Analysis and Visualization of Predicate Dependence on Formal Parameters and Global Variables

David Binkley and Mark Harman

Abstract—Empirical data concerning the qualitative and quantitative nature of program dependence is presented for a set of 20 programs ranging from 600 lines of code to 167,000 lines of code. The sources of dependence considered are global variables and formal parameters and the targets considered are a program's predicate nodes. The results show that as the number of formal parameters available to a predicate increases, there is a decrease in the proportion of these formal parameters which are depended upon by the predicate. No such correlation was found for global variables. Results from theoretical and actual computation time analysis indicate that the computation of dependence information is practical, suggesting that the analysis may be beneficial to several application areas. The paper also presents results concerning correlations that provide strong evidence that the global and formal dependence sources are independent of one another and that the numbers of globals and formals are independent of the size of the procedure that contains them. Finally, two visualization techniques for displaying dependence information are introduced. Illustrations show how these visualizations and predicate dependence analysis can assist in activities such as testing, comprehension, and evolution.

Index Terms—Dependence analysis, program slicing, program comprehension, software maintenance.

1 Introduction

Of the many software engineering applications supported by source code analysis, a large number are highly sensitive to dependences between program components. However, despite its importance, there is little published work on empirical assessment of the *level* of program dependence expected, nor how it varies within and between programs.

Results concerning dependence levels and trends are important for dependence-sensitive applications. Such applications can require considerable analysis effort (by both humans and machines). For example, human program comprehension is an inherently expensive task. Program slicing [8], [21], [42], [65], [75] has been used to reduce this expense, by focusing attention on a particular "slice" of the source code [12], [15], [17], [36], [51]. Slices are constructed using dependence analysis. The lower the dependence level, the smaller the slice. Using slicing can therefore reduce human analysis effort.

Automated analysis can be expensive too. For example, consider the automated search for branch-coverage adequate test data. Current tools do not integrate dependence analysis [47], [56], [67], [70]. However, the potential savings from dependence analysis are dramatic; *linear* dependence level reductions cause *exponential* search space reductions.

- *D. Brinkley is with Loyola College, Baltimore, MD 21210-2699. E-mail: binkley@cs.loyola.edu.*
- *M. Harman is with King's College London, Strand, London, WC2R 2LS, UK. E-mail: Mark@dcs.kcl.ac.uk.*

Manuscript received 19 Aug. 2003; revised 1 Apr. 2004; accepted 9 July 2004. Recommended for acceptance by L. Briand.
For information on obtaining reprints of this article, please send e-mail to: tse@computer.org, and reference IEEECS Log Number TSE-0112-0803.

This paper presents the results of an empirical investigation into the nature of program dependence for a set of 20 C programs. The goal of the study was to determine typical levels of dependence, to identify dependence trends, correlations and patterns, and to provide scalable visualization techniques for depicting the dependencies present in large programs.

The study focuses upon the extent to which a program's predicates depend upon two of their potential influencors —the global variables in scope and the formal parameters in the function or procedure which contains the predicate. Predicates are chosen because they allow dependence to be studied at a detailed and fine level of granularity, while results obtained refer to a semantically important aspect of the computation—the decision logic and control flow of the program.

The principal contributions of the paper can be summarized as follows:

- Results for dependence levels across 50,579 predicates indicate that a typical predicate depends upon approximately
  - 72 percent of formal parameters in the scope,
  - 48 percent of the global variables potentially used or defined by a call to the predicate's procedure, and
  - 2.4 percent of the global variables in scope.

  These results are important for dependence-sensitive activities such as test data generation [16], [22], [31], [52], [53], [56], [67], [70] (where the apparent size of the input space can be reduced by dependence analysis) and program comprehension

Name	Size KLOC	Predicates	Predicates per KLOC	Description
replace	0.6	84	140	Regular expression string replacement
copia	1.2	45	38	ESA signal processing code
compress	1.9	88	46	Data compression utility
which	5.4	136	25	Unix utility
barcode	5.9	365	62	Barcode generator
space	11.5	593	52	ESA ADL interpreter
ed	13.6	2627	193	Unix editor
prepro	14.8	619	42	ESA array pre-processing code
bc	16.8	432	26	Calculator
findutils	18.6	1819	98	File finding utilities
diffutils	19.8	2212	112	File comparing routines
flex2-4-7	21.2	911	43	BSD scanner generator (version 2.4.7)
flex2-5-4	21.5	1306	61	BSD scanner generator (version 2.5.4)
espresso	22.1	2497	113	logic simplification for CAD (from SPECmark)
ijpeg	28.2	2038	72	JPEG compressor (from SPECmark)
ntpd	45.6	3061	67	daemon for the network time protocol
a2ps	62.9	5984	95	postscript formatter
gnugo	81.7	5244	64	Gnu game player
sendmail	85.7	4991	23	Linux mail daemon
spice	167.0	15527	93	Digital circuit simulator
average	28.3	2529	73	

Fig. 1. Subject programs.

[6], [13], [33], [39], [55], [46] (where distinguishing true dependencies from apparent dependencies removes an unnecessary cognitive burden).

- The study provides evidence of a strong *inverse* correlation between the number of formal parameters visible to a predicate and the proportion actually used. There is no such correlation for globals.
- The paper investigates correlations between sources and targets of dependence, showing that the dependence sources considered are independent of one another and that, in general, the number of formals available to a predicate is not dependent upon the size of the function in which the predicate occurs.
- Visualizations are introduced for human analysis and interpretation of dependent proportions. The visualizations summarize dependence information about predicates, allowing trends and patterns to be identified.
- The paper illustrates ways in which the visualizations can be used to guide program comprehension and the human evaluation of software evolution. The paper also shows how the analysis can assist with reducing automated test data generation effort.
- Finally, the paper presents performance results from the implementation of the dependence analysis which show that the approach scales to medium/large programs, making it practical as a support for comprehension, test data generation, maintenance, and evolution.

The rest of the paper is organized as follows: Section 2 outlines the experimental setup and Section 3 describes the two visualizations used to view dependencies for large programs within a single graph. Sections 4 and 5 describe dependence ratios, levels and trends over all predicates studied, and correlations between them. Section 6 shows how the visualizations can help with program comprehension and evolution and how the predicate dependence analysis introduced in this paper can assist with automated test data generation. Section 7 considers threats to validity of the findings. Section 8 shows that dependence analysis is practical with regard to both its theoretical complexity and the actual computation time for the implementation used in the study. Section 9 presents related work. Finally, Section 10 summarizes the paper.

2 Experimental Setup

This section describes the experiential context and the experiment's instrumentation including the statistical tests used and the data collection process.

2.1 Experiment Context

The study concerns 20 C language subject programs that range in size from 600 LOC to 167K LOC. Three programs were obtained from the European Space Agency (ESA), while the remaining 17 consist of publicly available utilities and application programs. Fig. 1 presents summary information concerning the subjects.

2.2 Instrumentation

The study uses two statistical tests: Pearson's linear correlation and Spearman's rank correlation. Pearson's linear correlation is used to construct quantitative models of the data. Such correlations measure linear associations between variables. The output is a *correlation coefficient*, reported as the value R, and the coefficients of a linear *model*. The statistical significance of R can be summarized as

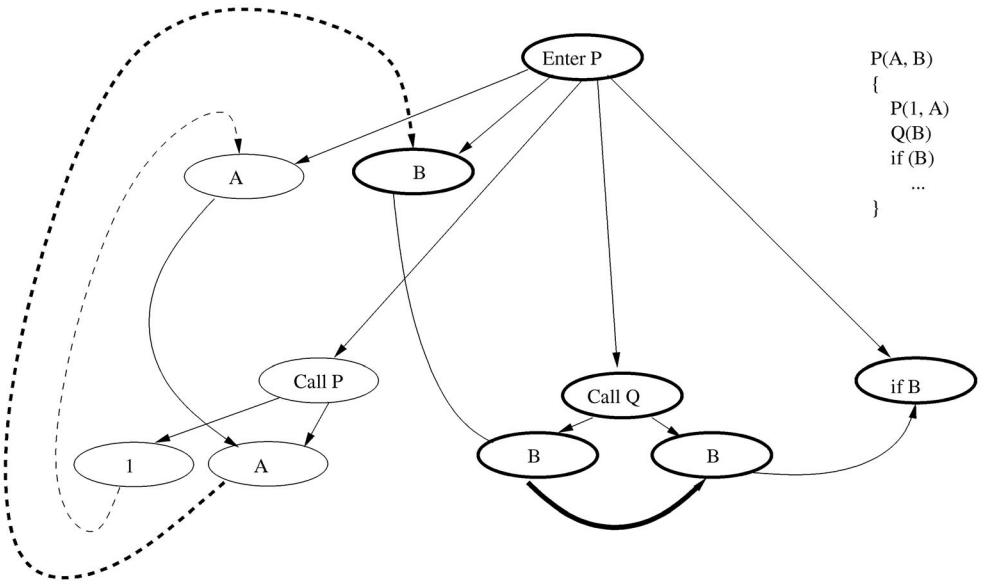

Fig. 2. Local slice computation example.

0.8 - 1.0 strong association, 0.5 - 0.8 moderate association, and 0.0 - 0.5 weak or no association. A negative value indicates an inverse correlation. For the effect of X, Y, and Z on A, the resulting model coefficients, m_i, belong to the linear equation

$$A = m_1X + m_2Y + m_3Z + b.$$

The statistical confidence associated with each coefficient is denoted by a *t*-score; a *t*-score of 1.94 or greater is considered significant. Spearman's rank correlation is also used. The Spearman test does not attempt to construct a linear correlation between the two variables, but merely gives a correlation value which indicates whether increases in one variable are accompanied by increases in the other.

The dependence data was collected by a modified version of the standard System Dependence Graph (SDG) backward interprocedural slicing algorithm [29]. An SDG represents a program as a collection of Procedure Dependence Graphs (PDGs), the nodes of which represent statements and predicates and the edges of which represent control and data-dependences between nodes. Interprocedural dependence edges connect PDGs to form the SDG. The final step in the construction of the SDG is the addition of "*summary*" edges at call-sites. These transitive dependence edges represent certain paths of dependence edges through called procedures from nodes representing initial parameter values to nodes representing final parameter values.

Slicing "at a node" in the SDG means finding all the statements which potentially affect the computation represented at that node. In terms of Weiser's definition of a slicing criterion, slicing at an SDG node n means slicing on (V, n), where the set of variables, V is the set of variables referenced at n.

The formal parameters and global variables that potentially affect a branch node b of procedure P can be computed by slicing with respect to the node for b and reporting the formal parameters and global variables whose nodes are in P's PDG and the slice. It turns out that this approach is inefficient because it needlessly considers components from other procedures. It is also a mild overapproximation in the presence of recursive procedures. Therefore, the slicer was modified to compute "local" slices [29], which only consider intraprocedural edges when slicing. Local slices safely identify those parameters that affect a predicate on a particular invocation of the procedure.

Example. Fig. 2 illustrates the time saving and increased precision resulting from the use of local slices. Only the relevant subgraph of `P`'s PDG is shown (e.g., the output of `P` is not shown). Consider the slice on "`if B`." The local slice used to gather predicate dependence includes the bold vertices. The time saving is illustrated by this slice as it traverses the summary edge (bold edge) associated with the call on procedure `Q`. By contrast, (pass 2 of) a full interprocedural slice would include elements from `Q` and incur the cost of discovering them. The increased precision is illustrated by the absence of formal parameter `A` in the local slice. A "full" interprocedural backward slice, would (during pass 1), traverse the bold interprocedural edge and eventually encounter the node for formal parameter `A`.

For completeness, some details of the operation of CodeSurfer when performing the data-flow analysis used to build SDGs are presented. First, the pointer analysis algorithm used by CodeSurfer is a modification of Andersen's algorithm [2]. The modification maintains the same precision while decreasing the time and space requirements [18]. The result is a flow and context insensitive points-to analysis that is almost as fast as Steensgaard's algorithm [63], but provides a more precise answer [61]. Unfortunately, flow or context sensitive points-to analysis is hard [26], [54], [77] and, at present, does not scale well. There is also empirical evidence that, for

program slicing, improved quality points-to sets do not have a significant impact on slice size [11], [49].

Second, it has been shown that structure-field treatment has a potential impact on slice size [76]. This is best understood by way of example. Consider the following C code:

```
struct file_ops
{
    int (*write_fn)();
    int (*read_fn)();
} file1;
```

Assume `write_fn` potentially points to the functions `write_ext2` and `write_nfs` and `read_fn` potentially points to the functions `read_ext2` and `read_nfs`. With structure fields expanded, the function call `(*file1.readfn)(···)` potentially calls `read_ext2` and `read_nfs`. With structure fields collapsed, all structure fields are treaded as a single variable and the call `(*file1.readfn)(···)` potentially calls `write_ext2` and `write_nfs` in addition to `read_ext2` and `read_nfs`. For programs with structures that contain pointers, and especially function pointers, this collapse can significantly increase slice size.

Arrays are passed to functions as whole structures. In common with other static analysis techniques, it is not possible (as it is in dynamic analysis) to consider individual array elements separately. Rather, an access to any array element is effectively treated as a may access to all elements.

