

A Framework for End-to-End Verification and Evaluation of Register Allocators

V. Krishna Nandivada¹
Fernando Magno Quintão Pereira²
Jens Palsberg²

¹ IBM India Research Laboratory, Delhi

² UCLA Computer Science Department, University of California, Los Angeles

Abstract. This paper presents a framework for designing, verifying, and evaluating register allocation algorithms. The proposed framework has three main components. The first component is MIRA, a language for describing programs prior to register allocation. The second component is FORD, a language that describes the results produced by the register allocator. The third component is a type checker for the output of a register allocator which helps find bugs. To illustrate the effectiveness of the framework, we present RALF, a tool that allows a register allocator to be integrated into the gcc compiler for the StrongARM architecture. RALF simplifies the development of register allocators by sheltering the programmer from the internal complexity of gcc. MIRA and FORD's features are sufficient to implement most of the register allocators currently in use and are independent of any particular register allocation algorithm or compiler. To demonstrate the generality of our framework, we have used RALF to evaluate eight different register allocators, including iterated register coalescing, linear scan, a chordal based allocator, and two integer linear programming approaches.

1 Introduction

1.1 Background

The register allocator is one of the most important parts of a compiler. Our experiments show that an optimal algorithm can improve the execution time of the compiled code by up to 250%. Although researchers have studied register allocation for a long time, many interesting problems remain. For example, in recent years PLDI (ACM SIGPLAN Conference on Programming Language Design and Implementation) has published several papers on register allocation [2004 (2 papers), 2005 (3 papers), 2006 (2 papers)]. While the essence of register allocation is well understood, developing a high-quality register allocator is nontrivial. In addition to understanding the register allocation algorithm, which can be complex, the developer must also know the internals of the compiler where the allocator will be implemented. For example, public domain compilers such as GCC [2] or SMLNJ [1] and compiler frameworks such as SUIF [16] or

SOOT [32] allow a programmer to implement a new register allocator. However, the programmer has to understand and work with their data structures which are complicated because register allocation affects both the machine specific and machine independent parts of the compilation process. One attempt to address this problem was Tabatabai et al.’s. [3] register allocation framework, implemented in the CMU C compiler. Their framework presents modules (e.g. graph construction, coalescing, color assignment, spill code insertion, etc) that different register allocators might need. However, if the allocator needs other mechanisms than those provided by the framework, the programmer must still deal with the internals of the CMU C compiler. A goal of our work is to completely shield the developer of register allocators from the internal complexities of a compiler.

Debugging register allocators is also a complicated task. Errors may surface in non-trivial ways; sometimes many instructions after the incorrect code. Moreover, the low-level nature of the machine code and its large size makes visual inspection of the register-allocator’s output tedious and error-prone. As a testimony of these difficulties, most recent publications in this field report only static data and not run-time measurements, let alone implementations in industrial compilers. Although static data (number of spills, number of registers used, etc) is important, it does not reflect the behavior of the register allocator in the presence of other optimizations and run time factors.

A few researchers have developed techniques for proving register allocators correct. Naik and Palsberg [22] proved the correctness of the ILP-based register allocator of Appel and George [6]. Ohori [27] designed a register allocation algorithm as a series of proof transformations which is correct by construction. Other researchers have shown how to validate the output of register allocators. Necula [25] presented a translation validation infrastructure for the `gcc` compiler that includes register allocation. Necula’s scheme treats the memory address of a spilled register as a variables, which allows reasoning about its live ranges, although relying on specific characteristics of the `gcc` compiler, such as addressing modes. Leroy [18] formally describes a technique to validate the output of graph coloring based register allocation algorithms. Basically, if the interference graph contains a pair of adjacent temporaries allocated to the same register, the verifier emits an error, otherwise it assumes that the code generated is correct. Andersson [5] and Pereira et al. [28] adopted similar approaches. Huang et al. [17] presented a more general approach which matches the live ranges of values in the original program against the live ranges of machine locations in the register-allocated program. A goal of our work is to use a type system to validate the output of register allocators, and to prove the soundness of the type system itself.

Annotations in the register-allocated code can help validation algorithms. Morrisett et al. [21] used type annotations to help guarantee memory safety, Necula and Lee [26] used more general annotations such as memory bounds, and Agat [4] used type annotations to validate the output of a register allocator.

A goal of our work is to validate the output of register allocators without extra annotations in the target code.

1.2 Our Results

We present a framework for designing, verifying, and evaluating register allocators. Our framework shields the developer from the internal complexities of a compiler, uses a type system to validate the output of register allocators, and does not rely on code annotations.

MIRA and FORD. Our framework centers around MIRA (Mathematical Intermediate representation for Register Allocation), a language for describing programs prior to register allocation, and FORD (FOrmat for Register allocation Directives), a language that describes the results produced by the register allocator. MIRA sources are abstract intermediate representations of programs immediately before the register allocation phase. MIRA descriptions contain architecture and program specific information. The former includes information such as number and classes of machine registers, number of caller-save registers, costs of loads and stores, etc. The latter consists of information such as the program’s control flow graph, use and definition sites of each variable, estimated usage frequency of each instruction, etc. In the context of our framework, the register allocator emits FORD directives that determine the mapping between registers and variables, the additional code inserted to handle spilling, etc. MIRA and FORD can accommodate many of the traditional register allocation algorithms and are simple enough to be easily used by the developer of register allocators.

Type system. We use a type system to verify that the output of a register allocator is correct. Our type system was inspired by Morrisett et al.’s type system for assembly language [21] and can be used with intermediate representations other than MIRA and FORD. A type correct program is guaranteed to have properties such as: (1) pseudos whose live ranges overlap are assigned to different registers, (2) live ranges of the same pseudo always reach a joint point assigned to the same register, and (3) a live register is not overwritten before it is used. We have found that typical errors in the implementation of a register allocator violate these properties. The type checker points out the locations of the register-allocated code where these properties fail. We have proved type soundness for our type system using the Twelf Meta-theorem prover [33]; our proof is available from the Ralf’s webpage <http://compilers.cs.ucla.edu/ralf>.

RALF. Our tool RALF (Register ALlocation Framework) allows a programmer to plug a new register allocator into gcc, without requiring the programmer to know any details of gcc’s implementation. RALF is an extra layer added on top of gcc, acting as a glue between gcc and the plugged-in register allocator. The new register allocator takes a MIRA program as input and gives a collection of FORD directives as output. The main objective of RALF is to be *simple*: when writing a register allocator compatible with our framework, the developer only has to write a program that translates MIRA to FORD. RALF treats the register allocator, which can be implemented in any language, as a black box whose only

purpose is to translate MIRA sources (provided by RALF) to FORD directives (which are fed back to RALF). Given an input program and the plugged-in register allocator, RALF generates a StrongARM binary executable using the new register allocator and the rest of gcc. RALF also implements our type checker for verifying the output of the register allocator. In addition to our own experiments with RALF, we have used RALF in an advanced compiler course at UCLA. As part of a class assignment, each student implemented a different register allocation algorithm to be used with the framework. More about these experiences can be found at RALF’s homepage <http://compilers.cs.ucla.edu/ralf>. Our current implementation of RALF compiles only C code to the StrongARM architecture; however, the MIRA and FORD description languages are designed to be general enough to fit other source languages and architectures. The RISC nature of the StrongARM architecture helps to keep RALF simple. Our implementation, which gets activated by different compiler switches, contains around 5000 lines of C code divided among 125 functions. The proposed framework is not intended for industrial implementations of register allocators, but for fast implementation and testing of research prototypes. The techniques used in our type system can be used also in a production compiler.

Comparison of eight register allocators. We have used RALF to implement and compare eight different registers allocators. The allocators tested range from classical algorithms, such as the usage-count based implementation [12], to novel approaches, such as register allocation via coloring of chordal graphs [28]. In addition of using static data, such as number of variables spilled, we compare the different algorithms by running the produced code on a StrongARM processor.

The remainder of this paper is organized as follows: Section 2 describes a simplified version of MIRA and FORD and characterizes the register allocation problem. Section 3 presents our type system, and Section 4 discusses our experimental results.

