

# Reducing Redundancies in Multi-revision Code Analysis

Carol V. Alexandru, Sebastiano Panichella and Harald C. Gall  
Software Evolution and Architecture Lab  
University of Zurich, Switzerland  
{alexandru,panichella,gall}@ifi.uzh.ch

**Abstract**—Software engineering research often requires analyzing multiple revisions of several software projects, be it to make and test predictions or to observe and identify patterns in how software evolves. However, code analysis tools are almost exclusively designed for the analysis of one specific version of the code, and the time and resources requirements grow linearly with each additional revision to be analyzed. Thus, code studies often observe a relatively small number of revisions and projects. Furthermore, each programming ecosystem provides dedicated tools, hence researchers typically only analyze code of one language, even when researching topics that should generalize to other ecosystems. To alleviate these issues, frameworks and models have been developed to combine analysis tools or automate the analysis of multiple revisions, but little research has gone into actually *removing* redundancies in multi-revision, multi-language code analysis. We present a novel end-to-end approach that systematically avoids redundancies every step of the way: when reading sources from version control, during parsing, in the internal code representation, and during the actual analysis. We evaluate our open-source implementation, LISA, on the full history of 300 projects, written in 3 different programming languages, computing basic code metrics for over 1.1 million program revisions. When analyzing many revisions, LISA requires less than a second on average to compute basic code metrics for all files in a single revision, even for projects consisting of millions of lines of code.

## I. INTRODUCTION

Static software analysis has a broad range of applications and is omnipresent both in software engineering research and practice. Professionals use software analysis towards goals such as enforcing coding guidelines, allocating resources and estimating future effort [1]–[4]. It also aids in development, for example by means of refactoring suggestions [3], [5]–[7] or bug detection [8], [9]. Practitioners most often apply their tools to the current version of their software, be it during development in an IDE or during continuous integration. Their tools often grow naturally out of needs within a software ecosystem and are tailored to the peculiarities of a specific language [10]–[13]. Practitioners also show growing interest in multi-language analyses [14], given that an increasing amount of software is written in multiple languages (like web applications and multi-platform mobile apps) [15]–[17].

Conversely, researchers use code analysis to discover patterns and anti-patterns in existing code [18]–[20], to learn more about software evolution [21]–[25] and to develop methods for change and bug prediction [3], [8], [9], [26]. In this context, tools developed for practical application are not always ideal,

because researchers expect high generalizability [27]–[33] and replicability [34] from their results, yet the analysis of large samples using existing methods is hard to automate for a number of reasons, which we discuss in the following section.

### A. Challenges

In the context of software engineering and evolution research, the selection of software projects suited for a particular investigation and the collection of code metrics are important steps toward large and replicable case studies. However, researchers need to overcome several technical and practical obstacles. First and foremost, software analysis for large projects is *costly in terms of time and resources* [29], [35], [36]. Moreover, nowadays software applications are written in multiple programming languages [15]–[17] and several issues with studying multi-language software have been identified by researchers and practitioners alike [15], [17], [37]. It is extraordinarily *hard to conduct generalizable code studies involving multiple languages* and attempts at solving this problem so far incur even higher performance penalties.

**Manual effort.** When the goal is to gain a broad overview on the evolution of many projects, the effort required to set up and run analyses for each individual project and revision presents a major problem. Tools which run on compiled applications are largely unsuitable, because the compiled version is rarely available for each revision and would need to be prepared manually. Automated build approaches so far exhibit limited effectiveness because projects tend to be unique in terms of the build environment and dependencies needed for compilation [35], [36], [38]–[40]. *Ideally, software evolution analysis tools should operate automatically for the most part.*

**Redundant analysis of unchanged source code.** The majority of existing code analysis approaches are not designed with a multi-revision use-case in mind [15], [17], [37]. They require one execution per revision, even if each commit makes only minute changes to the program. Researchers may be able to take advantage of tools that operate on the file level to only analyze files that change, but the ratio of changed to unchanged code within a single file is still very small and the majority of the analysis will still be redundant. *Ideally, a software evolution analysis tool should operate on the smallest possible delta between code revisions.*

**Source data duplication.** Most code analysis tools operate on files and folders [29], [35], [36]. While, even for large




Fig. 1. Method Count and Class Count for 4 releases of the facebook-csharp-sdk project.

projects, single revisions remain manageable, checking out thousands of revisions hinders effective parallelization. *Ideally, analyzing code from code repositories should occur in-place.*