For a predicate p, the *dependent proportion* refers to the proportion of *parameters* that p depends upon. For two predicates in the same function (i.e., two predicates having the same number of parameters available) a lower dependent proportion implies dependence on fewer parameters. In this definition, the term *parameter* refers to formal parameters and global variables. Global variables are further broken down into *visible globals*, those in scope at p, and *du-globals*, the subset of the visible globals potentially defined or used by a call to the function that contains p. This distinction between du-globals and visible globals is interesting because du-globals represent a first approximation to those globals upon which p depends. As will be seen, this first approximation is rather crude.

3 Visualization Techniques for the Dependent Proportion

Two visualization techniques are introduced to present dependence information. In order to visualize dependent proportions, three variables must be depicted:

- **Predicate Count**. The number of predicates summarized by a data point. Data points summarizing a larger number predicates carry more weight.
- **Max-Parameters**. The maximum number of parameters in scope (visible) at a predicate. Large maxima are more interesting because of the greater potential for savings from using dependence analysis.
- **Parameters-Used**. The number of parameters that affect a predicate according to the dependence analysis described in Section 2.2. The smaller this number the better for dependence-sensitive applications.

Program	Max Formals	Average Formals Used	Predicate Count
`replace`	1	1.0	8
	2	1.5	11
(*)	3	2.2	8
	4	3.6	47
	6	5.0	10
`prepro`	0	0.0	11
	1	1.0	137
	2	1.2	62
	3	1.1	176
	4	1.3	98
	5	1.1	53
	6	1.0	10
	7	1.6	18
	8	1.4	24
	9	1.0	4
	10	0.7	15
	11	1.0	5
	12	1.0	6

Fig. 3. Example formals data.

In these definitions "parameters" refers to formals and globals taken together. When only formals are being considered the terms *Max-Parameters* and *Parameters-Used* become *Max-Formals* and *Formals-Used*. Similar specialization is done when only globals are considered.

It is important to understand the difference between max-parameters and parameters-used. In the absence of any dependence analysis, *every* formal parameter and *every* global variable must be considered to potentially affect a predicate. Dependence analysis can reduce this number, but by its very nature, must be conservative: determining whether a formal definitely affects a predicate is undecidable. Therefore, a safe (conservative) approximation is computed; parameters-used excludes only those parameters that can be *guaranteed* not to affect a predicate.

3.1 Dependent Proportion Visualizations

To visualize the data, two graphical devices are used: the *Dependence Skyline* and the *Dependence Bubble Chart*. This section describes these visualizations and illustrates their interpretation for two of the smaller programs. Fig. 3 shows formal parameter dependence data for the programs `replace` and `prepro`. The column labeled "Predicate Count" indicates the number of predicates summarized in each row of the table. For example, consider the third row (marked "(*)" in Fig. 3). This row summarizes predicates in procedures with three formal parameters. Such predicates, by definition, are influenced by at most three formal parameters. As shown in the "Predicate Count" column, there are eight such predicates. For these predicates, the average number of formals depended upon (shown in the column labeled "Average Formals Used") is 2.2. Compare this row with the row immediately below it in the table. For predicates with four formal parameters in scope, an average of 3.6 actually affects the predicates. At first glance, this row seems to be telling a similar story: most of the formal

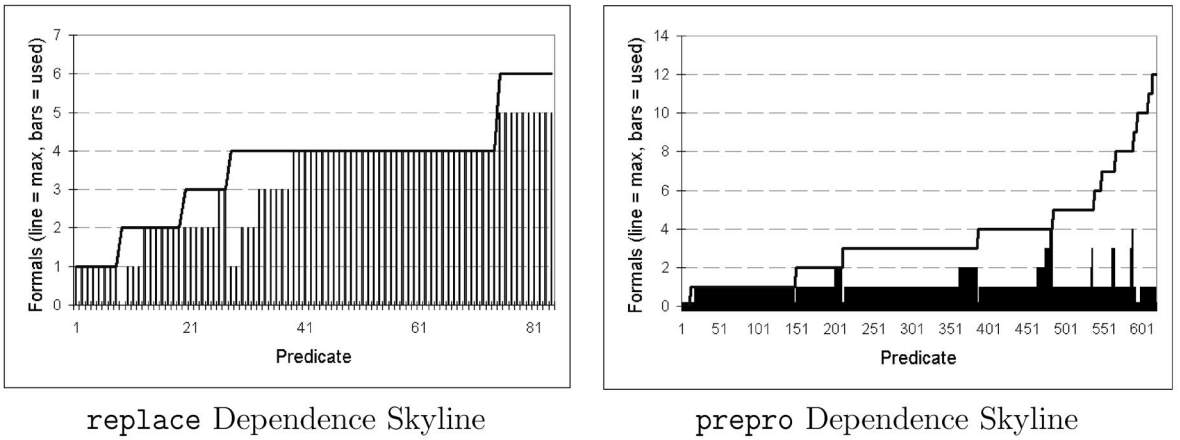

Fig. 4. Dependence skylines. Bars indicate actual predicate dependence. The horizon line running across the top of the dependence skyline is the "theoretical maximum"—the number of formals in scope. Predicates are ordered left-right by Max-Formals. The space between these two indicates ignore-able formal parameters.

parameters affect each predicate. However, Row 4 is more interesting than Row 3 because it summarizes almost six times as many predicates. This highlights the importance of the predicate count attribute.

The raw data summarized in Fig. 3 is shown in graphical form using dependence skylines in Fig. 4. A dependence skyline shows each predicate on the horizontal axis (sorted by max-parameters and, within that, sorted by parameters-used). The vertical axis denotes usage of parameters (both the theoretical maximum and the actual count revealed by dependence analysis). The "horizon line" which monotonically increases as it travels left to right across the top of the dependence skyline indicates the value of max-parameters, while the usage bars indicate the value of parameters-used.

For example, the data for predicates in procedures with three formal parameters (marked with "(*)" in Fig. 3) appears in the dependence skyline at points with X values 20 through 27 in the left of Fig. 4. Over this range, the horizon line (showing max-formals) is horizontal and represents all predicates with three visible formal parameters. The first six bars in this range have height 2, while the last two have height 3. (This gives the average height of 2.2 as shown in Fig. 3.) Each bar in a dependence skyline diagram summarizes a single static backward slice from a single predicate.

With a small example like `replace`, the dependence skyline replicates the tabular data in full detail. However, for a larger example, like `prepro`, the dependence skyline provides only a general *feel* for the raw data; it is unrealistic to separate out individual bars due to the large number of data points.

For larger data sets, dependence bubble charts can often better display the three dependent attributes. In such charts, the horizontal axis denotes max-parameters, while the vertical axis denotes parameters-used. Thus, the line $y = x$ represents the worst case (i.e., a predicate can reference no more than the visible parameters). For reference, this line is drawn as the solid line in all dependence bubble charts. Also drawn is a dashed line that represents the linear trend (computed using a linear least squares fit). Finally, bubble size represents predicate-count.

For example, consider `replace`'s dependence bubble chart for formal parameters shown on the left of Fig. 5. The largest bubble summarizes the data from Row 4 of the table in Fig. 3. Here, the average number of formals used by all predicates from `replace` with four formals in scope is just under 4. Recall from Fig. 3, the actual value is 3.6. Also, as can be quickly seen, this bubble is considerably larger than the others. In fact, it summarizes over half of the predicates in the program. The second largest bubble represents predicates in procedures with two formals (visually, these can be seen to depend upon an average of approximately 1.5 of the two available formal parameters).

Compare this with the `prepro` dependence bubble chart, which appears on the right of Fig. 5. Three important facts can be readily seen. First, divergence from the solid $y = x$ line very clearly highlights the *tremendous savings* which occur from dependence analysis of `prepro`. Second, the relative size of the bubbles indicates that most of this saving is to be expected in procedures with fewer than five formal parameters. Third, the dashed trend line indicates that formals-used is effectively constant (independent of the number of formals available). This trend line may have important implications for subsequent maintenance and evolution of `prepro`.

The dependence bubble charts and skylines for formal parameter dependence of the 20 programs studied are depicted in Figs. 6 and 7. The dependence bubble charts and skylines for du-global dependence of the 20 programs studied are depicted in Figs. 8 and 9. Results for visible global dependence are less interesting and so these are only summarized at the end of the next section.

4 Trends in Dependence Over All Predicates

One of the primary goals of this study is to determine levels, ratios, and trends in predicate dependence. This section reports key findings over all 50,579 predicates considered. Three related questions are asked that inquire about dependent proportion between predicates and first formals and then the two kinds of globals. The first question to be addressed is

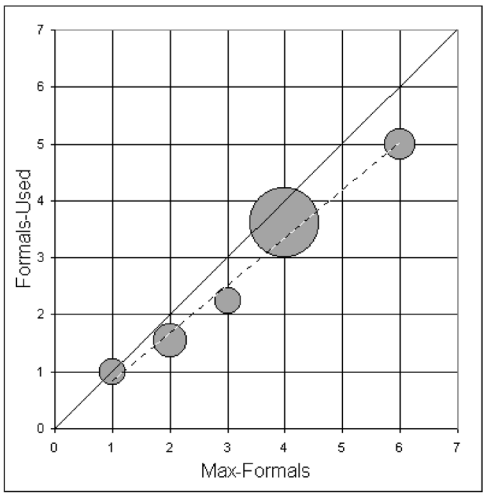

`replace` Dependence Bubble Chart

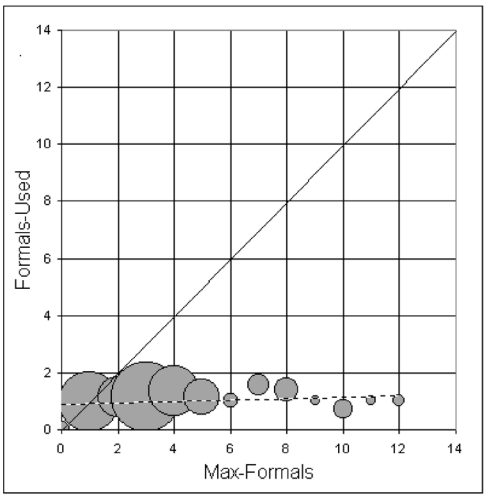

`prepro` Dependence Bubble Chart

Fig. 5. Dependence bubble charts. In these dependence bubble charts, bubble size indicates predicate-count. The solid line at 45^o indicates the maximum possible dependence (formals used = max formals). The further the bubbles are from this line, the greater saving analysis can produce. The dotted line indicates the linear trend (unweighted by bubble size).

> How does the dependent proportion change as the number of available formal parameters increases?

As seen in Section 4.2, the answer, put simply, is that

> As the number of available formals increases, the proportion of these depended upon drops.

This is an important and optimism-inducing finding. For example, impossibly large search spaces for test data generators may present more tractable problems than previously thought.

A similar question can be asked for globals. Sadly (but importantly), it cannot be said. Not for visible globals nor for du-globals. This provides some evidence for the assertion:

> Global variables are harmful.

That is, programs with a large number of global variables may present intractable problems for dependence-sensitive applications. There is a "silver lining" here: The percentage of visible globals used by a typical predicate is relatively low.

The data for formals is piecewise linear. Piecewise linear data can be analyzed by introducing a new variable that captures the pivot. For example, with a pivot of 10, this new variable has the value 0 for the first 10 observations and has the value $x - 10$ for all other observations. Thus, this new variable separates out the effect of the data after the pivot. In this case, the highest R-values are for pivots points of 8 and 10 ($R = 0.96$ and $R = 0.95$, respectively). There is little to choose between these two. Visually, the split starts at max-formals of 10, where the figures clearly show different dependence profiles.

For this reason, the following visualizations are split into two diagrams; one for values of max-formals less than 10 and one for max-formals at least 10. This piecewise linearity is starkly evident in the trend lines of the two dependence bubble charts. The following two equations separate out the two pieces. For values of max-formals below 10, the model is $formals\text{-}used = 0.14 + 0.68 \times max\text{-}formals$. For values of max-formals of at least 10, the model is $formals\text{-}used = 8.54 - 0.16 \times max\text{-}formals$. In essence, the number of formals-used increases linearly at roughly $0.68 \times max\text{-}formals$ until max-formals is 10, where formals-used begins a slight decline, but is rather independent of max-formals.

The decreasing dependence on formal parameters, can be readily appreciated visually in both dependence skylines and dependence bubble charts. The dependence skyline for all 50,579 predicates is shown in Fig. 10. The corresponding dependence bubble charts are shown at the top of Fig. 11.

When considering du-globals, the trend line for globals is just below the "$y = x/2$" line and reflects the 48 percent average usage (see the lower left of Fig. 11). The resulting Pearson model is $globals\text{-}used = 0.47 \times max\text{-}du\text{-}globals - 13$ with an R value of 0.83 (a strong linear association). For visible globals, the R value of 0.40 indicates there is no linear correlation. Thus, there are two possibilities. The more likely case is that max-visible-globals is simply not a good predictor of globals-used; however, it is it is possible (but unlikely considering the visualization of the data shown in the lower right of Fig. 11) that some nonlinear relationship exists. These two models for globals are important as they show that dependent proportions do not fall as the number of available globals increases.

