[2] G. Booch, *Software Components in Ada*, Benjamin Cummings, 1987.

[3] S. Chidamber and C. Kemerer, "A Metrics Suite for Object-Oriented Design," *IEEE Transactions on Software Engineering*, 1994, to appear.

[4] S. Conte, V. Shen and H. Dunsmore, *Software Engineering Metrics and Models*, Benjamin Cummings, 1986.

[5] D. Harms and B. Weide, "Copying and Swapping: Influences on the Design of Reusable Software Components," *IEEE Transactions on Software Engineering*, 17, 5 (May 1991), 424-435.

[6] J. Hollingsworth, B. Weide and S. Zweben, "Confessions of Some Used-Program Clients," *Proceedings of the Fourth Annual Workshop on Software Reuse*, Herndon, VA, November 1991.

[7] J. Hollingsworth, *Software Component Design-for-Reuse: A Language-Independent Discipline Applied to Ada*, Ph.D. dissertation, Dept. of Computer and Info. Sci., Ohio State University, Columbus, OH, Aug. 1992. Available by anonymous FTP from ftp.cis.ohio-state.edu in directory pub/tech-report/1993/TR01.

[8] W.R. LaLonde, "Designing Families of Data Types Using Exemplars," *ACM Trans. on Prog. Lang. and Systems*, 11, 2 (1989), 212-248.

[9] G. Milliken and D. Johnson, *Analysis of Messy Data, Vol. 1: Designed Experiments*, Van Nostrand Reinhold, 1984.

[10] R. Pressman, *Software Engineering: A Practitioner's Approach*, Third Edition, McGraw Hill, 1992.

[11] D. Ratkowski, M. Evans, and J.R. Alldredge, *Cross-Over Experiments*, Marcel Dekker, Inc., New York, 1993.

[12] E. Soloway and S. Iyengar, eds., *Empirical Studies of Programmers*, Ablex, 1986.

[13] B. Weide, W. Ogden and S. Zweben, "Reusable Software Components," in *Advances in Computers*, vol. 33, M.C. Yovits, ed., Academic Press, 1991, 1-65.

were trivial and these relatively trivial defects might be just as likely to be made when using direct implementation as they are when using layering. If the fraction of relatively trivial defects was high for many of the subjects, it then will be difficult to obtain statistical significance. With the small number of defects observed in these data, it is somewhat remarkable that we obtained *any* significant effects for this dependent variable.

A final item worth noting is that, had the analysis of the defect data for experiment 1 *not* corrected for overdispersion, it would falsely have concluded that there *was* a significant effect favoring layering. This illustrates how using the wrong statistical analysis in software engineering experiments such as these can mistakenly support the use of a particular technology even when the data do not really indicate such support.

6 Conclusion

The results of our experiments support the contention that, by using only a description of the functionality of and interfaces to existing components, new components can be developed with less effort than that required if the source code and representation data structures of the existing components are also used. In addition, it appears from our experiments that there certainly is no loss in the quality of the development process, at least in terms of the number of defects made during development, when the layering approach is used.

The empirical studies described in this paper illustrate interesting issues in the statistical analysis procedure, issues which, based on the authors' experiences, are not well-known to software engineering researchers. It is important that the proper analysis is used, lest the wrong conclusions be reached regarding the benefits of a particular method.

We are planning further experiments to more carefully analyze the defects made in the development of components such as those studied herein. It is important not only to characterize the kinds of defects observed, but also to provide if possible some cognitive explanation of these observations. We also are planning other studies to try to replicate the results reported herein. Studies such as these serve to provide a more sound and scientific basis for using (or not using) various software engineering methods.

Acknowledgment

The authors thank Rob Leighty of Ohio State University's Department of Statistics for his assistance with the statistical analyses.

References

[1] M. Aitkin, et al, *Statistical Modelling in GLIM*, Oxford Series Publications, 1989.