**Isolation of programming ecosystems.** Many code studies involve only one programming language (e.g., [15], [17], [24], [28], [30], [31], [37]). Every programming ecosystem comes with its own set of tools and any data gathering effort is duplicated with each additional language. Several successful approaches for performing studies on the evolution of software projects [21], [25], [41] and bug prediction [26], [42], [43] limit their applicability to a single programming language. *Ideally, software analysis approaches should be easily transferable to multiple programming languages.*

**Mismatch of concepts in different languages.** To analyze code in multiple languages using common tooling, researchers have devised meta-models as a common representation for source code [15], [17], [37]. These attempt to tie language-specific entities and structures to a common representation. However, each language exhibits unique features, hence accurate mappings are not always obvious and depend on the goals of each individual study or approach. *Ideally, code analysis tools should genericise over arbitrary structures and cross-or multi-language analyses should be enabled without having to adhere to a specific, pre-defined meta-model.*

**Low-resolution historical data.** Given all previous obstacles, the majority of code evolution studies not only use a small sample of projects in a single language, but also choose to analyze only a limited number of releases [16], [17], [44]. Considering that modern software development typically sees several commits per day for any given project [45], [46], analyzing only a few releases can paint an incomplete picture. For example, fig. 1 shows the Method Count and Class Count computed for 4 major releases of the facebook-csharp-sdk open source project. A preliminary interpretation of the graph suggests a rather steady rise and fall in the metrics of the source code as a whole. However, we cannot make any claims on how the software project evolved on a shorter time scale. We will revisit this example in our evaluation.

## B. Goal and Contributions

To presently conduct large scale code studies over the full history of many projects involving millions of program revisions, researchers need to either invest in resources for parallelization (e.g., [29]) or spend an excessive amount of time

running analyses (e.g., [20]). Software engineering research is yet lacking an approach that can capture the entire evolution of many software projects in different languages quickly and without excessive manual effort. The goal of this paper is to provide low-level groundwork towards overcoming these limitations. Specifically, we condense the challenges outlined in section I-A into two main research questions:

- **RQ<sub>1</sub>:** *How can redundancies in multi-revision program analysis be minimized?*
- **RQ<sub>2</sub>:** *How can we abstract over different programming languages in the context of high speed code analysis?*

Even though these are two distinct problems, imposing the restrictions of both problems onto a solution requires us to consider overlaps and mutual exclusions between the two. For example, performing multi-revision analyses for only one specific language is a very different goal from constructing a framework that accelerates code analysis for arbitrary languages, as it may allow the re-use of language-specific tools and data structures, such as a compilers and language-specific code representations. On the other hand, meta-models tend to be heavy-weight under most circumstances, whereas the high-speed requirement forces us to find a more light-weight solution. Ultimately, we believe that a generic solution is of greater benefit for the scientific community. That said, many of the techniques we present could be used independently.

## Our contributions can be summarized as follows:

- 1) We develop a novel algorithm for loading and analyzing a large number of different revisions of a project asynchronously and with minimum redundancy, enabling the concurrent analysis of thousands of revisions at minimal marginal cost for each additional revision.
- 2) We introduce light-weight entity mappings, a novel and flexible method for easily formulating ad-hoc mappings between language-specific entities and cross-language semantic concepts, without the requirement for an actual translation to a pre-defined meta-model.
- 3) We provide the research community with an open-source implementation of our approach, called LISA<sup>1</sup> (Lean Language-Independent Software Analyzer), as well as a replication package including 8.5 GB worth of fine-grained historical code measurements<sup>2</sup>.

## II. APPROACH

In the following sections, we describe several novel techniques which vastly reduce the redundancies inherent to analyzing multiple revisions of a program. This includes a multi-version graph representation of source code, an asynchronous parsing and deduplication algorithm and a highly flexible abstraction that enables the analysis of multiple programming languages using the same analysis instructions. We combine these techniques in LISA, although they could be applied independently to improve different aspects of existing applications.

<sup>1</sup><https://bitbucket.org/sealuzh/lisa>

<sup>2</sup><http://tiny.uzh.ch/C0>

### A. Multi-revision Graph Representation of Source Code

In this section, we describe a condensed graph representation which allows us to analyze the source code of thousands of revisions simultaneously.