Finally, consider the relationship between max and *dependent proportion* (i.e., percentage used). Scatter plots showing this relationship for formals, du-globals, and visible-globals appear in Fig. 12. Over all 20 programs, there is strong evidence for an inverse linear correlation between max-formals and dependence level. The resulting trend line, which has an R value of -0.833, is captured by $dependence\text{-}level = 86\ percent - 2.8\ percent \times max\text{-}formals$. While this linear equation better explains the data, the average dependent proportion for formals is about 72 percent. These linear equations were produced using unweighted averages. The fit is almost identical when weighted averages are used

$$(dependence\text{-}level = 85\ percent - 2.7\ percent \times max\text{-}formals.)$$

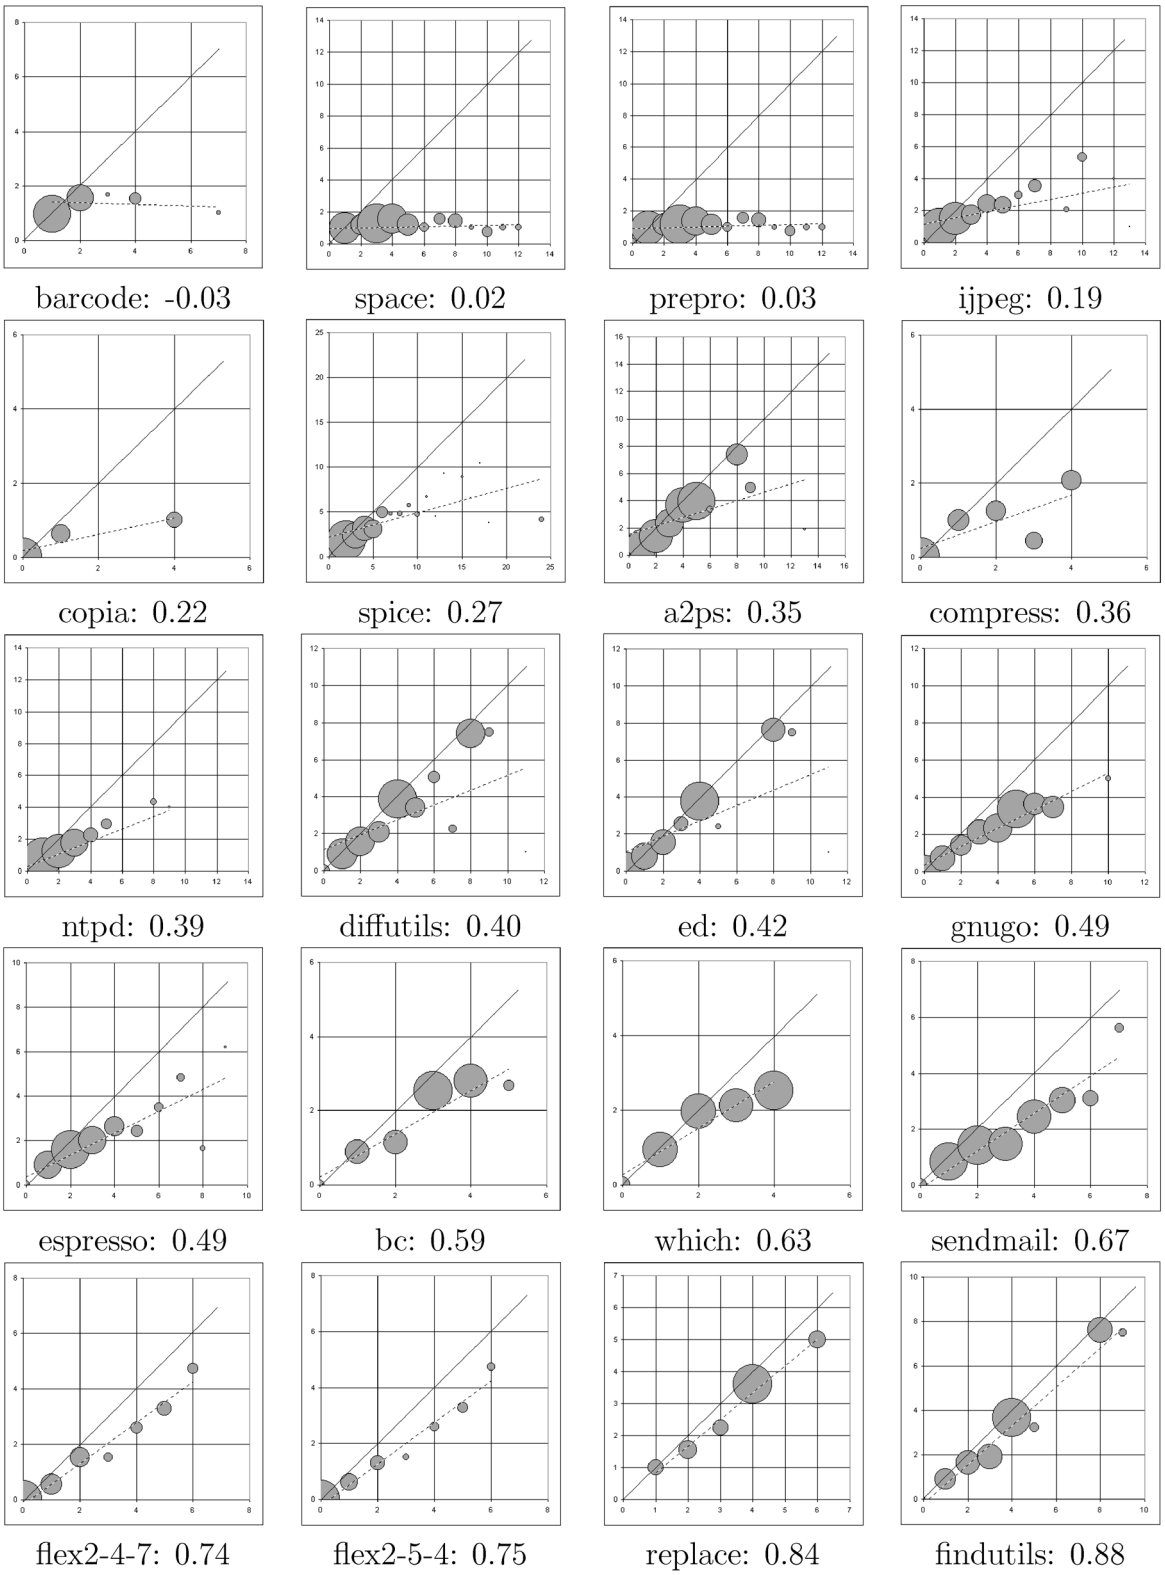

Fig. 6. Dependence bubble charts for formal parameter dependence. The number appearing after the name of each program is the slope of the dependence trend.

However, the R value for the weighted average is slightly lower, -0.66 rather than -0.833.

No linear correlations exist for globals. When considering du-globals (Fig. 12b), R has the values 0.080. For all visible globals (Fig. 12c), it has the value -0.27. Since there is no linear correlation, the averages become more interesting for globals than they are for formals (where a good linear fit was found). On average, 48 percent of du-globals and 2.4 percent of visible globals are used.

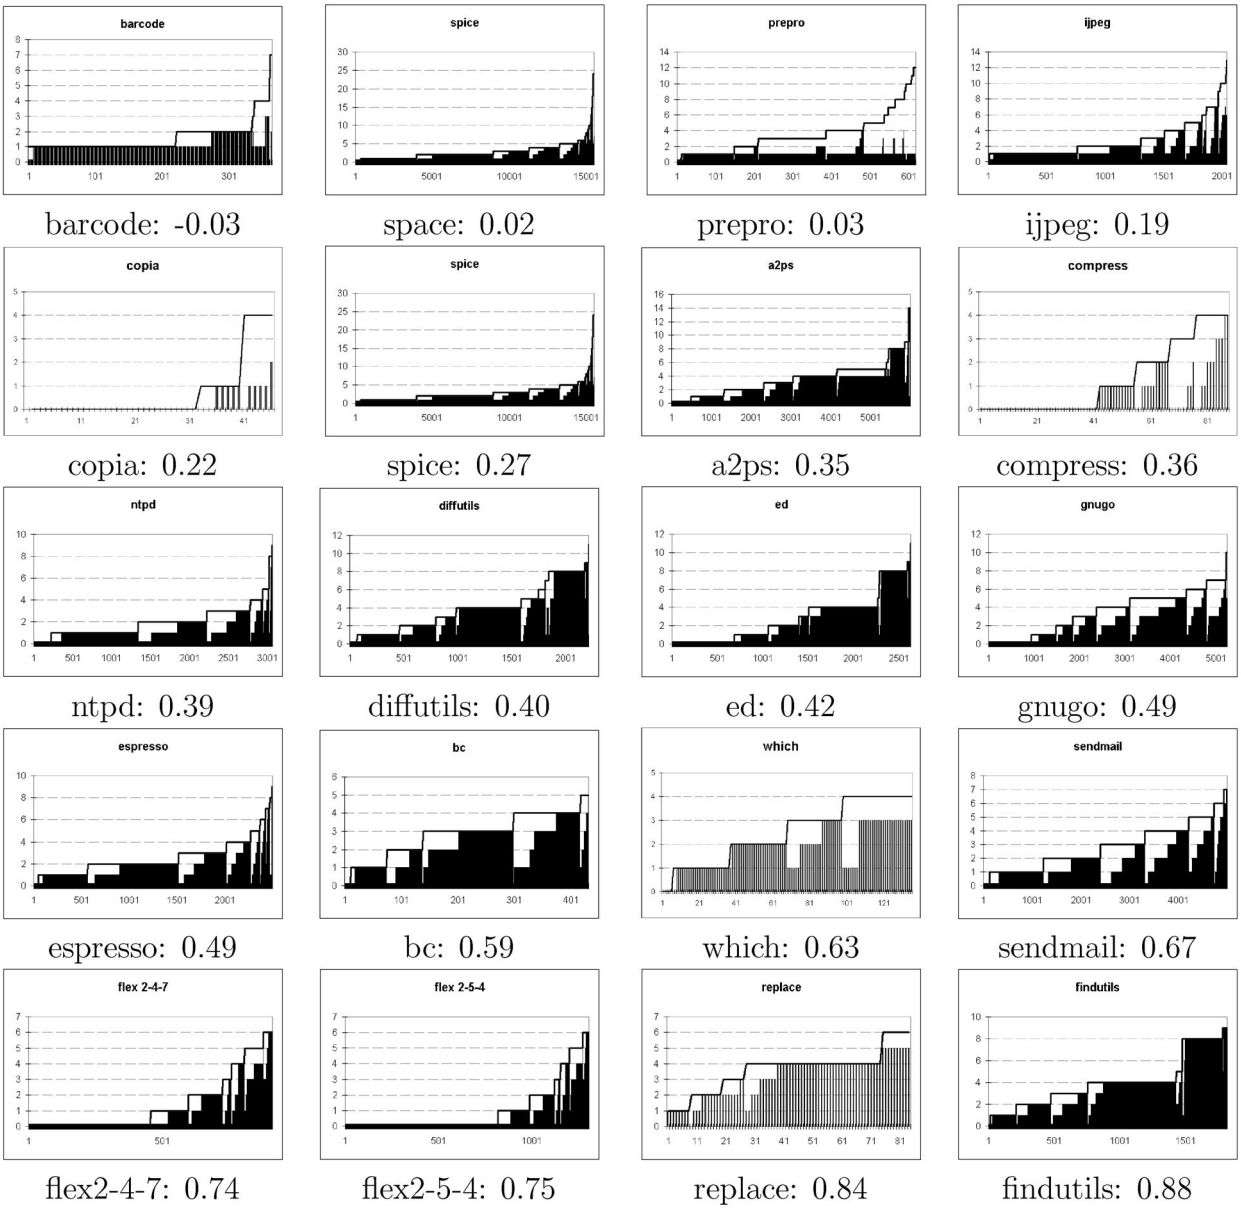

Fig. 7. Dependence skyline for formal parameter dependence. The number appearing after the name of each program is the slope of the dependence trend.

5 Correlations Between the Variables Studied

This section presents the results of various statistical tests to determine whether the values of interest in this study are correlated. The results tend to suggest that most of the variables of interest are not correlated to one-another. This is important, because it means that the results from Section 4 *really do* suggest that the nature of dependence of predicates upon formals is qualitatively and quantitatively different from the nature of dependence of predicates upon globals.

Tests were run for correlations between the numbers of nodes, edges, predicates, formals, and du-globals found in each procedure. Over all 20 programs studied, there were 7,191 procedures, so each test concerns 7,191 data points. Both Spearman Rank Correlations and Pearson Linear Correlations were used. The only significant correlations found are those where (as one would expect) increases in the size of the procedure led to a corresponding increase in the number of predicates and the number of du-globals. Interestingly, there is no evidence to suggest that increases in the size of a procedure are accompanied by corresponding increases in the number of formals.

Fig. 13 shows the R scores for Pearson and Spearman tests. Using Pearson's test, the only linear correlation (having an R score greater then 0.50) was between the number of nodes and edges in a PDG ($R = 0.714$). Because nodes and edges are correlated, the remaining tests are reported for nodes only. The values for edges are quite similar; however, as the number of edges includes some measure of the level of dependence in the procedure and not just its size, the number of nodes is a better measure of the size of the procedure.

