lookupBinding)
Fault: @ 1: (BBBW) S00C6 = DC fItLookupBinding ;S06
Fault: @ 1: (enterBinding)
Fault: @ 1: (BBBW) S00C5 = DC fItEnterBinding ;S05
Fault: @ 2: (blockSend)
Fault: @ 2: (BBBW) S00C2 = DC fItBlockSend ;S02
Fault: @ 2: (faultLimit0)
Fault: @ 2: (BBBW) S0088 = DC fItLimit ;LIMIT
Fault: @ 1: (allocObject)
Fault: @ 1: (BBBW) S00C4 = DC fItAllocObject ;S04
Fault: @ 1: (lookupBinding)
Fault: @ 1: (BBBW) S00C6 = DC fItLookupBinding ;S06
Fault: @ 1: (blockMove)
Fault: @ 1: (BBBW) S00C1 = DC fItBlockMove ;S01
Fault: @ 1: (faultLimit0)
Fault: @ 1: (BBBW) S0088 = DC fItLimit ;LIMIT
Fault: @ 1: (objectNode)
Fault: @ 1: (BBBW) S00D3 = DC fItObjectNode ;S13
Fault: @ 2: (faultXlate0)
Fault: @ 2: (BBBW) S008B = DC fItXLATE ;XLATE
Fault: @ 2: (lookupBinding)
Fault: @ 2: (BBBW) S00C6 = DC fItLookupBinding ;S06
Fault: @ 2: (enterBinding)
Fault: @ 2: (BBBW) S00C5 = DC fItEnterBinding ;S05
Fault: @ 3: (blockSend)
Fault: @ 3: (BBBW) S00C2 = DC fItBlockSend ;S02
Fault: @ 3: (faultLimit0)
```



```

MULH:      0  0.00%
NEQUAL:    0  0.00%
CARRY:     0  0.00%
HALT:      0  0.00%
  BZ:      0  0.00%
  LE:      0  0.00%
  BNZ:     0  0.00%
  LT:      0  0.00%

STOP:     2887  47.13%
Move:     1488  24.29%
ALU:       407   6.64%
Branch:   331   5.40%
Network:  330   5.39%
Field:    204   3.33%
DC:       143   2.33%
Fault:    142   2.32%
Assoc:    130   2.12%
NOP:       64   1.04%
Other:     0   0.00%

Foregnd:  3239  52.87%
Total:    6126

Fault Usage:
  objectNode:  21  26.25%
  faultXlate0: 14  17.50%
  lookupBinding: 11  13.75%
  faultCFut0:  10  12.50%
  faultLimit0:  8  10.00%
  blockSend:   4   5.00%
  allocObject:  4   5.00%
  enterBinding: 4   5.00%
  blockMove:   4   5.00%
  Total:      80

The xlate hit ratio is 109 out of 123 ( 88.62%).

376 words sent in 51 messages on priority 0.
Average message size: 7.37.
16.29 instructions/word (8.61 foreground instructions/word)
120.12 instructions/message (63.51 foreground instructions/message)
No priority 1 words sent.