**Graph representation.** We define a directed graph as an ordered pair  $G = (V, E)$ , consisting of a vertex set  $V$  and an edge set  $E$ . Each vertex  $v \in V$  is a tuple  $(i, M_v)$  consisting of a unique identifier  $i$  and a map  $M_v$  containing metadata on the vertex in the form of  $(k, m)$  key-value pairs. Every edge  $e \in E$  is a tuple  $(s, d, M_e)$  identified by the source and destination vertices  $s$  and  $d$  and can also contain a map  $M_e$  with metadata describing the edge. Such a graph is commonly used to model various representations of a program. For example, using a language grammar, one can create a concrete syntax tree (CST) for each file in a project, such that every vertex  $v$  represents a symbolic token in the original source code, connected to zero or more child vertices, populating the graph with disconnected trees (one for each file). In this case, vertices identify by their file name and syntax tree path, e.g., `/src/Main.java/CompilationUnit0` and store their literal token as metadata, while edges store no metadata. Similarly, one can build an abstract syntax tree (AST), where multiple tokens are interpreted and combined into higher-level entities. In that case, vertices also store additional metadata depending on the vertex type. For example, an AST vertex may represent a Java class and store metadata such as its name and visibility. A graph is also suitable for modeling a compiled program, in which case relationships between different parts of the program can be represented by additional edges containing metadata about the relationship, e.g., whether it describes attribute access or a method call. The metadata also serves as a container for additional data created during analysis. For example, a vertex representing a class may contain a method count as part of the metadata.

LISA uses Signal/Collect [47], a low-level graph framework, to store arbitrary graphs. Clients implement an interface with just two members: a definition of suffixes of file names that the client supports and a *parse* routine, which receives the content of a file and an agent for modifying the graph. Thus, any pre-existing parsers can be used in conjunction with LISA and any kind of revisional data can be represented. The data

contained in  $M_v$  and  $M_e$  is hence determined by the kind of graph that is loaded and by the analyses performed. Our implementation ships with adaptors for the JDK *javac* parser and for ANTLRv4. ANTLR is a parser generator for which existing grammars can be found for many languages<sup>3</sup>. Given an ANTLR grammar, it is possible to create a new LISA-compatible parser with only a few lines of boilerplate code.

**Asynchronous code-loading.** Differently from traditional tools, which check out individual revisions from the version control system to act on source code contained in files and folders, our implementation encourages direct, asynchronous access to the sources of multiple revisions. Our `SourceAgent` interface, which we implement for Git, mandates the preliminary creation of key-value pairs  $(p, B)$  for each file in any revision of the project, where  $p$  identifies a file by its path and  $B$  denotes a chronologically ordered sequence of sources for the file in different revisions. This way, each pair can be assigned to one of many parallel workers, which read the sources for a single file sequentially, while sources for different files can be read independently from each other. In Git, this extraction step can be performed in a single traversal of the Git graph database, without the need to ever check out a working copy of the code, as illustrated in fig. 2. That said, this technique will work with any version control system that allows direct access to individual revisions of files.

This approach has four distinct performance benefits over the traditional approach: Avoiding the checkout to a working directory first reduces the number of reads for each relevant file from 2 to 1, because each blob is only read once during parsing and never for a checkout, and it reduces the number of writes from 1 to 0, since blobs are never re-written to the file system. Second, it avoids the unnecessary checkout of any other assets which are not even relevant for the analysis. Third, it allows for the asynchronous, parallelized parsing of many revisions simultaneously, which is otherwise only possible by checking out multiple copies of the entire source code to different locations. Finally, files are only read in revisions where they underwent actual change.

**Multi-revision graph representation.** Typically, a single commit only changes a tiny part of an application. This means that the graph representations of two adjacent revisions overlap to a large degree and that most vertices are unchanged for many consecutive revisions (thousands, in practice). To avoid creating separate graphs for each revision, we extend our graph representation such that each vertex tuple  $(i, R, M_{vR})$  carries an additional set  $R$  denoting one or more *revision ranges*. A revision range  $r(s, e)$  consists of start and end revisions  $s$  and  $e$  and indicates in which revisions the vertex exists. Likewise, the metadata  $M_{vR}$  of a vertex now stores metadata for each revision range that the vertex represents. Edges on the other hand do not require any special treatment: whether or not an edge exists in a revision is determined by whether the two connecting vertices exist. In this paper, we will generally




Fig. 2. Contrary to the traditional method of checking out (i.e., reading and writing) *all* files in each revision and then parsing the source code from the filesystem (a), parsing the source code directly from a bare Git repository requires only a single read operation per *relevant* file (b).