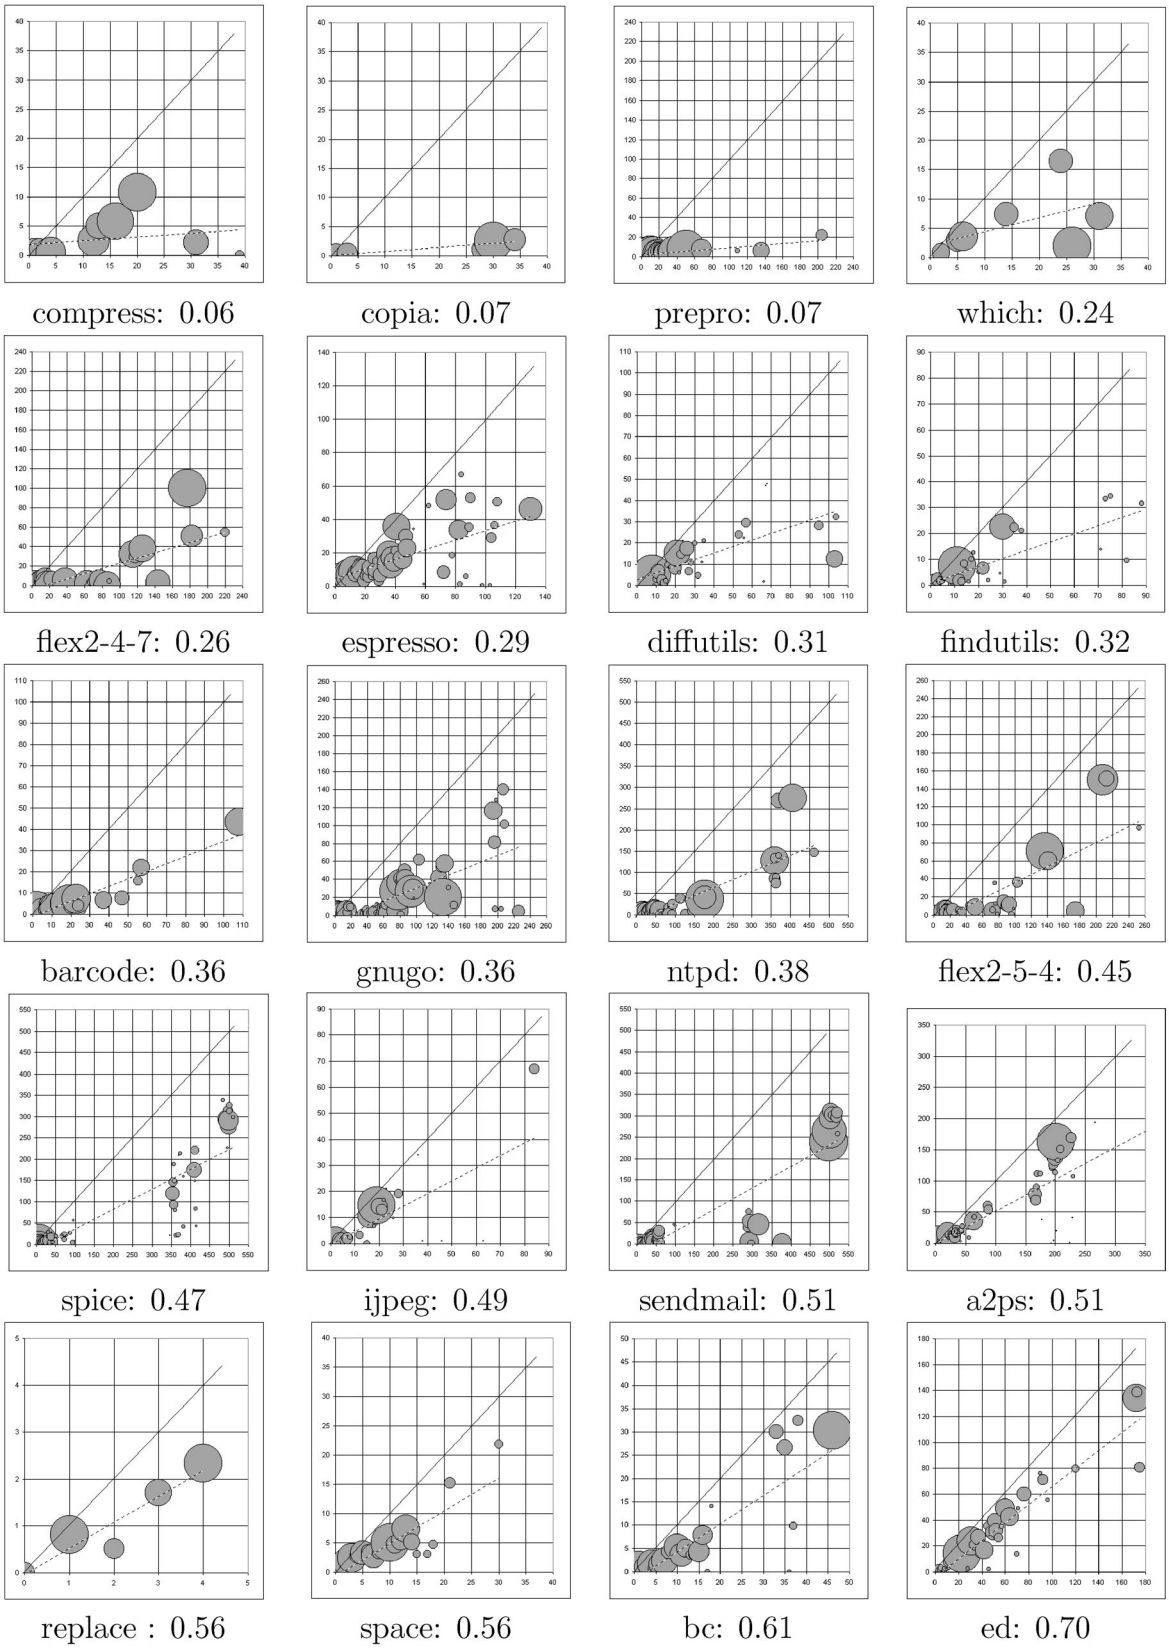

Fig. 8. Dependence bubble charts for du-global dependence. The number appearing after the name of each program is the slope of the dependence trend.

Using the Spearman test, in addition to the correlation between nodes and edges, three other correlations emerge. First, the number of predicates increases with the size of the procedure, but, as the Pearson R value is low, this relation is probably not a linear one. The remaining two relate the number of nodes to the number of du-globals and to the

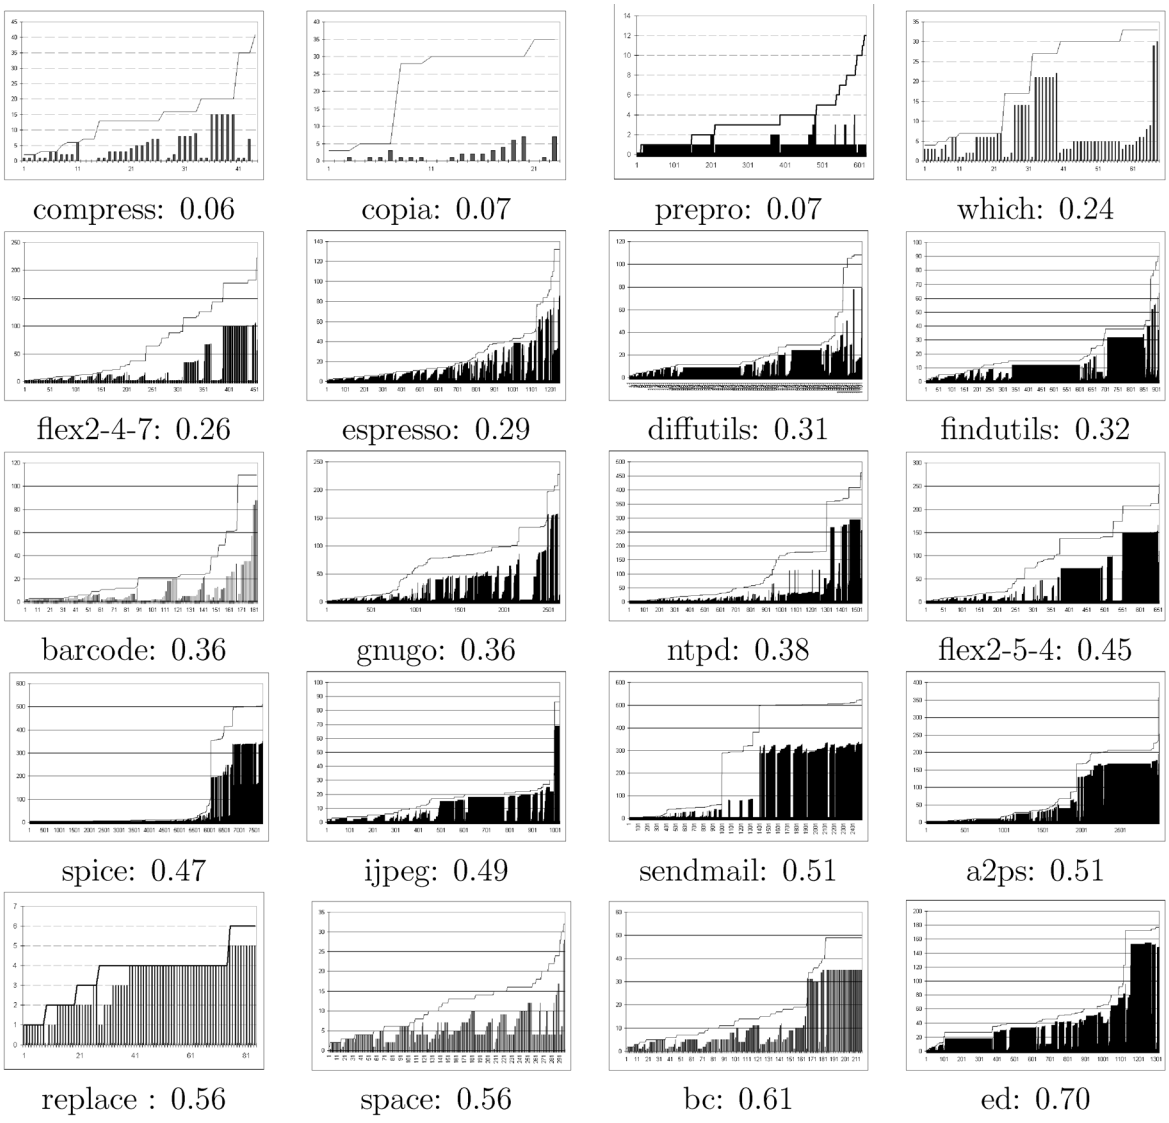

Fig. 9. Dependence skylines for du-global dependence. The number appearing after the name of each program is the slope of the dependence trend.

sum of the number of du-globals and the number of formals. In the final case, the number of du-globals being typically far larger than the number of formals dominates the statistics. The correlation with du-globals may seem surprising at first glance. However, du-globals are, by definition, "defined or used" in the procedure and, thus, roughly track its size.

Many of the tests considered in this section resulted in no correlation (neither strong nor weak being found). In many cases, the absence of a correlation (nonbold values in Fig. 13) is just as important as the presence of one. These results establish that the number of formal parameters is not correlated to the size of the procedure and that the number of formals and du-globals (the two primary dependence sources studied) are independent of one another. The observation of an absence of correlation between the number of formals in a procedure and the number of globals led to a set of further hypotheses about the relationship between the dependence sources available

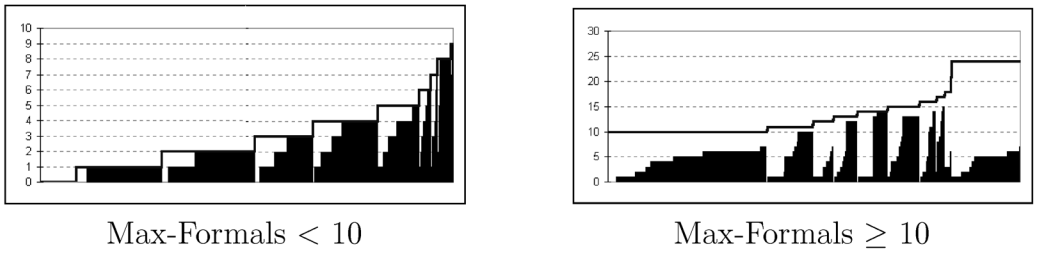

Fig. 10. Dependence skylines for all 50,579 predicates.

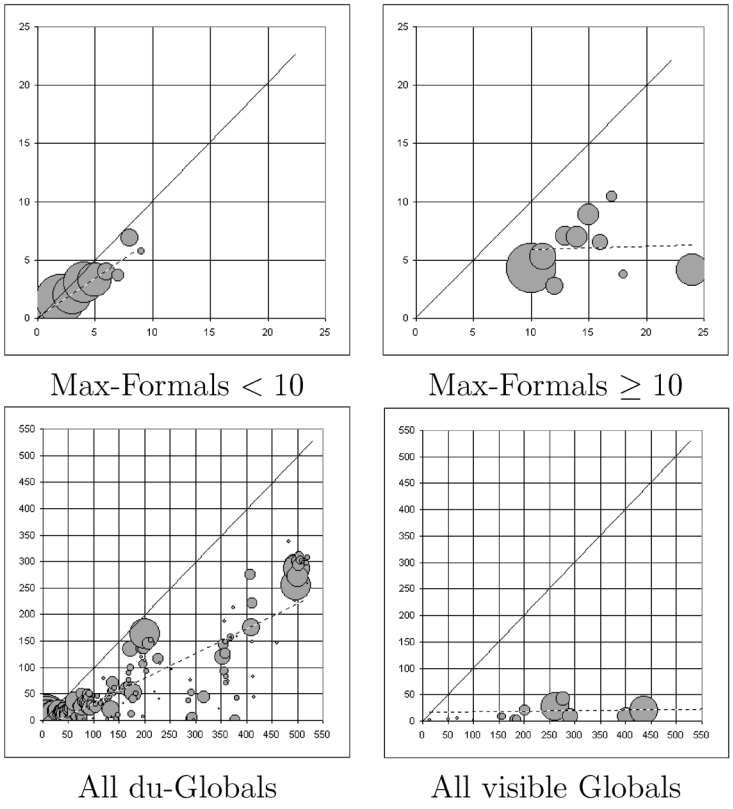

Fig. 11. Bubble charts for all 50,579 predicates. Note: scale of bubble size is different in each chart.

