Generation of efficient interprocedural analyzers with PAG*

Martin Alt Florian Martin

Universität des Saarlandes, Postfach 151150, 66041 Saarbrücken,
{alt|florian}@cs.uni-sb.de

Abstract. To produce high quality code, modern compilers use global optimization algorithms based on *abstract interpretation*. These algorithms are rather complex; their implementation is therefore a non–trivial task and error–prone. However, since they are based on a common theory, they have large similar parts. We conclude that analyzer writing better should be replaced with analyzer generation.
We present the tool PAG that has a high level functional input language to specify data flow analyses. It offers the specification of even recursive data structures and is therefore not limited to bit vector problems. PAG generates efficient analyzers which can be easily integrated in existing compilers. The analyzers are interprocedural, they can handle recursive procedures with local variables and higher order functions. PAG has successfully been tested by generating several analyzers (e.g. alias analysis, constant propagation) for an industrial quality ANSI-C and Fortran90 compiler.
Keywords: data flow analysis, specification and generation of analyzers, lattice specification, abstract syntax specification, interprocedural analysis, compiler construction

1 Introduction

Research in compiler generation has concentrated mostly on front end and lately on back end generation. The optimization phase has not received much attention. Only a few systems [21, 22, 25] for generation of analyzers were designed and built. All of them apply ad-hoc methods or heuristics if the language has subroutines. We present a new generative system for interprocedural analyses, PAG, that is able to produce analyzers which can be applied in several different compilers by instantiation of a well designed interface. The system is based on the theory of abstract interpretation. The philosophy of PAG is to support the designer of an analyzer by providing three languages for specifying the data flow problem, the abstract domains, and the compiler interface. This simplifies the construction process of the analyzers as well as the correctness proof and it results in a modular structure. The specifier is neither confronted with the implementation details of domain functionality nor with the traversal of the control flow graph or syntax tree nor with the implementation of suitable fixpoint algorithms. In the paper, we first briefly summarize the theory of the data flow analysis in section 2, then we introduce the structure of the generating system

* Partially funded by the **ESPRIT** Project #5399 (COMPARE)

in 3. Section 4 discusses the specification languages, followed by the presentation of our interprocedural solution mechanism in section 5. Some measurements are presented in section 6. Finally, section 7 summarizes the related works and exhibits some possible extensions.

2 Theoretical Background

The data flow analysis practiced nowadays was introduced mainly by [12] and refined by [11]. It is based on a *control flow graph (CFG)* that contains a node for every statement or basic block in a procedure and an edge for a possible flow of control. Furthermore we add a unique entry node s and exit node e and a labeling function that yields a syntax tree fragment for every node. A *data flow analysis problem (DFP)* or *instance of a data flow analysis framework* is a combination of such a graph with a complete lattice of values, called the *underlying lattice*, and a family of functions (one for each node). These functions express the local semantics and are therefore called *transfer functions*. If every transfer function is monotonic the problem is called a *monotone problem*. If they are even distributive it is a *distributive problem*.
To describe the solution of a data flow problem we'll first define the semantics of a path $\pi = n_1, \ldots, n_k$ in the *CFG*:

$$[\![\pi]\!] = \begin{cases} id & \text{if } \pi = \epsilon^1 \\ [\![(n_2, \cdots, n_k)]\!] \circ [\![n_1]\!] & \text{otherwise} \end{cases}$$

The desired solution of the *DFP* is union of the semantics of all paths applied to bottom, historically called the *meet over all paths solution*:

$$\mathbf{MOP}(n) = \bigsqcup \{[\![\pi]\!](\bot) \mid \pi \text{ is a path from } s \text{ to } n\}$$

for every node n of the *CFG*, where $\bot$ is the bottom element of the lattice. As the set of all paths from s to n is usually infinite, this solution is in general not computable. Therefore, the *minimal fixpoint solution* was introduced:

$$\mathbf{MFP}(n) = \begin{cases} [\![s]\!](\bot) & \text{if } n = s \\ [\![n]\!]\left(\bigsqcup \{\mathbf{MFP}(m) \mid m \text{ predecessor } n\}\right) & \text{otherwise} \end{cases}$$

Kam has proved in [11] that for every monotone data flow problem the **MFP** is greater (with respect to the ordering of the lattice) than the **MOP** solution, and therefore a save approximation. Moreover, if the *DFP* is distributive, both are equal. The interprocedural version of this theorem is presented in [13].

To solve a data flow problem, different iterative solution algorithms can be used which are guaranteed to terminate if the problem is monotone and the underlying lattice has no infinite ascending chains. Some of them, like the worklist algorithm, are presented in [9]. Other algorithms are the bounded fixpoint iteration from [18] or the higher order chaotic iteration sequences in [19]. If the chains in the lattice are infinite or if a speed up of the fixpoint operation is needed, widening and narrowing, explained in [5], is an appropriate method to solve the equation system.

[1] ϵ denotes the empty path

3 The System

So let's recall what's needed for a data flow analyzer:

1. a complete lattice D;
2. a meta transfer function which, if applied to a node of a CFG, yields a (monotone) transfer function from D to D;
3. a fixpoint algorithm with the appropriate data structures, e.g. an iteration algorithm and a working list;
4. a control flow graph, which is the input for the analyzer.

Some of these parts do not change much from one analyzer to another. To shorten the time of implementation and to simplify maintenance of the resulting system, we designed a special purpose programming language to describe (and create) such analyzers. By freeing the compiler designer from routine matters, he can focus on the crucial details. Our support language needs three specification components: the underlying lattice, the transfer functions and the solution method. This idea is not new. There have been a couple of experiments of generating analyzers, but only a few of them were successful. On one hand, the specification language should be powerful enough to cover a large class of problems, on the other hand the generated code must be sufficiently efficient if it is intended to work in a real compiler and not only to be a prototype for performing some tests.