2 A Simplified view of MIRA and FORD

Register allocation is the process of mapping a program \mathcal{M} that can use an unbounded number of variables, or *pseudo-registers*, to a program \mathcal{F} that must use a fixed (and generally small) number of *machine registers* to store data. In the remainder of this paper we use *register* in place of *machine register*, and *pseudo* for *pseudo-register*. If the number of registers is not sufficient to accommodate all the pseudos, some of them must be stored in memory; these are called *spilled* pseudos. Following the nomenclature normally used in the gcc community, we call the intermediate representation of programs \mathcal{M} the *Register Transfer Language (RTL)* and we use *Location Transfer Language (LTL)* to describe programs \mathcal{F} . These intermediate representations contain many details that are not relevant for the register allocation process. As shown in Figure 1, MIRA and FORD have been designed to constitute an interface between the register allocator and the RTL/LTL intermediate representations. They facilitate the development of




Fig. 1. Block diagram of our register allocation framework.

register allocation algorithms by hiding from the algorithm’s designer the details of the RTL/LTL representation that are not relevant to the register allocation process. In order to precisely characterize the register allocation problem, we will use a simplified version of MIRA and a simplified version of FORD in this section. We will call them sMIRA and sFORD respectively. For our purposes, a register allocator RA is a black box which takes an sMIRA program P_{sm} and a description of the target architecture, and produces an sFORD program P_{sf} . In this section, we will assume that the architecture specific information is a list of K machine registers $Regs$, and a set of caller save registers $CallerSave \subseteq Regs$. Thus, $P_{sf} = RA(P_{sm}, Regs, CallerSave)$. An actual register allocator would require more information, such as the class of each machine register, the cost of different operations, etc. Such information is given by the concrete specification of MIRA and FORD, which is presented in Appendix A in the full version of the paper, available from <http://compiler.cs.ucla.edu/ralf>.

2.1 Simplified MIRA

sMIRA programs are described by the grammar in Figure 2(a). We adopt an abstract representation of programs. All the operands not relevant to register allocation, such as constants, heap memory addresses, etc, are represented with the symbol \bullet . The only explicit operands are pseudos (p) and pre-colored registers (r, p). The latter are pseudos that have been assigned a fixed machine register due to architectural constraints. In this simplified presentation, we only use pre-colored register to pass parameters to function calls, and to retrieve their return value. Other operands, such as constants and memory references on the heap are abstracted out.

An sMIRA program is a sequence of *instruction blocks*. Each instruction block consists of an address label, represented by L , heading a sequence of instructions (I) followed by a jump. sMIRA programs can use an unbounded number of pseudo registers p . We do not distinguish the opcode of instructions, except *branches* and function calls. Branches affect the control flow, and *function calls* may cause caller save registers to be overwritten, once register allocation has been performed. A function call such as $(r_0, p_0) = \text{call } (r_1, p_1)..(r_s, p_s)$ uses pre-colored pseudos $(r_1, p_1)..(r_s, p_s)$ as parameters, and produces a return value in the pre-colored pseudo (r_0, p_0) . Notice that in sMIRA a call instruction does not contain a label; that is because we support only intra-procedural register allocation. An example of sMIRA program is given in Figure 2(b).

(Programs)	$P_{sm} ::= L_1 I_1; \dots; L_k I_k$		
(Code labels)	$L ::= L_1 \mid L_2 \mid \dots$		
(Instr. Sequence)	$I ::=$		
- (Jump)	$\text{jump } L$		
- (Sequence)	$i; I$		
(Pseudos)	$p ::= p_1 \mid p_2 \mid \dots$		
(Registers)	$r ::= r_1 \mid \dots \mid r_k$		
(Operands)	$o ::=$		
- (Constants)	\bullet		
- (Pseudos)	p		
- (Pre-colored)	(r, p)		
(Instructions)	$i ::=$		
- (Assignment)	$p = o$		
- (Ass. pre-col.)	$(r, p) = o$		
- (Cond. jump)	$\text{if } p \text{ jump } L$		
- (Function call)	$(r_0, p_0) = \text{call}$ $(r_1, p_1)..(r_s, p_s)$		

(a)	<table border="0"> <tr> <td>L_1</td> <td>1</td> <td>$p_0 = \bullet$</td> <td>L_2</td> <td>6</td> <td>$p_0 = p_0$</td> </tr> <tr> <td></td> <td>2</td> <td>$p_1 = \bullet$</td> <td></td> <td>7</td> <td>$(r_0, p_6) = \bullet$</td> </tr> <tr> <td></td> <td>3</td> <td>$p_2 = p_0$</td> <td></td> <td>8</td> <td>$(r_0, p_6) =$</td> </tr> <tr> <td></td> <td>4</td> <td>$\text{if } p_1 \text{ jump } l_3$</td> <td></td> <td></td> <td>$\text{call } (r_0, p_6)$</td> </tr> <tr> <td></td> <td>5</td> <td>$\text{jump } l_2$</td> <td></td> <td>9</td> <td>$p_2 = (r_0, p_6)$</td> </tr> <tr> <td></td> <td></td> <td></td> <td></td> <td>10</td> <td>$p_3 = p_0$</td> </tr> <tr> <td></td> <td></td> <td></td> <td></td> <td>11</td> <td>$p_4 = \bullet$</td> </tr> <tr> <td>L_3</td> <td>15</td> <td>$p_4 = p_0$</td> <td></td> <td>12</td> <td>$p_7 = p_4$</td> </tr> <tr> <td></td> <td>16</td> <td>$p_5 = p_2$</td> <td></td> <td>13</td> <td>$\text{if } p_3 \text{ jump } l_2$</td> </tr> <tr> <td></td> <td>17</td> <td>$\text{jump } \textit{exit}$</td> <td></td> <td>14</td> <td>$\text{jump } l_3$</td> </tr> </table>	L_1	1	$p_0 = \bullet$	L_2	6	$p_0 = p_0$		2	$p_1 = \bullet$		7	$(r_0, p_6) = \bullet$		3	$p_2 = p_0$		8	$(r_0, p_6) =$		4	$\text{if } p_1 \text{ jump } l_3$			$\text{call } (r_0, p_6)$		5	$\text{jump } l_2$		9	$p_2 = (r_0, p_6)$					10	$p_3 = p_0$					11	$p_4 = \bullet$	L_3	15	$p_4 = p_0$		12	$p_7 = p_4$		16	$p_5 = p_2$		13	$\text{if } p_3 \text{ jump } l_2$		17	$\text{jump } \textit{exit}$		14	$\text{jump } l_3$	(b)
L_1	1	$p_0 = \bullet$	L_2	6	$p_0 = p_0$																																																									
	2	$p_1 = \bullet$		7	$(r_0, p_6) = \bullet$																																																									
	3	$p_2 = p_0$		8	$(r_0, p_6) =$																																																									
	4	$\text{if } p_1 \text{ jump } l_3$			$\text{call } (r_0, p_6)$																																																									
	5	$\text{jump } l_2$		9	$p_2 = (r_0, p_6)$																																																									
				10	$p_3 = p_0$																																																									
				11	$p_4 = \bullet$																																																									
L_3	15	$p_4 = p_0$		12	$p_7 = p_4$																																																									
	16	$p_5 = p_2$		13	$\text{if } p_3 \text{ jump } l_2$																																																									
	17	$\text{jump } \textit{exit}$		14	$\text{jump } l_3$																																																									

Fig. 2. (a) Syntax of sMIRA programs. (b) Example of sMIRA program.

(Programs)	$P_{sf} ::= L_1 I_1; \dots; L_k I_k$		
(Code Labels)	$L ::= L_1 \mid L_2 \mid \dots$		
(Inst. seq.)	$I ::=$		
- (Jump)	$\text{jump } l$		
- (Sequence)	$i; I$		
(Pseudos)	$p ::= p_1 \mid p_2 \mid \dots$		
(Registers)	$r ::= r_1 \mid r_2 \mid \dots$		
(Mem. locs.)	$l ::= l_1 \mid l_2 \mid \dots$		
(Operands)	$o ::=$		
- (Constant)	\bullet		
- (Reg. bind)	(r, p)		
- (Mem. bind)	(l, p)		
(Instructions)	$i ::=$		
- (Assig.)	$(r, p) = o$		
- (Store)	$(l, p) = o$		
- (Cond. jump)	$\text{if } (r, p) \text{ jump } L$		
- (Func. call)	$(r_0, p_0) = \text{call}$ $(r_1, p_1)..(r_s, p_s)$		