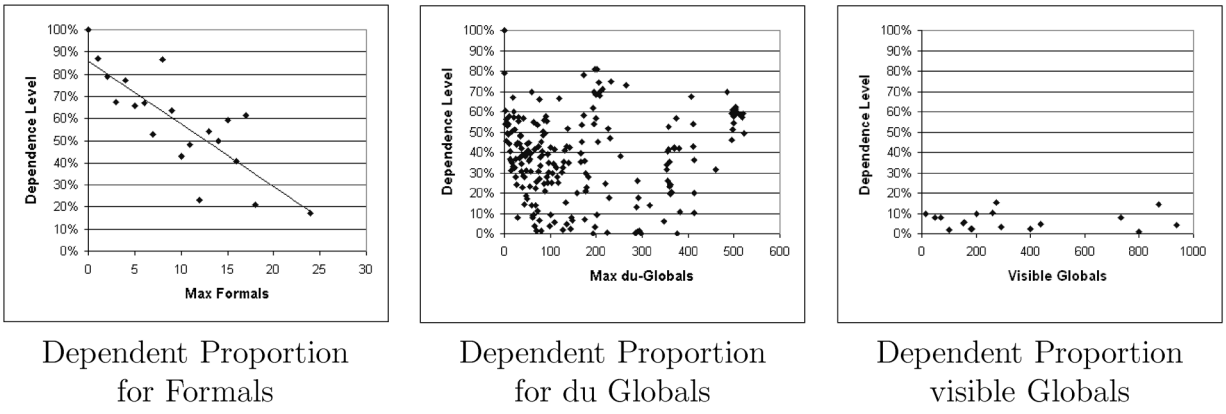

Fig. 12. Predicate dependence trends.

to a predicate (max-formals, max-globals, and max-visible-globals). No correlations were found, so a predicate which has a larger than average number of formals available to it does not necessarily have a larger than average number of du-globals available, etc. This provides evidence that these two dependence sources are independent of each other. The results over all predicates in all programs are shown in Fig. 14.

A test was also carried out for correlations between the number of formals and the number of du-globals available to a predicate in each program in isolation.[1] The results are shown in Fig. 15. Here, there does exist evidence of weak correlations for a few programs, but there is little compelling evidence for a trend.

In addition to asking whether the available numbers of each kind of dependence source are correlated, it is also sensible to ask whether the number actually used by a predicate are correlated. For example, it might be that, while the max-formals and the max-du-globals are not correlated, the number that the predicate uses are. If detected, this would suggest that predicates can be thought of as more or less dependent upon dependence sources *per se*, and that the distinction between particular *kinds* of dependence source is less important. Once again, however, both Spearman and Pearson tests show no correlations, yielding further evidence for the claim that the kinds of

1. The number of visible globals is essentially identical for all predicates in a program, so questions involving visible globals are not meaningful when considering each program independently.

X	vs.	Y	Pearson R Value	Spearman R Value
Nodes	vs.	Edges	**0.714**	**0.987**
Nodes	vs.	predicates	0.397	*0.663*
Nodes	vs.	formals	0.060	0.409
Nodes	vs.	du-globals	0.475	**0.840**
Nodes	vs.	vis-globals	0.063	0.198
Nodes	vs.	f+du-globals+vis-globals	0.086	0.321
Nodes	vs.	f+du-globals	0.476	**0.866**
Nodes	vs.	f+vis-globals	0.063	0.234
Nodes	vs.	du-globals+vis-globals	0.086	0.322
predicates	vs.	formals	0.117	0.374
predicates	vs.	du-globals	0.165	0.417
predicates	vs.	vis-globals	0.062	0.247
predicates	vs.	f+du-globals+vis-globals	0.070	0.301
predicates	vs.	f+du-globals	0.167	0.419
predicates	vs.	f+vis-globals	0.063	0.276
predicates	vs.	du-globals+vis-globals	0.070	0.304
formals	vs.	du-globals	0.064	0.247
formals	vs.	vis-globals	0.135	0.195
formals	vs.	du-globals+vis-globals	0.137	0.222
du-globals	vs.	vis-globals	0.177	0.059
du-globals	vs.	f+vis-globals	0.177	0.071
vis-globals	vs.	f+du-globals	0.179	0.096

Fig. 13. Pearson/Spearman correlations (bold values are strongly significant, italic values are weakly significant).

dependence source are independent of each other. The results are presented in Fig. 16.

The overall conclusion from this series of investigations into correlations between formals and globals is that the evidence all points to the finding that sources of dependence from formal parameters are wholly different in their nature and behavior from sources of dependence from globals. Combined with the findings from Section 4, it is reasonable to conclude that there is a qualitative and quantitative difference between formal and global dependence.

6 Applications for Dependence Visualizations

This section illustrates the application of the dependence analysis and the visualizations introduced in the paper. Applications that potentially benefit from dependence analysis and visulazations include program comprehension including understanding software evolution [6], [33], [39], [15], software surgery [17], testing including testability transformations and automatic structural testing, [16], [22], [52], [56], [70] array safety analysis, [12], and program flaw identification [67]. Three of these are illustrated in this section: Section 6.1 shows how the analysis of a predicate's dependence on formal parameters can be used to reduce test data generation effort in evolutionary testing. Sections 6.2 and 6.3 show the application of the visualizations as support and guidance for evolution and comprehension.

6.1 Application to Search-Based Test Data Generation

Search-based test data generation uses search techniques (such as genetic algorithms [27], [48]), to search for test data which meets some chosen test objective. The test objective is translated into a fitness function which guides the search, by determining when one test case is closer to meeting the objective than another.

More details about search-based test data generation can be found elsewhere [31], [32], [47], [50], [56], [57], [68], [71], [72]. For the work reported here, the important aspect of the technique is that it is search based. The space of possible inputs to the program or procedure constitutes the search space from which search-based testing attempts to locate a suitable test case. Consider, for example, the case of branch coverage. The test objective is to find an input that executes some branch of interest in the software under test. The search space is formed from the set of global variables and formal parameters available to the branch. For some chosen branch, controlled by a predicate p, not all of the search space will be relevant. Those global variables and formal parameters that cannot affect the value computed at p can be safely excluded from the search. In this way, the analysis

X	Y	Pearson R Value	Spearman R Value
Formals available	du-globals available	0.124	0.23
Formals available	visible globals available	0.008	0.27
du-globals available	visible globals available	0.0035	0.21

Fig. 14. Pearson/Spearman tests for dependence sources available to a predicate showing no correlation.

Program	Pearson R Value	Spearman R Value
a2ps	0.156	0.361
barcode	0.530	0.844
bc	0.274	0.561
compress	0.055	-0.070
copia	-0.396	0.524
diffutils	0.146	0.460
ed	0.248	0.764
espresso	0.347	0.745
findutils	0.354	0.537
flex	-0.047	0.110
gnugo	0.209	0.293
ijpeg	-0.137	-0.100
ntpd	0.090	0.263
prepro	0.181	0.648
replace	0.542	0.532
sendmail	0.326	0.389
space	0.628	0.750
spice	-0.036	-0.036
which	0.571	0.443

Fig. 15. Pearson/Spearman tests for dependence sources available to a predicate for each program in isolation.

used in the present paper can be used to augment search based testing by reducing the search space size.

To explore the potential of this possibility, a simple experiment was conducted, using the DaimlerChrysler Evolutionary Testing System [70]. The system uses genetic algorithms to search for test data that meet a variety of criteria. For this experiment, the system was configured to attempt branch coverage for the program shown in Fig. 17. The procedure takes, as input, 16 values. The first six (in variables `a` to `f`) affect the path through the procedure. The remainder (in the array `du`) do not. By varying the size of `du`, the effect of search space size can be investigated. In this case, the size of 0 and 10 provide the ratio 6 to 16, which approximates the average reduction for du-globals found in Section 3. The version with 16 input values has a larger search space, which potentially increases the test effort required to find the "right" input values to hit all branches. Predicate dependence analysis would determine that only the six formals `a` to `f` need be considered, thereby reducing the search space size and (hopefully) reducing test effort. To explore this potential, test data were generated for the 6 and 16 input versions of the procedure. Finally, observe that in this example, each predicate is dependent on two of the parameters. If, however, for example, the second predicate were nested within the first, then it would be dependent on four parameters because the slice would include the first predicate (via a control dependence edge).

Evolutionary algorithms are nondeterministic and contain some randomness in their behavior. To ensure that the results obtained were robust, data were collected from 10 separate trial runs of the test data generation system for each of the two versions of the program. The results are shown in Fig. 18. As can be seen, the optimized search (with the reduced search space) requires far fewer fitness evaluations. Here, the analysis of predicate dependence assists with the discovery of branch adequate test data generation. Even over all 10 runs, the worst-case effort for the optimized version (the tallest gray bar) is far better than the best case for the unoptimized version (the shortest white bar).

Of course, this initial set of results indicates merely that the approach is feasible and promising. A more detailed investigation of the application of predicate dependence analysis to search-based testing is the topic of ongoing and future work. However, the initial results are encouraging, particularly, since none of the current work on search-based testing [31], [47], [50], [56], [57], [60], [68] uses search-space reduction through source-code analysis.

6.2 Understanding Software Evolution

This section shows how the visualizations can be used to explore software evolution. It presents a case study in which two versions of the `flex` program were analyzed. These will be referred to as the "earlier version," `flex-2.4.7`, and the "later version," `flex-2.5.4`. In this section, the global variable bubble chart visualizations shown in Fig. 8 for each of the two versions will be used to explore some of the changes between the two versions. In this way, the section illustrates the way in that the bubble chart visualization might be used to support understanding of software evolution.

In general, looking at the bubbles in the global variable bubble charts of two versions of a program, a shift of bubbles to the right indicates that more globals are available to a typical predicate, while a shift upwards indicates that more globals (whether new or original) are being depended upon. An increase in bubble size indicates that predicates have been added.

Turning to the specific case of `flex`. A comparison the du-global bubble charts for the earlier and later versions shows both a rightward and an upward movement of bubbles. In general, tool support is required to identify comparable bubbles (those summarizing predicates of the same functions) with certainty. In practice, it can be done using only the visualizations with reasonable accuracy. For flex, comparable bubbles show a right shift of approximately 20-30 du-globals. For example, the earlier version's farthest right bubble is located at 220 on the horizontal axis (220 globals are available), while in the later version, the corresponding bubble is located at 253 (indicating 253 globals are available). Something has happened during the evolution of `flex` to increase the number of globals available to these predicates.

X	Y	Pearson R Value	Spearman R Value
Formals Used	du-globals Used	0.0081	0.002
Formals Used	visible globals Used	0.088	0.032
du-globals Used	visible globals Used	0.082	0.221

Fig. 16. Pearson/Spearman tests for dependence sources used by a predicate showing no correlation.

```
void testfunc(int a,int b, int c, int d, int e,int f, int du[10])
{
  if (a==100000 || b==100000)
    a++;
  if (c==100000 || d==100000)
    c++;
  if (e==100000 || f==100000)
    e++;
}
```

Fig. 17. Code used to experiment with the effect of dependence reduction.

Next, consider the two groups of four bubbles that have $x > 100$ and $y > 20$. That is, more than 100 globals available, but only 20 used. In moving from the earlier to the later version of `flex`, all four of these bubbles shift upwards (an indicator of increased reliance on globals).

Any such up-shift should be considered a warning that evolution has introduced potentially unwanted and undesirable data-dependencies. The bubbles direct attention toward the relevant sections of the code and the slices taken from it. For example, in the earlier version of `flex`, the slice on `if (performance_report > 0)` from function `readin` includes 32 LoC and 1 global (`performance_report`). By contrast, in the later version, the slice on this same statement includes 567 LoC and just over 100 globals. This is clearly a dramatic change in the dependence profile of the program and it merits further investigation.

The source code was analyzed to explore the reason for the increase in global dependence. It turns out that it had three principal causes. First, evolution of `flex` has added support for character class expressions; second, calls to standard I/O functions were replaced by calls to custom I/O functions; and, finally, there was an increase in command line option error checking. Centralization of I/O increases the interdependence between many of the globals variables. The consequence is that any predicates that depend upon some, tend to depend upon more of the rest. The increased error check ensures, *interalia*, that certain combinations of options are not permitted. Since these options are encoded in global variables, the increased error checking causes a dramatic increase in the number of globals, both available to and genuinely used by, the predicates. It also compounds the increased interdependence between globals.