3.1 Overall Structure

PAG has been designed to generate analyzers for compilation systems. The basic assumption is that such a compilation system offers an interface to the syntax tree and the control flow graph of the program to be analyzed. This representation of the intermediate structures can be found in nearly any compiler.

The interface to the control flow graph has to offer functions to walk over the graph, to get the identifier of a node, and to fetch the syntax tree which is the label of the node. To decouple the implementation of data structures of the graph from their logical functionality, the compiler designer has to write an interface which can be used for all generated analyzers working with this compiler. It is necessary to know the structure of the abstract syntax tree. It can be best defined by an extended tree grammar. From this grammar, the appropriate access and walk functions can be generated. Therefore PAG takes as input a file which describes the structure of the possible trees. Furthermore there are input files that construct the underlying lattice, the transfer functions and global parameters of the analyzer like the used fixpoint algorithm. All these input files are described in detail in section 4. An overview of the different input files to PAG is shown in figure 1.

4 The Specification Language

A PAG specification is divided in four parts: one for the definition of data types in the language DATLA, another for the description of the syntax trees, a third

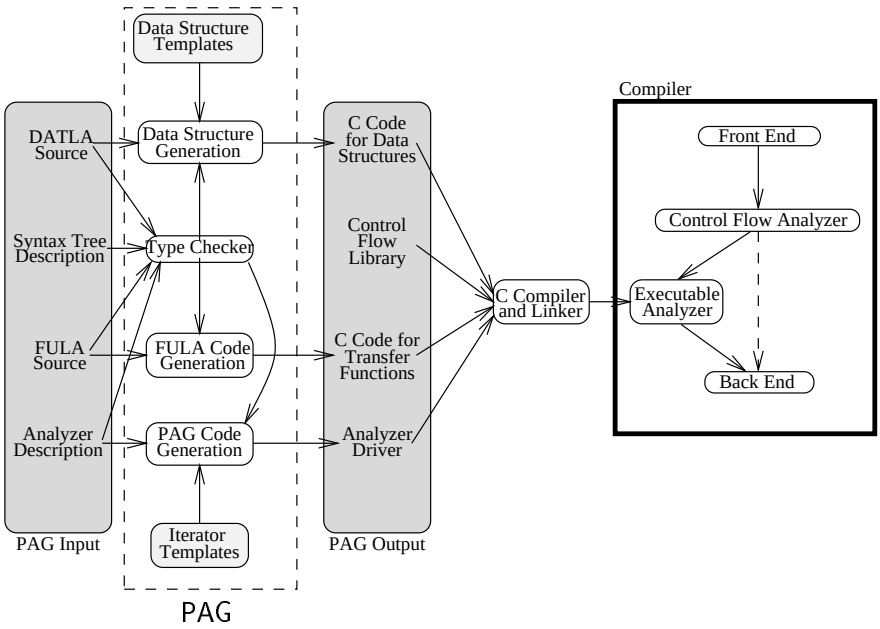

Fig. 1. The Structure of PAG

to specify the main structure of the analyzer, and a last to define the transfer functions in a language called FULA. This whole description is compiled into ANSI C code that performs the analysis outlined in section 5. The structure of these four parts will be shown in detail now. First, we present some notational conventions. We use `typewriter` style for reserved words, and *italic* style for nonterminal symbols. Among these 'v' is used for variables, 'p' for patterns, 'e' for expressions in general and 'b' for boolean valued ones. Finally '{}' is used for grouping and 'a^+' and 'a^*' for the usual repetitions. A complete example of an analyzer specification is given in appendix A. It is the description of an analyzer to detect assignments to live variables.

4.1 The Domain Specification Language

It is used to define the data types and data manipulation routines for the analyzer. Therefore the language is called DATLA, an acronym for data type definition language. We distinguish between sets and domains (complete lattices). They are constructed bottom up. This means that simple sets like numbers or enumerated sets build the base for more complicated sets like the set of all functions from numbers to truth values, which can be used as input for other set generating operators again. So we have some basic sets, and set operators. In addition there are operators to construct domains from sets, and to combine domains into new domains.

The specification is divided into two subparts: one for the set definitions

starting with the keyword `SET` and the other for the domain definitions with the keyword `DOMAIN`. Each definition is an equation with a single name that has to be defined on the left side and an operator applied to a couple of names on the right side. The predefined sets are up to now: `snum`: the set of signed integers, `unum`: the set of unsigned integers, `real`: the set of all floating point numbers, `bool`: the set of truth values `true` and `false`, `chr`: the character set (ASCII), and `string`: the set of all character sequences. This list can easily be extended even with user-defined types. These need to have implementations of every function the interface contains. Another possibility to generate a basic set is to enumerate a finite number of elements. Complex sets can be formed out of the basic ones through the following operators:

1. disjoint union of a finite number of sets: Disjointness is obtained by tagging the element of the sets participating in the union.
2. construction of the tuple space of a finite number of sets
3. power sets: building the set of all subsets of the original set
4. set of lists of elements of a set
5. set of functions from a set S_1 to a set S_2

In the second part of the specification, domains can be formed in one of the following ways:

1. enumeration of elements and of a (complete) partial order. The system checks whether it forms a lattice.
2. flattening of a set S: S is transformed into a domain by adding a least element $\bot$ and a greatest element $\top$. The elements of S are pointwise incomparable.
3. lifting of a domain D: new top and bottom elements are added which are greater respectively smaller than all other elements. The ordering of D is preserved.
4. building the power domain of a set or domain: the generated ordering is set inclusion.
5. construction of the tuple space of a finite number of domains. Tuples are ordered component–wise.
6. building the domain of functions from a set into a domain with point–wise ordering.