Source	df	Mean Sq.	f
Sequence	1	6624.8	0.13
Subject (within Sequence)	8	49544.4	
Treatment	1	69148.8	3.67
Sequence × Treatment	1	33784.2	1.79
Subject × Treatment (within Sequence)	8	18859.9	

Table 18: Experiment 3—Total Time Analysis

Treatment	Sequence				Overall	
	Layered First		Direct First			
	Mean	SD	Mean	SD	Mean	SD
Layered	3.00	2.45	1.00	1.22	2.00	2.11
Direct	5.00	4.85	2.80	2.49	3.90	3.81
Overall	4.00	3.77	1.90	2.08	2.95	3.15

Table 19: Experiment 3—Defect Data Summary

Total time, including that for testing, debugging and recoding, also was significantly better using layering when the component was an enhancement. This result held for both the simple and familiar queue, and the more complicated and less familiar partial map. The modification task provided no such significant effect, though the mean total time for layering was less than that for direct implementation. One possible explanation for this is that some defects made in the modification task tended to be nastier than those made in the enhancement task. It turned out that the average number of defects per subject was slightly higher for the enhancement task than for the modification task. If some of the defects for the modification task were, indeed, trickier, the time to debug and repair these defects would occupy a greater fraction of total time. Apparently, the *nature* of the defects was such that this debugging and repair time was not a function of the treatment, so the gain for layering in initial design and coding time is ameliorated when the debugging and repair time is added. Note that this could mean that trivial defects were just as trivial and nasty defects were just as nasty, whether layering or direct implementation was used.

No consistent effects in favor of layering were found for the defect data, though the mean number of defects was less for layering in each experiment. Here the small scale of the experiment may influence the results. Almost every subject had fewer than five defects. Some of these defects likely

Source	df	Mean Sq.	f
Sequence	1	2.57	0.84
Treatment	1	2.09	0.70
Sequence × Treatment	1	0.25	0.09
Error	16	2.98	

Table 20: Experiment 3—Defect Analysis (Poisson Regression)

Treatment	Sequence				Overall	
	Layered First		Direct First			
	Mean	SD	Mean	SD	Mean	SD
Layered	65.4	28.8	50.6	33.5	58.0	30.5
Direct	112.4	65.1	133.8	52.6	123.1	56.9
Overall	88.9	53.5	92.2	60.4	90.6	55.6

Table 15: Experiment 3—Initial Design/Coding Time Data Summary

Source	df	Mean Sq.	f
Sequence	1	54.5	0.02
Subject (within Sequence)	8	3609.0	
Treatment	1	21190.1	24.42*
Sequence $\times$ Treatment	1	1638.1	1.89
Subject $\times$ Treatment (within Sequence)	8	867.8	

Table 16: Experiment 3—Initial Design/Coding Time Analysis

the treatment effect, using all of the data, revealed a significant treatment effect in favor of layering (Tables 15 and 16).

The analysis of total time also revealed no significant sequence effect, nor a significant sequence by treatment interaction. However, the treatment effect was not significant (Tables 17 and 18).

The Poisson analysis of the defect data revealed no significant sequence effect nor a significant sequence by treatment interaction. After correcting for overdispersion, the test on treatment was not significant (Tables 19 and 20).

5.4 Discussion

For this set of experiments, we found that the use of layering consistently resulted in significantly faster time to complete the initial design and coding of the new component. This result held as the component changed from one with which the subjects were quite familiar to one with which the subjects were less familiar (and which was more complex as measured by the number of operations encapsulated in it and total lines of code). The result also held as the task changed from a component enhancement to a component modification.