<sup>3</sup>For example, <https://github.com/antlr/grammars-v4> contains grammar files for over 60 structured file formats.

denote a range as two dash-separated numbers in brackets, e.g., a revision range with start 10 and end 15 is denoted as  $\{10-15\}$ . No information is lost using this graph representation, as the structure of a single revision can still be identified simply by selecting only vertices whose revision ranges contain the desired revision. Even so, it saves a tremendous amount of space, because the majority of the graph remains unchanged for much of the entire evolution of a project.

### B. Low Redundancy Graph and Metadata

The interface provided by LISA for clients to populate the graph is transparent with regard to this multi-revision representation. When a client reads a file for a particular revision and adds a vertex with an identifier  $i$  to the graph, LISA transparently handles the case where a vertex with the same identifier already exists, and adds or extends a revision range inside the existing vertex to accommodate the new revision. This has several important consequences:

- 1) A vertex may contain differing metadata in different revisions. A simple example for this are literal tokens: There may be an “Integer” syntax vertex, whose metadata contains the number 5 as a literal in one revision, but 7 in another. In this case, two separate, adjacent ranges with different metadata are necessary to accurately represent both revisions using the same vertex.
- 2) If the same file was parsed in two adjacent revisions (e.g., 5 and 6), then it makes sense to merge the ranges  $\{5-5\}$  and  $\{6-6\}$  into a common range  $\{5-6\}$  for all vertices that underwent no change.
- 3) If a file has not undergone any changes in a revision, then, as per the previous section, it will not be parsed, leading to erroneous gaps in the revision ranges of a vertex. For example, if a file underwent changes in revisions 5 and 10, it will not be parsed in revisions 6 to 9, resulting in two ranges  $\{5-5\}$  and  $\{10-10\}$ . But the syntax nodes are actually present in-between these revisions, necessitating a range extension for a final configuration of  $\{5-9\}\{10-10\}$ , or even  $\{5-10\}$ , if the two parsed revisions contain the same metadata in the given vertex.
- 4) Finally, the metadata of a vertex for two adjacent revisions may also exhibit some overlap. Imagine a vertex representing a method during a computation which calculates the cyclomatic complexity and statement count of individual entities. Between two adjacent revisions, the method may have the same name and complexity, but a different statement count. In this case, two separate ranges are necessary to describe the vertex, yet most of the metadata is shared.

In the following section, we describe the data structures and algorithm used to efficiently build and store the revision-specific data in such a fashion that *neither vertices nor metadata are unnecessarily duplicated*.

**Sparse graph loading algorithm.** We explain the LISA range compression, shown in algorithm 1, by following the evolution of a single syntax tree vertex  $v$  of an exemplary project with just 10 revisions, as illustrated in table I.  $v$  is

parsed from a file  $f$  which Git stores as a blob for each revision where it underwent any change. The first row in table I enumerates the different revisions of the project. The letters in the second row denote unique metadata configurations that exist in  $v$ . In this example, the node has the same metadata in revisions 0 to 3, then the node is deleted in revision 4, then it reappears with different metadata in revisions 6 and then the metadata is again changed in revisions 8. The third row indicates whether the file  $f$  has been affected by a revision. In this example,  $f$  is parsed in revisions 0, 1, 3, 4, 5, 6 and 8. This implies that in revisions 2, 7 and 9 there were no changes in  $f$ , meaning that all vertices that existed in previous revisions (1, 6 and 8) are unchanged and still exist, even if the file wasn’t re-parsed. In revisions 4 and 5,  $f$  was parsed, but  $v$  did not appear in the source code of those revisions.

When a client adds a vertex to the graph and this vertex already exists, then LISA updates the vertex (UPDATEVERTEX in algorithm 1) by simply adding another single-revision range and queuing a range compression task leading up to the preceding revision. Hence, the range compression itself also runs asynchronously and “cleans up” behind the parser, compressing ranges added by the parser. For example, the parser may already have parsed the first 4 revisions (0 to 3), meaning that it will have created 3 single-revision ranges  $\{0-0\}$ ,  $\{1-1\}$  and  $\{3-3\}$  in vertex  $v$ . Since the metadata of  $v$  is

---

**Algorithm 1** LISA range compression algorithm. The *ranges* variable is a map from revision ranges to metadata and contains all metadata for a single vertex.  $\leftarrow$  assigns,  $\rightarrow$  references.