In contrast to other generators, PAG allows simultaneously recursive definition of sets and domains. In this way, for instance, the tree type can be easily expressed. But this and other features in DATLA result in the fact that infinite sets and domains, even with infinite chains, can be defined. Therefore the user has to make sure that only a finite part of the domain is used, or has to guarantee the termination of the analysis in another way. Environments are an example where only a finite part of some infinite domain is used. It is typical for program analysis that at each program point, some kind of information is stored for every variable that occurs in the program. This is usually modeled by a function from variables to the needed information. Obviously this domain has infinite chains because the set of variables is infinite. But in every program only a finite subset of variables is used, and so termination is guaranteed.

4.2 The Data Flow Description Language

The main part of the description of an analyzer is the specification of the data flow functions. For every node in the control flow graph, there has to be a function that transforms an incoming flow value into an outgoing one. This is expressed in a functional language FULA which was designed especially for that purpose. FULA programs can be compiled to efficient C code.

Overview: FULA is a first order functional language with eager evaluation. It has static scoping rules and the user defined types from DATLA. Interestingly, the language does not provide an explicit fixpoint operator. As we will see, the absence of an embedded fixpoint operator and the eager evaluation semantics do not restrict significantly the language constructs that PAG is able to analyze. In FULA every expression has a unique type that can be derived statically by a type inference algorithm. There are no implicit type casts in FULA. Any change of the type must be made explicit. For occurrences of a variable the corresponding definition is the syntactically innermost definition for that variable. Binding constructs are function definitions, `case` and `let` expressions.

Transfer Functions: The whole FULA source is split in two parts: one for definitions of auxiliary functions and one for the transfer functions. These transfer functions are written in a special notation: they don't need a name and are defined via patterns matching the labels of the control flow graph. They have an implicit parameter named `@` which is the incoming data flow value and they have to return a data flow value again.

Functions: There are two different types of functions in FULA. Firstly the functions defined in the language itself, and secondly those declared in DATLA. The latter are seen as datatypes and can be arguments to the first sort. This distinction is made because of the first order character of the language. So it is possible to write a (FULA) function, that takes a (DATLA) function as argument. For example a function that modifies the value of an environment for a given variable. Functions can only be defined with a name, which means that there is no lambda expression. Definitions are made by using patterns, and the cases can be spread over the whole specification. A single case looks like $f(p^*) = e$. For each function there can be an additional type definition. Nested functions are not allowed in FULA.

Patterns: They are used to define functions and in `case` expressions. Only *linear* patterns are usable, which means that the same variable is allowed only once per pattern. Pattern expressions may be nested. The following patterns are defined, while some static semantical restrictions, like type conditions, are to be followed:

1. constants: these can be elements of predefined sets like integers as well as user defined constants of enumerated sets or domains. Two special constants are $\bot$ and $\top$, the bottom and top elements of all lattices.
2. variables, but each variable once per nested pattern.
3. empty list and empty set. There are two types of empty lists: those defined in DATLA and those introduced by the syntax tree (see 4.4). They are notated as `[]` or `[!]` and {} for the empty set.

4. cons patterns for both types of lists: p`:`p.
5. tuples: (`,`$\cdots$). Notice that a tuple pattern matches always if the sub patterns match.
6. wild card: it is denoted by an underbar `_` and matches every input.

Expressions can be:

1. constants (built-in and enumerated)
2. if expressions: `if` b `then` e_1 `else` e_2 `endif` for conditionals. This can be seen as the only non eager construct in DATLA, because e is evaluated first, and then e_1 or e_2 but never both.
3. let expressions: `let` $\{\ v_i = e_i;\}^+$ `in` e to introduce a number of variables v_i local to the expression e. The eagerness results in the fact that every e_i is evaluated before e is evaluated.
4. case expressions: `case` $\{v_j\}^+$ `of` $\{\{p_j\}_i^+ \Rightarrow e_i;\ \}^+$ `endcase` gives the possibility to examine the structure of one or more expressions bound to the variables v_j. The result of this expression is the value of the first e_i for which all p_j's match the values of the v_j's
5. function application:both kinds of functions DATLA and FULA can be applied to zero or more expressions. The number of expressions applied to must correspond to the arity of the function because functional expressions that would result if a function is applied to more or less arguments than defined, are not allowed in a first order language.
6. print expressions: `print` (e_0) `++` e is equivalent to the expression e but with the side effect that e_0 is printed out.
7. built in function application (pre- and infix notation). These are functions that are generated for certain data types. All these functions are listed below (for a more refined documentation see [16]):
   - for every type: equal `=` and not equal `!=`
   - for every domain additionally: `lub`, `glb`, $\sqsubseteq$, $\sqsupseteq$, $\sqsubset$, $\sqsupset$
   - `unum`, `snum` ,`real`: mathematical operations like: plus `+`, minus `-`, times `*`, integer division `/`, modulo `%`, amount $|\ .\ |$; bitwise operations: and `&`, or `|`, xor `^`; comparison functions: greater $>$, greater equal $>=$, less $<$, less equal $<=$; cast functions: `snum`, `unum`, `chr`, `real`
   - `bool`: and `&&`, or `||`, not `!`
   - `chr`: cast functions `snum`, `unum`, `real`
   - `string`: concatenation `.`, character selection [.]
   - flatted, lifted domains: `lift`, `drop`
   - lists: cons `:`
   - sums: `is-`*type_name*, `down-`*type_name*, `up-`*type_name* to lift an element of a sum component into the sum and vice versa.
   - tuples: select `!`, tuple construction $(e_1, \cdots, e_n)$
   - power sets: add an element `~`, subtract an element `#`, member test `??`
   - functions: application .{ . }, value changing .[. \ .], creation with default element [$-$ $>$.]