Treatment	Sequence				Overall	
	Layered First		Direct First			
	Mean	SD	Mean	SD	Mean	SD
Layered	157.8	99.4	112.0	83.6	134.9	89.9
Direct	193.2	112.9	311.8	327.4	252.5	239.2
Overall	175.5	102.0	211.9	248.7	193.7	185.9

Table 17: Experiment 3—Total Time Data Summary

Source	df	Mean Sq.	f
Sequence	1	123873.8	3.37
Subject (within Sequence)	8	36725.2	
Sequence × Treatment	1	60280.2	12.59*
Subject × Treatment (within Sequence)	8	4786.8	
Treatment (first period only)	1	119163.0	14.35*
Error	8	8302.4	

Table 12: Experiment 2—Total Time Analysis

Treatment	Sequence				Overall	
	Layered First		Direct First			
	Mean	SD	Mean	SD	Mean	SD
Layered	0.80	1.79	0.80	0.84	0.80	1.32
Direct	2.60	1.67	6.60	2.61	4.60	2.95
Overall	1.70	1.89	3.70	3.56	2.70	2.96

Table 13: Experiment 2—Defect Data Summary

Source	df	Mean Sq.	f
Sequence	1	4.24	2.37
Treatment	1	16.52	9.23*
Sequence × Treatment	1	0.79	0.44
Error	16	1.79	

Table 14: Experiment 2—Defect Analysis (Poisson Regression)

Treatment	Sequence				Overall	
	Layered First		Direct First			
	Mean	SD	Mean	SD	Mean	SD
Layered	69.8	16.6	119.0	103.8	94.4	74.7
Direct	127.6	29.6	273.2	131.1	200.4	118.0
Overall	98.7	37.9	196.1	138.0	147.4	110.4

Table 9: Experiment 2—Initial Design/Coding Time Data Summary

Source	df	Mean Sq.	f
Sequence	1	47433.8	3.63
Subject (within Sequence)	8	13082.6	
Sequence $\times$ Treatment	1	11616.2	7.86*
Subject $\times$ Treatment (within Sequence)	8	1477.5	
Treatment (first period only)	1	41371.6	18.87*
Error	8	2912.0	

Table 10: Experiment 2—Initial Design/Coding Time Analysis

the first period data, revealed that the layering approach required significantly less effort ($f = 18.87$, $F_{1,8,.95} = 5.32$).

The analysis of total time gave results similar to that of initial design and coding time. Tables 11 and 12 show the data and analysis; the treatment test using only the first period data was significant in favor of layering ($f = 14.35$, $F_{1,8,.95} = 5.32$). Again, this replicated the results of Experiment 1.

The Poisson analysis of the defect data showed no significant effect either for sequence or sequence by treatment interaction. The treatment test, after correcting for overdispersion, showed a significant difference in favor of layering (i.e., the layering approach gave rise to significantly fewer defects). This result differed from that in Experiment 1, where no significant effect was observed (see Tables 13 and 14 for details).

5.3 Experiment 3 — Partial Map Modification

In the analysis of initial design and coding time for the partial map modification experiment, there was no significant effect for sequence nor for the sequence by treatment interaction. The analysis of

Treatment	Sequence				Overall	
	Layered First		Direct First			
	Mean	SD	Mean	SD	Mean	SD
Layered	113.0	62.9	160.6	147.3	136.8	109.7
Direct	191.0	40.5	458.2	236.1	324.6	212.9
Overall	152.0	64.6	309.4	242.9	230.7	190.9

Table 11: Experiment 2—Total Time Data Summary

Source	df	Mean Sq.	f
Sequence	1	3287.1	0.30
Subject (within Sequence)	16	11087.7	
Sequence × Treatment	1	33878.8	7.16*
Subject × Treatment (within Sequence)	16	4748.5	
Treatment (first period only)	1	13735.8	7.81*
Error	16	1759.6	

Table 6: Experiment 1—Total Time Analysis

Treatment	Sequence				Overall	
	Layered First		Direct First			
	Mean	SD	Mean	SD	Mean	SD
Layered	0.89	0.93	1.33	1.41	1.11	1.18
Direct	2.67	3.57	4.22	3.99	3.44	3.76
Overall	1.78	2.69	2.78	3.26	2.28	2.99

Table 7: Experiment 1—Defect Data Summary

Source	df	Mean Sq.	f
Sequence	1	1.38	0.48
Treatment	1	7.84	2.72
Sequence × Treatment	1	<0.01	<0.01
Error	32	2.88	

Table 8: Experiment 1—Defect Analysis (Poisson Regression)