---

```

1: procedure UPDATEVERTEX( $v, rev, meta$ )
2:   CREATERANGE( $v.ranges, rev, meta$ )
3:   queue(COMPRESSRANGES( $v.ranges, v.mark, rev-1$ ))
4: procedure CREATERANGE( $ranges, rev, meta$ )
5:    $new \leftarrow (\{rev, rev\} \rightarrow meta)$ 
6:    $ranges \leftarrow ranges + new$ 
7: procedure COMPRESSRANGES( $ranges, mark, end$ )
8:   if  $end = null$  then ▷ Case 1
9:      $mark \leftarrow ranges.head$ 
10:  else
11:     $next \leftarrow ranges.get(end + 1)$ 
12:    if  $mark = null \ \& \ next = null$  then ▷ Case 2
13:      doNothing
14:    else if  $mark = null \ \& \ next \neq null$  then ▷ Case 3
15:       $mark \leftarrow next$ 
16:    else if  $mark \neq null \ \& \ next = null$  then ▷ Case 4
17:       $extended \leftarrow (\{mark.start, end\} \rightarrow mark.meta)$ 
18:       $ranges \leftarrow ranges - mark + extended$ 
19:       $mark \leftarrow null$ 
20:    else if  $mark \neq null \ \& \ next \neq null$  then
21:      if  $mark.meta = next.meta$  then ▷ Case 5
22:         $merged \leftarrow (\{mark.start, next.end\} \rightarrow mark.meta)$ 
23:         $ranges \leftarrow ranges - mark - next + merged$ 
24:         $mark \leftarrow merged$ 
25:      else if  $mark.meta \neq next.meta$  then ▷ Case 6
26:         $extended \leftarrow (\{mark.start, end\} \rightarrow mark.meta)$ 
27:         $ranges \leftarrow ranges - mark + extended$ 
28:         $mark \leftarrow next$ 

```

---

TABLE I  
METADATA OF AN EXEMPLARY GRAPH VERTEX.

Revision	0	1	2	3	4	5	6	7	8	9			
Content	A	A	A	A			B	B	C	C			
Touched	✓	✓		✓	✓	✓	✓		✓				
Revision	0	1	2	3	4	5	6	7	8	9	End	Mark	Case
0	A										-	-	1
1	A	A									0	{0-0}	5
2	A											{0-1}	
3	A		A								2	{0-1}	5
4			A								3	{0-3}	4
5			A								4	-	2
6			A				B				5	{6-6}	3
7			A				B					{6-6}	
8			A				B		C		7	{6-6}	6
9			A				B		C			{8-8}	
final			A				B		C		9	{8-8}	4

the same in all these revisions, including revision 2, where no changes were made to the file at all, the goal of the algorithm is to combine these ranges into a single range  $\{0-3\}$ .

We now step through the execution of the compression algorithm for each revision of the example vertex. Note that for each possible *case* in the algorithm, table I contains the case number in the last column. Also note that “marking” a range simply means saving a reference to that range as a temporary property of the vertex. This *mark* indicates whether a revision range is “open-ended” in the sense that it may be extended in the next revision. After parsing revision 0, the compression algorithm simply needs to *mark* the first existing range in  $v$  (case 1). After parsing revision 1, it finds a *marked* range and an existing range at revision  $end+1 = 1$ . It compares the two ranges, finds that they are equal (case 5), merges them and *marks* the resulting range. Revision 2 saw no changes in  $f$ , so  $v$  remains unchanged by the parser. Note however that in reality, the code represented by  $v$  continues to exist, hence the revision range needs to be extended. This happens in revision 3, where  $f$  was modified again. The algorithm finds the *marked* range and a range at  $end+1 = 3$  containing the same metadata, prompting a merge (case 5). In revision 4, there exists a *marked* range, but there is no range at  $end+1 = 4$ . A merge is not necessary, so it just *unmarks* the existing *marked* range (case 4). In revision 5, there is no *marked* range and no *next* range, so nothing happens (case 2). This simply means that even though  $f$  had undergone changes, they did not affect  $v$ , which does not exist in  $end$  or in  $next$ . In revision 6, there is no *marked* range, but there is a *next* range, so it is *marked* (case 3). Revision 7 saw no changes. In revision 8, there is a *marked* range and a *next* range, but they are not equal, so the *marked* range is extended until  $end$  and then the *next* range is *marked* instead (case 6). Revision 9 saw no changes. In the end, a final pass is made through all vertices in the graph, extending any *marked* ranges to include

the last revision. This is necessary because these *marked* ranges previously ended with the revision where the file was last *parsed*, although these vertices actually exist until the end of the project history. In this example, the *marked* range is extended to include the last revision (case 4).