4.3 Analyzer Description

In order to generate analyzers, some additional declarations are necessary:

1. `direction`: values can be `forward` or `backward`. This specifies whether the data flow is along the edges of the control flow graph or in the opposite direction.
2. `carrier`: the name of the domain that is used for the analysis; i.e. the type of the flow values.
3. `combine`: the name of a FULA function, that merges information which comes over different edges. This is usually the function `lub` for the least upper bound.
4. `init`: the initial value that is associated with every control flow node. Usually the neutral element of the combine function ($\perp$).
5. `init_start`: an initialization value for the start node (the end node for a backward problem).

4.4 The Interface Description

This part of the specification is used to define the structure of the syntax tree which is constructed from the source program by the compiler. With this definition, the PAG compiler is able to generate access functions for the tree, following certain naming conventions. This is needed because the generated analyzer must be able to walk over the syntax tree to determine for instance the instruction given on a control flow node. The form of the tree is introduced as a tree grammar with two additions. Firstly, there is a notation to introduce lists of nonterminals in order to gain more simplicity in notation and handling. Secondly we have a notation *nonterminal* `==` *simple_type* to identify a class of nonterminals with a built in type. For such a rule, PAG generates functions to cast an element of *nonterminal* into an element of *simple_type*. Generally the nonterminals of the syntax tree and lists of them are considered as types that are usable in FULA. The concrete syntax is quite obvious, and is skipped here with a pointer to the PAG reference manual [16].

Another part of the interface is that to the control flow graph. But this interface is quite simple and fixed irrespective of the programming language. So the control flow graph is abstracted by a library that has to be supported by the compiler. It includes functions to fetch the successors or ancestors of a node, or to access the corresponding syntax tree fragment.

5 Interprocedural Analysis

In the previous sections, we have focused our interest on intraprocedural analysis. Yet good programming style requires a lot of small procedures that are used at many program points. So it is a must not only to look inside every procedure on its own but at the whole program in order to achieve good analysis results. As already mentioned, the number of problems with interprocedural analysis is large: to achieve excellent results very much time is needed, and fast algorithms

do not find out everything possible. So one has to find a compromise between time consumption and precision.

5.1 Fundamental Algorithms

There are different methods to handle procedures in data flow analysis. We shortly summarize some techniques:

- Non–recursive procedures can be inlined, and the intraprocedural algorithm can be applied. This is only applicable for small programs because the transformed program may grow exponentially.
- Procedure calls may be considered as ordinary statements that make all available information invalid. Or destroy only all information based on those objects that may be modified by that call; this assumes a previous analysis phase that computes for each procedure the set of potentially modified variables (data flow problem over call graph).
- Use a two-level algorithm: first compute an abstract form of the procedure (called *effect* or *jump* in [8]), which maps flow information from the beginning of a procedure to its end; then do standard iteration and use the related effect function at any procedure call; the computation of the effect functions brings up new restrictions on the underlying lattice, as they cannot be computed in general or even if they can, the representation may grow exponentially ([2]) This method is especially useful if the lattices are finite. Many restricted versions of the constant propagation relay on that, e.g. constant copy propagation [14] or even demand-driven versions [7].
- Do an analysis within the procedure body: one variant of this method summarizes the effects of all calls at the beginning of a procedure and another keeps different calls separate. This approach is also known as *call string* approach and is described in [20].

5.2 Our Approach

Our approach tries to achieve the following goals:

- to guarantee termination for every input and every analyzer that terminates intraprocedurally;
- to keep the additional effort of the compiler designer for the interprocedural analysis low;
- time and space complexity should be as low as possible to guarantee the practical usability of the generated analyzer;
- the results from the interprocedural analysis should be as good as possible.

To get a precise solution (this means the **MOP** solution) for the interprocedural case for all problems that are terminating intraprocedurally is not possible. As stated in [20] this is only possible if the domain is finite or the DFP is distributive or if there is no recursion in the program. All three cases restrict our mentioned goals too much. The approach we have chosen is similar to the call string method, because it offers a flexible technique that can niftily trade time and space for precision (see section 5.4). One can observe that in most 'every day live' programs

calls to a procedure from different call sites[2] are made with distinct values. And in addition to that, a large number of function calls are non–recursive. So the different call sites are worth the effort to analyze them separately. Furthermore it can be useful to keep deeper levels of calls separated. This will be clarified by the following example:

```
static int i;                r() {                    q() {
main() {                        q();                     i += 1;
   i = 1;                    }                        }
   r();
   r();
}
```

If we separate in a constant propagation analyzer (compare [23]) every call chain of length up to two, we can figure out that the call of `q` in `r` is made twice with different values. The separation of the calls is achieved by introducing arrays or vectors of data flow values for each node of the control flow graph instead of using only one value per node. In the different array elements we can store the data that comes from different call sites. If there are dynamically more calls than there are fields in the vector (which has a static length) they will be merged together with the (specified) combination operation.

After the analysis has finished, an extra advantage of this vector is that it can be used in different ways: for a simple optimization inside the procedure the meet of all elements can be used because that is the flow value which is valid for all incarnations of the procedure. But it can also be used to determine if it is useful to specialize the procedure for certain values.

5.3 Formal Description

Definition1 (Program Representation).
We represent a program P with the procedures $P_0, \ldots, P_n$, where P_0 is the main procedure as a **super graph** $G^* = (N^*, E^*, s_0, e_0)$. G^* consists of a number of intraprocedural control flow graphs $G_0, \ldots, G_n$ which represent the procedures of P. In difference to the standard CFGs, every call in G^* is represented by three nodes: a call site node, a return site node, and a local node. The call node has one edge to the entry node of the called procedure and another to the local node. The return node has the exit node of the called procedure and the local node as a predecessor.