(a)	<table border="0"> <tr> <td>L_1</td> <td>1</td> <td>$(r_1, p_0) = \bullet$</td> <td>L_2</td> <td>6</td> <td>$(r_1, p_0) = (r_1, p_0)$</td> </tr> <tr> <td></td> <td>2</td> <td>$(r_0, p_1) = \bullet$</td> <td></td> <td>7</td> <td>$(r_0, p_6) = \bullet$</td> </tr> <tr> <td></td> <td>3</td> <td>$(r_1, p_2) = (r_1, p_0)$</td> <td>\uparrow</td> <td>$(l_1, p_0) = (r_1, p_0)$</td> </tr> <tr> <td></td> <td>\uparrow</td> <td>$(l_0, p_2) = (r_1, p_2)$</td> <td></td> <td>8</td> <td>$(r_0, p_6) = \text{call } (r_0, p_6)$</td> </tr> <tr> <td></td> <td>4</td> <td>$\text{if } (r_1, p_1) \text{ jump } l_3$</td> <td>$\downarrow$</td> <td>$(r_1, p_0) = (l_1, p_0)$</td> </tr> <tr> <td></td> <td>5</td> <td>$\text{jump } l_2$</td> <td></td> <td>9</td> <td>$(r_2, p_2) = (r_0, p_6)$</td> </tr> <tr> <td></td> <td></td> <td></td> <td>\uparrow</td> <td>$(l_0, p_2) = (r_2, p_2)$</td> </tr> <tr> <td>L_3</td> <td>15</td> <td>$(r_1, p_4) = (r_1, p_0)$</td> <td></td> <td>10</td> <td>$(r_2, p_3) = (r_1, p_0)$</td> </tr> <tr> <td></td> <td>\downarrow</td> <td>$(r_0, p_2) = (l_0, p_2)$</td> <td></td> <td>11</td> <td>$(r_0, p_4) = \bullet$</td> </tr> <tr> <td></td> <td>16</td> <td>$(r_0, p_5) = (r_0, p_2)$</td> <td></td> <td>12</td> <td>$(r_0, p_7) = (r_0, p_4)$</td> </tr> <tr> <td></td> <td>17</td> <td>$\text{jump } \textit{exit}$</td> <td></td> <td>13</td> <td>$\text{if } (r_1, p_3) \text{ jump } l_2$</td> </tr> <tr> <td></td> <td></td> <td></td> <td></td> <td>14</td> <td>$\text{jump } l_3$</td> </tr> </table>	L_1	1	$(r_1, p_0) = \bullet$	L_2	6	$(r_1, p_0) = (r_1, p_0)$		2	$(r_0, p_1) = \bullet$		7	$(r_0, p_6) = \bullet$		3	$(r_1, p_2) = (r_1, p_0)$	\uparrow	$(l_1, p_0) = (r_1, p_0)$		\uparrow	$(l_0, p_2) = (r_1, p_2)$		8	$(r_0, p_6) = \text{call } (r_0, p_6)$		4	$\text{if } (r_1, p_1) \text{ jump } l_3$	\downarrow	$(r_1, p_0) = (l_1, p_0)$		5	$\text{jump } l_2$		9	$(r_2, p_2) = (r_0, p_6)$				\uparrow	$(l_0, p_2) = (r_2, p_2)$	L_3	15	$(r_1, p_4) = (r_1, p_0)$		10	$(r_2, p_3) = (r_1, p_0)$		\downarrow	$(r_0, p_2) = (l_0, p_2)$		11	$(r_0, p_4) = \bullet$		16	$(r_0, p_5) = (r_0, p_2)$		12	$(r_0, p_7) = (r_0, p_4)$		17	$\text{jump } \textit{exit}$		13	$\text{if } (r_1, p_3) \text{ jump } l_2$					14	$\text{jump } l_3$	(b)
L_1	1	$(r_1, p_0) = \bullet$	L_2	6	$(r_1, p_0) = (r_1, p_0)$																																																																		
	2	$(r_0, p_1) = \bullet$		7	$(r_0, p_6) = \bullet$																																																																		
	3	$(r_1, p_2) = (r_1, p_0)$	\uparrow	$(l_1, p_0) = (r_1, p_0)$																																																																			
	\uparrow	$(l_0, p_2) = (r_1, p_2)$		8	$(r_0, p_6) = \text{call } (r_0, p_6)$																																																																		
	4	$\text{if } (r_1, p_1) \text{ jump } l_3$	\downarrow	$(r_1, p_0) = (l_1, p_0)$																																																																			
	5	$\text{jump } l_2$		9	$(r_2, p_2) = (r_0, p_6)$																																																																		
			\uparrow	$(l_0, p_2) = (r_2, p_2)$																																																																			
L_3	15	$(r_1, p_4) = (r_1, p_0)$		10	$(r_2, p_3) = (r_1, p_0)$																																																																		
	\downarrow	$(r_0, p_2) = (l_0, p_2)$		11	$(r_0, p_4) = \bullet$																																																																		
	16	$(r_0, p_5) = (r_0, p_2)$		12	$(r_0, p_7) = (r_0, p_4)$																																																																		
	17	$\text{jump } \textit{exit}$		13	$\text{if } (r_1, p_3) \text{ jump } l_2$																																																																		
				14	$\text{jump } l_3$																																																																		

Fig. 3. (a) Syntax of sFORD programs (b) Example of sFORD program.




Fig. 4. (a-e) Examples of errors due to wrong register allocation.

2.2 Simplified FORD

A register allocator produces sFORD programs, which are represented by the grammar in Figure 3(a). Operands in sFORD are bindings of pseudos (p) to *machine locations*. In our representation, a machine location can be either a physical register (r), or a memory address (l). In addition to calls and branches, we distinguish loads “ $= (l, p)$ ”, and stores “ $(l, p) =$ ”, because these instructions are used to save and restore spilled values. Notice that we make an explicit distinction between *code labels*, represented by L , and data labels (stack locations), represented by l . In the sFORD representation, caller-save registers can be overwritten by function calls; thus, the register allocator must guarantee that pseudos that are alive across function calls are not mapped to caller-save registers.

Let P_{sm} be the sFORD program in Figure 3(b), and consider an architecture where $\text{Regs} = \{r_0, r_1, r_2\}$ and $\text{CallerSave} = \{r_0, r_1\}$. Let RA be a hypothetical register allocator, such that $P_{sf} = RA(P_{sm}, \text{Regs}, \text{CallerSave})$ is the program in Figure 3. In our example, RA has found the following register assignment: (r_1, p_0) , (r_0, p_1) , (r_1, p_3) , (r_1, p_4) , (r_0, p_5) , (r_0, p_6) and (r_0, p_7) . The pseudo p_2 has been spilled due to the high register pressure between instructions 11 and 12; its memory location is given by the label l_0 . Furthermore, pseudo p_0 has been spilled to memory location l_1 because it is stored in the caller save register r_0 and is live across a function call.

3 Type Checking

Inaccuracies in the implementation of a register allocation algorithm may result in different types of errors in the sFORD program. In Figure 4 we illustrate five different errors that can be produced by a flawed register allocator. In Fig. 4(a), p_0 was defined in register r_0 at instruction 1, but it is expected to be found in register r_1 when used in instruction 2. In Fig. 4(b) register r_0 is overwritten in instruction 2 while it contains the live pseudo p_0 . Fig. 4(c) describes a similar situation, but in this case a memory location is overwritten while the value it holds is still alive. In Fig. 4(d), we assume that r_1 is a caller save register. In this case, pseudo p_1 may have its location overwritten during the execution of the function call in instruction 2. Finally, in Fig. 4(e) the value of p_0 , stored in register r_0 may be overwritten, depending on the path taken during the execution

of the program. This last error is particularly elusive, because its consequences might not surface during the testing of the target program.

In order to guarantee that the values used in the sMIRA program are preserved in the sFORD representation, we use a type checker inspired by [21]. The basic data used in our type system are machine locations, and the type of a machine location is the pseudo-register that it stores. In our case, the register allocator annotates each definition or use of data with its type. A definition of a machine location, e.g. $(r_i, p_j) = \bullet$ corresponds to declaring r_i with the type p_j . Let (r_i, p_j) be an annotated machine location. Intuitively, every time this machine location is used, e.g. $(r, p) = (r_i, p_j)$, its annotated type corresponds to the type that can be discovered by a type inference engine if the sFORD program is correct. For instance, the program in Figure 4(a) is incorrect because p_0 , the type of r_1 in the second instruction cannot be inferred.