Finally, all four bubbles that show significant up-shift also show a predicate count increase (depicted as a bubble size increase in the chart). For example, in `yyparse`, to track the beginning-of-line (bol), the action "`YY_USER_ACTION`" from each `switch` case in the earlier version is replaced, in the later version, by

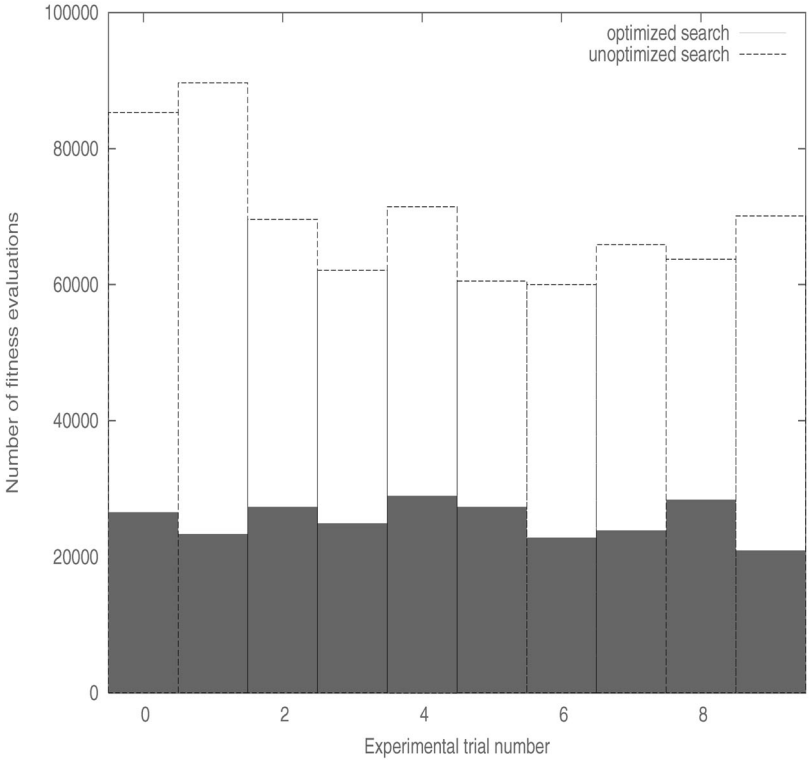

Fig. 18. Predicate analysis to optimize search-based test data generation.

function	# Nodes	# Predicates	# Formals	# du-Globals
go_..._compression	220	2	0	46
go_..._compression_and_decompression	258	1	0	63
main	301	0	2	75
parse_args	1,352	45	2	84

Fig. 19. Functions in ijpeg with largest numbers of du-globals available.

```
#define YY_RULE_SETUP \
  if ( yyleng > 0 ) \
    yy_current_buffer->yy_at_bol = \
         (yytext[yyleng - 1] == '\n'); \
  YY_USER_ACTION
```

This results in the insertion of a new predicate into each switch case.

Adding command line checking and character classes were clearly a move to improve the functionality of flex. The visualizations draw attention to these changes, but whether they are justified will depend upon the perceived importance of the improved functionality. However, the centralization of I/O is more of a perfective maintenance activity. It does not aim to improve functionality, but to reorganize the software, presumably with the goal of assisting further evolution in some way. The use of the bubble charts draws attention the impact of this change and raises questions about its justification.

6.3 Using the Visualizations to Assist Comprehension

In addition to the patterns seen over multiple version of a program, bubble charts provide visual clues as to abnormalities within a program. In this section, two features that stand out in bubble charts are explored in more detail. These features, the existence of outliers and gaps in the occurrence of bubbles on the horizontal axis, can aid in the understanding potential trouble spots within the code.

Outliers are evident in the global dependence bubble charts of ijpeg, barcode, gnugo, and flex2-5-4 (See Fig. 8). The outliers are collections of predicates (all having access to the same number of globals). However, all the outlier predicates tend to come from a single function; thus, in the following discussion "outlier function f" is often used in place of "outlier predicates from function f."

To begin with, ijpeg has a clear outlier in its dependence on globals (the small bubble on the top right). This bubble is from predicates for which the value of max-du-globals is large. The top four functions from ijpeg, when sorted by max-globals, are presented in Fig. 19. The outlier comes from predicates in the function parse_args, which has 45 predicates with 84 globals available to them. All of these predicates test for command-line options. Each is contained in a for loop that depends (transitively) on 67 globals. As each predicate is dependent on the for loop, each predicate depends on 67 globals. A fragment of the code is presented in Fig. 20. Interestingly, the other three functions from Fig. 19 are from the same file as parse_args. These three functions all have very low dependence and contain very few predicates (0, 1, and 2); thus, they produce the three very small bubbles (dots) near the horizontal axis.

The second example with an outlier is barcode. Barcode's main function has 108 du-globals available and 33 predicates. It looks like an outlier because its bubble is much further to the right than the others. Further investigation reveals that the call tree for barcode is broad and shallow. The function main (transitively) calls all functions. Therefore, it potentially uses all global variables. Because the call-tree is broad and shallow, no other function potentially uses such a large number of globals.

However, the main function appears on the trend line for dependence. Its du-globals are those of its called functions. It therefore has more du-globals available to it, but depends upon a similar proportion of these to the average for all the other functions. This can be contrasted with the outlier for ijpeg that, not only had more globals available to it, but contains predicates with a higher-than-average dependent proportion.

The gnugo program has three bubbles that denote sets of predicates with about 200 available du-globals, but that use fewer than 10 of these. These are the three bubbles that appear on the right hand side of the bubble chart, near the horizontal axis. Each comes from a different file, but each is an outlier for a similar reason; all call most of the program, initialize a large number of globals, and contain predicates that are largely unaffected by the initialized variables. Data

```
parse_args
{
    ...
    for(; argv[i]; i++) {
        ...
        else if( !strcmp( argv[i], "-difference.histogram" ) ) {
             histogram_flag = atoi( argv[++i] );
        }
   ...
}
```

Fig. 20. Fragment of code from ijpeg corresponding to outlier.

function	# Nodes	# Predicates	# Formals	# du-Globals
main	3,809	47	2	226
do_move	1,069	8	4	204
play_test	940	12	2	197

Fig. 21. Functions in gnugo with largest numbers of du-globals available.

for the three functions, corresponding to the three outlier bubbles are presented in Fig. 21.

For example, the function main contains 320 lines of code, 180 of that form a switch statement, used to process command-line arguments. A fragment of code that illustrates this is contained in Fig. 22. This accounts for most of the direct global uses. The remaining lines of main are primarily concerned with determining the "play mode" for the game.

The rest of the globals that main may use or modify come from called functions. For example, main directly calls function play_test (also an outlier), which potentially uses 197 globals. This function is a test harness that calls most of the rest of the program. Likewise, the function do_move calls most of the rest of the program.

By contrast to barcode, gnugo has longer transitive call chains in its call graph. This creates several functions, (main, do_move, and play_test), which potentially access a large number of globals. However, unlike barcode (and ijpeg), the globals being initialized do not affect the initialization process logic (i.e., its predicates). This explains why the three bubbles are "unusually close" to the horizontal axis.

The final outlier bubble chart is for flex2-5-4. Here, there is one bubble on the "far right" (near $x = 17$) that touches the horizontal axis. This is an outlier, likely to draw the viewer's attention, since it contains a large number of predicates potentially affected by a large number of globals, but that turn out to be affected by very few of these "potential" globals. The bubble corresponds to the function make_tables, which has 174 globals available to it and 56 predicates. This function merely outputs information including the values of many global variables. However, these variables do not affect any of the 56 predicates (which are used to determine that data to output). Furthermore, the predicates are all linear (un-nested) and simple, for example,

```
if ( C_plus_plus ).
```

In addition to outliers, a number of the bubble charts for global variables (see Fig. 8) show gaps in the occurrence of bubbles along the horizontal axis. This occurs, for example, in the global variable bubble charts for copia, findutils, bc, and a2ps. Examination of the code reveals that these were all caused by the occurrence of a set of mutually recursive functions. For example, program bc has a gap between the values 18 and 33 on the horizontal axis. There are six bubbles to the right of the gap, which have varying numbers of du-globals available to them. The data for these six functions and the others in bc are presented in Fig. 23.

The call graph for these six functions is depicted in Fig. 24. In bc, the function main calls try_file. As a result, it has available to it, all the du-globals of try_file together with a few of its own. Mutually recursive functions have broadly the same number of du-globals available to them since they are all in the transitive closure of each one's call graph. The reader may wonder why, therefore, the three mutually recursive functions dc_func, dc_eval_an, and dc_evalstr do not have an *identical* set of du-globals. They do not because some of them contain killing assignments to one of the potential du-globals. These killing assignments occur before the value of the potential du-global is referenced. In such situations, the dependence algorithm used in this study does not include them in the set of du-globals (a safe approximation since no predicate can depend on the initial value of such a global).

In all the cases studied, the presence of gaps on the horizontal axis was caused by the existence of mutual recursion among a set of functions. Of course, the fact that a program contains mutual recursion can also be determined, quite naturally, from the call graph of the program. However, the bubble charts allow for greater comprehension of the impact of such mutual recursion upon dependence. Consider the six mutually recursive functions in bc for an example. The size and vertical location of the bubbles indicates the number of predicates involved and the average dependence of these upon the set of du-globals. For example, dc_func has a large number of predicates (69) and they tend to depend heavily upon the set of du-globals. By comparison, dc_eval_an and try_file have very few predicates (1 and 4, respectively) and these depend upon very few of the du-globals available to them. This indicates, *interalia*, that dc_eval_an and try_file will be easier to generate test data for than dc_func.

7 Threats to Validity

For an experiment not involving human subjects, only two potential threats to validity need to be considered. These are

```
/* Now weed through all of the command line options.  */
 while ((i=getopt_long(argc, argv, "-ab:B:d:D:Egh::H:l:L:o:p:r:fsStTvwm:",
      long_options, NULL)) != EOF)
  {
    switch (i) {
      case 'T' : printboard++; set_opt_display_board(1); break;
      case 't' :  ++verbose; break;
      case 'a' :  allpats = 1; break;
...
```

Fig. 22. Fragment of code from gnugo for command line processing.

function	# Nodes	# Predicates	# Formals	# du-Globals
dc_binop	211	5	2	13
dc_triop	275	7	2	14
dc_out_num	85	2	4	15
out_num	557	20	3	15
dc_getnum	481	20	3	16
dc_print	95	2	4	17
dc_printal	70	1	1	18
··· gap ···				
dc_evalfil	498	11	1	33
dc_evalstr	332	12	1	35
dc_eval_an	149	1	1	36
try_file	180	4	1	37
main	472	6	2	38
dc_func	2315	69	3	46

Fig. 23. Analysis of the gap in the bubble chart for bc global dependence.

threats to external validity and internal validity. External validity, sometimes refereed to as selection validity, is the degree to which the findings can be generalized. In this experiment, it is possible that the selected programs are not representative of programs in general and thus results from the experiment do not apply to "typical" programs. The wide variety of programs used in the study helps to mitigate this concern. As most of these are open-source programs, it remains possible that nonopen source programs will exhibit different behavior.

Internal validity is the degree to which conclusions can be drawn about the causal effect of the independent variable on the dependent variable. In this experiment, potential threats come from inappropriate statistical tests or violations of statistical assumptions, inaccurate underlying analysis, and the degree to which the variables used in the study accurately measure the concepts they purport to measure (a form of construct validity). The statistical tests used were chosen based on past experiments and guidance of those trained in statistics. This serves to mitigate the possibility of an inappropriate test being employed. The analysis assumes that program slicing can determine the parameters used by a predicate. While, in theory, slicing does this, in practice, there are potential implementation

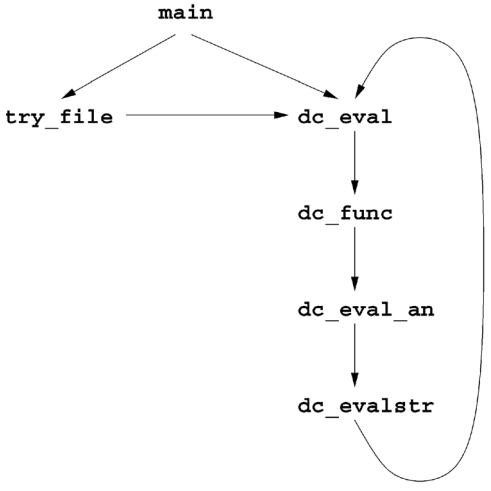

Fig. 24. Call graph of mutually recursive functions from bc.

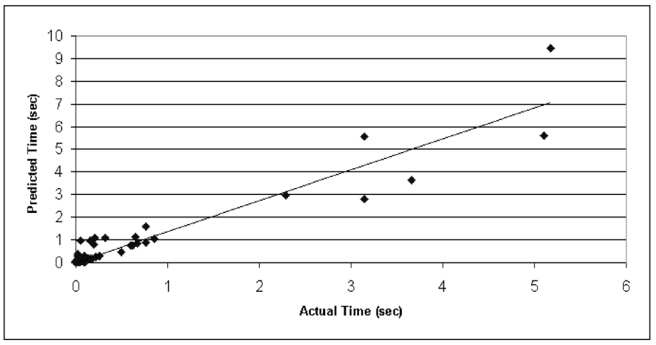

Fig. 25. Actual computation time versus predicted computation time per function.

issues. For example, the slicing tools used might contain an error, or employ imprecise analyses (e.g., imprecise data-flow analysis or imprecise points-to analysis). To mitigate these concerns, the tools used were all mature and well tested. These concerns could be further addressed by replication of the experiment using alternative analyses (e.g., Steensgaard's pointer analysis) and alternative slicing tools (e.g., Sprite [49] or Unravel [44]).