Figure 2 is an example of a super graph. The local node between call and return node is useful to model the behavior of local variables: the compiler designer can specify that after the local node the flow value is constantly bottom[3]. Then the implicitly performed combine operation before the return node will always yield the value returned from the procedure. But in the constant propagation for example it is also possible to set all global variables at the local node and all variables local to the call at the procedure exit to bottom. Then the local

[2] i.e. different places in the source code
[3] if the combine function is specified as a meet

variables arrive at the return node directly from the call and all other variables are coming from the procedure exit. This is correct if none of the local variables can be modified (due to aliasing) in the procedure. Therefore the concept of the local node results in a large variety of possibilities to handle local values for the specifier of the analyzer; the interface to the abstract syntax tree allows even to access information that has been computed by other compiler parts; e.g. alias graphs.

There are two problems with the construction of a super graph: if a call is made to a function variable the called procedure is not known. Then, there must be an edge to the entry of every (potentially called) procedure. The second one arises if the called procedure is not known because it's not in the source code like any kind of library function. Then we introduce a dummy procedure with which the unknown call is connected. The analyzer designer has to specify worst case assumptions for that dummy procedure, or can do better if he knows the behavior of the library function. Thus, PAG supports unlike all other generators a specification mechanism for the semantics of libraries.

In a super graph we can carry out a standard intraprocedural data flow analysis. So we don't have any additional expenses for computing the effect of the functions. As a refinement we should use only interprocedural valid paths in the graph. A path is called a *valid interprocedural path* if it contains only matching pairs of call and return nodes of a procedure like a correctly braced expression. Exact definitions can be found in [14, 15].

Now we can introduce the discrimination of the data of different call sites by assigning not only a single value of the lattice with every control flow node, but an array of values. It's clear that the length of these arrays in every procedure should be the same for each of its control flow nodes, and at least one.

Definition 2 (Graph Extension).
For a super graph G^*, we define the triple $(G^*, \mathbf{Arity}, \mathbf{map}_c)$ to be a *graph extension* if **Arity** is a function that maps every node $n \in N^*$ to a natural number, the length of the data flow array, and $\mathbf{map}_c$ is a family of functions for every call site $c =$ call P where $\mathbf{map}_c : [1..\mathbf{Arity}(c)] \rightarrow [1..\mathbf{Arity}(\text{entry } P)]$ maps every position of the data flow array of the calling procedure to a position in the array of P.

Instead of considering paths between nodes of the super graph, we use now paths between the elements of the data flow arrays. So a new graph with pairs of nodes and vector positions as nodes can be defined.

Definition 3 (Extended Super Graph).
For a graph extension $(G^*, \mathbf{Arity}, \mathbf{map}_c)$ we call the graph $G^*_E = (N^*_E, E^*_E, s^*_E, e^*_E)$ *extended super graph*, with $N^*_E = \{(n, i) \mid n \in N^* \text{ and } i \in [1..\mathbf{Arity}(n)]\}$, $s^*_E = (s^*, 1)$, $e^*_E = (e^*, 1)$, and $((n_1, i_1), (n_2, i_2)) \in E^*_E$ iff n_2 is an entry node and $i_2 = \mathbf{map}_{n_1}(i_1)$, or n_1 is an exit node and $i_1 = \mathbf{map}_c(i_2)$ with c being the call belonging to n_2, or $i_2 = i_1$.

An extended super graph with its extended paths is shown in figure 2. The interprocedural analysis can be performed by applying an intraprocedural fixpoint

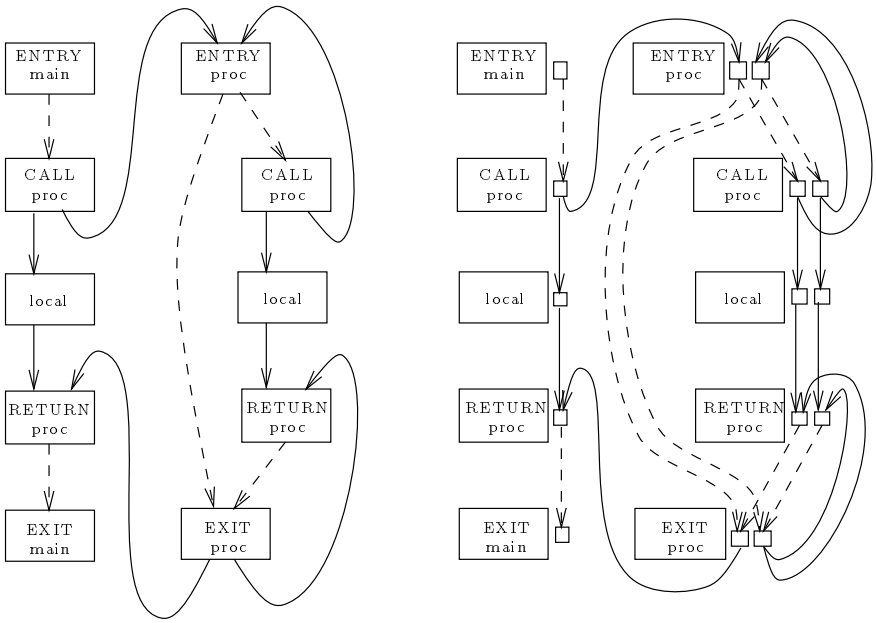

Fig. 2. A super graph and its extended super graph

algorithm to the extended super graph. The merging of the data at different call sites is done automatically, if the $\mathbf{map}_c$ are chosen appropriately: In figure 2, the two edges of the call to proc from itself are leading into the same data flow element at the head of proc. So a meet of the two values of the call site will be calculated if the second element of the vector at the entry of proc is needed. The corresponding element from the exit node will be duplicated and propagated to the return node inside proc. To solve a data flow problem with extended graphs, it is necessary to find suitable map and arity functions. With these it is possible to tune the analysis: the higher the arities the better the precision we can achieve, but the more time and space is needed.

5.4 Mappings

The task to select a mapping is mainly to find a compromise between time/space complexity and preciseness of the analysis. So we will explain now some methods to calculate pairs of arity and mapping.