3.1 Operational semantics of sFORD programs

We define an abstract machine to evaluate sFORD programs. The state M of this machine is defined in terms of a tuple with four elements: (C, D, R, I) . If M is a program state and we have M' such that $M \rightarrow M'$, then we say that M can *take a step*. A program state M is *stuck* if M cannot take a step. A program state M *goes wrong* if $\exists M' : M \rightarrow^* M'$ and M' is stuck. I is defined in Figure 3(a); the code heap C , data heap D , and register bank R are defined below.

$$\begin{aligned} \text{(Code Heap)} \quad C &::= \{L_1 = I_1, \dots, L_k = I_k\} \\ \text{(Data Heap)} \quad D &::= \{l_1 = p_1, \dots, l_m = p_m\} \\ \text{(Register Bank)} \quad R &::= \{r_1 = p_1, \dots, r_n = p_n\} \\ \text{(Machine State)} \quad M &::= (C, D, R, I) \end{aligned}$$

The evaluation rules for our abstract machine are given in Figure 5. Rules 1, 2 and 3 are to evaluate the operands of sFORD. The assignment statement (Rule 4) modifies the mapping in the register bank, and the store statement (Rule 5) modifies the mapping in the data heap. The result of a conditional branch has no importance in our representation. Therefore, an instruction such as $\text{if } (r, p) \text{ jump } v$ is evaluated non-deterministically by either Rule 6 or Rule 7. We conservatively assume that a call instruction changes the contents of all the caller save registers, and defines a register with the return value (Rule 8). In order to simulate the effects of a function call on the caller-save registers we define the erasing function “ \diamond ” below. We augment the set of pseudo-registers, e.g. $\{p_1, p_2, \dots\}$ with \perp . This pseudo will be used as the type of non-initialized registers and we assume that it is not defined in any instruction of the original sMIRA program.

$$\begin{aligned} \diamond : R \times X &\mapsto R' \\ (R \diamond X)(r) &= \perp \text{ if } r \in X, \text{ else } R(r) \end{aligned}$$

In Figure 5 the set X is replaced by the set of caller-save registers. We draw the attention of the reader to the premises of rules 4 to 8, which ensure that a

location used by an instruction indeed contains the pseudo that is expected by that instruction. For $(r_1, p_1) = (r_0, p_0)$, the premise of Rule 4 ensures that r_0 is holding the value of p_0 when i is executed.

$$D, R \vdash \bullet \quad (1)$$

$$D, R \vdash (r, p), \text{ if } R(r) = p \wedge p \neq \perp \quad (2)$$

$$D, R \vdash (l, p), \text{ if } D(l) = p \wedge p \neq \perp \quad (3)$$

$$\frac{D, R \vdash o}{(C, D, R, (r, p) = o; I) \rightarrow (C, D, R[r \mapsto p], I)} \quad (4)$$

$$\frac{D, R \vdash o}{(C, D, R, (l, p) = o; I) \rightarrow (C, D[l \mapsto p], R, I)} \quad (5)$$

$$\frac{D, R \vdash (r, p)}{(C, D, R, \text{if } (r, p) \text{ jump } L; I) \rightarrow (C, D, R, I')} \quad C_{cond} \quad (6)$$

$$C_{cond} = L \in \text{domain}(C) \wedge C(L) = I'$$

$$\frac{D, R \vdash (r, p)}{(C, D, R, \text{if } (r, p) \text{ jump } L; I) \rightarrow (C, D, R, I)} \quad (7)$$

$$\frac{\forall (r_i, p_i), 1 \leq i \leq s, C, D, R \vdash (r_i, p_i)}{(C, D, R, (r_0, p_0) = \text{call } (r_1, p_1), \dots, (r_s, p_s); I) \rightarrow (C, D, (R \diamond \text{callerSave})[r \mapsto p], I)} \quad (8)$$

$$(C, D, R, \text{jump } L) \rightarrow (C, D, R, I) \quad \text{if } C_{jump} \quad (9)$$

$$C_{jump} = L \in \text{domain}(C) \wedge C(L) = I$$

Fig. 5. Operational Semantics of sFORD programs

Operands

$$\vdash \bullet : \text{Const} \quad (10)$$

$$\frac{\Gamma(r) = p \quad p \neq \perp}{\Gamma \vdash (r, p) : p} \quad (11)$$

$$\frac{\Delta(l) = p \quad p \neq \perp}{\Delta \vdash (l, p) : p} \quad (12)$$

Instructions

$$\frac{\Delta; \Gamma \vdash o : t \quad p \neq \perp}{\Psi \vdash (r, p) = o : (\Gamma \times \Delta) \mapsto (\Gamma[r : p] \times \Delta)} \quad (13)$$

$$\frac{\Delta; \Gamma \vdash o : t \quad p \neq \perp}{\Psi \vdash (l, p) = o : (\Gamma \times \Delta) \mapsto (\Gamma \times \Delta[l : p])} \quad (14)$$

$$\frac{\Gamma \vdash (r, p) : p \quad \Psi \vdash L : (\Gamma' \times \Delta') \quad (\Gamma \times \Delta) \leq (\Gamma' \times \Delta')}{\Psi \vdash \text{if } (r, p) \text{ jump } L : (\Gamma \times \Delta) \mapsto (\Gamma \times \Delta)} \quad (15)$$

$$\frac{\forall (r_i, p_i), 1 \leq i \leq s, \Gamma \vdash (r_i, p_i) : p_i \quad p_0 \neq \perp}{\Psi \vdash (r_0, p_0) = \text{call } (r_1, p_1), \dots, (r_s, p_s) : (\Gamma \times \Delta) \mapsto ((\Gamma \diamond \text{callerSave})[r : p] \times \Delta)} \quad (16)$$

Instruction sequences

$$\frac{\Psi \vdash L : (\Gamma' \times \Delta') \quad (\Gamma \times \Delta) \leq (\Gamma' \times \Delta')}{\Psi \vdash \text{jump } L : (\Gamma \times \Delta)} \quad (17)$$

$$\frac{\Psi \vdash i : (\Gamma \times \Delta) \mapsto (\Gamma' \times \Delta') \quad \Psi \vdash I : (\Gamma' \times \Delta')}{\Psi \vdash i; I : (\Gamma \times \Delta)} \quad (18)$$

Bank of Registers

$$\frac{\forall r \in \text{domain}(\Gamma). \vdash R(r) : \Gamma(r)}{\Psi \vdash R : \Gamma} \quad (19)$$

Data Heap

$$\frac{\forall l \in \text{domain}(\Delta). \vdash D(l) : \Delta(l)}{\Psi \vdash D : \Delta} \quad (20)$$

Code Heap

$$\frac{\forall L \in \text{domain}(\Psi). \Psi \vdash C(L) : \Psi(L)}{\vdash C : \Psi} \quad (21)$$

Machine states:

$$\frac{\vdash C : \Psi \quad \vdash D : \Delta \quad \vdash R : \Gamma \quad \Psi \vdash I : (\Gamma' \times \Delta') \quad (\Gamma \times \Delta) \leq (\Gamma' \times \Delta')}{\vdash (C, D, R, I)} \quad (22)$$

Fig. 6. Type System of sFORD

3.2 Typing Rules

We define the following types for values:

$$\text{value types } t ::= p \mid \text{Const}$$

We define three typing environments:

$$\begin{aligned} \text{(Code heap type)} \quad \Psi &::= \{L_1 : (\Gamma_1 \times \Delta_1), \\ &\quad \dots, L_k : (\Gamma_k \times \Delta_k)\} \\ \text{(Register bank type)} \quad \Gamma &::= \{r_1 : p_1, \dots, r_m : p_m\} \\ \text{(Data heap type)} \quad \Delta &::= \{l_1 : p_1, \dots, l_n : p_n\} \end{aligned}$$

Operands that have no effect on the register allocation process are given the type `Const`. The type of a machine location (register or memory address) is determined by the pseudo that is stored in that location. The environment Γ contains the types of the machine registers, and the typing environment Δ contains the types of locations in the data heap. We will refer to Γ and Δ as *location environments*. The environment Ψ determines the type of each instruction block. The type of an instruction sequence is given by the minimum configuration of the bank of registers and data heap that the sequence must receive in order to be able to execute properly. We consider instructions as functions that modify the location environments, that is, an instruction i expects an environment $(\Gamma \times \Delta)$ and returns a possibly modified environment $(\Gamma' \times \Delta')$. We define an ordering on location environments as follows:

$$\Gamma \leq \Gamma' \quad \text{if } \forall r, r : p \in \Gamma' \text{ then } r : p \in \Gamma \quad (23)$$

$$\Delta \leq \Delta' \quad \text{if } \forall l, l : p \in \Delta' \text{ then } l : p \in \Delta \quad (24)$$

$$(\Gamma \times \Delta) \leq (\Gamma' \times \Delta') \quad \text{if } \Gamma \leq \Gamma' \wedge \Delta \leq \Delta' \quad (25)$$

The type rules for sFORD programs are given in Fig. 6. According to rule 11, the type of a register binding such as (r, p) is the temporary p , but, only if p is the type of r in the Γ environment, otherwise it does not type-check. Similarly, Rule 12 determines the type of memory bindings. Rules 13, 14 and 16 change the location environment. The ordering comparison in the premises of Rules 15 and 17 is necessary to guarantee that all the registers alive at the beginning of an instruction block have well defined types.