@0..3)

```

Figure 5-14. MDPSim Transcript

This transcript shows a MDPSim session in which the user loads the rangesum assembly code and calls the sum function with the argument 10 on a 2x2x1-node J-Machine with COSMOS using only internal memory (-msize 0x1000). Since watching faults was enabled, MDPSim prints each fault encountered at each MDP as it is running. The fault message gives the number of the MDP on which the fault occurred, the number of the fault vector, and the name of the fault; the {BBBW} is additional MDPSim breakpoint and watchpoint information. Finally, after 1544 steps the answer 55 is produced and displayed.

The dynamic instruction statistics for the run are also shown. About half of the time is spent distributing the functions to all of the nodes; the second time sum is called with the argument 10, it only takes 893 ticks to produce the answer (a tick is the time it takes every node to execute one instruction; MDPSim assumes that every instruction runs in the same amount of time).

Chapter 6. Debugging

Optimist II, Cosmos, and the Concurrent Smalltalk applications are large programs, and debugging them is an important consideration. I will not discuss the process of debugging Optimist II itself; standard Common Lisp and CLOS techniques such as building firewalls and providing print routines for important data structures were used.

The primary approach to debugging MDP code I took is prevention. I made sure that the Cosmos design was sound before running it. The criticality criteria were very helpful in avoiding re-entrancy and double fault problems. Nevertheless, while the prevention approach was successful on Cosmos itself, it cannot be the sole debugging method used on the Concurrent Smalltalk programs. Instead, a combination of debugging means at various levels has been provided.

Debugging Concurrent Smalltalk Code

The first line of defense is the Optimist II compiler itself. The compiler will complain when it detects errors such as incorrect function argument counts or bad types, if types are declared.

The second line of defense is the interpreter in the Optimist II compiler. The interpreter can be used to run Concurrent Smalltalk programs before they are downloaded into MDPSim or onto a J-Machine. The interpreter provides nearly complete checking of Concurrent Smalltalk programs, so it should catch most of the remaining bugs. However, the interpreter will not catch bugs which occur only on large data sets, nor will it find Cosmos's or the Optimist II code generator's bugs.

Debugging MDP Code on MDPSim

Debugging becomes considerably more difficult once the code is in assembly language form. Fortunately, Cosmos does include some facilities for debugging Concurrent Smalltalk programs.

The third line of defense is comprised of the safety features built into the MDP architecture. Type and bounds checking were extremely valuable when debugging Cosmos, as they catch most common type errors when they happen and prevent runaway programs from doing too much damage to the machine state. Without these facilities debugging Cosmos and Concurrent Smalltalk programs could have been intractable.

The fourth line of defense consists of safety checks built into a number of critical places in Cosmos. These checks include:

- A check in the CFUT handler that distinguishes real cfutures from uninitialized variables, together with the initialization of memory and globals to values that will cause CFUT faults.
- Checks in the XLATE and INVADR handlers for references to primitive, nonexistent, or deleted objects. Without these checks, such references would generate messages that wander about the J-Machine forever.
- A check in the RETURN handler to make sure that the context was expecting the value that was returned. This check catches the extremely elusive bug of replying to the same continuation twice, as the second reply message may overwrite a variable in the context after it has been reallocated to a completely unrelated function. The bug will be caught eventually, even if the second function stores a cfuture into the same context location, because then there will still be two replies to the same context location, and the cycle will repeat itself. Of

recommended to try to find bugs in the earlier steps of the compilation process because the tools at those levels are more robust and informative (but not as faithful to the J-Machine).

Although Cosmos includes many checks for the common Concurrent Smalltalk programming errors, Cosmos does not protect itself from itself—it does not detect corruption in its data structures. Fortunately, segmentation by the MDP ensures that those data structures could only be corrupted by Cosmos itself, as well-compiled Concurrent Smalltalk programs cannot reference data outside their segments. Cosmos was mainly debugged by design, with only minor debugging necessary once the operating system was written.

MDPSim also helps in debugging MDP code by providing watchpoints, breakpoints, the HALT instruction, hazard detection, and determinism, which allows any bug to be reproduced.

Chapter 7. Performance Measurements

Both Cosmos and the code output by Optimist II were optimized for speed. This chapter presents some measurements that determine just how fast compiled Concurrent Smalltalk runs on a J-Machine. Both theoretical derivations and real measurements are presented and compared. Both calculations indicate that the average grain size (the ratio of useful instructions executed to messages sent) for running Concurrent Smalltalk on a J-Machine is between 50 and 70 instructions, and the average number of instructions executed per method is about 100 instructions. This is a pity if the average method only performs a few instructions' worth of real computation, yet, since Cosmos and the code output by Optimist II are already heavily optimized, it does not seem likely that incremental changes will reduce these numbers much further.

In addition to the above figures, various other statistics are presented. The static and dynamic instruction use frequencies were collected to identify areas in which the MDP's hardware performance could be improved; no major surprises were found there. These frequencies indicate that the MDP spends an average of about 2 cycles per instruction; this number increases to 4 if slow external DRAM is used to hold the user program and data.

Finally, the network load is analyzed. The network should not become saturated until more than 343 MDPs are put together; if a larger J-Machine is to be built, either the network will have to be made faster, the operating system slower, or considerable attention will have to be paid to locality.

Table 7-2. Selected User Action Instruction Counts