Treatment	Sequence				Overall	
	Layered First		Direct First			
	Mean	SD	Mean	SD	Mean	SD
Layered	99.1	41.8	57.2	21.3	78.2	38.7
Direct	94.0	44.3	180.4	65.4	137.2	70.1
Overall	96.6	41.9	118.8	79.0	107.7	63.4

Table 3: Experiment 1—Initial Design/Coding Time Data Summary

Source	df	Mean Sq.	f
Sequence	1	4466.7	1.84
Subject (within Sequence)	16	2434.7	
Sequence $\times$ Treatment	1	37056.3	20.73*
Subject $\times$ Treatment (within Sequence)	16	1787.5	
Treatment (first period only)	1	6609.7	14.09*
Error	16	469.1	

Table 4: Experiment 1—Initial Design/Coding Time Analysis

= 163.4 min., mean for direct first = 182.6 min., $f = 0.30$, $F_{1,16,.95} = 4.49$), there was a significant sequence by treatment interaction ($f = 7.16$, $F_{1,16,.95} = 4.49$), and the first period data showed that the layering approach was significantly faster ($f = 7.81$, $F_{1,16,.95} = 4.49$). Tables 5 and 6 contain the relevant data and ANOVA, respectively, for total time.

The Poisson analysis of the defect data revealed no significant sequence effect (mean for layering first = 1.78 , mean for direct first = 2.78 , $f = 0.48$, $F_{1,32,.95} = 4.16$) nor a significant sequence by treatment interaction ($f < 0.01$, $F_{1,32,.95} = 4.16$). Tables 7 and 8 contain the relevant statistics. There was no significant treatment effect after correcting for overdispersion ($f = 2.72$, $F_{1,32,.95} = 4.16$).

5.2 Experiment 2 — Partial Map Enhancement

The analysis of initial design and coding time for the partial map enhancement was similar to that for the queue enhancement. There was no significant sequence effect, but there was a significant sequence by treatment interaction (Tables 9 and 10). The test for the treatment effect, using only

Treatment	Sequence				Overall	
	Layered First		Direct First			
	Mean	SD	Mean	SD	Mean	SD
Layered	145.1	81.4	102.8	47.6	123.9	68.2
Direct	181.8	111.0	262.3	102.4	222.1	111.5
Overall	163.4	96.3	182.6	112.8	173.0	103.8

Table 5: Experiment 1—Total Time Data Summary

that the mean and variance are equal. The well-known vast differences in subjects makes this assumption unlikely to be met in our experiments. When the true variance is higher than that assumed by the Poisson model, a phenomenon called "overdispersion" is present. The statistical analysis therefore must deal not only with non-normality, but also with overdispersion. Without correcting for overdispersion, for example, the test for significance of the treatment (approach) may falsely indicate significance. Fortunately, there are statistical procedures to deal with this, too [1].

5 Results

In each of the analyses that follow, the significance tests are done using the standard alpha level of .05, though a case might be made that a higher alpha level (e.g., .10) would be reasonable. In the first experiment, the subjects already were very familiar with the component in question (the queue). They had seen, and likely used several times, the standard linked representation structure used in this experiment. Hence, the simplicity and familiarity of the component itself might mask true effects due to the treatment. Prior to the second and third experiments, the partial map component used in these experiments had been studied by the subjects in an exercise in which they did an implementation using hashing. A hashing implementation was, in fact, used as the representation of the partial map component provided in experiments two and three (though it likely would not have used exactly the same code as did any of the subjects). Hence, we felt that, if our experiments were biased at all, they were biased in favor of our *not* getting a significant treatment effect when in fact one was present. Therefore, our judgment was that, if the experiment showed a significant treatment effect, it was not likely to be spurious.

5.1 Experiment 1 — Queue Enhancement

The analysis of variance for this experiment, using initial design and coding time, revealed no significant effect for sequence (mean for layering first = 96.6 min., mean for direct first = 118.8 min., $f = 1.84$, critical $F_{1,16,.95} = 4.49$), but did indicate a significant sequence by treatment interaction ($f = 20.73$, $F_{1,16,.95} = 4.49$). The means and standard deviations for each of the four cells are shown in Table 3, while the relevant components of the ANOVA table are shown in Table 4. In the ANOVA tables, statistically significant f values (at the .05 level) are indicated by an asterisk.

Following the approach outlined in the previous section, the treatment effect was tested using only the first period data, with the result that the layering approach required significantly less effort ($f = 14.09$, $F_{1,16,.95} = 4.49$).

When total time was used as the dependent variable, a situation similar to that for initial design and coding time was observed. There was no significant effect for sequence (mean for layering first

4 Statistical Analysis Issues

The analysis of variance (ANOVA) model for these experiments allows us to decide if there is a significantly different effect, on each dependent variable, of the layering approach versus the direct implementation approach (in statistical terminology, this is the "treatment effect" in the experiment). In addition, the model allows us to decide if there is a significant difference in the two possible sequences (orders) in which these two approaches were employed. Finally, it allows us to assess if the treatment and the sequence interact; that is, we can test if there is a different effect of the approach depending on the order in which the two approaches were employed. Since each subject did the same task twice (once for each of the two approaches), but did so in only one of the two possible sequences, in statistical terminology the subjects are nested within sequences. Thus, this model is sometimes called a "nested factorial" design. In the model, the variance associated with subjects (nested in sequences) is used to test for the sequence effect, while the variance of the subject by treatment interaction is used to test for the treatment effect and the treatment by sequence interaction. Some texts call this latter variance the "time error term" [9].

However, the normal nested factorial analysis of variance may not be appropriate for these experiments, for two reasons.

1. An important concern is the potential that a subject's having done a task once will affect performance on the second attempt at the task (even though a different approach is being used). This possibility of a "carryover effect," as it is known in statistics, requires that care be taken when deciding if there really is an effect due to the approach used.

   A common and accepted way of dealing with this issue is first to test if there is a significant sequence effect and/or a significant interaction effect. If there is neither, then the normal nested factorial analysis of variance is used to test for the effect of the treatment (i.e., the approach). If either the sequence effect or the interaction effect is significant, however, the treatment effect is tested using only the data for the first period in the sequence. That is, for those subjects who did the layering approach first, only their data for the layering approach is used; for those who did the direct implementation first, only their data for the direct implementation approach is used. The error estimate used for this test is a function of the subject error and time error terms [9].

2. The defect data typically will comprise small integer values, whose distribution does not satisfy the normality assumptions required of standard analysis of variance. Defect data may be modeled more accurately by a Poisson-like distribution, and there are statistical procedures to do Poisson (regression) analysis, allowing tests for the significance of the sequence, treatment, and sequence by treatment interaction effects. However, a Poisson distribution would imply

Treatment	Sequence	
	Layered First	Direct First
Layered	9 subj. in Exp. 1	9 subj. in Exp. 1
Direct	5 in Exp. 2 and 3	5 in Exp. 2 and 3

Table 2: Factors in the Experimental Design

improve the subjects' abilities to do the direct implementation version of the enhancement, relative to their ability to do the layered version. Ten subjects participated in this experiment.

The third experiment was a "modification of functionality" experiment using the partial map component. Given the basic partial map package used in the second experiment, the task was to create a component that encapsulated the concept of an *almost constant map*, which maps all but a finite number of domain elements to a "default" range value. Other than `Initialize`, `Finalize`, and `Swap`, the operations required by this package are as follows.

- `Reset` Reset the map to a constant function in which all domain elements map to a specified default value.
- `Get_Default` Return the default value, to which most of the elements are mapped.
- `Swap_Range_Value` Modify the range value associated with a given domain element.
- `Remove_Any_Anomaly` Make an arbitrary domain element, presently not mapping to the default value, instead map to the default value.
- `Test_If_Anomaly` Test if a given domain value maps to the default value.
- `Test_If_Constant` Test if all domain values map to the default value.

In this experiment, we felt that the subjects would be creating a rather unfamiliar component using what, by now, would be a somewhat familiar component. The same ten subjects who participated in the second experiment also participated in this third experiment.

To mitigate the differential effects of subjects on task completion, in each experiment each subject performed the task using *both* the layering and direct implementation approaches. The order in which the two approaches were followed was assigned randomly, subject to a counterbalancing with respect to the experience level of the subjects (i.e., we didn't want more experienced subjects doing the tasks in one order while the less experienced subjects did the tasks in the other order). Thus, as illustrated in Table 2, each of the experiments was a two-period crossover design [9, 11].

understood unbounded queue component. In addition to the "universal" package operations of `Initialize`, `Finalize`, and `Swap` [13, 5, 7], the basic Ada queue package provided the standard `Enqueue`, `Dequeue`, and `Is_Empty` operations. The enhancement task was to add operations to `Copy` a queue, `Clear` (i.e., empty) a queue, `Append` one queue to another, and `Reverse` a queue. The representation structure used for the queue was a standard linked data structure with pointers to the front and rear. Eighteen subjects participated in this experiment. Since the functionality and representation structure used in this package were so familiar to the subjects — both from classical data structures and a previous lab exercise — we felt that it would be difficult to obtain significant differences due to the approach used. Preliminary results from this experiment were reported in [6].

The second experiment was a replication of the first experiment using a more complex component that encapsulated a "partial map." This component allows a client to create an associative mapping which is a partial function from an arbitrary domain type to an arbitrary range type; it is useful in table processing applications. In addition to `Initialize`, `Finalize`, and `Swap`, the basic package provided the following operations.

- `Make_Defined` Assign a given range value to a given domain element (add an association to the map).
- `Make_Undefined` Undefine a given, presently defined domain element (i.e., remove a particular association from the map).
- `Make_Any_One_Undefined` Undefine some presently defined domain element, chosen arbitrarity and returned by the implementation (remove an arbitrary association from the map).
- `Test_If_Defined` Test if a given domain element is defined.
- `Test_If_Any_One_Defined` Test if there are *any* defined elements in the map.

The enhancement task for this experiment was to add operations that would `Display` (print) the map, `Clear` it (making every domain element undefined, according to the map), `Combine` two partial maps (assuming there were no inconsistently defined mappings in these two maps), and `Remap` all domain elements that map to one range value so they instead map to another range value. The representation chosen for the partial map was a hash table. The subjects were taught about this package in a previous lab assignment, where they were asked to implement the package using the hash table representation. Nevertheless, in comparison with the unbounded queue package, we felt that both the specification of the partial map package (i.e., the description of the functionality of the operations) and the representation of the data type were more complicated. Furthermore, we felt that the experience with the representation structure gained in the earlier lab could only

3 Overview of the Experiments

Our experiments were conducted as part of a course on the subject of "Software Components in Ada." This course was offered in two separate quarters, once in the summer of 1991 and the second time in the fall of 1992. In each case, the students in the course were graduate and upper-division undergraduate students majoring in computer science. Several of the students were full-time employees in the computing field.

The lectures heavily emphasized the trade-offs evident with principles of encapsulation, abstraction and layering, and presented a detailed engineering discipline for designing, formally specifying, and correctly and efficiently implementing Ada generic packages exporting abstract data types. Several programming assignments illustrated main points from the lectures, and the experiments were conducted in conjunction with some of these assignments.

Our main independent variable for the experiments was the type of reuse approach used by the subjects, and the dependent variables of interest were effort and quality of the reuse activity. To assess effort, we had the subjects keep careful records of the time spent in the initial designing/coding, testing, and debugging/recoding phases of the task. These times were recorded individually for each operation exported by the component under investigation. Both the initial designing/coding time and the total time to complete the task (including the testing and debugging/recoding time) were of interest to this study.

To assess quality, for each operation of the component the subjects recorded the number of defects they needed to fix. Only those defects that caused run-time failures during testing were included.

We emphasized to the students the importance of being internally consistent in keeping and reporting the data, and stressed that grades in the course would have nothing to do with the reported numbers. We audited the information provided by the students through post-assignment interviews, and the programs submitted by the students were tested by the instructor to ensure (as well as possible) that no lingering defects remained. We were somewhat skeptical of the data on a per-operation basis, but were confident that the aggregate information provided by the students was accurate.

Several factors other than reuse approach might influence the results on the dependent variables. One of these, the nature of the activity (i.e., enhancement of functionality versus modification of functionality) was mentioned in the previous section. Other important factors include the component involved in the reuse activity, the subjects' familiarity with the component specification and representation, and the abilities of the subjects. In planning the experiments, we attempted to deal with each of these issues.

The first experiment was an "enhancement of functionality" exercise using a simple, well-

Component Devel. Approach	Type of Reuse	Information Used	Understanding Required
Direct Impl.	White box (open,transparent)	Coding details, Representation data structures	Purpose/functionality and implementation of existing components
Layered	Black Box (closed,opaque)	Functional Specification Interface description	Purpose/functionality of existing components

Table 1: Comparison of Component Development Approaches

"open," or "transparent" reuse. A second approach to developing the desired functionality is to use only the specifications of the functionality of, and the interfaces to, the existing components. We call this the "layering" approach because the new functionality is provided by components built on top of the existing encapsulated units. Other names for this approach might be "black box," "closed," or "opaque" reuse (see Table 1).

Both approaches require that the implementor understand the purpose of and functionality provided by the existing components. But the first approach also requires an understanding of the *implementation* of the existing components, while the second approach does not. Intuitively, this suggests that black-box reuse demands of the implementor a lower cognitive load, and hence should reduce the *effort* required, at least for the initial design and coding of the new components.

The expected effect on *quality* of layering versus direct implementation is somewhat less clear. In this paper, we will use the term "quality" in the sense of "correctness," though we know that quality has other dimensions, too. Since the need to work with, and possibly misuse, the existing implementation is eliminated with black box strategies, one might expect better quality (i.e., fewer defects) from components developed using layering. On the other hand, since with layering one has available only the operations provided by the existing components and not the underlying representation data structures, the algorithms used in the layering approach might differ from those used in the direct implementation approach. This use of different algorithmic approaches may give rise to different distributions of defects. Even if the same algorithmic approaches are used in the two methods, most of the defects may occur in the attempt to synthesize the algorithm from the existing operations, rather than in the manipulation of representations of existing data structures.

Both of the tasks described above, that of adding functionality and that of modifying functionality, were studied in our experiments. The experiments are described in detail in the following section.

emerging software technologies.

In this paper, we report on three experiments designed to assess the effect, on effort and quality, of designing and developing software components by taking full advantage of existing abstractions, compared with an approach that allows seeing and using information about the implementation of the existing abstractions when designing and developing new components. Each of the experiments is a controlled study [4], and in each study there are interesting aspects to the statistical analysis required to determine if there is a significant difference between the two approaches.

The next section gives some additional background to motivate the experiments, while Section 3 describes each experiment in detail. Section 4 discusses important issues that affect the statistical analyses appropriate for these experiments. Section 5 discusses the results.

2 Background

Controlled studies are relatively rare in software engineering, and this often is attributed to the prohibitive cost of repeating large systems development tasks so that the effect of using a particular method can be examined relative to a baseline. Most attempts at doing a controlled analysis have been done with small sets of programmers (frequently students) on very small projects. While the applicability of the results of such studies to "real" systems is not obvious, studies of this nature have proven useful in the past. In some cases, they have shown that a particular method may not have the desired effect, or they have helped us to understand better the various factors that may influence the effect of a method (see, e.g., [12]). If the tasks performed by subjects in these experiments indeed are a subset of the tasks performed in the development of larger systems, and if the results of these experiments can be replicated, then the experiments can offer useful information about pieces of the complex process of software development.

What kinds of tasks might be typical of large systems development, and therefore worth investigating via controlled experimental studies? Two possibilities are 1) adding functionality to an existing component, and 2) building a component that has *similar* functionality to that of an existing component. In the former case, all that is desired is a new set of operations, but the existing component is of the right kind to provide these operations; the existing component's set of capabilities just needs to be enhanced. In the latter case, a closely related component might be available, but the desired component must provide somewhat *different*, rather than merely *additional*, functionality.

In both types of tasks, one can imagine developing the components in at least two ways. One way is to "directly implement" the desired functionality using the coding details of existing components. This involves making use of information about the data structures used to represent the objects encapsulated by the existing components. Some people refer to this type of reuse as "white box,"

1 Introduction

It is well-known that software systems typically exceed their expected development and maintenance costs. While there are many perceived reasons for this, one reason is that the components of these systems tend to be developed nearly from scratch, instead of being predominantly reuses of existing components [10].

The ability to successfully reuse components in new systems depends on the components being properly encapsulated. That way, new components can be "layered" on top of them, taking advantage of the abstract notions already encapsulated and avoiding reimplementation of these abstractions. For many years, programming languages have provided various means of encapsulating data structures and their associated operations, and the means of layering new components using these encapsulated components. Supposedly, there are productivity and quality gains to be had by following this layering technique.

However, currently we do not see in software systems the widespread use of encapsulation to layer the system's components. This is true even in popular component libraries. For example, in [2], Booch represents a map abstraction as a hash table using chaining for collision resolution. But he codes from scratch the lists that represent chains. He does not reuse his list package. Studies of object-oriented class hierarchies often have found that these "hierarchies" in fact are very flat [3], indicating that most components are not really built on top of existing components.

One possible reason for this lack of layering is that the original components are not well-designed, so that 1) the components needed in the next application are not quite those that are currently available, and 2) the available components cannot easily be converted (by layering) into those that *are* needed. Features and standard use of programming languages, such as Ada's restriction on the mode for parameters to functions and mixed use of private and limited private types, also inhibit one's ability to compose components [7].

Another reason may be that it is not obvious that, even if there were a useful component available, it would be better to reuse it by respecting its encapsulation than it would be to develop the required new functionality by accessing the underlying representation of the component. Performance is most often mentioned as the basis for this belief, although in our experience the performance penalty typically is minimal if the proper abstract functionality is encapsulated in the component. Language features such as code inheritance actually encourage violation of encapsulation while appearing to support layering [8].

We know of no controlled studies showing that the use of layering and encapsulation improves the cost or quality of component development, based on such measures as time to design/develop or number of defects in the resulting product. While abstract and anecdotal arguments have some place in technology assessment, it is vital that sound empirical support be obtained for the use of

The Effects of Layering and Encapsulation on Software Development Cost and Quality

Stu Zweben, Steve Edwards, Bruce Weide
The Ohio State University *

Joe Hollingsworth
Indiana University Southeast †

April 1994

Abstract

Software engineers often espouse the importance of using abstraction and encapsulation in developing software components. They advocate the "layering" of new components on top of existing components, using only information about the functionality and interfaces provided by the existing components. This layering approach is in contrast to a "direct implementation" of new components, utilizing unencapsulated access to the representation data structures and code present in the existing components. However, there is no empirical evidence that indicates whether the layering approach improves the productivity or quality of the development activity.

We discuss three controlled experiments designed to gather such empirical evidence. The results support the contention that layering significantly reduces the effort required to build new components. Furthermore, the quality of the components, in terms of the number of defects introduced during their development, is at least as good using the layered approach.

We also discuss some interesting issues related to the statistical analysis of experiments such as these.

Index Terms: empirical study, encapsulation, software components, abstract data types, software development, software reuse.

Research supported in part by the National Science Foundation under grants CCR-9111892 and CCR-9311702, and by ARPA under contract F30602-93-C-0243, monitored by the U.S. Air Force Materiel Command, Rome Laboratories, ARPA order number A714.

*Department of Computer and Information Science, 2036 Neil Avenue, Columbus, OH 43210

†Department of Computer Science, 4201 Grant Line Rd., New Albany, IN 47150