1. In the simplest case, the arity of each procedure is one, and the mapping functions are the identity. So the information of each call to a procedure is mixed together.
2. Another simple way is to count the number of calls to a procedure p in the program text (this is the number of control flow edges in the super graph that are leading to the entry of p) and to take this number as the arity of p. Of course one must choose the arity of P_0 as one although it is formally not

called. Then the $\mathbf{map}_c$ functions project all elements of the tuple at the call site to a single fixed position in the array of the called procedure.
The effect of these functions is that the meet of the data flow information at every call site in the program is kept separately in the called procedure p, but if there are further calls from p to some other procedure q the information is mixed up with the flow values of all other calls to p. In practice, it has shown that the arity of some procedures calculated with this method becomes quite large in real programs (up to 100). So it seems to be reasonable not to use higher values for the arities in large programs, due to space and time restrictions.

3. To be more precise, it is necessary to keep deeper levels of the call separated. If there are for example two different call sites c_1 and c_2 of a procedure p and p calls q we would like to have as many fields in the array of q as there are fields at c_1 plus the number of fields that are at c_2 in which we can map the resulting information of c_1 and c_2. So the arity of a procedure should be the sum of the arities of all call sites. As induction base we have **Arity**(entry P_0) = 1. But this works only if the call graph is acyclic which means that the program is non–recursive.
   So we have to find the strongly connected components (SCC, see [17]) in the call graph G and consider them as a single procedure in a call graph G'. To this acyclic call graph G' we can apply the method described for the non–recursive case, with the additional rule that the arity of a SCC is the arity calculated for the non–recursive case multiplied by a constant k which reflects the fact that there can be many recursive calls inside the SCC. Afterwards, we expand the melted nodes again and assign to every procedure the arity of the compound node.
   The resulting mapping for calls leading to a procedure outside any SCC or calls from outside into a SCC is simple because every element in the call site vector has a corresponding one in the vector at the procedure entry. For the other mappings, different ways can be chosen. An example is shown in figure 3.
4. If the underlying lattice D is finite a mapping can be constructed that results in the precise solution of the DFP. Call chains of length up to $|D|^2 * K$ (K is the number of call sites in the program) must be analyzed separately.

Summary
So let's shortly summarize the advantages and disadvantages of our approach:

— we are usually not precise, but this can be achieved by using an appropriate mapping if the domain is finite.
— complex lattices and large programs result in analyzers that use a lot of memory and time; to overcome the space problems, we added a garbage collector for the flow elements (FULA is functional and therefore has a copy semantics).
\+ every intraprocedural *DFP* can be extended to an interprocedural analyzer without further restrictions of the domain.

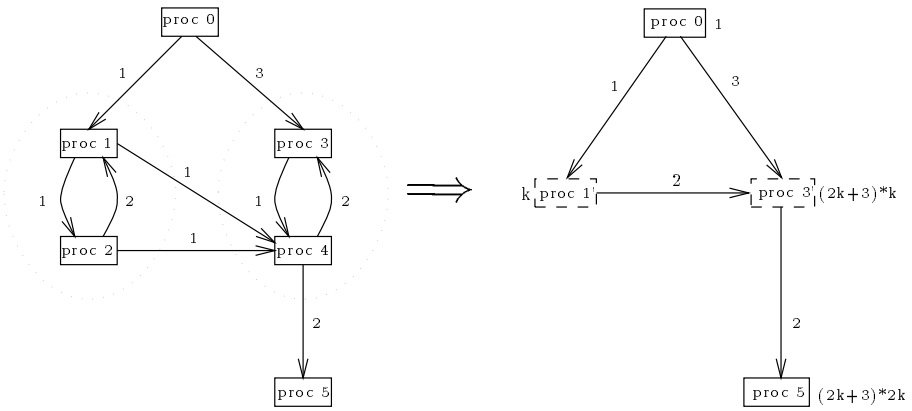

Fig. 3. The calculation of the arity in method 3. The numbers on the edges mean the numbers of call sites, and the numbers at the nodes are the calculated arities, e.g. the edge `proc 2` to `proc 1` means that there are two call sites of `proc 1` in `proc 2`.

+ additional specification is only needed for parameters, return values and local variables.
+ tuning is possible through the choice of arity and mappings.
+ the method yields detailed results about the different calls.

6 Practical Results

PAG generated analyzers have successfully been tested in a large compiler system with ANSI-C and Fortran90 front ends as well as with different back ends (including SPARC) in the **ESPRIT** project COMPARE. We have specified the following data flow problems:

- copy constant propagation: only assignments of the form `x := c` where `c` is a constant, and `x := y` are taken into account.
- linear constant propagation: here additionally, statements of the form `x := c*y + d` are considered, where `c` and `d` are constants.
- full constant propagation: every right side is taken into account. Note that the lattice for this analyzer is infinite and the transfer functions are not distributive.
- alias analysis: computes *must* and *may* aliases
- liveness: a well known bit vector problem
- dominator: similar to liveness

These analyzers run together with handwritten and generated compiler phases (called *engines*) in the **CoSy** compilation model that has been introduced by [1]. In figure 4 we have listed the size of the different specifications and analyzers. The specification of the syntax tree is the same for all analyzers as well as a large part of the hand coded interface listed in the last row. As there are some differences in

	Lines of Specification				Generated	Interface
	DATLA	FULA	syntax	$\sum$	Code (kB)	(lines)
copy constant propagation	233	12	144	389	268	1283
linear constant propagation	400	14	144	558	309	1251
full constant propagation	628	12	144	784	371	1283
alias analysis	203	10	144	357	215	1084
live variables	95	10	144	249	222	1054
dominator analysis	22	3	144	169	147	1084

Fig. 4. The size of specification and the generated analyzer

the subsequent treatment of the computed data, they are not fully identical. We measured the practical relevance of the 2. mapping of section 5.4 by applying the full and the copy constant propagation to several 'real' programs. The results are shown in figure 5. We have also generated analyzers for a completely different compiler which takes a subset of Pascal as input.