Action	Instruction Count	Description
Function or Method Call	=11+nargs. May be higher if arguments must be touched or lower if many SEND2s are used.	Call a function or a method. The time does not include the CFUT fault or reply time.
Reply with Suspend	8-10	Return a reply to the caller.
Primitive	1-4 for instructions and up to 400 or more for system calls.	Perform a primitive operation such as an addition or a conditional.

Nargs is the number of arguments sent in the application message.

time above 30 instructions per function or method invocation is spent in the operating system¹.

Analysis

A juxtaposition of the main figures from Tables 7-1 and 7-2 reveals that a typical program will spend about 70% of its active time in the operating system and 30% of the time in user code. Furthermore, the program will take about 100 instructions per function invoked, except for tail-forwarded functions which will only take about 25 instructions each. About 20 extra instructions should be added for each method dispatch that the compiler is unable to optimize out. To derive these estimates the following system of accounting is used: the work ascribed to a function invocation consists of all work needed to call the function on the originating node plus all work needed to dispatch the function on the called node, but not including the work done by the called function to call other functions.

Standard Invocations

Each non-tail-forwarded function invocation requires the processing of an ObjectNode call, a function message send, a reply message, and optionally a cfuture fault on the originating node, and a function dispatch and a reply on the called node. Assuming that the average function call has two arguments, the total operating system work for the above activity is:

$$\begin{aligned}
 & \text{ObjectNode} + \text{ApplyFunction} \\
 & + c(\text{CFUT fault} + \text{restarting Reply}) + (1-c)(\text{non-restarting Reply}) \\
 & = 9 + 4 + 69c + 12(1-c) \\
 & = 25 + 57c \text{ instructions.}
 \end{aligned}$$

c is the probability that a cfuture will be referenced before being replaced by the returned value. This probability can vary over a wide range depending on the branching factor of the program call graph. *c* is 1.0 for a recursive factorial program and 0.5 for a recursive fibonacci or rangesum program. If a branching factor between 1 and 2 is assumed, *c* will be somewhere between 0.5 and 1.0; suppose it is 0.75, which results in 68 instructions executed in the operating system per function invocation.

The total user code work is

$$\text{Function call} + \text{Primitives called by function} + \text{Reply with Suspend.}$$

The time spent executing primitives will vary greatly depending on the application; 10 instruction seems reasonable for most cases, although it will be higher if the user program calls Divide or allocates objects. Substituting this number and the average number of arguments yields a total user code work of

¹Tail-forwarded calls are cheaper because the net cost of a tail-forwarded call is one call and no return, which is about 15 user code instructions.

$13 + 10 + 9 = 32$ instructions.

Thus, the total amount of work taken to process one function invocation is 100 instructions, out of which about 10 instructions (the primitives) could be construed as being "useful" work and the rest overhead. This figure does not include any object migration or XLATE miss overhead. These results should not be interpreted as implying that an MDP running Cosmos has a performance 10 times slower than a comparable processor in a sequential computer because sequential computers also have a considerable function calling and parameter passing overhead.

Tail-Forwarded Invocations

Tail-forwarded applications are considerably more efficient. Using the accounting method outlined above results in ascribing

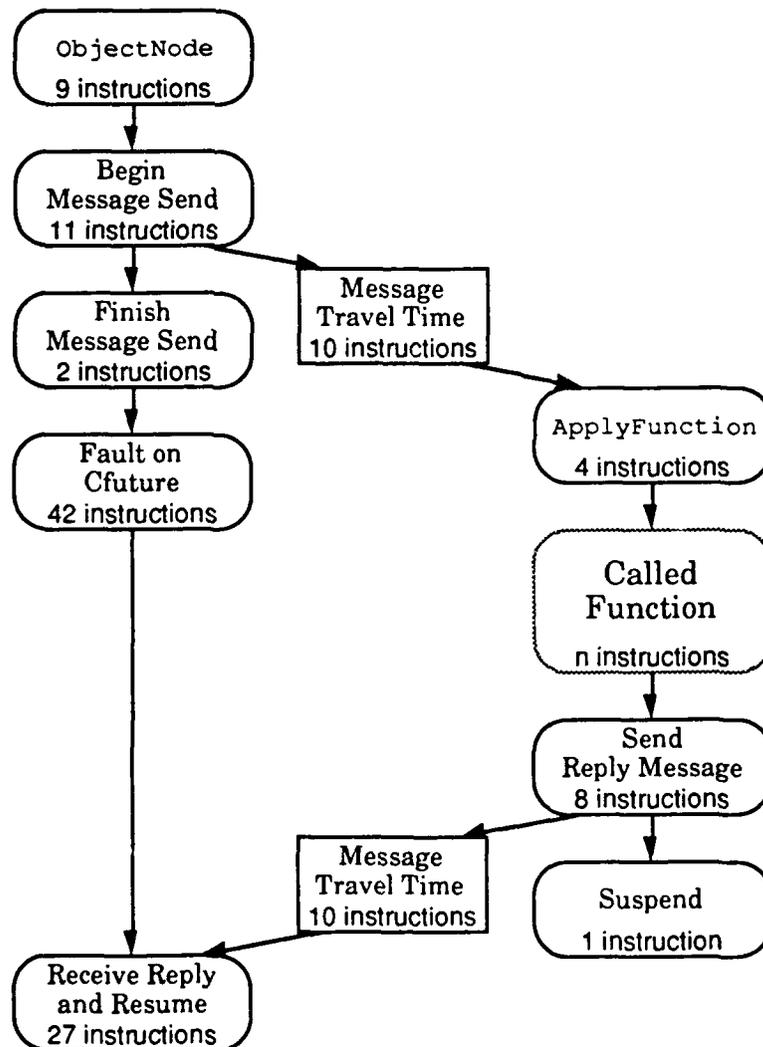


Figure 7-1. Function Invocation Latency

The latency of the network is estimated at about 10 instruction times (20 cycles) to send a message between two randomly chosen nodes on a 4096-node machine.

If n is the time taken by the called function, the latency of invoking a function is $9+11+10+4+8+10+27+n = 79+n$ instructions unless the called function takes fewer than 12 instructions, in which case the latency is $9+11+2+42+27 = 91$ instructions.

ObjectNode + ApplyFunction = 13 instructions

operating system overhead and

Function call + Primitives called by function = 23 instructions

user code work. The total work done is 36 instructions, out of which again 10 instructions is "useful" work.

Latency

The preceding analysis calculated the total amount of work needed per function invocation in a program, which determines throughput on a fully loaded system in which each processor is busy; however, another important component of performance is latency. It turns out that the latency of a function invocation can be lower than the amount of work done by the function invocation because two processors (the caller and the callee) can execute much of the function invocation in parallel.

Assuming no other activity in the system, a non-tail-forwarded function invocation will consist of the caller sending a message to the callee. Then the callee evaluates the function, while the caller takes a cfuture fault (or calls another function, but this won't matter). Unless the called function is very short, the caller will finish the cfuture fault processing and then idle before it gets the reply message from the callee. Finally, the callee replies to the caller, which restarts the calling process.

As can be seen in Figure 7-1, the latency of a function call is 79 instructions in addition to the time taken to execute the function; if the function takes fewer than 12 instructions to execute, the overall latency is 91 instructions. These numbers are less than the total amount of work done by the system (104 instructions).

Summary

The results above indicate that the number of instructions needed to process a function invocation for Cosmos running on a J-Machine should be about 100 instructions, with the notable exception of tail-forwarded functions, which require only about 36 instructions. The instruction counts may be higher if many primitive calls are made or if the operating system faults often.

7.2. Measurements

Grain Size and Machine Load

To attempt to measure the J-Machine's performance and grain size, I ran several programs, including factorial (Figure 7-2); rangesum as listed in Chapter 5; rangesum2 (Figure 7-3), which is a version of rangesum which builds and traverses a data structure; and sort (Figure 7-4), which generates and sorts an array of n pseudo-random numbers using the Batcher parallel sort technique described on page 112 of [28].