8 Theoretical and Actual Time Complexity

The previous three sections show the significant benefits possible using dependence analysis. However, these benefits can only be realized if it is possible to compute the dependence information in reasonable time. Fortunately, this is possible. This section describes the predicted and actual computational effort required to compute the dependence information presented in Sections 3 and 4. The algorithms takes as input and SDG. The construction and complexity of building an SDG are described elsewhere [29], [58], [18].

The algorithmic complexity of the analysis can be given in terms of the size of a procedure, S, and the number of predicates, P. The complexity of a brute force algorithm is $O(P \times S)$. This algorithm takes P slices, each having linear ($O(S)$) complexity. It is possible to reduce this to $O(S)$ assuming $O(1)$ bit vector operations, by propagating bit vectors over a topologically sorted DAG produced by first finding strongly connected components (SCCs) in the subgraph reachable from predicate nodes [64].

The actual data was collected on a PII 450 with 386Mb memory using an implementation of the brute force algorithm. The average time per predicate was 7.6ms. Statistical modeling of the computation time shows a tight fit with that predicted by its complexity. Using the sum of the number of edges and nodes in a PDG as the size measure, S, yields the following model ($R = 0.99995$)

$$analysis\ time\ (s) = 0.17 \times P \times S.$$

Furthermore, the t-score for the coefficient of $P \times S$ is 7,881, which, being much greater than 1.94, is statistically significant. This indicates an excellent fit as seen the scatter plot of actual and predicted time shown in Fig. 25. Note that most of the points in the plot occur right on the line with values less then 0.020.

			Time			
			All Predicates (seconds)		per Predicate (ms)	
Program	KLOC	Predicates	Total	Analysis	Total	Analysis
replace	0.6	84	0.10	0.01	1.19	0.12
copia	1.2	45	4.12	0.14	91.56	3.11
compress	1.9	88	0.22	0.01	2.50	0.11
which	5.4	136	0.43	0.02	3.16	0.15
barcode	5.9	365	1.34	0.06	3.67	0.16
space	11.5	593	2.06	0.10	3.97	0.17
ed	13.6	2627	11.34	5.60	4.32	2.13
prepro	14.8	619	2.67	0.10	4.31	0.16
bc	16.8	432	1.64	0.08	3.80	0.19
findutils	18.6	1819	10.58	6.43	5.82	3.53
diffutils	19.8	2212	11.59	6.83	5.24	3.09
flex2-4-7	21.2	911	4.38	0.96	4.81	1.05
flex2-5-4	21.5	1306	16.86	8.47	12.91	6.47
espresso	22.1	2497	11.68	0.84	4.62	0.34
ijpeg	28.2	2038	7.45	0.44	3.62	0.22
ntpd	45.6	3061	33.81	4.75	11.05	1.55
a2ps	62.9	5984	761.79	707.82	127.30	118.29
gnugo	81.7	5244	32.29	2.75	6.16	0.52
sendmail	85.7	4991	1248.98	122.74	250.21	24.60
spice	167.0	15527	1401.72	59.30	90.27	3.82
average	28.3	2529	121.90	43.25	20.52	7.64

Fig. 26. Computation times.

Fig. 26 shows computation times per program and per predicate for each of the 20 programs. The table separates out from the total running time, which is dominated by reading in the SDG, the time to perform just the dependence analysis. Programs `a2ps`, and to a lesser extent `sendmail`, are outliers in the per-predicate costs because they contain rather large functions. For example, the largest function, from `a2ps`, contains 34,252 nodes, which is more than the total number of nodes for some of the mid-sized programs.

9 Related Work

The analysis reported here was first presented by the authors in a preliminary version of this paper [9]. The analysis is a variable dependence analysis based upon a modified form of a standard program slicing algorithm [29]. Slicing was introduced by Weiser [73] in 1979 as a code extraction technique to determine, statically, the set of statements that potentially affect a set of variables at some chosen program point. It has since been extended to allow for dynamic analysis [35] and conditioned analysis [13]. There are several surveys on slicing in the literature, which cover slicing applications and techniques [8], [66] and empirical results collected about slicing [11].

There has been considerable previous work concerning program slice size and trade-offs between algorithmic precision and slice size [3], [10], [24], [37], [43], [49], [55], [62], [69], [79]. To some extent, these studies are also concerned with program dependence, albeit, indirectly, due to the way in which slices are constructed from a model of program dependence [29], [75]. However, the work reported here is the first large-scale study of predicate dependence.

The initial work on program slicing envisaged that slices would be an end product, presented to the user, to assist with comprehension [74] and debugging [45]. Much of the subsequent work on slicing has concerned the extension of slicing algorithms to cover language features such procedures [28], jumps [1], [5], [14], [20], [38] and pointers [43], [49], [54], [55], and to improve the efficiency of slicing algorithms [3], [30].

The work reported in the present paper follows an emerging trend in slicing research to use slicing analysis as a *means to an end rather than an end in itself*. For instance, previous work has shown that slicing can be used to construct metrics for assessing cohesion [7] and coupling [23]. Slicing has also been used as a support for software testing [19], [22], [25], reengineering [33], [34], [40], virus detection [41], and security [78]. In all these applications, though the slices do influence the overall results, they are "invisible" to the slice user.

There has also been previous work on using slicing and dependence analysis in visualization tools. For instance, Ball and Eick [4] developed a tool called SeeSlice, which allows the user to see the entire text of the program at once and to view the relative directness of dependence (in terms of distance from the slice node in the SDG). Balmas [6], developed a tool for dependence visualization, which allows the user to view the whole system at a high level of abstraction and to zoom in on particular aspects of dependence. Rilling [59] uses a hyperbolic tree (a kind of "fish eye view") to allow dependence information to be displayed in such as way that the size of the view corresponds to the focus of attention; less important details appear at the edge of the tree, where they are barely visible, while important dependencies appear much bigger.

All these visualization techniques are concerned with managing the sheer quantity of information present in a large system when one wishes to display all dependencies. The work reported here differs from this previous work on dependence visualization because it does not aim to show the dependency relationships themselves. Rather, it shows the *impact* of the dependencies upon the predicates of the program. The goal of the bubble chart is to allow

observations to be made at a high level about the effect of dependence upon the predicates (and thereby the control logic) of the program.

10 Summary and Future Work

This paper presented results from a study of 50,579 predicates in 20 C programs ranging from 600 lines of code to 167,000 lines of code. In general, it found that a typical predicate depends upon 72 percent of the formal parameters in the scope, 48 percent of the global variables potentially used or defined by a call to the predicate's procedure, and 2.4 percent of the global variables in scope.

The study discovered a strong inverse correlation between an increase in the number of formal parameters visible to a predicate and the proportion actually depended upon. There is no such correlation for either kind of global variable. The results of analysis for correlations provide strong evidence that the global and formal dependence sources are independent of one another and that the number of formals is independent of the size of the procedure that contains them.

It was shown how predicate dependence analysis can be useful as a supporting analysis for search based software testing. The analysis allows the search space to be reduced with a consequent decrease in the effort required to generate test data. More work is needed to provide a full evaluation of this potential application, but the initial results are encouraging.

The paper also introduced two visualizations for human analysis and interpretation of dependent proportions: bubble charts and skyline diagrams. It was shown that the bubble chart visualization can help with comprehension, by directing attention to code features that cause interesting properties in the visualizations, such as "gaps" in dependence and outlier bubbles in the bubble charts. It was also shown, for example, that bubble charts are useful in supporting the analysis of change between different releases of a system.

Future work is required to explore these possibilities fully and to consider other possible applications of the analysis techniques used. For instance, consider the situation where a function has an unusually large number of formals, with an unusually low dependence of predicates upon the values of these formals. Such a function would show up as an outlier on the formal parameter dependence bubble chart. It might be that this suggests a possibility for refactoring; the function needs to be split, with subsets of the original set of formals in each subfunction. The dependence analysis might be a guide, both to the identification of candidates for refactoring and to the determination of "cut points." Finally, future work will consider the correlations between the results of dependence analysis and other indicators of qualities of the software such as maintainability, reliability, and testability.

Acknowledgments

All three referees provided detailed comments and suggestions that helped to improve the presentation and content of the paper. DaimlerChrysler AG Berlin kindly made available their evolutionary test data generation tool, and André Baresel of DaimlerChrysler performed the experiments, for which results were reported in Section 6.1. The authors also wish to thank GrammaTech Inc (http://www.grammatech.com) for providing CodeSurfer, Gregg Rothermel for providing the `space` program, Chris Morel for his statistical guidance, and Erin Ptah who assisted with the visualizations. Mark Harman is supported, in part, by EPSRC grants GR/R98938, GR/M58719, GR/M78083, and GR/R43150 and by two development grants from DaimlerChrysler, Berlin. Dave Binkley is supported by the US National Science Foundation grant CCR-0305330.

References

[1] H. Agrawal, "On Slicing Programs with Jump Statements," *Proc. ACM SIGPLAN Conf. Programming Language Design and Implementation,* pp. 302-312, June 1994.

[2] L.O. Andersen, "Program Analysis and Specialization for the C Programming Language," PhD thesis, DIKU, Univ. of Copenhagen, May 1994.

[3] D.C. Atkinson and W.G. Griswold, "Effective Whole-Program Analysis in the Presence of Pointers," *Proc. ACM SIGSOFT Sixth Int'l Symp. Foundations of Software Eng. (FSE '98),* pp. 46-55, Nov. 1998.

[4] T. Ball and S.G. Eick, "Visualizing Program Slices," *Proc. Symp. Visual Languages,* A.L. Ambler and T.D. Kimura, eds., pp. 288-295, Oct. 1994.

[5] T. Ball and S. Horwitz, "Slicing Programs with Arbitrary Control-Flow," *Proc. First Conf. Automated Algorithmic Debugging,* P. Fritzson, ed., pp. 206-222, 1993.

[6] F. Balmas, "Using Dependence Graphs as a Support to Document Programs," *Proc. Second IEEE Int'l Workshop Source Code Analysis and Manipulation,* pp. 145-154, Oct. 2002.

[7] J.M. Bieman and L.M. Ott, "Measuring Functional Cohesion," *IEEE Trans. Software Eng.,* vol. 20, no. 8, pp. 644-657, Aug. 1994.

[8] D.W. Binkley and K.B. Gallagher, "Program Slicing," *Advances in Computing, Volume 43,* M. Zelkowitz, ed., pp. 1-50, 1996.

[9] D.W. Binkley and M. Harman, "An Empirical Study of Predicate Dependence Levels and Trends," *Proc. 25th IEEE Int'l Conf. Software Eng. (ICSE 2003),* pp. 330-339, May 2003.

[10] D.W. Binkley and M. Harman, "Results from a Large-Scale Study of Performance Optimization Techniques for Source Code Analyses Based on Graph Reachability Algorithms," *Proc. IEEE Int'l Workshop Source Code Analysis and Manipulation (SCAM 2003),* pp. 203-212, Sept. 2003.

[11] D.W. Binkley and M. Harman, "A Survey of Empirical Results on Program Slicing," *Advances in Computers,* vol. 62, pp. 105-178, 2004.

[12] D.W. Binkley, M. Harman, L.R. Raszewski, and C. Smith, "An Empirical Study of Amorphous Slicing as a Program Comprehension Support Tool," *Proc. Eighth IEEE Int'l Workshop Program Comprehension (IWPC 2000),* pp. 161-170, June 2000.

[13] G. Canfora, A. Cimitile, and A. DeLucia, "Conditioned Program Slicing," *Information and Software Technology, special issue on program slicing,* M. Harman and K. Gallagher, eds., vol. 40, pp. 595-607, 1998.

[14] J.-D. Choi and J. Ferrante, "Static Slicing in the Presence of Goto Statements," *ACM Trans. Programming Languages and Systems,* vol. 16, no. 4, pp. 1097-1113, July 1994.

[15] A. DeLucia, A.R. Fasolino, and M. Munro, "Understanding Function Behaviours through Program Slicing," *Proc. Fourth IEEE Workshop Program Comprehension,* pp. 9-18, Mar. 1996.

[16] R. Ferguson and B. Korel, "The Chaining Approach for Software Test Data Generation," *ACM Trans. Software Eng. and Methodology,* vol. 5, no. 1, pp. 63-86, Jan. 1996.

[17] K.B. Gallagher, "Evaluating the Surgeon's Assistant: Results of a Pilot Study," *Proc. Int'l Conf. Software Maintenance,* pp. 236-244, Nov. 1992.

[18] Grammatech Inc., The codesurfer slicing system, 2002.

[19] M. Harman and S. Danicic, "Using Program Slicing to Simplify Testing," *Software Testing, Verification and Reliability,* vol. 5, no. 3, pp. 143-162, Sept. 1995.

[20] M. Harman and S. Danicic, "A New Algorithm for Slicing Unstructured Programs," *J. Software Maintenance and Evolution,* vol. 10, no. 6, pp. 415-441, 1998.

[21] M. Harman and R.M. Hierons, "An Overview of Program Slicing," *Software Focus,* vol. 2, no. 3, pp. 85-92, 2001.

[22] M. Harman, L. Hu, R.M. Hierons, J. Wegener, H. Sthamer, A. Baresel, and M. Roper, "Testability Transformation," *IEEE Trans. Software Eng.,* vol. 30, no. 1 pp. 3-16, Jan. 2004.

[23] M. Harman, M. Okunlawon, B. Sivagurunathan, and S. Danicic, "Slice-Based Measurement of Coupling," *Proc. 19th ICSE Workshop Process Modelling and Empirical Studies of Software Evolution,* R. Harrison, ed., May 1997.