7 PAG and Beyond

As far as we know, only a few analyzer generators have been implemented. Many papers have discussed analyzer generators from a theoretical point of view, but not many serious efforts have been undertaken to produce practical implementations. Projects we know of are the following:

1. Tjiang and Hennessy [21] presented an analyzer generator *Sharlit* which uses *path compression* for evaluating the path functions. The user has to give simplification rules and has to support the implementation of the flow values and graph routines in low level C code. There is no possibility of automatic generation of the domain code and no automatic handling of procedures.
2. Venkatesch and Fischer [22] presented a tool *Spare* which is based on the abstract interpretation framework. It is developed mainly for testing purpose of some abstract interpretations. Moreover it has no built in facilities to handle a function mechanism.
3. Yi and Harrison [25] developed an abstract interpretation based analyzer generator. It has the ability to tune the analysis by using a restricted form of narrowing. But it allows only domains of finite cardinality.

We have shown in this paper that PAG is able to generate analyzers that can analyze large programs. Specifying an analysis like the copy propagation shortens the time of implementation to something between one and two hours if the interface is already there (which is needed only once per compilation system). One of the authors was able to adapt an already existing specification of constant propagation to include alias information within three man hours.

For the future we plan to investigate additional things:

- new fixpoint algorithms: we plan to implement some other fixpoint algorithms and investigate a hybrid approach, where we can mix different algorithms and keep the correctness and termination properties.

program	*procs*	*nodes*	*objects*	IPC_C	IPC_F	CON_C	CON_F	$\#\text{obj}_C$	$\#\text{obj}_F$	$\#\text{s}_C$	$\#\text{s}_F$
cdecl	32	3209	835	7	7	1010	1010	2/80	3/80	357	357
ed	46	3448	784	7	9	1481	1509	7/59	1159	138	138
fft	7	638	195	3	3	439	439	0/19	0/19	62	62
flex	128	11974	2301	49	49	4926	4980	59/355	59/336	1072	1072
flops	2	640	139	9	9	5820	6209	90/47	161/61	240	1105
gzip	99	8332	1740	11	11	4353	4358	36/250	37/280	1148	1148
heap	3	361	94	8	8	1432	1432	13/34	13/34	55	55
linpack	12	1196	278	6	6	587	610	28/51	32/55	490	490
ratfor	51	3200	762	0	0	934	934	6/54	6/54	323	323
twig	81	4525	1085	6	6	1332	1332	3/64	3/64	375	375
xmodem	30	4460	1157	38	38	4849	4963	6/164	12/166	520	526
clinpack	12	1257	285	6	6	668	691	28/51	32/55	490	490
cloops	126	3212	620	127	131	13341	14436	56/177	60/179	1377	1388
dhry	13	662	239	8	8	279	286	2/13	3/14	99	99
search	3	89	34	2	2	189	189	7/8	7/8	9	9
whet	7	492	121	3	3	397	448	16/44	40/49	104	104

Practical results measured for the copy constant propagation (X_C) and the full constant propagation(X_F).

IPC: the number of constant objects at the procedure headers
$\sum_{entry\ p} |\{o|(\sqcup_{i=1}^{Arity(p)} flow(p)[i])(o) \neq \bot, \top\}|$

CON: the sum of the number of constant object at all nodes
$\sum_{flownode\ n} |\{o|(\sqcup_{i=1}^{Arity(n)} flow(n)[i])(o) \neq \bot, \top\}|$

#obj: x/y, x the number of constant source code variables,
y the number of constant temporaries (frontend introduced)

#s: foldings of subexpressions

Fig. 5. Practical Measurements

- reduction of the space consumption: the vectors of the data flow elements will be large, if the program mainly consists of procedure calls. In that case most of the entries at procedure headers are the same; thus, a melting of these elements may result in less precise data flow information, however, with decreased time and space consumption. This can be generalized by replacing the equality with a distance notion. This distance can be inductively defined on the structure of the lattices. The user can specify a concrete distance d and during the analysis elements e_1, e_2, within this distance, are replaced by the result of the (specified) combination function $combine(e_1, e_2)$. Then the mappings of section 5.4 are dynamically redefined.
- automatically reasoning over PAG specifications: an interesting research direction is the analysis of PAG specifications. Using an eager first-order functional language for specifying the transfer functions allows automatic proof of some properties like monotonicity.