None of the programs in Fig. 4 type-checks. In Fig. 4(a), r_1 is not declared with type p_0 , thus the premise of Rule 11 is not satisfied. In Fig. 4(b), the type of r_0 , before the execution of instruction 3, is p_1 , not p_0 , as expected by the type annotation. Again, Rule 11 is not satisfied. Fig. 4(c) presents a similar case, but using a memory location instead of a register: the type of l_0 in instruction 3 is not p_0 , as expected, but p_1 . The type of the used operand would not satisfy the premise in Rule 12. In Fig. 4(d), r_1 has type \perp before the execution of instruction 3, which is different from the expected type p_1 . Finally, in Fig. 4(e) $\Psi(L_2) = ([r_0 : p_0] \times \Delta)$, but the type of instruction 4 is $([r_0 : p_1] \times \Delta) \mapsto ([r_0 : p_1] \times \Delta)$. This would not satisfy the inequality in Rule 17.

3.3 Type Soundness

We state the lemmas and theorems that constitute our soundness proof. Our soundness proof assumes that the sMIRA program has no uses of undefined pseudos. If the sFORD program type-checks, then it preserves the values alive in the original sMIRA code. Proofs can be found in the extended version of this paper.

Theorem 1. (Preservation) *If $\vdash M$, and $M \rightarrow M'$, then $\vdash M'$.*

Lemma 1. (Canonical Values) *If $\vdash C : \Psi, \vdash D : \Delta$ and $\vdash R : \Gamma$ then:*

1. *If $\Psi \vdash L : (\Gamma \times \Delta)$, then $L \in \text{domain}(C)$, $C(L) = I$ and $\Psi \vdash I : (\Gamma \times \Delta)$.*
2. *If $\Delta; \Gamma \vdash o : p$, then $o = r$, or $o = l$. If $o = r$, then $r \in \text{domain}(R)$, else if $o = l$, then $l \in \text{domain}(D)$.*
3. *If $\Delta; \Gamma \vdash o : \text{Const}$, then $o = \bullet$.*

Theorem 2. (Progress) *If $\vdash M$, then M is a final state, or there exists M' such that $M \mapsto M'$.*

Corollary 1. (Soundness) *If $\vdash M$, then M cannot go wrong.*

We have checked the proof using Twelf [33].

3.4 Preservation of callee-save registers

Our type system is intra-procedural, and it assumes that a function call preserves callee-save registers. This must be verified for each procedure, after its type-checking phase, when every instruction has a well know type. If we assume that a machine register is either caller-save or callee-save, this verification step can be done via a simple test. Let L_0 be the label of the first instruction in the procedure, and let L_e be an exit point. Let $\Psi(L_0) = (\Gamma_0 \times \Delta_0)$, and let $\Psi(L_e) = (\Gamma_e \times \Delta_e)$. Callee-save registers are preserved if $\forall L_e, \Gamma_0 \diamond \text{CallerSave} \geq \Gamma_e \diamond \text{CallerSave}$.

4 Experimental Results

Figure 1 presents a high-level block diagram of RALF. RALF interfaces the transformation between the RTL and LTL intermediate representations used by gcc. RALF's front end consists mainly of gcc's parser, gcc's optimization phases, and code to produce a MIRA program from a RTL program. The back end consists mostly of a type checker, code for producing LTL instructions, and gcc's code generation engine. RALF interacts with the implementation of a register allocator via ordinary ASCII files. Given a RTL program P , RALF translates P into a MIRA ASCII program \mathcal{M} , which is then fed to the plugged-in register allocator. The register allocator outputs a set of FORD directives \mathcal{F} , which are then given back to RALF. RALF checks that $(\mathcal{M}, \mathcal{F})$ is a correct mapping via the type system described in Section 3, applies the directives \mathcal{F} on the original

RTL program, and generates gcc’s LTL code. RALF’s back end does not need the original MIRA file in order to produce the LTL program. When producing the RTL code, RALF inserts loads and stores for callee save registers at the entrance and exit of each function. (A smart register allocator might decide to take on that responsibility itself and RALF has an option for that.)

We have tested RALF with eight different register allocators: (1: gcc -O2) the allocator present in the gcc compiler, which has two main phases: (a) aggressive register allocation for local variables within basic blocks, (b) conservative allocation for the whole function. (2: Naive) The *naive* register allocator, which spills all the pseudos. (3: UBC) usage count based register allocator [12], (4: IRC) iterated register coalescing [14], (5: Chordal) register allocation via coloring of chordal graphs [28], (6: LS) linear scan [29], (7: RA) integer linear program (ILP) [24], (8: SARA) stack location allocation combined with register allocation (SARA) [24].

Five of the register allocators have been implemented in Java (2, 3, 4, 5 and 6). Algorithms 7 and 8 have been implemented in AMPL [11]. The interface provided by RALF is extremely simple, and most of the code used to parse and output MIRA/FORD files could be reused among the different implementations. Table 7 compares the size of each implementation; (J) stands for Java, and (A) for AMPL. We do not compare the execution speed of the allocators, because they have been implemented in different languages.

RA	#LOC	
	RA	Interface
Naive	48 (J)	773 (J)
UBC	2766 (J)	773 (J)
IRC	3538 (J)	773 (J)
Chordal	4134 (J)	773 (J)
LS	385 (J)	1100 (J)
RA	495 (A)	298 (A)
SARA	731 (A)	400 (A)

Fig. 7. Comparison between different register allocators plugged on RALF

We have plugged each of the eight algorithms into RALF and then tested the produced code on a StrongARM/XScale processor, with 64MB SDRAM, and no cache. We have drawn our benchmark programs from a variety of sources. We chose these benchmarks in part because, contrary to the most traditional ones (e.g. SPEC), they can be publicly distributed on RALF’s website, which can be accessed at <http://compiler.cs.ucla.edu/ralf>.

- Stanford Benchmark suite: a collection of seven programs that test recursive calls and array indexing.
- NetBench [20]: url is a network related benchmark that implements HTTP based switching; md5 is a typical cryptographic algorithm.

Bench	LOC	RTL	gcc-O2		LS		RA		SARA		Chordal		IRC		UCB		Naive	
			mem	csr	mem	csr	mem	csr	mem	csr	mem	csr	mem	csr	mem	csr	mem	csr
Stanford	307	1082	20	81	171	107	22	63	24	69	34	134	70	100	44	133	854	0
yacr2	3979	10838	1078	289	3035	335	1003	123	1109	142	2121	357	1957	314	2200	361	8181	0
ft	2155	3218	299	130	538	151	225	87	230	106	371	162	561	100	360	169	2184	0
c4	897	40948	187	123	715	301	176	145	179	151	416	170	453	145	394	170	3531	0
mm	885	3388	386	92	2494	68	375	116	380	92	591	93	648	93	687	93	2590	0
url	652	1264	102	54	313	16	120	56	120	58	155	61	201	51	203	63	860	0
md5	790	3464	519	110	1869	433	500	120	500	120	570	228	697	174	580	229	2714	0

Fig. 8. Compile time statistics.




Fig. 9. Comparison of different register allocators using execution time of benchmarks as the metric.

- Pointer-intensive benchmark [8]: This benchmark suite is a collection of pointer-intensive benchmarks. Yacr2 is an implementation of a maze solver and Ft is an implementation of a minimum spanning tree algorithm.
- c4 and mm are taken from the comp.benchmarks USENET newsgroup at <http://www.cs.wisc.edu/~thomas/comp.benchmarks.FAQ.html>. c4 is an implementation of the connect-4 game, and mm is matrix multiplication benchmark.

The number of lines of C code (LOC) and the number of instructions in the RTL for these benchmarks are presented in Figure 8. These benchmarks are non-floating point programs. (We had to edit few of the programs to remove some code that uses floating point operations; we did so only after ensuring that the code with floating point operations is not critical to the behavior of the program.) For each benchmark, Figure 8 presents two static compile time statistics: the number of memory accesses (mem) due to spill/reload instructions, and the number of callee save registers (csr) used by the register allocator (leads to more memory accesses).