[24] M.J. Harrold and N. Ci, "Reuse-Driven Interprocedural Slicing," *Proc. 20th Int'l Conf. Software Eng.,* pp. 74-83, Apr. 1998.

[25] R.M. Hierons, M. Harman, and S. Danicic, "Using Program Slicing to Assist in the Detection of Equivalent Mutants," *Software Testing, Verification and Reliability,* vol. 9, no. 4, pp. 233-262, 1999.

[26] M. Hind, "Pointer Analysis—Haven't We Solved This Problem Yet?" *Program Analysis for Software Tools and Eng. (PASTE '01),* June 2001.

[27] J.H. Holland, *Adaption in Natural and Artificial Systems.* MIT Press, 1975.

[28] S. Horwitz, T. Reps, and D.W. Binkley, "Interprocedural Slicing Using Dependence Graphs," *Proc. ACM SIGPLAN Conf. Programming Language Design and Implementation,* pp. 25-46, June 1988.

[29] S. Horwitz, T. Reps, and D.W. Binkley, "Interprocedural Slicing Using Dependence Graphs," *ACM Trans. Programming Languages and Systems,* vol. 12, no. 1, pp. 26-61, 1990.

[30] S. Horwitz, T. Reps, and M. Sagiv, "Demand Interprocedural Dataflow Analysis," *Proc. SIGSOFT'95 Third ACM SIGSOFT Symp. Foundations of Software Eng. (FSE '95),* pp. 104-115, Oct. 1995.

[31] B.F. Jones, H.-H. Sthamer, and D.E. Eyres, "Automatic Structural Testing Using Genetic Algorithms," *The Software Eng. J.,* vol. 11, pp. 299-306, 1996.

[32] B.F. Jones, D.E. Eyres, and H.H. Sthamer, "A Strategy for Using Genetic Algorithms to Automate Branch and Fault-Based Testing," *The Computer J.,* vol. 41, no. 2, pp. 98-107, 1998.

[33] R. Komondoor and S. Horwitz, "Semantics-Preserving Procedure Extraction," *Proc. 27th ACM SIGPLAN-SIGACT Symp. Principles of Programming Languages (POPL-00),* pp. 155-169, Jan. 2000.

[34] R. Komondoor and S. Horwitz, "Effective Automatic Procedure Extraction," *Proc. 11th IEEE Int'l Workshop Program Comprehension,* May 2003.

[35] B. Korel and J. Laski, "Dynamic Program Slicing," *Information Processing Letters,* vol. 29, no. 3, pp. 155-163, Oct. 1988.

[36] B. Korel and J. Rilling, "Program Slicing in Understanding of Large Programs," *Proc. Sixth IEEE Int'l Workshop Program Comprenhesion (IWPC '98),* pp. 145-152, 1998.

[37] J. Krinke, "Evaluating Context-Sensitive Slicing and Chopping," *Proc. IEEE Int'l Conf. Software Maintenance (ICSM 2002),* pp. 22-31, Oct. 2002.

[38] S. Kumar and S. Horwitz, "Better Slicing of Programs with Jumps and Switches," *Proc. Fifth Int'l Conf. Fundamental Approaches to Software Eng. (FASE 2002),* pp. 96-112, 2002.

[39] S. Kusumoto, A. Nishimatsu, K. Nishie, and K. Inoue, "Experimental Evaluation of Program Slicing for Fault Localization," *Empirical Software Eng.,* vol. 7, pp. 49-76, 2002.

[40] A. Lakhotia and J.-C. Deprez, "Restructuring Programs by Tucking Statements Into Functions," *Information and Software Technology Special Issue on Program Slicing,* M. Harman and K. Gallagher, eds., vol. 40, pp. 677-689, 1998.

[41] A. Lakhotia and P. Singh, "Challenges in Getting Formal with Viruses," *Virus Bull.,* Sept. 2003.

[42] L.D. Larsen and M.J. Harrold, "Slicing Object-Oriented Software," *Proc. 18th Int'l Conf. Software Eng.,* pp. 495-505, 1996.

[43] D. Liang and M.J. Harrold, "Efficient Points-To Analysis for Whole-Program Analysis," O. Nierstrasz and M. Lemoine, eds., *Proc. Foundations on Software Eng.,* 1999.

[44] J.R. Lyle, D.R. Wallace, J.R. Graham, K.B. Gallagher, J.P. Poole, and D.W. Binkley, "Unravel: A CASE Tool to Assist Evaluation of High Integrity Software, Volume 1: Requirements and Design," Technical Report NISTIR 5691, US Department of Commerce, Technology Administration, Nat'l Inst. of Standards and Technology, Computer Systems Laboratory, Gaithersburg, Md., 20899, 1995.

[45] J.R. Lyle and M. Weiser, "Automatic Program Bug Location by Program Slicing," *Proc. Second Int'l Conf. Computers and Applications,* pp. 877-882, 1987.

[46] S. Mancoridis, B.S. Mitchell, Y.-F. Chen, and E.R. Gansner, "Bunch: A Clustering Tool for the Recovery and Maintenance of Software System Structures," *Proc. IEEE Int'l Conf. Software Maintenance,* pp. 50-59, 1999.

[47] C.C. Michael, G. McGraw, and M.A. Schatz, "Generating Software Test Data by Evolution," *IEEE Trans. Software Eng.,* vol. 27, no. 12, pp. 1085-1110, Dec. 2001.

[48] M. Mitchell, *An Introduction to Genetic Algorithms.* MIT Press, 1996.

[49] M. Mock, D.C. Atkinson, C. Chambers, and S.J. Eggers, "Improving Program Slicing with Dynamic Points-To Data," *Proc. Tenth ACM SIGSOFT Symp. Foundations of Software Eng. (FSE '02),* W.G. Griswold, ed., pp. 71-80, Nov. 2002.

[50] F. Mueller and J. Wegener, "A Comparison of Static Analysis and Evolutionary Testing for the Verification of Timing Constraints," *Proc. Fourth IEEE Real-Time Technology and Applications Symp. (RTAS '98),* pp. 144-154, June 1998.

[51] J. Ning, A. Engberts, and V. Kozaczynski, "Automated Support for Legacy Code Understanding," *Comm. ACM,* vol. 37, no. 5, pp. 50-57, 1994.

[52] A.J. Offutt, "An Integrated System for Automatically Generating Test Data," *Proc. First Int'l Conf. Systems Integration,* R.T. Ng et al., eds., pp. 694-701, Apr. 1990.

[53] A.J. Offutt, Z. Jin, and J. Pan, "The Dynamic Domain Reduction Approach to Test Data Generation," *Software Practice and Experience,* vol. 29, no. 2, pp. 167-193, Jan. 1999.

[54] A. Orso, S. Sinha, and M.J. Harrold, "Effects of Pointers on Data Dependences," *Proc. Ninth IEEE Int'l Workshop Program Comprehension (IWPC '01),* pp. 39-49, May 2001.

[55] A. Orso, S. Sinha, and M.J. Harrold, "Incremental Slicing Based on Data-Dependences Types," *Proc. IEEE Int'l Conf. Software Maintenance (ICSM 2001),* pp. 158-167, Nov. 2001.

[56] R.P. Pargas, M.J. Harrold, and R.R. Peck, "Test-Data Generation Using Genetic Algorithms," *The J. Software Testing, Verification and Reliability,* vol. 9, pp. 263-282, 1999.

[57] H. Pohlheim and J. Wegener, "Testing the Temporal Behavior of Real-Time Software Modules Using Extended Evolutionary Algorithms," *Proc. Genetic and Evolutionary Computation Conf.,* W. Banzhaf et al., eds., vol. 2, p. 1795, July 1999.

[58] T. Reps, S. Horwitz, M. Sagiv, and G. Rosay, "Speeding Up Slicing," *Proc. ACM Foundations of Software Eng. (FSE '94),* pp. 11-20, Dec. 1994.

[59] J. Rilling and S.P. Mudur, "On the Use of Metaballs to Visually Map Source Code Structures and Analysis Results Onto 3D Space," *Proc. 10th Working Conf. Reverse Eng.,* pp. 42-52, Oct. 2002.

[60] A. Schultz, J. Grefenstette, and K. Jong, "Test and Evaluation by Genetic Algorithms," *IEEE Expert,* vol. 8, no. 5, pp. 9-14, 1993.

[61] M. Shapiro and S. Horwitz, "The Effects of the Precision of Pointer Analysis," *Lecture Notes in Computer Science,* vol. 1302, pp. 16-34, 1997.

[62] S. Sinha, M.J. Harrold, and G. Rothermel, "System-Dependence-Graph-Based Slicing of Programs with Arbitrary Interprocedural Control-Flow," *Proc. 21st Int'l Conf. Software Eng.,* pp. 432-441, May 1999.

[63] B. Steensgaard, "Points-To Analysis in Almost Linear Time," *Proc. Conf. Record of the 23rd ACM SIGPLAN-SIGACT Symp. Principles of Programming Languages (POPL '96),* pp. 32-41, Jan. 1996.

[64] R.E. Tarjan, "Testing Flow Graph Reducibility," *J. Computer and System Sciences,* vol. 9, no. 3, pp. 355-365, Dec. 1974.

[65] F. Tip, "A Survey of Program Slicing Techniques," Technical Report CS-R9438, Centrum voor Wiskunde en Informatica, Amsterdam, 1994.

[66] F. Tip, "A Survey of Program Slicing Techniques," *J. Programming Languages,* vol. 3, no. 3, pp. 121-189, Sept. 1995.

[67] N. Tracey, J. Clark, and K. Mander, "Automated Program Flaw Finding Using Simulated Annealing," *Proc. Int'l Symp. Software Testing and Analysis,* pp. 73-81, Mar. 1998.

[68] N. Tracey, J. Clark, and K. Mander, "The Way Forward for Unifying Dynamic Test-Case Generation: The Optimisation-Based Approach," *Proc. Int'l Workshop Dependable Computing and Its Applications (DCIA),* pp. 169-180, Jan. 1998.

[69] G.A. Venkatesh, "Experimental Results from Dynamic Slicing of C Programs," *ACM Trans. Programming Languages and Systems,* vol. 17, no. 2, pp. 197-216, Mar. 1995.

[70] J. Wegener, A. Baresel, and H. Sthamer, "Evolutionary Test Environment for Automatic Structural Testing," *Information and Software Technology Special Issue on Software Eng. Using Metaheuristic Innovative Algorithms,* vol. 43, no. 14, pp. 841-854, 2001.

[71] J. Wegener, K. Grimm, M. Grochtmann, H. Sthamer, and B.F. Jones, "Systematic Testing of Real-Time Systems," *Proc. Fourth Int'l Conf. Software Testing Analysis and Rev. (EuroSTAR 96),* 1996.
[72] J. Wegener, H. Sthamer, B.F. Jones, and D.E. Eyres, "Testing Real-Time Systems Using Genetic Algorithms," *Software Quality,* vol. 6, pp. 127-135, 1997.
[73] M. Weiser, "Program Slices: Formal, Psychological, and Practical Investigations of an Automatic Program Abstraction Method," PhD thesis, Univ. of Michigan, Ann Arbor, 1979.
[74] M. Weiser, "Programmers Use Slicing When Debugging," *Comm. ACM,* vol. 25, no. 7, pp. 446-452, July 1982.
[75] M. Weiser, "Program Slicing," *IEEE Trans. Software Eng.,* vol. 10, no. 4, pp. 352-357, 1984.
[76] S.H. Yong, S. Horwitz, and T. Reps, "Pointer Analysis for Programs with Structures and Casting," *Proc. SIGPLAN 99 Conf. Programming Language Design and Implementation,* pp. 91-103, May 1999.
[77] J.-s. Yur, B.G. Ryder, and W.A. Landi, "An Incremental Flow- and Context-Sensitive Pointer Aliasing Analysis," *Proc. 21st Int'l Conf. Software Eng.,* pp. 442-452, May 1999.
[78] X. Zhang and R. Gupta, "Hiding Program Slices for Software Security," *Proc. First Ann. IEEE/ACM Int'l Symp. Code Generation and Optimization,* pp. 325-336, Mar. 2003.
[79] X. Zhang, R. Gupta, and Y. Zhang, "Precise Dynamic Slicing Algorithms," *Proc. 25th IEEE Int'l Conf. and Software Eng. (ICSE 2003),* pp. 319-329, May 2003.

David Binkley is a professor of computer science at Loyola College in Maryland. He has worked on slicing since 1988. He was instrumental in the development of the Wisconsin Program Slicing System, through his seminal PhD work on the System Dependence Graph (SDG) Slicing algorithm. From 1993 to 1999, he worked on the Unravel slicing system for the National Institute of Standards and Technology (NIST) and, from 2000 to 2001, he worked full time on the CodeSurfer commercial slicing system for GrammaTech Inc. More recently, he has worked on amorphous slicing and empirical studies of slicing; work funded by the US National Science Foundation.

Mark Harman is a professor of software engineering and head of the Software Engineering Group at the Department of Computer Science, King's College, London. He has worked extensively on program slicing, transformation, and testing, and, more recently, he founded the field of search-based software engineering. He is funded by the UK Engineering and Physical Sciences Research Council (EPSRC) and DaimlerChrysler, Berlin. He is program chair for several conferences and workshops including the 20th International Conference on Software Maintenance and the Search-Based Software Engineering Track of the Genetic and Evolutionary Computation Conference.