```
(defun fact (n)
  (if (zero? n)
      1
      (* n (fact (- n 1)))))
```

Figure 7-2. Factorial Program

```
(defclass pair (object)
  car
  cdr)

(defun cons (x y):pair
  (put-car-cdr (new pair) x y))

(defmethod put-car-cdr pair (x y):pair
  (cset car x)
  (cset cdr y)
  self)

(defun make-countlist (low:integer high:integer)
  (if (> low high) (halt))
  (if (= low high)
      low
      (let ((middle (/ (+ low high) 2)))
        (cons (make-countlist low middle)
              (make-countlist (+ middle 1) high))))))

(defmethod reduce pair (op:funct)
  (op (reduce car op) (reduce cdr op)))

(defmethod reduce integer (op:funct)
  self)

(defun add (x y)
  (+ x y))

(defun reduce-add (tree)
  (reduce tree add))

(defmethod ramp integer ()
  (make-countlist 0 self))

(defmethod rangesum2 integer ()
  (reduce-add (ramp self)))
```

Figure 7-3. Rangesum2 Program

This program exercises several Concurrent Smalltalk object facilities such as allocating objects and traversing trees.

Concurrent Smalltalk on the Message-Driven Processor

```
(defclass distarray (distobj)
  value)

(defmethod initialize distarray (low,high:integer f:funct)
  (if (= low high)
    (cset (get-value (co group low)) (f low))
    (clet ((middle (/ (+ low high) 2)))
      (concurrently
        (initialize group low middle f)
        (initialize group (+ middle 1) high f))))))

(defun make-distarray (n modulus)
  (clet ((da (new distarray n)))
    (initialize da 0 (- n 1) (lambda (x) (mod (* x x x) modulus))
      da)))

(defmethod sort-exchanges distarray (low,high,p,r,d:integer)
  (if (<= low high)
    (if (= low high)
      (clet ((low2 (+ low d)))
        (clet ((v1 (get-value (co group low)))
              (v2 (get-value (co group low2))))
          (if (> v1 v2)
            (concurrently
              (cset (get-value (co group low)) v2)
              (cset (get-value (co group low2)) v1))))))
      (clet ((middle (/ (+ low high) 2)))
        (concurrently
          (sort-exchanges group low middle p r d)
          (sort-exchanges group (+ middle 1) high p r d))))))

(defmethod sort-q distarray (p,q,r,d:integer)
  (sort-exchanges group 0 (- (logical-limit self) (+ d 1)) p r d)
  (if (<> p q)
    (sort-q group p (/ q 2) p (- q p))))

(defmethod sort-p distarray (half,p:integer)
  (sort-q group p half 0 p)
  (if (> p 1)
    (sort-p group half (/ p 2))
    group))

(defmethod sort distarray ()
  (clet ((half (ash 1 (- (integer-length (- (logical-limit self) 1)) 1))))
    (sort-p group half half)))

(defun sort-distarray (n modulus)
  (sort (make-distarray n modulus)))
```

Figure 7-4. Sort Program

Sort-distarray, given the values of n and $modulus$, sorts an array of n pseudo-random numbers. The i th pseudo-random number is equal to $i^3 \bmod modulus$. The Batcher sort algorithm is used, as presented on page 112 of [28].

Measurements were done on a 4-node and a 16-node simulated J-Machine. The results of the trials are summarized in Table 7-3.

The grain size is the third number in the working instructions executed column. The time to process one function invocation is approximately twice the grain size unless tail-forwarding is used extensively. Except for sorting 4 numbers and the trivial factorial case, the results indicate function invocation times of between 81 and 162 instructions, which means that the estimate of 100 in the previous section was about right. Many of the functions in the sort sample program are tail-forwarded, so the average function invocation time for that example is less than twice the grain size. In addition, the sort program has a grain size higher than predicted in the previous section. This is probably due to frequent calls to the multiplication, division¹, and `co` primitives as well as to distribution of large code objects; the grain size does decrease for larger input values.

¹A division by 2 is just a single `ASH` instruction, but the division in `make-distarray` requires a complete `Divide` call.

Comparison with Dataflow

Ellen Spertus made a few performance numbers available for her implementation of dataflow on the J-Machine [34]. I compared her timings with those obtained by Optimist II/Cosmos on the same examples. The program used was the factorial function listed in Figure 7-5.

The dataflow interpreter took 431 steps to compute the factorial of 4. The Concurrent Smalltalk version of the factorial program took 725 steps to execute from a cold start but only 265 steps from a hot start. The dataflow interpreter allocates code statically and references absolute addresses, so every timing is effectively a hot start. The dataflow interpreter took 628 steps to compute three factorials of 4 in parallel, while the Concurrent Smalltalk code took 399 steps to complete the task. Thus, for this simple example the Concurrent Smalltalk/Optimist II/Cosmos combination is faster than dataflow, but not by much. However, Concurrent Smalltalk is more dynamic than the current dataflow system in [34].