References

1. M. Alt, U. Aßmann, and H. van Someren. Cosy Compiler Phase Embedding with the Cosy Compiler Model. In *CC'94*, LNCS 786 Springer, 1994.
2. D. Callahan, K. D. Cooper, K. Kennedy, and L. Torczon. Interprocedural constant propagation. In *SIGPLAN '86*, volume 21, June 1986.
3. P. Cousot and R. Cousot. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *POPL'77*, pages 238–252, 1977.
4. P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *POPL'79*, pages 269–282, January 1979.
5. P. Cousot and R. Cousot. Abstract interpretation frameworks. *Journal of Logic Computation*, 2(4):511–547, 1992.
6. D. Dhamdhere, B. Rosen, and F. K. Zadeck. How to analyze large programs efficiently and informatively. In *SIGPLAN PLDI '92*, volume 27, June 1992.
7. E. Duesterwald, R. Gupta, and M. L. Soffa. Demand-driven computation of interprocedual data flow. In *POPL '95*, January 1995.
8. D. Grove and L. Torczon. Interprocedural Constant Propagation: A Study of Jump Function Implementations. In *ACM SIGPLAN PLDI '93*, 1993.
9. M. Hecht. *Flow Analysis of Computer Programs*. North Holland, 1977.
10. N. Jones and S. Muchnick. Even simple programs are hard to analyze. In *POPL'75*, pages 106–118, January 1975.
11. J. Kam and J. Ullman. Monotone data flow analysis frameworks. *Acta Informatica*, 7:305–317, 1977.
12. G. Kildall. A unified approach to global program optimization. In *POPL'73*, pages 194–206, October 1973.
13. J. Knoop and B. Steffen. The interprocedural coincidence theorem. In *CC'92*, LNCS 641, pages 125–140., 1992.
14. T. Reps, M. Sagiv, and S. Horwitz. Precise interprocedural dataflow analysis with application to constant propagation. In *TAPSOFT'95, Arhus, Denmark*, 1995.
15. Florian Martin. Die Generierung von Datenflußanalysatoren. Master's thesis, Universität des Saarlandes, 1995.
16. Florian Martin. *PAG Reference Manual*. Universität des Saarlandes, 1995.
17. Kurt Mehlhorn. *Data Structures and Algorithms 2: Graph Algorithms and NP-Completeness*. Springer Verlag, 1984. ISBN 3-540-13641-X.
18. H.R. Nielson and F. Nielson. Bounded fixed point iteration. In *ACM POPL'92*, pages 71–82, Albuquerque, NM, January 1992.
19. Mads Rosendahl. Higher-order chaotic iteration sequences. In *PLILP'93, Tallinn, Estonia*, LNCS 714, pages 332–345. Springer-Verlag, 1993.
20. M. Sharir and A. Pnueli. Two approaches to interprocedural data flow analysis. In *Program Flow Analysis: Theory and Applications*, chapter 7, Prentice-Hall, 1981.
21. S. Tjiang and J. Hennessy. Sharlit – A tool for building optimizers. In *PLDI'92*.
22. G.V. Venkatesch and C. N. Fischer. Spare: A development environment For Program Analysis Algorithms. In *IEEE Transactions on Software Engineering*, 1992.
23. M. N. Wegman and F. K. Zadeck. Constant propagation with conditional branches. In *ACM POPL'85*, pages 291–299, New Orleans, LA, January 1985.
24. Reinhard Wilhelm and Dieter Maurer. *Compiler Design*. International Computer Science Series. Addison–Wesley, 1995.
25. K. Yi and W.L. Harrison III. Automatic generation and management of interprocedural program analyses. In *ACM SIGPLAN-SIGACT*, January 1993.

A Specification of Liveness

```
DOMAIN
varset = set(unum)
vars   = lift(varset)

PROBLEM livevar
direction : backward
carrier   : vars
init      : bot
init_start: lift(bot)
combine   : comb

TRANSFER

Assign(e1,e2):
   minus(@,def(e1)) lub use(e1) lub use(e2);

Evaluate(exp): @ lub use(exp);

If(exp): @ lub use(exp);

FuncCall(_,pars,_): @ lub use_list(pars);

Call(_,pars): @ lub use_list(pars);

default: @;

SUPPORT
comb(a,b) = a lub b;

// Subtracts list of unums from set of unums
minus:: vars, [unum] -> vars;
minus(l_set,list) =
 let x <= l_set;
 in lift( case list of
               a:as => x # a;
                 [] => x;
          endcase);

// Calculate used objects of expressions
use_list([!]) = lift({});
use_list(x::xs) = use(x) lub use_list(xs);

// Calculate used objects of an expression
use::EXPR -> vars;
use(Content(ObjectAddr(obj)))
   = lift({} ~ id(obj));
use(Content(e)) = lift(top);
use(Member(e,_)) = use(e);
use(Unop(_,e)) = use(e);
use(Binop(_,e1,e2)) = use(e1) lub use(e2);
use(_) = lift({});

// Return defined objects
def(ObjectAddr(obj)) = id(obj):[];
def(_) = [];

// Get the identifier of an object
id(obj) = unum(val-INT(
  case obj of
    DataGlobal(id,_) => id;
    Local(id,_,_) => id;
    Parameter(id,_) => id;
    Register(id,_,_) => id;
    _ => error("Unknown object");
  endcase));


SYNTAX

START : STMT;

STMT: Evaluate(Expr:EXPR)
  | BeginProc(Params:OBJ*,Locals:OBJ*)
  | Assign(Lhs:EXPR,Rhs:EXPR)
  | Call(Proc:EXPR,Params:EXPR*)
  | FuncCall(Proc:EXPR,Params:EXPR*,Res:EXPR)
  | CallRet(Params:OBJ*,Locals:OBJ*)
  | EndProc(Params:OBJ*,Locals:OBJ*)
  | Goto(Target:EXPR)
  | Return(Value:EXPR,Next:EXPR)
  | If(Cond:EXPR)
  | EndFuncCall(Res:EXPR)
  | EndCall();

EXPR: IntConst(Value:UNIV_INT)
  | RealConst(Value:UNIV_REAL)
  | BoolConst(Value:BOOL)
  | ObjectAddr(Obj:OBJ)
  | Content(Addr:EXPR)
  | Subscript(Base:EXPR,Index:EXPR)
  | Member(Base:EXPR,Field:OBJ*)
  | Unop(Op:UNOP,Value:EXPR)
  | Binop(Op:BINOP,Left:EXPR,Right:EXPR);

UNOP: Cast()
  | Abs()
  | Neg()
  | Not();

BINOP: Plus()
  | Diff()
  | Mult()
  | And()
  | Or()
  | Div()
  | Quo()
  | Mod();

OBJ: DataGlobal(id:INT,IsVolatile:BOOL)
  | Local(id:INT,IsVolatile:BOOL,Proc:OBJ)
  | Parameter(id:INT,IsVolatile:BOOL)
  | Register(id:INT,IsVolatile:BOOL,Reg:INT);

BOOL == bool;
NAME==str;
INT==snum;
UNIV_INT == snum;
UNIV_REAL == real;
UNIV_ADDRESS == unum;
```