The chart in Figure 9 compares the execution times of the programs produced by each of the register allocators. The execution times have been normalized

against the time obtained by programs compiled with gcc at the -O2 optimization level. It can be noted that, although the gcc algorithm is heavily tuned for the StrongARM architecture, Chordal, UBC, and IRC present comparative performances. Also Chordal and IRC’s performances are similar, which confirms the results found by Pereira et al [28]. Figure 9 suggests an upper limit on the gains that any register allocator can make. Even in the most extreme case, the code generated by the naive allocator is worse by a factor of 2.5 (as compared to the optimal solution found by the ILP based allocator). An important point is that most of these benchmarks deal with structures and arrays that require compulsory memory accesses, and it seems that these accesses overshadow the spill cost and hence such a small (2.5 times) improvement. Another conclusion is that the obtained execution time given by optimal solutions (SARA and RA), that run in worst-case exponential time, is not much lower than that obtained by polynomial time heuristics. It should be pointed that our experiments compare specific implementations of the allocators, and are run on a specific target machine; however, it was our objective to be as faithful as possible to the original description of each algorithm.

5 Conclusion

We have presented a framework that facilitates the development of register allocation algorithms. Our framework consists of the two description languages MIRA and FORD, plus a type system. Our framework is easy to use, and versatile enough to support a wide variety of register allocation paradigms. In order to validate this claim, we have developed RALF, an implementation of our framework for the StrongARM architecture, and used it to compare eight different register allocators.

Our framework has several limitations which may be overcome in future work: FORD does not permit the register allocator to modify the control flow graph of the target program; the grammar of MIRA supports only intra-procedural register allocation; the validation algorithm does not handle bitwidth aware register allocations [30]; and FORD does not support the concept of re-materialization. The implementation of RALF itself has the limitations that it targets only the ARM architecture, and RALF does not handle pseudos of type float or double (we have opted for this restriction to keep the implementation simple). So far we have experimented with only the C front-end of gcc; RALF can be seamlessly used with any language supported by gcc.

RALF, our benchmarks, and our Twelf proof are publicly available from <http://compilers.cs.ucla.edu/ralf>. Along with RALF we provide a small collection of tools that we have used during the design and testing of register allocators, such as a data flow visualizer. As a personal testimony, we can say that the use of the type system plus the abstractions provided by MIRA and FORD have saved us significant time and effort during the implementation of several register allocators.

References

1. Standard ML of New Jersey. 2000. [http:// www. smlnj. org/](http://www.smlnj.org/)
2. GNU C compiler. 2005. [http:// gcc. gnu. org](http://gcc.gnu.org)
3. Ali-Reza Adl-Tabatabai, Thomas Gross, and Guei-Yuan Lueh. Code reuse in an optimizing compiler. In *OOPSLA*, pages 51–68. ACM Press, 1996.
4. Johan Agat. Types for register allocation. *Lecture Notes in Computer Science*, 1467:92–111, 1997.
5. Christian Andersson. Register allocation by optimal graph coloring. In *CC*, pages 34 – 45. Springer, 2003.
6. Andrew W Appel and Lal George. Optimal spilling for cisc machines with few registers. In *PLDI*, pages 243–253. ACM Press, 2001.
7. Andrew W. Appel and Lal George. Optimal spilling for CISC machines with few registers. In *PLDI*, pages 243–253, 2001.
8. Todd M. Austin, Scott E. Breach, and Gurindar S. Sohi. Efficient detection of all pointer and array access errors. In *PLDI*, pages 290–301, 1994.
9. G. J. Chaitin. Register allocation and spilling via graph coloring. *SIGPLAN Notices*, 17(6):98–105, June 1982.
10. K.M. Elleithy and E.G. Abd-El-Fattah. A genetic algorithm for register allocation. In *Ninth Great Lakes Symposium on VLSI*, pages 226–, 1999.
11. Robert Fourer, David M. Gay, and Brian W. Kernighan. *AMPL A modeling language for mathematical programming*. Scientific Press, 1993. [http:// www. ampl. com](http://www.ampl.com).
12. R. A. Freiburghouse. Register allocation via usage counts. *Commun. ACM*, 17(11):638–642, 1974.
13. Changqing Fu and Kent Wilken. A faster optimal register allocator. In *MICRO*, pages 245–256. IEEE Computer Society Press, 2002.
14. Lal George and Andrew W. Appel. Iterated register coalescing. *TOPLAS*, 18(3):300–324, May 1996.
15. David W. Goodwin and Kent D. Wilken. Optimal and near-optimal global register allocations using 0-1 integer programming. *SPE*, 26(8):929–968, August 1996.
16. Mary W. Hall, Jennifer-Ann M. Anderson, Saman P. Amarasinghe, Brian R. Murphy, Shih-Wei Liao, Edouard Bugnion, and Monica S. Lam. Maximizing multi-processor performance with the SUIF compiler. *IEEE Computer*, 29(12):84–89, 1996.
17. Yuqiang Huang, Bruce R. Childers, and Mary Lou Soffa. Catching and identifying bugs in register allocation. In *SAS*. Springer, 2006.
18. Xavier Leroy. Formal certification of a compiler back-end or: programming a compiler with a proof assistant. In *POPL*, pages 42–54. ACM Press, 2006.
19. Guei-Yuan Lueh, Thomas Gross, and Ali-Reza Adl-Tabatabai. Global register allocation based on graph fusion. In *Languages and Compilers for Parallel Computing*, pages 246–265, 1996.
20. G. Memik, B.Mangione-Smith, and W.Hu. Netbench: A benchmarking suite for network processors. *IEEE International Conference Computer-Aided Deisgn*, November 2001.
21. Greg Morrisett, David Walker, Karl Crary, and Neal Glew. From system F to typed assembly language. *TOPLAS*, 21(3):527–568, 1999.
22. Mayur Naik and Jens Palsberg. Compiling with code size constraints. *Transactions on Embedded Computing Systems*, 3(1):163–181, 2004.

23. V. Krishna Nandivada and Jens Palsberg. Efficient spill code for SDRAM. In *CASES*, pages 24–31, 2003.
24. V. Krishna Nandivada and Jens Palsberg. Sara: Combining stack allocation and register allocation. In *CC*, pages 232–246, 2005.
25. George C. Necula. Translation validation for an optimizing compiler. In *PLDI*, pages 83–95. ACM, 2000.
26. George C. Necula and Peter Lee. The design and implementation of a certifying compiler. In *PLDI*, pages 333–344, 1998.
27. Atsuchi Ohori. Register allocation by proof transformation. *Science of Computer Programming*, 50(1-3):161–187, 2004.
28. Fernando MQ Pereira and Jens Palsberg. Register allocation via coloring of chordal graphs. In *ASPLAS*, 2005.
29. Massimiliano Poletto and Vivek Sarkar. Linear scan register allocation. *TOPLAS*, 21(5):895–913, 1999.
30. Sriraman Tallam and Rajiv Gupta. Bitwidth aware global register allocation. In *POPL*, pages 85–96, 2003.
31. Omri Traub, Glenn H. Holloway, and Michael D. Smith. Quality and speed in linear-scan register allocation. In *PLDI*, pages 142–151, 1998.
32. Raja Vallee-Rai, Phong Co, Etienne Gagnon, Laurie Hendren, Patrick Lam, and Vijay Sundaresan. Soot - a Java bytecode optimization framework. In *CASCON*, 1999.
33. Frank Pfenning and Carsten Schrmann. System description: Twelf - a meta-logical framework for deductive systems. In (CADE-16), pages 202-206. Springer-Verlag

A Concrete Specification of MIRA and FORD

In this section we describe the concrete syntax of MIRA and FORD. MIRA includes all the data in an sMIRA program, but augments it with architecture specific information. On the other hand, a FORD program is considerably different than the sFORD program that it describes. As we show in Section A.2, a FORD program is a set of mappings that describe how pseudos are mapped to registers and memory locations. It does not include informations about the control flow graph of the sFORD program, because it can be inferred from the MIRA description (or the RTL program that MIRA is abstracting).