As a result of the compression, instead of storing metadata for each revision (eight in case of  $v$ ) only three ranges with different metadata remain (visible in the final row of table I). In practice, revision ranges tend to comprise a much larger number of revisions. If a file is added to a project in an early revision, it can be that the majority of vertices within exhibit the same metadata for thousands of subsequent revisions.

**Range splitting and shared metadata.** After loading the graph, the initial metadata for a vertex will likely be contained in a small number of revision ranges. However, once the graph is being analyzed (as detailed in section II-C), the metadata in different ranges may start to diverge to a certain degree. For example, a “Class” node might exist in revisions 5 to 280 resulting in a single all-encompassing revision range  $\{5-280\}$ , but its attribute count could be computed as 4 in the first 100 revisions and as 7 in the remaining revisions. This means that the revision range  $\{5-280\}$  needs to be split into  $\{5-105\}\{106-280\}$ , and that separate metadata needs to be stored for those two revision ranges. LISA performs these splits dynamically during the computation whenever necessary. However, storing completely separate metadata for each revision range would again introduce significant data duplication: in our current example, only the attribute count will be different for the two ranges. Hence, we store metadata in referentially transparent, immutable data structures called *case classes*, which are essentially ‘copy-on-write’, meaning that while two distinct metadata containers exist after the split, all their data is actually shared, except for the attribute count.

### C. Multi-revision Code Analysis

Signal/Collect, the graph framework used in LISA, operates under a computational paradigm where vertices in the graph communicate to form an emergent computation. More specifically, each vertex can *signal* other vertices on outgoing edges to transmit arbitrary messages, *collect* incoming signals and *modify* its own state. We ask the reader to refer to [47] for an in-depth description of this paradigm.

To formulate analyses in LISA, the user defines a *data structure* for analysis data (which LISA integrates into the metadata of a vertex), a *start* function governing how and where in the graph the first signals are emitted for this particular analysis, and an *onCollect* function, which determines how incoming signals are processed. Both these functions are *side effect free* and return ‘new’ metadata to reflect changes (however, referential transparency of case classes in Scala prevents unnecessary data duplication). An exemplary implementation for calculating cyclomatic complexity [48] can be found online<sup>4</sup>. The data structure used by this analysis (defined on line 15) contains a boolean to determine whether

<sup>4</sup><http://tiny.uzh.ch/ED>

Listing 1  
DEMO.JAVA

```
1 public class Demo {
2     public void run() {
3         for (int i = 1; i < 100; i++) {
4             if (i % 3 == 0 || i % 5 == 0) {
5                 System.out.println(i);
6             }
7         }
8     }
9 }
```

or not to persist the computed value for a given vertex and a single integer value to hold the actual complexity value. The start function causes an `MccPacket` to be sent out from each *leaf* vertex in the graph (lines 19-23). The value contained in the packet is 1 by default, or 2 if the leaf vertex itself is already a node that increases the complexity of the program. In regular programs, having a leaf node contribute to the complexity is impossible (e.g., there cannot be an ‘If’ statement with no truth check beneath it in the tree), but as we will see in the next section, LISA filters irrelevant nodes during parsing, such that the entire tree below a branching statement may be omitted from the graph. The `onCollect` function accumulates incoming values according to the definition of cyclomatic complexity (line 28) and once all child values have been received, it sends the complexity of the local vertex to its parent (closure starting at line 32). A transient value is computed and stored in the metadata of each vertex in the graph as a consequence of this algorithm, but the analyst would likely want to know the complexity only for specific kinds of nodes. The determination, whether the `mcc` value for a given vertex should be persisted is done on line 34, where the `persist` property is set `true` if the vertex is a class, method or file. The user could easily add other vertex types (e.g., blocks or closures) if needed.