```
(defun fact (n)
  (if (<= n 1)
      1
      (* n (fact (- n 1)))))
```

Figure 7-5. Factorial Program used in Dataflow

Network Load

As seen in Table 7-3, the network loading is usually between one word every 8 instructions and one word every 20 instructions, with the earlier figure dominating as the J-Machine utilization approaches 100%. If an average MDP instruction length is taken to be 2.0 cycles, this implies that a program could inject words into the network as fast as one word every 16 cycles on every MDP.

Suppose that we run one of the above programs on a J-Machine organized as a $k \times k \times k$ mesh. Let $N = k \times k \times k$ be the number of nodes. To a first-order approximation, the capacity of the network is $3N$ half-word-hops/cycle¹, or $1.5N$ word-hops/cycle. Assuming random sources and destinations, a message will have to travel an average of $k/3$ nodes on each of the three dimensions, so the expected distance the message has to travel is $3k/3 = k$ nodes. Hence, the network's theoretical capacity is the delivery of $1.5N/k = 1.5k^2$ words per cycle. On the other hand, the program offers $N/16$ words/cycle to the network, which means that unless locality is exploited or the program slowed down, there will be an upper bound on the size of the J-Machine which can run Cosmos.

A mesh loaded at about 30% of its theoretical capacity should be able to route messages without excessive delays [32]. To calculate the maximum k , set

$$0.3 \times 1.5k^2 = k^3/16$$

$$k = 7.2.$$

Thus, the network should not become a critical resource until a J-Machine with over $7^3 = 343$ nodes is built. If the network routing speed is doubled, network loading should not be problematic until the J-Machine exceeds $14^3 = 2744$ nodes. On the other hand, should the Cosmos operating system be sped up somehow, the critical size might fall below 343 nodes. Serious attention to locality will have to be paid if a J-Machine larger than a few hundred nodes is built; conversely, if only a small J-Machine is built, it may not be adequate for testing algorithms for exploiting locality because almost any algorithm will work.

¹The J-Machine network can transmit half a word between every pair of adjacent MDPs on every cycle.

Table 7-6. Memory Access Frequencies

Operating System memory usage:

Reads: 394430 (0.15/instruction, 0.27/working instruction)
 Writes: 152756 (0.06/instruction, 0.10/working instruction)
 Fetches: 2295682 (0.88/instruction, 1.56/working instruction)

Heap memory usage:

Reads: 152262 (0.06/instruction, 0.10/working instruction)
 Writes: 138807 (0.05/instruction, 0.09/working instruction)
 Fetches: 317299 (0.12/instruction, 0.22/working instruction)

Total memory usage:

Reads: 546692 (0.21/instruction, 0.37/working instruction)
 Writes: 291563 (0.11/instruction, 0.20/working instruction)
 Fetches: 2612981 (1.00/instruction, 1.78/working instruction)

3.48 cycles/working instruction

1.87 cycles/working instruction without external RAM

The numbers above indicate the number of memory references (reads, writes, and fetches) done to the operating system (everything except the heap) and heap areas of memory by Sort running on 16 MDPs with an input of 100. The numbers for the other sample programs are similar. The cycles per instruction figures were calculated by adding the instruction frequencies from Table 7-5 weighted by the instruction times together with the memory usage frequencies weighted by memory access times.

The 4096-word internal memory contains all of the operating system data and code and a small portion of the heap (about 2100 words). The rest of the heap (65536 words) lies in slow external memory. When running on a real J-Machine, the sort program will achieve somewhere between 1.87 and 3.48 cycles per working instruction depending on how much of the program and data resides in the internal memory portion of the heap.

Considering that internal memory read, write, and fetch times average 1, 0, and 1/8 cycles¹, respectively, while external memory read, write, and fetch times are 6, 5, and 3 cycles², respectively, a loss of only a factor of two in performance by placing the user program and data in external memory is surprisingly low. The reason for such a low cycles-per-working-instruction figure when the user program and data are in external memory is the high Cosmos overhead. The MDP spends most of its time executing Cosmos code, which decreases the cycles-per-working-instruction number from what it would otherwise have been. For the same reason, changes that would reduce Cosmos overhead at the expense of user program size are undesirable in most cases.

¹The write time is 0 because it is absorbed by the execution of the WRITE instruction—WRITE does not require any extra cycles when writing to memory as opposed to a register. Eight instructions can be fetched in one cycle for an effective fetch time of 1/8 cycle per instruction; the branch instruction cycle counts already include the overhead for fetching the next set of instructions.

²Two instructions are fetched at a time from external memory in 6 cycles, for an effective fetch time of 3 cycles per instruction.

7.3. Conclusion

Context Switching Performance

A large component of the current operating system overhead time is the time taken to save and restore contexts, especially in the CFUT fault handler. One possibility to increase the speed of the CFUT fault handler is to not save data registers and not copy the message upon a CFUT fault [11]. Not saving data registers would reduce the fault handler's time by 4 instructions¹, while not copying the message would reduce it by 6 more instructions. However, these gains would come at a price—the size of the object code would increase because the compiler could not effectively allocate variables to registers; it is not clear whether the savings in the operating system overhead would outweigh the increased time spent executing user code, especially if the user code lies in external DRAM, while the operating system lies in fast internal SRAM.

Summary

Both the derived and measured data indicate that the grain size for running Concurrent Smalltalk on the J-Machine is 50 to 70 instructions. Since most functions involve two messages (one apply message and one return message), the average number of instructions needed to process a function call is between 100 and 140; actually, it is probably closer to 100 because of tail forwarding.

When running entirely from internal memory, the MDP executes one instruction about every two cycles; if user programs and data have to be accessed from external memory, that count increases to about four cycles per instruction. The network load was calculated assuming a fast program (two cycles per instruction) injecting messages into the network at the fastest observed rate (one word every eight instructions) and utilizing 100% of the J-Machine's processors. If the messages are sent randomly under the above conditions, the J-Machine network will saturate when a J-Machine with over 343 MDPs is built. Of course, most programs will not be as fast, but some crafted library routines could impose network loads as high as indicated above. To prevent network saturation, either the network will have to be made faster, the program slower, or some means of exploiting locality invented.

¹The reduction would be 8 instructions if the data registers did not have to be restored by the reply handler; however, it is difficult for the reply handler to distinguish the cases in which it has to restore registers because some unanticipated fault like overflow happened from the cases in which it doesn't; the extra instructions needed to make this decision would make this optimization not worthwhile.

Chapter 8. Future Evolution

Although working Concurrent Smalltalk programs have been demonstrated, the Concurrent Smalltalk programming system is by no means complete. Some suggestions for improvements were discussed throughout the previous chapters—more optimizations could be added to the compiler, distributed objects could be distributed more uniformly, and storage used by free BRAT entries and free standard contexts could be placed back into the heap's free storage pool.

Nevertheless, the possible modifications are by no means limited to the minor ones listed there. The Concurrent Smalltalk programming system is still an evolving research and demonstration vehicle, and many issues still have to be addressed before it becomes a truly general-purpose system. This chapter lists these issues together with potential approaches for addressing them.

The first section lists features that were left out of the Concurrent Smalltalk implementation that are desirable in a full system. These features are useful in many specialized applications, but the system can work without them.

The second section lists the resource management concerns raised by the implementation of Cosmos. These concerns include load balancing, garbage collection, name space reuse, fanout bottlenecks, and parallelism control. A few ideas are suggested about handling the fanout bottleneck and parallelism control problems, but many of these issues are still in the research stage.

The third section outlines a few changes that could be made to the MDP architecture that would improve the performance of Cosmos and compiled Concurrent Smalltalk programs.

Comments

Comments may be placed anywhere in source files. A comment starts with a semicolon (;) and is terminated by the end of the line. Comments are treated as if they were line breaks by the reader.

(define name name)

Top-level Primitive

This primitive defines the first name as an alias for the object specified by the second name. For example, if the second name refers to a global, after this primitive is executed, both names will refer to the same global.

(undef name)

Top-level Primitive

This primitive removes the top-level definition of name, if any. It should be used with caution, as it is possible to bring the system into an inconsistent state using undef.

cset, concurrently, nconcurrently, parallel, nparallel, block, loop, future, or lazy-future statement, let or clet bindings, or some other statement that permits parallel execution without synchronization.

(reply expression) **Macro**
 (reply (continuation expression)*) **Primitive**

The first variant of `reply` evaluates expression and sends its value to continuation. Execution then proceeds with the next statement of the current method, if any. Reply is not strict—it may reply a future or a cfuture. The value of a reply statement is nil.

The second variant of `reply` is used to return values to named continuations. The `reply` takes an even number of arguments; within each pair, the first argument is the continuation name and the second one its value.

(return expression) **Macro**
 (return (continuation expression)*) **Macro**

Return is equivalent to a `reply` followed by an `exit`—the values of the expressions are sent to the caller, and the execution of the method or function terminates subject to the caveats in the `exit` statement description.

(return-value-expected?) :boolean **Function**
 (return-value-expected? continuation) :boolean **Function**

`Return-value-expected?` returns true if the caller of the method or function is expecting a `reply` for continuation (or continuation if continuation is not specified). It is not guaranteed to return false otherwise, so an implementation that always returns true is acceptable.

The associative restricted selectors allow an arbitrary number of arguments; they compile into pairwise invocations of the corresponding methods. The grouping order is not specified.

Methods declared with restricted selectors should not have side effects.

The identities in Table A-3 have been carefully selected to allow efficient implementation of primitive operations without sacrificing functionality. Some identities have been omitted on purpose. For example, * does not have to be commutative in general, nor does (* a 0) have to equal 0. Not requiring these identities allows * to be used to multiply quaternions and matrices.

The restricted selectors not, and, or, and xor may not be distinguishable from lognot, logand, logor, and logxor on all implementations. Redefining these should be avoided; if they must be redefined, only one set should be redefined.

Table A-3. Restricted Selectors

not and or xor lognot logand logor logxor
< <= > >= = <>
neg + - reverse-- * // reverse-// mod reverse-mod
ash reverse-ash integer-length

Table A-4. Identities among Primitive Methods

- + is associative and commutative.
- 0 is an identity for +.
- (- a b) = (reverse-- b a).
- (- a b) = (+ a (neg b)).
- (neg (neg a)) = a.
- * is commutative with scalar constants and associative.
- 1 is an identity for *.
- (* a -1) = (neg a).
- (* a 2^e) = (ash a e).
- (// a (neg b)) = (neg (// a b)).
- (// a 2^e) = (ash a -e).
- (// a b) = (reverse-// b a).
- (mod a (neg b)) = (neg (mod a b)).
- (mod a b) = (reverse-mod b a).
- (ash a b) = (reverse-ash b a).
- (ash 0 a) = 0.
- (ash a 0) = a.
- (not (not a)) = a.
- and, or, and xor are associative and commutative.
- (and a false) = false.
- (and a true) = a.
- (or a false) = a.
- (or a true) = true.
- (xor a false) = a.
- (xor a true) = (not a).
- (lognot (lognot a)) = a.
- logand, logor, and logxor are associative and commutative.
- (logand a 0) = 0.
- (logand a -1) = a.
- (logor a 0) = a.
- (logor a -1) = -1.
- (logxor a 0) = a.
- (logxor a -1) = (lognot a).
- (< a b) = (not (>= a b)).
- (> a b) = (not (<= a b)).
- (= a b) = (not (<> a b)).
- (< a b) = (> b a).
- (<= a b) = (>= b a).
- (= a b) = (= b a).
- (<> a b) = (<> b a).

(class-kind? o:object c:class):boolean

Method

(class-member? o:object c:class):boolean

Method

Class-kind? returns a boolean value that specifies whether the given object is an instance of the given class or one of its subclasses. **Class-member?** is just like **Class-kind?** except that it returns true only if the object is a direct instance of the given class.

(subclass? c1, c2:class):boolean

Method

Subclass? returns true if c1 is a subclass of c2 and false otherwise.

Appendix E. Optimist II Listing

This listing has been removed due to space constraints. For a copy of the source, please send mail to waldemar@ai.mit.edu or billd@ai.mit.edu. A slightly older, printed copy of the source can also be found in the original Master's thesis version of this document.