A.1 MIRA syntax

The MIRA language describes RTL programs. One of our major concerns when designing MIRA was to hide most of the compiler specific details from the programmer, while including the information required by different register allocators. The language format of MIRA is similar to the language used for the data input in AMPL [11]. A program in MIRA is specified in terms of sets, such as the set of instructions (*Insts*), and integer parameters, such as the cost of a load instruction (*LoadCost*). Some of the sets used in MIRA/FORD can have quadratic time and this may impact on the compilation time; however, the main objective of our framework is not to allow fast execution of a specific compiler, but to facilitate the development and comparison of different register allocators. A MIRA description has two main parts: an *architecture specific* section, and a *program specific* section.

The Architecture Specific Section The architecture specific section of MIRA describes the bank of registers available in the target machine, and the costs associated with particular instructions, such as loads and stores. Given a computer architecture, e.g X86, PowerPC, StrongARM, etc, we let *NRegs* denote the number of machine registers available. A class of machine registers is a set of mutually assignable registers. We denote the number of register classes of a particular architecture by *NClasses*. A machine register is a member of at least one class. MIRA uses the following sets to describe architecture-specific information:

$$\begin{aligned} \text{Regs} &\subseteq \{1..N\text{Regs}\} \\ \text{Classes} &\subseteq \{1..N\text{Classes}\} \\ \text{RegXClass} &\subseteq \text{Regs} \times \text{Classes} \\ \text{RegGroup} &\subseteq \text{Regs} \times X, \quad X \subseteq \text{Regs} \\ \text{RegSize} &\subseteq \text{Regs} \times N \\ \text{CallerSave} &\subseteq \text{Regs} \\ \text{MoveCost} &\in N \\ \text{LoadCost} &\in N \\ \text{StoreCost} &\in N \end{aligned}$$

The classes of each register are determined by pairs in the set *RegXClass*. When an assignment to one register name can affect the value of another, we say that

such register names *alias*. Aliased registers are part of the same *register group*. For example, in the X86 architecture, the first 8 bits of register **AX** also belong to register **AL**; thus, we say that **AX** and **AL** are in the same register group. We abstract the grouping of registers by the **RegGroup** parameter. Registers that can be overwritten by function calls are called *caller-save*. In MIRA these are members of the **CallerSave** set. Finally, the costs of load, store and move instructions, in CPU cycles, are given by the parameters **LoadCost**, **StoreCost** and **MoveCost** respectively. These parameters are useful in the implementation of spilling heuristics. We denote the set of natural numbers by N .

The Program Specific Section We denote the number of instructions in the input RTL program by **NInsts**, and the number of pseudos by **NPseudos**. The following definitions are used to describe program specific information:

$$\begin{aligned}
\text{Insts} &\subseteq \{1 \dots \text{NInsts}\} \\
\text{UsageFreq} &\subseteq \text{Insts} \times N \\
\text{Pseudos} &\subseteq \{1 \dots \text{NPseudos}\} \\
\text{MemLoc} &\subseteq \{1 \dots \text{NMemLoc}\} \\
\text{PseudoXClass} &\subseteq \text{Pseudos} \times \text{Classes} \\
\text{Def} &\subseteq \text{Insts} \times \text{Pseudos} \\
\text{Use} &\subseteq \text{Insts} \times \text{Pseudos} \\
\text{PreColoreds} &\subseteq \text{Regs} \times \text{Pseudos} \\
\text{PrevInst} &\subseteq \text{Insts} \times \text{Insts} \\
\text{Calls} &\subseteq \text{Insts} \\
\text{Moves} &\subseteq \text{Insts} \\
\text{PsSize} &\subseteq \text{Pseudos} \times N
\end{aligned}$$

The set **Insts** represents the instructions, and the set **Pseudos** represents the pseudo-registers in the RTL program. **MemLoc** is the set of memory locations available for the register allocator to store spilled pseudos and **NMemLoc** is the total number of memory locations. By default, MIRA declares a number of memory locations equal to the number of pseudo-registers. The set **UsageFreq** gives the usage frequency of each instruction, that is, an estimate of how many times that instruction will be executed. This parameter can be estimated statically, or via dynamic profile. The class of each pseudo is defined by the product set **PseudoXClass**. The set **Def** relates a pseudo-register to the instructions where it is defined, and the set **Use** relates a pseudo to the instructions where it is used. Pre-colored registers are given in the set **PreColoreds**. These registers permit to model many idiosyncrasies of the target architecture, such as calling conventions and access of global locations. If $(p, r) \in \text{PreColoreds}$, then pseudo-register p must be assigned machine register r . The control flow graph is determined by pairs in the set **PrevInst**; $(i, j) \in \text{PrevInst}$ indicates that instruction i is followed by instruction j . The set **Calls** distinguish instructions that are function calls, and thus, can overwrite the caller-save registers. Move instructions are determined by the set **Moves**. Good register allocators could assign the same machine register to both operands of a copy, so that it can be coalesced. The set **PsSize** gives

the size of each pseudo in bits; for example, if pseudo p has size of 8 bits then $(p, 8) \in \text{PsSize}$.

We omit the description of some parameters that deal with particular register banks (dealing with newer architectures), load-multiple instructions, store-multiple instructions (enables faster memory access), known-loads (load instructions that are already present), known-stores (store instructions that are already present), and inversions ([23]). The full MIRA grammar is publicly available from the website <http://compilers.cs.ucla.edu/ralf>.

An example of a MIRA program The MIRA description produced for the smIRA program in Figure 2(b) is shown below. (For readability, we use labels to distinguish numbers that denote different entities, such as instructions (i), registers (r), pseudos (p), etc. These labels do not exist in an actual MIRA program).

```

/* Architecture specific information */
Regs = {r0, r1, r2}
Classes = {c0}
RegXclass = {(r0,c0) (r1,c0) (r2,c0)}
RegSize = {(r0, 32) (r1, 32) (r2, 32)}
RegGroup = {}
CallerSave = {r0, r1}
LoadCost = 35
StoreCost = 35
MoveCost = 1
/* Program specific information */
Insts = {i1, i2, i3, i4, i5, i6, i7, i8, i9,
i10, i11, i12, i13, i14, i15, i16, i17}
UsageFreq = {(i1, 1),(i2, 1),(i3, 1),(i4, 1),(i5, 1),
(i6, 10),(i7, 10),(i8, 10),(i9, 10),(i10, 10),(i11, 10),
(i12, 10),(i13, 10),(i14, 1),(i15, 1),(i16, 1),(i17, 1)}
Pseudos = {p0, p1, p2, p3, p4, p5, p6, p7}
PsSize = {(p0, 32),(p2, 32),(p3, 32)
(p4, 32) (p5, 32),(p6,32),(p7,32)}
MemLoc = {l1, l2, l3, l4, l5, l6, l7}
PseudoXClass = {(p0, c0),(p1, c0),(p2, c0)
(p3, c0),(p4, c0),(p5, c0),(p6, c0),(p7, c0)}
Def = {(i1, p0),(i2, p1),(i3, p2),(i6, p0),
(i7, p6),(i8, p6),(i9, p2),(i10, p3),(i11, p4),
(i12, p7),(i15, p4),(i16, p5)}
Use = {(i3, p0),(i4, p1),(i6, p0),(i8, p6),
(i9, p6),(i10, p0),(i12, p4),(i13, p3),(i15, p0),
(i16, 2)}
PreColoreds = {(p6, r0),(p7, r0)}
PrevInst = {(i1, i2),(i2, i3),(i3, i4),(i4, i5),(i4, i15),
(i5, i6),(i6, i7),(i7, i8),(i8, i9),(i9, i10),(i10, i11),
(i11, i12),(i12, i13),(i13, i14),(i13, i6),(i14, i15),
(i15, i16),(i16, i17)}

```

```

Calls = {i8}
Moves = {i3, i6, i9, i10, i15, i16}

```

If, instead of the hypothetical machine used in the previous example, the target machine were X86, the architecture specific section of the MIRA program would be the code below. Notice how the `RegGroup` set shows that `AX` is contained within `EAX` and `AL` is contained in both `AX` and `EAX`.

```

/* Machine specific information */
RegisterClasses = {8bits, 16bits, 32bits}

RegXClass = {(AL,8bits) (BL,8bits) (CL,8bits) (DL,8bits)
(AX,16bits) (CX,16bits) (DX,16bits) (BX,16bits)
(EAX,32bits) (ECX,32bits) (EDX,32bits) (EBX,32bits)}
/* We omit AH, BH, CH, DH due to space constraints */

RegGroups = {(AL, {AX, EAX}), (AX, {EAX}),
(BL, {BX, EBX}), (BX, {EBX})
(CL, {CX, ECX}), (CX, {ECX})
(DL, {DX, EDX}), (DX, {EDX})}

CallerSave = {AL AX EAX BL BX EBX}

RegSize = {(AL, 8), (BL, 8), (CL, 8), (DL, 8), (AX, 16),
(BX, 16), (CX, 16), (DX, 16), (EAX, 32), (EBX, 32),
(ECX, 32), (EDX, 32)}