Once an analysis is completed, the analysis results are persisted using a user-provided persistence strategy, for example to store values in a database for further processing. LISA ships with a persister that will dump all data, for which persistence has been enabled, into a sparse CSV file, where each line corresponds to a vertex in a given range and each column corresponds to one kind of data. Where data has not been computed for a specific vertex, the cell remains empty.

Since the computations are executed directly on the range-compressed graph, calculations are executed only once for a revision range of any particular vertex and outgoing signals are also attached to a revision range. The revision range of a vertex may be split, as described in section II-A, if an incoming signal concerns a partial range of the receiving vertex. In this fashion, the number of computations necessary to compute metrics for individual commits is vastly reduced.

#### D. Multi-language Analysis Genericism

Programming languages share certain concepts and many code-related analyses can be expressed for different languages. And while the concrete syntax tree representation can vary greatly for different languages and parsers, the *relative* structure can be fairly similar. For example, when comparing the

Listing 2  
A VERTEX LABEL MAPPING FOR JAVA.

```
1 object AntlrJavaParseTree extends Domain {
2     override val mapping = Map(
3         'method -> Set("MethodDeclaration"),
4         'branch -> Set("IfStatement", "ForStatement",
5             "WhileStatement", "CatchClause",
6             "SwitchLabel", "TryStatement",
7             "DoStatement", "ConditionalExpression")
8     )
9 }
```

ASTs parsed from a Java and a Python program, the exact sequence, labeling and nesting of vertices leading from a root node to the leaf nodes of a method may differ greatly. However, structural features (where the vertices are located relative to each other), in the context of code analysis (and not program compilation) are very similar: A method/function vertex that has a class vertex as a parent is contained in that class; or: two `ifs` on the same level represent 4 possible paths through the local scope, while one `if` nested in another creates only 3 possible paths. This means that we primarily need to explicitly capture *entities*, while retaining the *structure* as offered by the parser. Furthermore, and especially in the context of large scale code analysis, the representation of the source code should capture only the minimum necessary for a particular analysis. Consider we wanted to compute the cyclomatic complexity of the example in listing 1. It contains 140 AST vertices when parsed using ANTLR, yet most of them are entirely irrelevant towards the complexity metric.

The solution we propose to solve this problem is notably simple: It requires a one-way many-to-many relation  $T \rightarrow L$  from the domain of entity types  $T$  required for a particular analysis, to the co-domain of parser-specific labels  $L$  used by a particular language. In LISA, this relation is simply expressed as a map from symbols to sets of labels. Listing 2 shows an example which is sufficient for the complexity analysis, which notably doesn't contain any language-specific labels. It specifies how the entities mentioned by the analysis (`'method` and `'branch`) can be identified in Java. Note that the analysis also checks for other potential symbols (`'class` and `'unit`), but specifying these is not mandatory. Applied to the `Demo` program, this mapping populates the graph with a mere 5 vertices connected in a straight line: one vertex each for the file `Demo.java` itself, the method `run`, the `for` loop, the `if` statement and the `||` operator. All other vertices (such as numbers and other operators) are automatically ignored at the parsing step. When applying `MccAnalysis` to this graph, the result (a complexity of 4, which is persisted only for the vertex matching the `'method` symbol) is still correct.

Note that it is still possible to write language-specific analyses by specifying the explicit label used by the parser instead of a mapped symbol, in case a language-specific structure needs to be analyzed.

#### E. Approach Summary

Compared to traditional approaches, LISA avoids redundancies at every juncture of multi-version analysis: It analyzes code taken directly from the Git database and only re-reads

files containing changes, while user-defined mappings vastly reduce the number of vertices required to perform analyses on the loaded source code. Both code and metadata present in multiple revisions is stored only once and the computations themselves are also executed only once for each revision range at the vertex/subtree level, avoiding the expensive re-computation of data at the file level. Finally, the computed data is selectively persisted to keep the results manageable for further analysis. Furthermore, the user-defined mappings not only reduce the size of the graph, but also make analyses transferable to other languages, which LISA can easily support via existing ANTLR grammars or through dedicated parsers.