```

LB_Done:      MOVE     FIP,IP
fltLookupBinding = IP:abs:|fault|unchecked|LookupBinding<<offsetN

-----
:| Delete a binding of R1 in the BRAT.  Halt if no such binding existed.
:| The purgeBinding entry point also purges the binding from the xlate table.
-----
:|
:|Call: deleteBinding
:|Call: purgeBinding
:|
:|In:   R1     Key.
:|
:|Criticality 5.
:|
:|Alters R0.
:|

PurgeBinding:  MOVE     NIL,R0           :Criticality 6. Purge the object's binding from the XLATE table.
               ENTER   R1,R0
DeleteBinding:  MOVE     R2,FOP0        :Criticality 6. Save R2 and R3.
               MOVE     R3,FOP1
               ROT      R1,-BRATLenLog*4,R2
               XOR     R1,R2,R2         :Calculate the hash code for R1.
               ROT     R2,-BRATLenLog*2,R0 :The hash code is the XOR of the four bytes of R1,
               XOR     R2,R0,R2         :the same as the XLATE hash code.
               ROT     R2,-BRATLenLog,R0
               XOR     R2,R0,R2
               MOVE     BRATLength-1,R0
               AND     R2,R0,R2         :R2 contains a hash code between 0 and BRATLength-1.
               DC      BRATstart-2
DB_Next:       ADD     R2,R0,R2
               ADD     R2,2,R0
               MOVE     [R0,A0],R2      :Follow the linked list of BRAT entries starting with
               BNIL   R2,"DB_Halt     :the one in R0. Leave if R0 is NIL.
               EQ     R1,[R2,A0],R3    :Compare the key against R1.
               BF     R3,"DB_Next     :Check the next entry if it doesn't match.
               ADD     R2,2,R2         :Otherwise delete this entry.
               MOVE     [R2,A0],R3
               MOVE     R3,[R0,A0]
               MOVE     [BRATFree,A0],R3
               MOVE     R3,[R2,A0]
               SUB     R2,2,R2
               MOVE     R2,[BRATFree,A0]
               MOVE     FOP1,R3
               MOVE     FOP0,R2
               MOVE     FIP,IP
DB_Halt:       HALT     haltBRATDelete

fltPurgeBinding = IP:abs:|fault|unchecked:PurgeBinding<<offsetN
fltDeleteBinding = IP:abs:|fault|unchecked>DeleteBinding<<offsetN

```


Concurrent Smalltalk on the Message-Driven Processor

```

:#####
:##
:## Control Manager ##
:##
:#####

```

```

-----
:| Execute an Apply, ApplyFunction, or ApplySelector message.
-----

```

```

Apply:      MOVE    [applyFuncnt,A3],R1    ;Criticality 0. Get the funct.
            CHECK   R1,TAG0,R2           ;If it has tag 0, assume it is a selector.
            BT      R2,*ApplySelector
            CHECK   R1,ID,R2             ;If it has tag ID, assume it is a function.
            BT      R2,*ApplyFunction
            HALT    haltApply            ;Otherwise the message was invalid.

ApplyFunction: MOVE [applyFuncnt,A3],R0    ;Criticality 0. Get the function.
                XLATE R0,objectXLATE,A0
                DC    IP:oFunctionCode<<offsetN ;Start executing at the second word of the function.
                MOVE  R0,IP

ApplySelector: MOVE [applyReceiver,A3],R0  ;Criticality 0. Get the receiver.
                PROBE R0,R1                ;Probe it, hoping it is an ID or DID.
                BNIL R1,*AS_Miss
                MOVE  R1,A2                ;If so, point A2 to the instance object.
                MOVE  R0,ID2
                MOVE  [objectHeader,A2],R0 ;Extract the class from the object header.
                WTAG  R0,INT,R0
                ROT   R0,-hdrClassN,R0
                AND   R0,hdrClassM,R0
AS_1:         MOVE  [applyFuncnt,A3],R1    ;Get the selector.
                CALL  lookupMethodU       ;R0 now contains INT: class.
                DC    IP:oFunctionCode<<offsetN ;Go execute the method.
                XLATE R2,objectXLATE,A0
                MOVE  R0,IP