```

A.2 FORD directives

The following sets are used in FORD specifications:

$$\begin{aligned}
RDef &\subseteq \text{Insts} \times \text{Pseudos} \times \text{Regs} \\
RUse &\subseteq \text{Insts} \times \text{Pseudos} \times \text{Regs} \\
Loads &\subseteq \text{Insts} \times \text{Pseudos} \times \text{Regs} [\times \text{MemLoc}] \\
Stores &\subseteq \text{Insts} \times \text{Pseudos} \times \text{Regs} [\times \text{MemLoc}] \\
Moves &\subseteq \text{Insts} \times \text{Regs} \times \text{Regs}
\end{aligned}$$

The sections `RDef` and `RUse` describe the mapping between pseudos and machine registers. A tuple such as (i, p, r) in `RUse` means that pseudo p must be located in register r when used at instruction i . The same tuple in `RDef` would state that pseudo p must be located in register r when defined in instruction i . Load instructions can be specified in the `Loads` section. This section consists of tuples of the form $(i, p, r [l])$, meaning that pseudo p is loaded in register r before instruction i . Optionally, the register allocator may specify the memory label l from which the data should be retrieved. Store information is given in the `Stores` section. This section consists of tuples of the form $(i, p, r [l])$, signifying that pseudo p is stored from register r after instruction i . The optional l

determines the location where the data should be stored. Normally the register allocator does not make assumptions about the exact memory address where spilled values are stored. This should be the default approach of any allocator based on MIRA/FORD. However, our framework allows the allocator to specify a memory label to retrieve/access data. This may be useful, for instance, to coalesce memory references. The register allocator can insert move instruction between registers, for example, to split live ranges of pseudos in order to avoid spilling. Such instructions are declared in the `Moves` section of the FORD file. If $(i, r_1, r_2) \in \text{Moves}$, then r_2 is moved to r_1 before instruction i .

We omit the description of some sections that deal with spill code using load-multiple and store-multiple instructions (enables faster memory access), bitwise operations (to handle exchanging between registers, bit packing etc) and inversions ([23]). The full grammar is at <http://compilers.cs.ucla.edu/ralf>.

An example of a FORD program The register allocation directives for the program in Fig. 3(b) is described by the FORD specification below:

```
RDef = {(i1,p0,r1),(i2,p1,r0),(i3,p2,r1),(i6,p0,r1),
(i7, p6, r0),(i8, p6, r0),(i9, p2,r2),(i10,p3,r2),
(i11,p4,r0),(i12,p7,r0),(i15,p4,r1),(i16,p5,r0)}
RUse = {(i3, p0, r1),(i4, p1, r1),(i6, p0, r1),
(i8, p6, r0),(i9, p6, r0),(i10,p0,r1),(i12,p4,r0),
(i13, p3, r1),(i15, p0, r1),(i16, p2, r0)}
Loads = {(i8, p0, r1, 11),(i15, p2, r0, 10)}
Stores = {(i3,p2,r1,10),(i7,p0,r1,11),(i9,p2,r2,10)}
Moves = {}
```

A.3 Table verification

Before running the type verification engine on an SFORD program, we perform a sequence of simple checks on it, which verify if the sets produced by the register allocator in the FORD file are consistent. They can be executed in $O(1)$ time via simple table inspection. For instance, if $(i, p, r) \in \text{RDef}$, then $(i, p) \in \text{Def}$, $r \in \text{Regs}$, and if $(p, c) \in \text{PseudoXClass}$, then there exists $(r, c) \in \text{RegXClass}$. In this section we give four examples of table checks that can be performed on top of the MIRA/FORD specification.

Unique definition each pseudo p defined at instruction i is assigned a unique register r :

$$\begin{aligned} &\forall i \in \text{Insts} : \forall p \in \text{Pseudos} : \\ &(i, p) \in \text{Def} \Rightarrow \exists r \in \text{Regs} : (i, p, r) \in \text{RDef} \\ &\wedge \nexists s \in \text{Regs} : (i, p, s) \in \text{RDef} \wedge s \neq r \end{aligned}$$

Unique use each pseudo p used at instruction i is assigned a unique register r :

$$\begin{aligned} &\forall i \in \text{Insts} : \forall p \in \text{Pseudos} : \\ &(i, p) \in \text{Use} \Rightarrow \exists r \in \text{Regs} : (i, p, r) \in \text{RUse} \\ &\wedge \nexists s \in \text{Regs} : (i, p, s) \in \text{RUse} \wedge s \neq r \end{aligned}$$

Class matching A pseudo must be assigned a register of its class.

$$\begin{aligned} & \forall i \in \text{Insts} : \forall p \in \text{Pseudos} : \forall r \in \text{Regs} : \\ & (i, p, r) \in \text{RUse} \vee (i, p, r) \in \text{RDef} \Rightarrow \\ & (\exists c \in \text{Classes} : (r, c) \in \text{RegXClass} \wedge (p, c) \in \text{PseudoXClass}) \end{aligned}$$

Pre-colored integrity Pre-colored pseudos cannot be assigned registers by the register allocator. Their mapping is fixed by the architecture.

$$\begin{aligned} & \forall i \in \text{Insts} : \forall p \in \text{Pseudos} : \forall r \in \text{Regs} : \\ & (i, p, r) \in \text{RUse} \vee (i, p, r) \in \text{RDef} \Rightarrow \\ & (\nexists s : (s, p) \in \text{PreColoreds}) \end{aligned}$$

A.4 Case study: Naive Register Allocator

```
function NaiveRegAlloc
input: MIRA  $M$ ,
output: FORD  $F$ 
  for each instruction  $i \in M.\text{Insts}$  do
    for each  $(i, p_1) \in M.\text{Use}$ 
      add  $(i, p_1, r_0)$  to  $F.\text{Loads}$ 
      add  $(i, p_1, r_0)$  to  $F.\text{RUse}$ 
    for each  $(i, p_2) \in M.\text{Def}$ 
      add  $(i, p_2, r_0)$  to  $F.\text{RDef}$ 
      add  $(i, p_2, r_0)$  to  $F.\text{Stores}$ 
```

Fig. 10. Pseudo code for naive register allocator.

In order to give a glimpse of how a register allocator implemented on top of RALF looks, we show the so called naive register allocator. This algorithm loads each pseudo before each use and stores it back after each definition. The algorithm is outlined in Figure 10. This version of the naive algorithm assumes that there exists at most one use per instruction, what is in accordance with the grammar in Figure 2(a). The actual implementation of this algorithm, used by RALF, handles three-address code, where each instruction has at most two uses. The only instruction that uses more than one register is `call`, but the used registers are pre-colored and cannot be assigned locations. Our example assumes an architecture that only has one class of registers; thus, r_0 can be used as the target location of any pseudo.

B Generality of the Framework

The MIRA description language provides the information required by most of the register allocation algorithms currently in use. In order to validate this claim,

	Alg	I/P/L/U/D/Pr	F	J	C	M	X
A	[28]	✓		✓	✓	✓	
	[14]	✓		✓	✓	✓	✓
	[9]	✓				✓	✓
	[19]	✓		✓	✓	✓	
B	Naive	✓					
	[29]	✓				✓	
	[31]	✓				✓	
	[12]	✓		✓	✓		
C	[7]	✓		✓	✓	✓	✓
	[13]	✓		✓	✓	✓	
	[15]	✓		✓	✓	✓	✓
	[24]	✓		✓	✓	✓	
D	[30]	✓		✓	✓		✓
	[10]	✓		✓	✓	✓	

Table 1. Data structures used by each register allocation algorithm.

we show the data required by a number of different register allocators in Table 1. We categorize the data in a format similar to the input interface described in section A.1: I:Insts, P:Pseudos, and L:MemLoc, U:Use, D:Def, Pr:PrevInst, C:Calls, M:Moves and F:UsageFreq. The X column represents features not supported by RALF. We have separated the register allocators into four categories: (A) Graph coloring based, (B) Control flow based, (C) ILP based, and (D) others. It can be noticed that the data provided by our intermediate representations accommodates the necessities of most of the register allocators; however the framework has limitations, which will be addressed as future works: