
Survey of Research on Software Clones

Rainer Koschke¹

Universität Bremen, Germany

Summary. This report summarizes my overview talk on software clone detection research. It first discusses the notion of software redundancy, cloning, duplication, and similarity. Then, it describes various categorizations of clone types, empirical studies on the root causes for cloning, current opinions and wisdom of consequences of cloning, empirical studies on the evolution of clones, ways to remove, to avoid, and to detect them, empirical evaluations of existing automatic clone detector performance (such as recall, precision, time and space consumption) and their fitness for a particular purpose, benchmarks for clone detector evaluations, presentation issues, and last but not least application of clone detection in other related fields.

After each summary of a subarea, I am listing open research questions.

1 Software Redundancy, Code Cloning, and Code Duplication

There are different forms of redundancy in software. Software comprises both programs and data. In the data base community, there is a clear notion of redundancy that has lead to various levels of normal forms. A similar theory does not yet exist for computer programs.

In computer programs, we can also have different types of redundancy. We should note that not every type of redundancy is harmful. For instance, programming languages use redundant declarations so that a compiler is able to check consistency between declarations and their uses. Also, at the architectural level, n-version programming is a strategy in which redundancy is purposefully and consciously used to implement reliable systems.

Sometimes *redundant* is used also in the sense of *superflous* in the software engineering literature. For instance, Xie and Engler show that superflous (they call them redundant) operations such as idempotent operations, assignments that are never read, dead code, conditional branches that are never taken, and redundant NULL-checks can pinpoint potential errors [1, 2, 3].

Redundant code is also often misleadingly called *cloned* code in the literature—although that implies that one piece of code is derived from the

other one in the original sense of this word. According to the Merriam-Webster dictionary, a *clone* is one that appears to be a copy of an original form. It is a synonym to *duplicate*. Although cloning leads to redundant code, not every redundant code is a clone. There may be cases in which two code segments that are no copy of each other just happen to be similar or even identical by accident. Also, there may be redundant code that is semantically equivalent but has a completely different implementation.

There is no agreement in the research community on the exact notion of redundancy and cloning. Ira Baxter’s definition of clones expresses this vagueness:

Clones are segments of code that are similar according to some definition of similarity. —Ira Baxter, 2002

According to this definition, there can be different notions of similarity. They can be based on text, lexical or syntactic structure, or semantics. They can even be similar if they follow the same pattern, that is, the same building plan. Instances of design patterns and idioms are similar in that they follow a similar structure to implement a solution to a similar problem.

Semantic difference relates to the observable behavior. A piece of code, *A*, is semantically similar to another piece of code, *B*, if *B* subsumes the functionality of *A*, in other words, they have “similar” pre and post conditions.

Unfortunately, detecting such semantic redundancy is undecidable in general although it would be worthwhile as you can often estimate the number of developers of a large software system by the number of hash table or list implementations you find.

Another definition of redundancy considers the program text: Two code fragments form a redundancy if their program text is similar. The two code fragments may or may not be equivalent semantically. These pieces are redundant because one fragment may need to be adjusted if the other one is changed. If the code fragments are executable code, their behavior is not necessarily equivalent or subsumed at the concrete level, but only at a more abstract level. For instance, two code pieces may be identical at the textual level including all variable names that occur within but the variable names are bound to different declarations in the different contexts. Then, the execution of the code changes different variables. Figure 1 shows two textually identical segments in the line range of 4–6 and 10–12, respectively. The semantic difference is that the first segment sets a global variable whereas the second one a local variable. The common abstract behavior of the two code segments is to iterate over a data structure and to increase a variable in each step.

Program-text redundancy is most often the result of *copy&paste*; that is, the programmer selects a code fragment and copies it to another location. Sometimes, these programmers are forced to copy because of limitations of the programming language. In other cases, they intend to reuse code. Sometimes these clones are modified slightly to adapt them to their new environment or

```

1  int sum = 0;
2
3  void foo(Iterator iter){
4      for (item = first(iter); has_more(iter); item = next(iter)){
5          sum = sum + value (item);
6      }
7  }
8  int bar(Iterator iter){
9      int sum = 0;
10     for (item = first(iter); has_more(iter); item = next(iter)){
11         sum = sum + value (item);
12     }
13 }

```

Fig. 1. Example of code clones.

purpose. Several authors report on 7-23% code duplication [4, 5, 6]; in one extreme case even 59% [7].

Clearly, the definition of redundancy, similarity, and cloning in software is still an open issue. There is little consensus in this matter. A study by Walenstein et al. [8], for instance, reports on differences among different human raters for clone candidates. In this study, clones were to be identified that ought to be removed and Walenstein et al. gave guidelines towards clones worthwhile being removed. The human raters of the clones proposed by automated tools did rarely agree upon what constitutes a clone worth to be removed. While the sources of inter-rater difference could be the insufficient similarity among clones or the appraisal of the need for removal, the study still highlights that there is no clear consensus yet, even for task-specific definitions of clones.

Our community should try to develop a general notion of redundancy, similarity, and cloning, and then identify more task-oriented categorizations of clones. Other research areas have similar difficulties in defining their fundamental terms. For instance, the architecture community debates the notion of *architecture* and the community of object-oriented programming the notion of *object*. Typically, these heated discussions cool down after a while when no consensus can be reached and no new arguments can be made. So, should we just stop discussing this matter, too? To some extent, these fundamental terms define the field. So it is important to clarify them. It is difficult, for instance, to create benchmarks to evaluate automatic clone detectors if it is unclear what we consider a clone.

2 Types of Clones

Program-text clones can be compared on the basis of the program text that has been copied. We can distinguish the following types of clones accordingly:

- **Type 1** is an exact copy without modifications (except for whitespace and comments).
- **Type 2** is a syntactically identical copy; only variable, type, or function identifiers have been changed.
- **Type 3** is a copy with further modifications; statements have been changed, added, or removed.

Baker further distinguishes so called parameterized clones [9], which are a subset of type-2 clones. Two code fragments A and B are a parameterized clone pair if there is a bijective mapping from A 's identifiers onto B 's identifiers that allows an identifier substitution in A resulting in A' and A' is a type-1 clone to B (and vice versa).

While type-1 and type-2 clones are precisely defined and form an equivalence relation, the definition of type-3 clones is inherently vague. Some researchers consider two consecutive type-1 or type-2 clones together forming a type-3 clone if the gap in between is below a certain threshold of lines [4, 23]. Another precise definition could be based on a threshold for the Levenshtein Distance, that is, the number of deletions, insertions, or substitutions required to transform one string into another. There is no consensus on a suitable similarity measure for type-3 clones yet.

The above simple classification is still very rough. Balazinska et al. introduced a more refined classification for function clones [10] as described in Table 1. This classification makes sense for selecting a suitable strategy for clone removal. For instance, the design pattern *TemplateMethod* may be used to factor out differences in the types used in different code fragments or the design pattern *Strategy* can be used to factor out algorithmic differences [11, 12]. Furthermore Balazinska et al. argue that each class is associated with a different risk in clone removal.

Kasper et al.'s classification is the most elaborated classification to date [13, 14] (cf. Table 2). The first level is a hint about the distance of clones. An argument can be made (although there is no empirical study on this hypothesis) that it is likely that clones between files are more problematic than within the same file as that it is more likely to overlook the former clones when it comes to consistent changes. The second decision distinguishes which syntactic units are copied. The third gives the degree of similarity and the fourth may be used to filter irrelevant or spurious clones.

There are still several open fundamental and terminological questions in software redundancy as follows:

- What are suitable definitions of similarity for which purpose?

<ul style="list-style-type: none"> • difference in method attributes (<code>static</code>, <code>private</code>, <code>throws</code>, etc.) • single-token difference in function body <ul style="list-style-type: none"> – further distinction into type of token: <ul style="list-style-type: none"> – called method – parameter type – literal – ... • token-sequence difference in function body <ul style="list-style-type: none"> – one unit (expression or statement) differs in token sequence – two units – more than two units
--

Table 1. Classification by Balazinska et al.

<ol style="list-style-type: none"> 1. At first level, distinguish clones within the same or across different files 2. then, according to type of region: <ul style="list-style-type: none"> • functions • declarations • macros • hybrids (in more than one of the above) • otherwise (among typedefs, variable declarations, function signatures) 3. then, degree of overlap or containment 4. then, according to type of code sequence: <ul style="list-style-type: none"> • initialization clones (first five lines) • finalization clones (last five lines) • loop clones (60% overlap of bodies) • switch and if (60% overlap of branches) • multiple conditions: several switch and if statements • partial conditions: branches of switch/if are similar
--

Table 2. Classification by Kasper et al. [13, 14]

- What other categorizations of clones make sense (e.g., semantics, origins, risks, etc.)?
- What is the statistical distribution of clone types?
- Are there correlations among orthogonal categories?
- Which strategies of removal and avoidance, risks of removal, potential damages, root causes, and other factors are associated with these categories?
- Can we create a theory of redundancy similar to normal forms in databases?

3 The Root Causes for Code Clones

A recent ethnographic study by Kim and Notkin [15] has shed some light on why programmers copy and paste code. By observing programmers in their daily practice they identified the following reasons.

Sometimes programmers are simply forced to duplicate code because of limitations of the programming language being used. Analyzing these root causes in more detail could help to improve the language design.

Furthermore, programmers often delay code restructuring until they have copied and pasted several times. Only then, they are able to identify the variabilities of their code to be factored out. Creating abstract generic solutions in advance often leads to unnecessarily flexible and hence needlessly complicated solutions. Moreover, the exact variabilities may be difficult to foresee. Hence, programmers tend to follow the idea of extreme programming in the small by not investing too much effort in speculative planning and anticipation.

Systems are modularized based on principles such as information hiding, minimizing coupling, and maximizing cohesion. In the end—at least for systems written in ordinary programming languages—the system is composed of a fixed set of modules. Ideally, if the system needs to be changed, only a very small number of modules must be adjusted. Yet, there are very different change scenarios and it is not unlikely that the chosen modularization forces a change to be repeated for many modules. The triggers for such changes are called *cross-cutting concerns*. For instance, logging is typically a feature that must be implemented by most modules. Another example is parameter checking in defensive programming where every function must check its parameters before it fulfills its purpose [16]. Then copy&paste dependencies reflect important underlying design decisions, namely, cross-cutting concerns.

Another important root cause is that programmers often reuse the copied text as a template and then customize the template in the pasted context.

Kasper et al. have investigated clones in large systems [17]. They found what they call *patterns of cloning* where cloning is consciously used as an implementation strategy. In their case study, they found the following cloning patterns:

Forking is cloning used to bootstrap development of similar solutions, with the expectation that evolution of the code will occur somewhat independently, at least in the short term. The assumption is that the copied code takes a separate evolution path independent of the original. In such a case, changes in the copy may be made that have no side effect on the original code.

Templating is used as a method to directly copy behavior of existing code but appropriate abstraction mechanisms are unavailable. It was also identified as a main driver for cloning in Kim and Notkin's case study [15]. Templating is often found when a reused library has a relatively fixed protocol (that is, a required order of using its interface items) which manifests as

laying out the control flow of the interface items as a fixed pattern. For instance, the code in Figure 1 uses a fixed iteration scheme for variable `iter`.

Customization occurs when currently existing code does not adequately meet a new set of requirements. The existing code is cloned and tailored to solve this new problem.

Very likely other more organizational aspects play a role, too. Time pressure, for instance, does not leave much time to search for the best long-term solution. Unavailable information on the impact of code changes leads programmers to create copies in which they make the required enhancement; such changes then are less likely to affect the original code negatively. Inadequate performance measures of programmers' productivity in the number of lines of code they produce neither invite programmers to avoid duplicates.

Although the two studies by Kim and Notkin as well as Kapsner and Godfrey are important first contributions, this area certainly requires more similar studies about root causes for cloning. Other potential reasons should be investigated:

- insufficient information on global change impact
- badly organized reuse process (type-4 clones)
- questionable productivity measures (LOCs per day)
- time pressure
- educational deficiencies, ignorance, or shortsightedness
- intellectual challenges (e.g., generics)
- professionalism/end-user programming (e.g., HTML, Visual Basic, etc.)
- development process ([18]: XP yields less clones?)
- organizational issues, e.g., distributed development organizations

Identifying the root causes would help us to fight the reasons, not just the symptoms, for instance, by giving feedback for programming language design.

4 Consequences of Cloning

There are plausible arguments that code cloning increases maintenance effort. Changes must be made consistently multiple times if the code is redundant. Often it is not documented where code has been copied. Manual search for copied code is infeasible for large systems. Furthermore during analysis, the same code must be read over and over again, then compared to the other code just to find out that this code has already been analyzed. Only if you make a detailed comparison, which can be difficult if there are subtle differences in the code or its environment, you can be sure that the code is indeed the same. This comparison can be fairly expensive. If the code would have been implemented only once in a function, this effort could have been completely avoided.

For these reasons, code cloning is number one on the stink parade of bad smell by Fowler [19]. But there are also counter arguments. In Kapsler and Godfrey’s study [17], code cloning is a purposeful implementation strategy which may make sense under certain circumstances. Moreover, it is not clear when you have type-3 clones whether the unifying solution would be easier to maintain than several copies with small changes. Generic solutions can become overly complicated. Maintainability can only be defined in a certain context with controlled parameters. That is, a less sophisticated programmer may be better off maintaining copied code than a highly parameterized piece of code. Moreover, there is a risk associated with removing code clones [20]. The removal requires deep semantic analyses and it is difficult to make any guarantees that the removal does not introduce errors. There may be even organizational reasons to copy code. Code cloning could, for instance, be used to disentangle development units [20].

The current debate lacks empirical studies of the costs and benefits of code cloning. To my knowledge, there is only one empirical study that explores the interrelationship of code cloning and maintainability. Monden et al. [21] analyzed a large system consisting of about 2,000 modules written in 1 MLOC lines of Cobol code over a period of 20 years. They used a token-based clone detector to find clones that were at least 30 lines long. They searched for correlations of maximal clone length with change frequency and number of errors. They found that most errors were reported for modules with clones of at least 200 lines. They also found still many errors in modules with shorter clones up to 50 lines. Yet, interestingly enough, they found the lowest error rate for modules with clones of 50 to 100 lines. Monden et al. have not further analyzed why these maintainability factors correlate in such a way with code cloning.

Chou et al. [22] investigated the hypothesis that if a function, file, or directory has one error, it is more likely that it has others. They found in their analysis of the Linux and OpenBSD kernels that this phenomenon can be observed most often where programmer ignorance of interface or system rules combines with copy-and-paste. They explain the correlation of bugs and copy-and-paste primarily by programmer ignorance, but they also note that – in addition to ignorance – the prevalence of copy-and-paste error clustering among different device drivers and versions suggests that programmers believe that “working” code is correct code. They note that if the copied code is incorrect, or its placed into a context it was not intended for, the assumption of goodness is violated.

Li et al. [23] use clone detection to find bugs when programmers copy code but rename identifiers in the pasted code inconsistently. On average, 13 % of the clones flagged as copy-and-paste bugs by their technique turned out to be real errors for the systems *Linux kernel*, *FreeBSD*, *Apache*, and *PostgreSQL*. The false positive rate is 73 % on average, where on average 14 % of the potential problems are still under analysis by the developers of the analyzed systems.

The empirical studies described here show interesting results. Most certainly, more empirical research along these lines should be conducted on the relation of cloning and quality attributes of software.

5 Clone Evolution

There are a few empirical studies on the evolution of clones, which describe some interesting observations. Antoniol et al. propose time series derived from clones over several releases of a system to monitor and predict the evolution of clones [24]. Their study for the data base system *mSQL* showed that the prediction of the average number of clones per function is fairly reliable. In another case study for the Linux kernel, they found that the scope of cloning is limited [25]. Only few clones can be found across subsystems; most clones are completely contained within a subsystem. In the subsystem *arch*, constituting the hardware architecture abstraction layer, newer hardware architectures tend to exhibit slightly higher clone rates. The explanation for this phenomenon is that newer modules are often derived from existing similar ones. The relative number of clones seems to be rather stable, that is, cloning does not occur in peaks. This last result was also reported by Godfrey and Tu who noticed that cloning is common and steady practice in the Linux kernel [26]. However, the cloning rate does increase steadily over time. Li et al. [23] observed for the Linux kernel in the period of 1994 to 2004 that the redundancy rate has increased from about 17 % to about 22 %. They observed a comparable behavior for FreeBSD. Most of the growth of redundancy rate comes from a few modules, including *drivers* and *arch* in Linux and *sys* in FreeBSD. The percentage of copy-paste code code increases more rapidly in those modules than in the entire software suite. They explain this observation by the fact that Linux supports more and more similar device drivers during this period.

Kim et al. analyzed the clone genealogy for two open-source Java systems using historical data from a version control system [27]. A clone genealogy forms a tree that shows how clones derive in time over multiple versions of a program from common ancestors. Beyond that, the genealogy contains information about the differences among siblings.

Their study showed that many code clones exist in the system for only a short time. Kim et al. conclude that extensive refactoring of such short-lived clones may not be worthwhile if they likely diverge from one another very soon. Moreover, many clones, in particular those with a long lifetime that have changed consistently with other elements in the same group cannot easily be avoided because of limitations of the programming language.

Clearly, as already mentioned above, more empirical studies along these lines are required. These studies should take industrial systems into account, too, as it is unclear to which extent these current observations can be attributed to the nature of open-source development. It would also be interest-

ing to investigate what the degree of cloning tells about the organization or development process. For instance, a study by Nickell and Smith reports that the extreme programming projects in their organization produce significantly fewer clones [18].

These studies should be target at the following questions (among others):

- How do clones evolve in industrial systems?
- What does their evolution tell about the development organization?
- What are reliable predictors?

6 Clone removal

If you do want to remove clones, there are several way to do so. There are even commercial tools, such as *CloneDr*¹ by *Semantics Design* to automatically detect and remove clones. Cloning and automatic abstraction and removal could even be a suitable implementation approach as hinted by Ira Baxter:

Cloning can be a good strategy if the you have the right tools in place. Let programmers copy and adjust, and then let tools factor out the differences with appropriate mechanisms. –Ira Baxter, 2002

In simple cases, you can use functional abstraction to replace equivalent copied code by a function call to a newly created function that encapsulates the copied code [28, 29]. In more difficult cases, when the difference is not just in the variable names that occur in the copied code, one may be able to replace by macros if the programming languages comes with a preprocessor. A preprocessor offers textual transformations to handle more complicated replacements. If a preprocessor is available, one can also use conditional compilation. As the excessive use of macros and conditional compilation may introduce many new problems, the solution to the redundancy problem may be found at the design level. The use of design patterns is an option to avoid clones by better design [11, 12]. Yet, this approach requires much more human expertise and, hence, can be less automated. Last but not least, one can develop code generators for highly repetitive code.

In all approaches, it is a challenge to cut out the right abstractions and to come up with meaningful names of generated functions or macros. Moreover, it is usually difficult to check the preconditions for these proposed transformations—be they manual or automated.

Empirical investigations of costs and benefits of clone removal are needed so that informed refactoring decisions can be made:

- What is the relation of clone types to quality attributes?
- Is there a relevance ranking of clone types for removal?
- How are clone types associated with suitable removal techniques and what are their benefits, costs and risks?

¹ Trademark of Semantic Designs, Inc.

7 Clone Avoidance

Rather than removing clones after the offense, it may be better to avoid them right from the beginning by integrating clone detection in the normal development process. Lague et al. identify two ways to integrate clone detection in normal development [6].

It can be used as *preventive control* where the code is checked continuously—for instance, at each check-in in the version control system or even on the fly while the code is edited—and the addition of a clone is reported for confirmation.

A complementary integration is *problem mining* where the code currently under modification is searched in the rest of the system. The found segments of code can then be checked whether the change must be repeated in this segment for consistency.

Lague et al. assessed the benefits of integrating clone detection in normal development by analyzing the three-year version history of a very large procedural telecommunication system [6]. In total, 89 million of non-blank lines (including comments) were analyzed, for an average size of 14.83 million lines per version. The average number of functions per release was 187,000.

Problem mining is assessed by the number of functions changed that have clones that were not changed; that is, how often a modification was missed potentially. Preventive control is assessed by the number of functions added that were similar to existing functions; that is, the code that could have been saved.

It is interesting to note, that—contrary to their expectations—they observed a low rate of growth in the number of overall clones in the system, due to the fact that many clones were actually removed from the system.

They conclude from their data that preventive control would help to lower the number of clones. Many clones disappeared only long after the day they came into existence. Early detection of clones could lead to taking this measure earlier.

They also found that problem mining could have provided programmers with a significant number of opportunities for correcting problems before end-user experienced them. The study indicates a potential for improving the software system quality and customer satisfaction through an effective clone management strategy.

The open issues are:

- What is the best integration of clone detection in the normal development process?
- In particular, what are the benefits and costs of possible integrations?
- What are reliable cloning indicators to trigger refactoring actions?
- How can we help organizations that already have many clones? How do we manage these clones?

8 Clone Detection

While there is an ongoing debate as to whether remove clones, there is a consensus about the importance to at least detect them. Clone avoidance during normal development, as described in the previous section, as well as making sure that a change can be made consistently in the presence of clones requires to know where the clones are. Manual clone detection is infeasible for large systems, hence, automatic support is necessary.

Automated software clone detection is an active field of research. This section summarizes the research in this area. The techniques can be distinguished at the first level in the type of information their analysis is based on and at the second level in the used algorithm.

8.1 Textual comparison

The approach by Rieger et al. compares whole lines to each other textually [7]. To increase performance, lines are partitioned using a hash function for strings. Only lines in the same partition are compared. The result is visualized as a dotplot, where each dot indicates a pair of cloned lines. Clones may be found as certain patterns in those dotplots visually. Consecutive lines can be summarized to larger cloned sequences automatically as uninterrupted diagonals or displaced diagonals in the dotplot.

Johnson [30, 31] uses the efficient string matching by Karp and Rabin [32, 33] based on fingerprints.

Marcus et al. [34] compare certain pieces of text, namely, identifiers using latent semantic indexing, a technique from information retrieval. The idea here is to identify fragments in which similar names occur as potential clones.

8.2 Token comparison

Baker's technique is also a line-based comparison. Instead of a string comparison, the token sequences of lines are compared efficiently through a suffix tree. First, each token sequence for a whole line is summarized by a so called *functor* that abstracts of concrete values of identifiers and literals [4]. The functor characterizes this token sequence uniquely. Assigning functors can be viewed as a perfect hash function. Concrete values of identifiers and literals are captured as parameters to this functor. An encoding of these parameters abstracts from their concrete values but not from their order so that code fragments may be detected that differ only in systematic renaming of parameters. Two lines are clones if they match in their functors and parameter encoding.

The functors and their parameters are summarized in a suffix tree, a trie that represents all suffixes of the program in a compact fashion. A suffix tree can be built in time and space linear to the input length [35, 36]. Every branch in the suffix tree represents program suffixes with common beginnings, hence, cloned sequences.

Kamiya et al. increase recall for superfluous different, yet equivalent sequences by normalizing the token sequences [37].

Because syntax is not taken into account, the found clones may overlap different syntactic units, which cannot be replaced through functional abstraction. Either in a preprocessing [38, 39, 40] or post-processing [41] step, clones that completely fall in syntactic blocks can be found if block delimiters are known. Preprocessing and postprocessing both require some syntactic information—gathered either lightweight by counting tokens opening and closing syntactic scopes or island grammars [42] or a full-fledged syntax analysis [40].

8.3 Metric comparison

Merlo et al. gather different metrics for code fragments and compare these metric vectors instead of comparing code directly [6, 5, 43, 44]. An allowable distance (for instance, Euclidean distance) for these metric vectors can be used as a hint for similar code. Specific metric-based techniques were also proposed for clones in web sites [45, 46].

8.4 Comparison of abstract syntax trees (AST)

Baxter et al. partition subtrees of the abstract syntax tree of a program based on a hash function and then compare subtrees in the same partition through tree matching (allowing for some divergences) [47]. A similar approach was proposed earlier by Yang [48] using dynamic programming to find differences between two versions of the same file.

Suffix trees—central to token-based techniques following Baker’s idea—can also be used to detect sequences of identical AST nodes. In the approach by Koschke et al. [49], the AST nodes are serialized in preorder traversal, a suffix tree is created for these serialized AST nodes, and the resulting maximally long AST node sequences are then cut according to their syntactic region so that only syntactically closed sequences remain.

8.5 Comparison of program dependency graphs

Token-based techniques and syntax-based techniques depend upon the textual order of the program. If programmers modify the order of the statements in copied code (for instance, to camouflage plagiarism) or use code cloning as in the templating implementation strategy (see Section 3), where the basic skeleton of an algorithm is reused and then certain pieces are adjusted to the new context, the copied code will not be found.

Yet, the order cannot be changed arbitrarily without changing the meaning of the program. All control and data dependencies must be maintained. A program dependency graph [50] is a representation of a program that represents only the control and data dependency among statements. Clones may be

identified as isomorphic subgraphs in a program dependency graph [51, 52]. Because this problem is NP hard, the algorithms use approximative solutions.

8.6 Other techniques

Leitao [53] combines syntactic and semantic techniques through a combination of specialized comparison functions that compare various aspects (similar call subgraphs, commutative operators, user-defined equivalences, transformations into canonical syntactic forms). Each comparison function yields an evidence that is summarized in an evidence-factor model yielding a clone likelihood. Walter et al. [54] and Li et al. [55] cast the search for similar fragments as a data mining problem. Statement sequences are summarized to item sets. An adapted data mining algorithm searches for frequent item sets.

9 Comparison of Clone Detection Algorithms

The abundance of clone detection techniques calls for a thorough comparison so that we know the strength and weaknesses of these techniques in order to make an informed decision if we need to select a clone detection technique for a particular purpose.

Clone detectors can be compared in terms of recall and precision of their findings as well as suitability for a particular purpose. There are several evaluations along these lines based on qualitative and quantitative data.

Bailey and Burd compared three clone and two plagiarism detectors [56]. Among the clone detectors were three of the techniques later evaluated by the subsequent study by Bellon and Koschke [57], namely, the techniques by Kamiya [37], Baxter [47], and Merlo [43]. For the latter technique, Bailey used an own re-implementation; the other tools were original. The plagiarism detectors were JPlag [58] and Moss [59].

The clone candidates of the techniques were validated by Bailey, and the accepted clone pairs formed an oracle against which the clone candidates were compared. Several metrics were proposed to measure various aspects of the found clones, such as scope (i.e., within the same file or across file boundaries), and the findings in terms of recall and precision.

The syntax-based technique by Baxter had the highest precision (100%) and the lowest recall (9%) in this experiment. Kamiya's technique had the highest recall and a precision comparable to the other techniques (72%). Interestingly, Merlo's metric-based technique showed the least precision (63%).

Although the case study by Bailey and Burd showed interesting initial results, it was conducted on only one relatively small system (16 KLOC). However, because the size was limited, Bailey was able to validate all clone candidates.

A subsequent larger study was conducted by Bellon and Koschke [60, 57]. Their likewise quantitative comparison of clone detectors was conducted for

4 Java and 4 C systems in the range of totalling almost 850 KLOC. The participants and their clone detectors evaluated are listed in Table 3.

Participant	Tool	Comparison
Brenda S. Baker [4]	<i>Dup</i>	Token
Ira D. Baxter [47]	<i>CloneDr</i>	AST
Toshihiro Kamiya [37]	<i>CCFinder</i>	Token
Jens Krinke [51]	<i>Duplix</i>	PDG
Ettore Merlo [43]	<i>CLAN</i>	Function Metrics
Matthias Rieger [7]	<i>Duploc</i>	Text

Table 3. Participating scientists

Table 4 summarizes the findings of Bellon and Koschke’s study. Row *clone type* lists the type of clones the respective clone detector finds (for clone types, see Section 2). The next two rows qualify the tools in terms of their time and space consumption. The data is reported at an ordinal scale - -, -, +, + + where - - is worst (the exact measures can be found in the paper to this study [60, 57]). Recall and precision are determined as in Bailey and Burd’s study by comparing the clone detectors’ findings to a human oracle. The same ordinal scale is used to qualify the results; exact data are reported in the paper [60, 57].

	Baker	Baxter	Kamiya	Krinke	Merlo	Rieger
Clone type	1, 2	1, 2	1, 2, 3	3	1, 2, 3	1, 2, 3
Speed	+ +	-	+	--	+ +	?
RAM	+	-	+	+	+ +	?
Recall	+	-	+	-	-	+
Precision	-	+	-	-	+	-

Table 4. Results from the Bellon and Koschke study.

The difference in precision of Merlo’s approach in this comparison to the study by Bailey and Burd can be explained by the fact that Merlo compared not only metrics but also the tokens and their textual images to identify type-1 and type-2 clones in the study by Bellon and Koschke.

While the Bailey/Burd and Bellon/Koschke studies focus on quantitative evaluation of clone detectors, other authors have evaluated clone detectors for the fitness for a particular maintenance task. Rysselberghe and Demeyer [61] have compared text-based [7, 62], token-based [4], and metric-based [43] clone detectors for refactoring. They compare these techniques in terms of suitability (can a candidate be manipulated by a refactoring tool?), relevance (is there a priority which of the matches should be refactored first?), confidence (can

one solely rely on the results of the code cloning tool, or is manual inspection necessary?), and focus (does one have to concentrate on a single class or is it also possible to assess an entire project?). They assess these criteria qualitatively based on the clone candidates produced by the tools. Figure 2 summarizes their conclusions.

critierion	most suitable technique
suitability	metric-based
relevance	no difference
confidence	text-based
focus	no difference

Fig. 2. Assessment by Rysselberghe and Demeyer

Bruntink et al. use clone detection to find cross-cutting concerns in C programs with homogeneous implementations [63]. In their case study, they used *CCFinder*—Kamiya’s [37] tool evaluated in other case studies, too—one of the Bauhaus² clone detectors, namely *ccdimpl*, which is a variation of Baxter’s technique [47], and the PDG-based detector *PDG-DUP* by Komondoor [52]. The cross-cutting concerns they looked for were error handling, tracing, pre and post condition checking, and memory error handling. The study showed that the clone classes obtained by Bauhaus’ *ccdimpl* can provide the best match with the range checking, null-pointer checking, and error handling concerns. Null-pointer checking and error handling can be found by *CCFinder* almost equally well. Tracing and memory error handling can best be found by *PDG-DUP*.

The open issues in evaluation clone detectors is to overcome the following limitations of current benchmarks:

- single oracle (until recently)
 - differences among different human raters for clone candidates [8] when clones ought to be removed.
- yes/no decision rather than degree of confidence
- clones length measured as lines rather than tokens
- current benchmark insists on contiguous lines/tokens
- current benchmark is based on clone pairs rather than clone classes

There is also an organizational aspect that our community needs to address. Once we agreed on a benchmark, benchmarking should become standard procedure of the community. Every new technique should be compared against this benchmark prior to publication.

² <http://www.axivion.com>

10 Clone Presentation

Because there is typically a huge amount of clones in large systems and these clones differ in various attributes (type, degree of similarity, length, etc.), presentation issues of clone information is critical. This huge information space must be made accessible to a human analyst. The analyst needs a holistic view that combines source code views and architectural views.

There have been several proposals for clone visualization. Scatter plots are two-dimensional charts where all software units are listed on both axes [64, 7, 65]. There is a dot if two software units are similar. The granularity of software units may differ. It can range from single lines to functions to classes and files to packages and subsystems. Visual patterns of cloning may be observed by a human analyst. A problem with this approach is scalability for many software units and the order of the listed software units as this has an impact on the visual patterns. While there is a “natural” order for lines (i.e., lexical order) within a file, it is not clear how to order more coarse-grained units such as functions, files, and packages. Lexical order or their names is in most cases as arbitrary as random order.

Johnson [66] proposes Hasse diagrams for cloning between sets of files so that one can better see whether code has been copied between files, which is possibly more critical than cloning within a file. A Hasse diagram (named after a German mathematician) is used to draw a partial order among sets as an acyclic graph. Directed arcs connect nodes that are related by the order relation and for which no other directed path exists.

In Johnson’s context, each match of a block of text identifies a range of characters (or lines) from two (or more) files. For each subset of files, one can total the number of characters that the matching process has discovered to match between the given set of files. A subset of files forms a node, if the files have non-zero matches. The inclusion between subsets of files yields the edges.

Rieger et al. [67] propose to use Michele Lanza’s polymetric views [68] to visualize various aspects of clones in one view. A polymetric view is again based on the graph metaphor and representation where a node represents a software unit and an edge a cloning relation. Visually, additional information can be attached to the graph by the degree of freedom for the position (X/Y in the two-dimensional space), color of nodes and edges, thickness of edges, breadth and width of nodes. Propose a fixed set of metrics combinations to be mapped onto graphical aspects to present the clone information from different perspective for different tasks.

Beyond polymetric views, Rieger et al. [67] propose a variation of tree maps to show the degree of cloning along with the system decomposition. Tree maps display information about entities with a hierarchical relationship in a fixed space (for instance, the whole system on one screen) where the leaves of the hierarchy contain a metric to be visualized. Each inner node aggregates the metric values of its descendants. Each node is represented through a piece of

the available space. The space of a descendent node is completely contained in space of its ancestor. There is no overlap in space for nodes that are not in an ancestor/descendant relation. This is how the hierarchy is presented. Essentially the hierarchy is projected into the two dimensional space seen from the root of the hierarchy. In order to show the hierarchy clearly, the space of each node appears as rectangle where the direction of subdivision of nested nodes is alternated horizontally and vertically at each level. The space of each rectangle is proportional to the metric.

This visualization was originally proposed by Ben Shneiderman in the early 1990s to show space consumption of a hard disk with a hierarchical file system. While space is used very efficiently, problems arise when the hierarchy is deeply nested.

There is a large body of research on information visualization in general and software visualization in particular that we have not yet explored for clone visualization. In order to understand which visualization works best for which purpose, we need more systematic empirical research.

11 Related Fields

Clone detection has applications in other fields and—vice versa—ideas from related fields can be reused for clone detection.

Bruntink et al. for instance, use clone detectors to search for code that could be factored out as aspects using an aspect-oriented language [16, 63]. They identify error handling, tracing, pre and post condition checking, and memory error handling. Although they used classic clone detectors that were not designed for this particular purpose, the clone detectors appeared to be helpful. Classic clone detectors try to find similar code—similar in terms of their program text. The implementations of an aspect, on the other hand, are often very heterogeneous and are similar only at a more semantic level. For instance, precondition checking tests each parameter of a function for certain criteria. At the implementation level, functions differ in the order and type of parameters so that checks are generally different in the program text.

The code compaction community for that try to minimize the memory footprint of programs for small devices use very similar algorithms to identify redundant code that could be compressed [69].

The detection of plagiarism faces similar but even worse problems as clone detection [59, 58, 70, 71, 72, 73, 74]. In plagiarism cases, people try to camouflage their copy, which makes it more difficult to detect the plagiarism. In order to reuse classic clone detectors for plagiarism, we would need to reduce programs to a normal form for comparison. This normalization, on the other hand, could lead to false positives.

Another application of clone detection is the comparison of versions or variants of software systems. While versions derive from each other, variants have a common ancestor. In both cases, they are very similar. In software

evolution research, where information of software units is observed over time or versions, respectively, it is necessary to map the software entities of one version to those of the other version in order to carry over the information. This problem is called the *origin analysis* [75]. The same problem needs to be solved when two software variants are to be compared or merged [76]. Relying solely on names of these units for this analysis may be misleading if a refactoring like *renaming* has taken place [19]. Also, the refactoring *extract method* moves statements from one function to create a new function. Clone detection can help to establish a mapping between two versions or variants of a program. Several authors have used clone detection techniques or at least a code similarity measure to determine this mapping [77, 78, 79, 26, 80, 81, 82].

The difference of comparing versions or variants to detecting clones is that the task here is to map a code entity onto only one or at least a small set of candidates in the other system, the comparison is only between systems (clones within the same version or variant are irrelevant), cloning is the rule rather than the exception as the two versions or variants overlap to a very high degree, the focus is on the differences rather than the similarities, and the comparison should tolerate renaming and all refactorings that move entities around such as pull-up field, move method, etc.

As there are similar related fields, the interesting questions here are “What can clone detection learn from other fields” and “What can other fields learn from clone detection?”

References

1. Xie, Y., Engler, D.: Using redundancies to find errors. In: Proceedings of the 10th ACM SIGSOFT symposium on Foundations of software engineering, ACM Press (2002) 51–60
2. Xie, Y., Engler, D.: Using redundancies to find errors. SIGSOFT Softw. Eng. Notes **27**(6) (2002) 51–60
3. Xie, Y., Engler, D.: Using redundancies to find errors. IEEE Computer Society Transactions on Software Engineering **29**(10) (2003) 915–928
4. Baker, B.S.: On finding duplication and near-duplication in large software systems. In Wills, L., Newcomb, P., Chikofsky, E., eds.: Second Working Conference on Reverse Engineering, Los Alamitos, California, IEEE Computer Society Press (1995) 86–95
5. Kontogiannis, K., Mori, R.D., Merlo, E., Galler, M., Bernstein, M.: Pattern matching for clone and concept detection. Automated Software Engineering **3**(1/2) (1996) 79–108
6. Lague, B., Proulx, D., Mayrand, J., Merlo, E., Hudepohl, J.: Assessing the benefits of incorporating function clone detection in a development process. In: International Conference on Software Maintenance. (1997) 314–321
7. Ducasse, S., Rieger, M., Demeyer, S.: A Language Independent Approach for Detecting Duplicated Code. In: International Conference on Software Maintenance. (1999) 109–118

8. Walenstein, A., Jyoti, N., Li, J., Yang, Y., Lakhota, A.: Problems creating task-relevant clone detection reference data. In: Working Conference on Reverse Engineering, IEEE Computer Society Press (2003)
9. Baker, B.S.: A program for identifying duplicated code. In: Computer Science and Statistics 24: Proceedings of the 24th Symposium on the Interface. (1992) 49–57
10. Balazinska, M., Merlo, E., Dagenais, M., Lague, B., Kontogiannis, K.: Measuring clone based reengineering opportunities. In: IEEE Symposium on Software Metrics, IEEE Computer Society Press (1999) 292–303
11. Balazinska, M., Merlo, E., Dagenais, M., Lague, B., Kontogiannis, K.: Partial redesign of java software systems based on clone analysis. In: Working Conference on Reverse Engineering, IEEE Computer Society Press (1999) 326–336
12. Balazinska, M., Merlo, E., Dagenais, M., Lague, B., Kontogiannis, K.: Advanced clone-analysis to support object-oriented system refactoring. In: Working Conference on Reverse Engineering, IEEE Computer Society Press (2000) 98–107
13. Kapsner, C., Godfrey, M.W.: Toward a taxonomy of clones in source code: A case study. In: Evolution of Large Scale Industrial Software Architectures. (2003)
14. Kapsner, C., Godfrey, M.: A taxonomy of clones in source code: The reengineers most wanted list. In: Working Conference on Reverse Engineering, IEEE Computer Society Press (2003)
15. Kim, M., Bergman, L., Lau, T., Notkin, D.: An ethnographic study of copy and paste programming practices in OOPL. In: International Symposium on Empirical Software Engineering, IEEE Computer Society Press (2004) 83–92
16. Bruntink, M., van Deursen, A., Tourwe, T., van Engelen, R.: An evaluation of clone detection techniques for crosscutting concerns. In: International Conference on Software Maintenance. (2004) 200–209
17. Kapsner, C., Godfrey, M.W.: "clones considered harmful" considered harmful. In: Working Conference on Reverse Engineering. (2006)
18. Nickell, E., Smith, I.: Extreme programming and software clones. In: Working Conference on Reverse Engineering, IEEE Computer Society Press (2003)
19. Fowler, M.: Refactoring: improving the design of existing code. Addison Wesley (1999)
20. Cordy, J.: Comprehending reality: Practical challenges to software maintenance automation. In: International Workshop on Program Comprehension, IEEE Computer Society Press (2003) 196–206
21. Monden, A., Nakae, D., Kamiya, T., Sato, S., Matsumoto, K.: Software quality analysis by code clones in industrial legacy software. In: IEEE Symposium on Software Metrics. (2002) 87–94
22. Chou, A., Yang, J., Chelf, B., Hallem, S., Engler, D.R.: An empirical study of operating system errors. In: Symposium on Operating Systems Principles. (2001) 73–88
23. Li, Z., Lu, S., Myagmar, S., Zhou, Y.: Copy-paste and related bugs in large-scale software code. IEEE Computer Society Transactions on Software Engineering **32**(3) (2006) 176–192
24. Antoniol, G., Casazza, G., Penta, M.D., Merlo, E.: Modeling clones evolution through time series. In: International Conference on Software Maintenance, IEEE Computer Society Press (2001) 273–280
25. Antoniol, G., Villano, U., Merlo, E., Penta, M.: Analyzing cloning evolution in the linux kernel. Information and Software Technology **44**(13) (2002)

26. Godfrey, M., Tu, Q.: Growth, evolution and structural change in open source software. In: Workshop on Principles of Software Evolution. (2001)
27. Kim, M., Sazawal, V., Notkin, D., Murphy, G.C.: An empirical study of code clone genealogies. In: European Software Engineering Conference and Foundations of Software Engineering (ESEC/FSE). (2005)
28. Fanta, R., Rajlich, V.: Removing clones from the code. *Journal on Software Maintenance and Evolution* **11**(4) (1999) 223–243
29. Komondoor, R., Horwitz, S.: Eliminating duplication in source code via procedure extraction. Technical report 1461, UW-Madison Dept. of Computer Sciences (2002)
30. Johnson, J.H.: Identifying redundancy in source code using fingerprints. In: Conference of the Centre for Advanced Studies on Collaborative research, IBM Press (1993) 171–183
31. Johnson, J.H.: Substring matching for clone detection and change tracking. In: International Conference on Software Maintenance, IEEE Computer Society Press (1994) 120–126
32. Karp, R.M.: Combinatorics, complexity, and randomness. *Communications of the ACM* **29**(2) (1986) 98–109
33. Karp, R.M., Rabin, M.: Efficient randomized pattern-matching algorithms. *IBM Journal Research and Development* **31**(2) (1987) 249–260
34. Marcus, A., Maletic, J.: Identification of high-level concept clones in source code. In: International Conference on Automated Software Engineering. (2001) 107–114
35. McCreight, E.: A space-economical suffix tree construction algorithm. *Journal of the ACM* **32**(2) (1976) 262–272
36. Baker, B.S.: Parameterized Pattern Matching: Algorithms and Applications. *Journal Computer System Science* **52**(1) (1996) 28–42
37. Kamiya, T., Kusumoto, S., Inoue, K.: CCFinder: A Multi-Linguistic Token-based Code Clone Detection System for Large Scale Source Code. *IEEE Computer Society Transactions on Software Engineering* **28**(7) (2002) 654–670
38. Synytskyy, N., Cordy, J.R., Dean, T.: Resolution of static clones in dynamic web pages. In: Workshop on Website Evolution. (2003) 49–56
39. Cordy, J.R., Dean, T.R., Synytskyy, N.: Practical language-independent detection of near-miss clones. In: Conference of the Centre for Advanced Studies on Collaborative research, IBM Press (2004) 1–12
40. Gitchell, D., Tran, N.: Sim: a utility for detecting similarity in computer programs. In: SIGCSE '99: The proceedings of the thirtieth SIGCSE technical symposium on Computer science education, ACM Press (1999) 266–270
41. Higo, Y., Ueda, Y., Kamiya, T., Kusumoto, S., Inoue, K.: On software maintenance process improvement based on code clone analysis. In: International Conference on Product Focused Software Process Improvement. Volume 2559 of Lecture Notes In Computer Science., Springer (2002) 185–197
42. Moonen, L.: Generating robust parsers using island grammars. In: Proceedings of the Working Conference on Reverse Engineering, IEEE Computer Society Press (2001) 13–22
43. Mayrand, J., Leblanc, C., Merlo, E.: Experiment on the automatic detection of function clones in a software system using metrics. In: International Conference on Software Maintenance. (1996) 244–253

44. Kontogiannis, K., DeMori, R., Bernstein, M., Galler, M., Merlo, E.: Pattern matching for design concept localization. In: Working Conference on Reverse Engineering, IEEE Computer Society Press (1995) 96–103
45. Di Lucca, G., Di Penta, M., Fasolino, A.: An approach to identify duplicated web pages. In: International Computer Software and Applications Conference. (2002) 481–486
46. Lanubile, F., Mallardo, T.: Finding function clones in web applications. In: European Conference on Software Maintenance and Reengineering. (2003) 379–386
47. Baxter, I.D., Yahin, A., Moura, L., Sant’Anna, M., Bier, L.: Clone Detection Using Abstract Syntax Trees. In Koshgoftaar, T.M., Bennett, K., eds.: International Conference on Software Maintenance, IEEE Computer Society Press (1998) 368–378
48. Yang, W.: Identifying syntactic differences between two programs. *Software-Practice and Experience* **21**(7) (1991) 739–755
49. Koschke, R., Falke, R., Frenzel, P.: Clone detection using abstract syntax suffix trees. In: Working Conference on Reverse Engineering, IEEE Computer Society Press (2006)
50. Ferrante, J., Ottenstein, K., Warren, J.: The program dependence graph and its use in optimization. *ACM Trans. on Prog. Lang. and Sys.* **9**(3) (1987) 319–349
51. Krinke, J.: Identifying Similar Code with Program Dependence Graphs. In: Working Conference on Reverse Engineering. (2001) 301–309
52. Komondoor, R., Horwitz, S.: Using slicing to identify duplication in source code. In: Proc. Int. Symposium on Static Analysis. (2001) 40–56
53. Leitao, A.M.: Detection of redundant code using r2d2. In: Workshop Source Code Analysis and Manipulation, IEEE Computer Society Press (2003) 183–192
54. Wahler, V., Seipel, D., von Gudenberg, J.W., Fischer, G.: Clone detection in source code by frequent itemset techniques. In: Workshop Source Code Analysis and Manipulation. (2004) 128–135
55. Li, Z., Lu, S., Myagmar, S., Zhou, Y.: Cp-miner: A tool for finding copy-paste and related bugs in operating system code. In: Operating System Design and Implementation. (2004) 289–302
56. Bailey, J., Burd, E.: Evaluating clone detection tools for use during preventative maintenance. In: Workshop Source Code Analysis and Manipulation, IEEE Computer Society Press (2002) 36–43
57. Bellon, S., Koschke, R.: Comparison and evaluation of clone detection tools. *IEEE Computer Society Transactions on Software Engineering* (2006) submitted for publication.
58. Prechelt, L., Malpohl, G., Philippsen, M.: Jplag: Finding plagiarisms among a set of programs. Technical report, University of Karlsruhe, Department of Informatics (2000)
59. Schleimer, S., Wilkerson, D.S., Aiken, A.: Winnowing: local algorithms for document fingerprinting. In: Proceedings of the SIGMOD. (2003) 76–85
60. Bellon, S.: Vergleich von Techniken zur Erkennung duplizierten Quellcodes. Diploma thesis, no. 1998, University of Stuttgart (Germany), Institute for Software Technology (2002)
61. Van Rysselberghe, F., Demeyer, S.: Evaluating clone detection techniques from a refactoring perspective. In: International Conference on Automated Software Engineering. (2004)

62. Rieger, M.: Effective Clone Detection Without Language Barriers. Dissertation, University of Bern, Switzerland (2005)
63. Bruntink, M., van Engelen, R., Tourwe, T.: On the use of clone detection for identifying crosscutting concern code. *IEEE Computer Society Transactions on Software Engineering* **31**(10) (2005) 804–818
64. Church, K.W., Helfman, J.I.: Dotplot: A program for exploring self-similarity in millions of lines for text and code. *Journal of American Statistical Association, Institute for Mathematical Statistics and Interface Foundations of North America* **2**(2) (1993) 153–174
65. Ueda, Y., Kamiya, T., Kusumoto, S., Inoue, K.: Gemini: Maintenance support environment based on code clone analysis. In: *IEEE Symposium on Software Metrics*, IEEE Computer Society Press (2002) 67–76
66. Johnson, J.H.: Visualizing textual redundancy in legacy source. In: *Conference of the Centre for Advanced Studies on Collaborative research*, IBM Press (1994) 32
67. Rieger, M., Ducasse, S., Lanza, M.: Insights into system-wide code duplication. In: *Working Conference on Reverse Engineering*, IEEE Computer Society Press (2004) 100–109
68. Lanza, M., Ducasse, S.: Polymetric views - a lightweight visual approach to reverse engineering. *IEEE Computer Society Transactions on Software Engineering* **29**(9) (2003) 782–795
69. Cooper, K., McIntosh, N.: Enhanced code compression for embedded risc processors. In: *Proc. ACM Conf. on Programming Language Design and Implementation*, ACM Press (1999) 139–149
70. Flammia, G.: On the internet, software should be milked, not brewed. *IEEE Expert* **11**(6) (1996) 87–88
71. Manber, U.: Finding similar files in a large file system. In: *Proceedings of the Winter Usenix Technical Conference*. (1994) 1–10
72. Jankowitz, H.T.: Detecting plagiarism in student pascal programs. *Computer Journal* **1**(31) (1988) 1–8
73. Madhavji, N.H.: Compare: a collusion detector for pascal. *Techniques et Sciences Informatiques* **4**(6) (1985) 489–497
74. Grier, S.: A tool that detects plagiarism in pascal programs. *SIGSCE Bulletin* **13**(1) (1981)
75. Tu, Q.: On navigation and analysis of software architecture evolution. Master's thesis, University of Waterloo (1992)
76. Hunt, J.: Extensible, Language Aware Differencing and Merging. Dissertation, University of Kaiserslautern (2001)
77. Tu, Q., Godfrey, M.W.: An integrated approach for studying architectural evolution. In: *International Workshop on Program Comprehension*, IEEE Computer Society Press (2002) 127–136
78. Zou, L., Godfrey, M.: Detecting merging and splitting using origin analysis. In: *Working Conference on Reverse Engineering*, IEEE Computer Society Press (2003)
79. Godfrey, M., Tu, Q.: Evolution in open source software: A case study. In: *International Conference on Software Maintenance*, IEEE Computer Society Press (2000)
80. Van Rysselberghe, F., Demeyer, S.: Reconstruction of successful software evolution using clone detection. In: *International Workshop on Principles of Software Evolution*. (2003) 126–130

81. Xing, Z., Stroulia, E.: Refactoring detection based on umldiff change-facts queries. In: Working Conference on Reverse Engineering, IEEE Computer Society Press (2003) 263–274
82. Xing, Z., Stroulia, E.: Analyzing the evolutionary history of the logical design of object-oriented software. IEEE Computer Society Transactions on Software Engineering **31**(10) (2005) 850–868