AS_Miss:     CALL  typeOf                  ;Call the rea. class-extraction routine.
            BR    *AS_1

msgApply = Apply<<offsetN
msgApplyFunction = ApplyFunction<<offsetN
msgApplySelector = ApplySelector<<offsetN

```

```

:*****
:## Initialization ##
:*****

;-----
;| Initialize the MDP.
;-----

InitializeMDP: DC ADDR:invalid ;Clear the user address and ID registers.
MOVE R0,A0
MOVE R0,A1
MOVE R0,A2
MOVE R0,A3
MOVE R0,A0B
MOVE R0,A1B
MOVE R0,A2B
MOVE R0,A3B
MOVE R0,A0B
MOVE R0,A1B
MOVE R0,A2B
MOVE R0,A3B
MOVE N1,R0
MOVE R0,ID0
MOVE R0,ID1
MOVE R0,ID2
MOVE R0,ID3
MOVE R0,ID0B
MOVE R0,ID1B
MOVE R0,ID2B
MOVE R0,ID3B
MOVE R0,ID0B
MOVE R0,ID1B
MOVE R0,ID2B
MOVE R0,ID3B
MOVE -1,R1 ;Clear all globals to CFUT:-1.
WTAG R1,CFUT,R1 ;R1 contains CFUT:-1.
DC 64
IMDP_ClrGlobals: SUB R0,1,R0
MOVE R1,[R0,AC]
BNZ R0,"IMDP_ClrGlobals

DC ADDR:Queue1Start<<baseN ;Initialize the queues.
MOVE R0,QH1
DC ADDR:Queue1Start<<baseN (Queue1End-Queue1Start-1)
MOVE R0,QBM
DC ADDR:Queue0Start<<baseN
MOVE R0,QH0
DC ADDR:Queue0Start<<baseN (Queue0End-Queue0Start-1)
MOVE R0,QBM

MOVE N11,R2 ;R2 contains N11.
DC ADDR:XlateStart<<baseN (XlateEnd-XlateStart-1)
MOVE R2,[LimitOverride,A0] ;Initialize LimitOverride.
MOVE R0,TBM ;Initialize the xlate table.
DC ADDR:XlateStart<<baseN
MOVE R2,[FastContextQueue,A0] ;Initialize FastContextQueue.
MOVE R0,A2
IMDP_ClrXlate: DC XlateEnd-XlateStart
SUB R0,1,R0 ;Clear every entry in the table to NIL.
MOVE R2,[R0,A2]
BNZ R0,"IMDP_ClrXlate
MOVE R2,[BRATFree,A0] ;Initialize BRATFree.
DC ADDR:BRATStart<<baseN
MOVE R0,A2
IMDP_ClrBrat: DC BRATEnd-BRATStart ;Clear the BRAT.
SUB R0,1,R0
MOVE R2,[R0,A2]
BNZ R0,"IMDP_ClrBrat
DC FixedHeapStart ;Initialize the heap.
MOVE R0,[HeapStart,A0]
MOVE R0,[FirstFree,A0]
MOVE R0,R3
DC HeapEnd
MOVE R0,[LastFree,AC]
IMDP_ClrHeap: IF FASTSIM
MOVE R1,[R3,A0] ;Clear the heap to CFUT:-1.
ADD R3,1,R3
GE R3,R0,R2
BF R2,"IMDP_ClrHeap
END
MOVE NNR,R2 ;Initialize RandomSeed, and SerialNode.
MOVE R2,[RandomSeed,A0]
DC nodeMask
MOVE R0,[NodeMask,A0]
AND R2,xM,R3 ;Calculate this node's serial number from the NNR value.
ROT R2,-yN,R0
AND R0,yM,R0
ROT R0,xL,R0
OR R3,R0,R3
ROT R2,-zN,R0
AND R0,zM,R0
ROT R0,x1+y1,R0
OR R3,R0,R3
MOVE R3,[SerialNode,AC]
DC ID: (nFastContexts-1)<<serialN
OR R0,R2,R0 ;Initialize LastObjectID and NextDistobjID.
MOVE R0,[LastObjectID,A0]
MOVE R0,R0
MOVE R0,[NextDistobjID,A0]
MOVE R0,[nFastContexts,R3] ;Make nFastContexts fast contexts.
IMDP_MakeFast: SUB R3,1,R3
MOVE R3,[Temp:NITM Context,A0] ;Save the number of fast contexts yet to be made.
ROT R3,serialN,R1
MOVE NNR,R3 ;Put the node number into the context ID.
OR R1,R3,R1

```


- [40] Yokote, Yasuhiko, and Tokoro, Mario. "Experience and Evolution of Concurrent-Smalltalk." *Proceedings of the 1987 Object-Oriented Programming Systems, Languages, and Applications Conference*, October 1987.
- [41] Wills, D. Scott. "Multi-Model Execution on a Fine Grain Message Passing Substrate." MIT Concurrent VLSI Architecture Memo 20, January 1989.
- [42] Wills, D. Scott. "Pi: A Parallel Architecture Interface for Multi-Model Execution." MIT Artificial Intelligence Laboratory Technical Report 1245, July 1990.
- [43] Wulf, William M., Johnsson, Richard K., Weinstock, Charles B., Hobbs, Steven O., and Geschke, Charles, M. *The Design of an Optimizing Compiler*. American Elsevier, New York, 1975.