### III. EVALUATION

In RQ1, we ask how redundancy can be avoided when analyzing multiple revisions of a project. Our solution, as summarized in section II-E, combines multiple novel techniques and we want to know how effective they are. For RQ2, we want to know if we can analyze code from multiple programming languages using the same underlying framework and analysis formulations. We want to gauge the practicality of light-weight mappings, our proposed solution. We frame our evaluation in a practical research context, by attempting to answer a fundamental question with regard to software evolution analysis: *How many releases of a project have to be analyzed to gain an accurate picture of its evolution?*

Our evaluation is structured as follows. In section III-A, we conduct a large-scale study over the full history of 300 projects written in 3 different languages. Based on this, we discuss the performance of our approach and the effectiveness of its redundancy removal techniques in section III-B and also discuss the practicality of light-weight mappings. In section III-C, we present the results of our artifact study. Finally, we discuss threats to our study in section III-D.

We provide a comprehensive replication package, including all scripts used in the evaluation as well as all the resulting data (8.5 GB worth of analysis results) online<sup>5</sup>.

#### A. Code Study

Using research reports and statistics on popular Java and JavaScript projects [49], [50] as well as GitHub’s “most starred” query option, we selected 100 Java, C# and JavaScript projects, ignoring projects that do not contain an application per se (such as tutorials and demo projects) and preferring projects with a larger number of commits (for a median of 2956 commits for Java, 2051 for C# and 1097 for JavaScript). We proceeded to formulate analyses to compute the number of classes, methods, method parameters, variables and statements, the cyclomatic complexity, control flow nesting depth and number of distinct control flow paths, and to detect the Brain-Method [51] code smell. We also compute the number of direct children and total number of vertices beneath each vertex as a proxy for the size of the code base. We fed LISA with existing ANTLR grammars for all three languages and proceeded to

define suitable mappings to match the entities relevant to our analyses to the ones used in the ANTLR grammars. Even though LISA supports analyzing multiple languages within the same computation, we wanted to observe its performance characteristics on a language-by-language basis, so that we can assess the impact of the used parser on overall performance.

We ensured the correctness of the formulated analyses (and thus the range compression algorithm) by creating a sample project for each language which contains a large number of code combinations varying across multiple revisions, manually calculating the expected code metrics and confirming that they match the results procured through LISA. Then we ran the tool on the Git URLs of the 300 projects to compute the metrics for all revisions connected to the Git HEAD and persisted those metrics at the file, class and method level.

LISA’s hardware requirements scale with the size of the project and the number of revisions (or rather the amount of actual change between revisions). Up to thousands of revisions and 100’000s of LOC, commodity hardware with 4-16GB of RAM is sufficient. However, analyzing big projects, such as Mono, containing over 100k revisions and over 5M LOC, requires more RAM. To accommodate the analysis of such projects, we used a Google Compute Engine instance with 24 cores and 158GB of RAM to analyze all projects.

#### B. Evaluation of the Approach

The figures shown in table II indicates that the redundancy reduction techniques we apply are extremely effective. We briefly revisit the challenges outlined in section I-A:

**Source data duplication.** LISA reads directly from bare Git. Even though this may seem a simple matter, to our knowledge, no existing *code* analysis tool takes advantage of this technique. Not only does it avoid checkouts, which would be costly both in terms of time and storage space, but it also enables LISA to parse the source code of many releases in parallel, further speeding up the analysis.

**Redundant analysis of unchanged code.** LISA only reads the bare minimum necessary to still capture a complete representation for analysis. It’s crucial to note that table II shows the number of files and lines *actually parsed* by LISA, which are orders of magnitude *smaller* than the volume contained in all analyzed releases added together. To give a concrete example: analyzing **all 106 160 revisions** of the `mono` project, LISA only had to read 160 422 521 lines of code. This is merely **40 times** the number of lines of code contained in the most recent revision alone (3 914 887 LOC). Likewise, the 2.5 billion lines of code LISA parsed also represent only a fraction of the total code volume actually contained in all revisions of all projects.

**Manual effort.** Formulating the analyses (198 lines of Scala for the 11 metrics in our study) constitutes the lion share in terms of manual effort, while writing the language mappings is a matter of minutes. One simply reads the grammar and selects the right labels for different entities. To analyze a project, LISA only requires the Git URL. Since it acts solely on source code, the analyst need not bother with libraries, build tools and other hindrances. Thus, LISA allows for short turn-around

<sup>5</sup><http://tiny.uzh.ch/C0>

TABLE II

AMOUNT OF CODE ANALYZED AND ANALYSIS DURATIONS IN THE STUDY.

		Java	C#	JavaScript
Projects		100	100	55
Revisions analyzed	total	646 261	489 764	204 301
	smallest	<sup>1</sup> 715	<sup>2</sup> 394	<sup>3</sup> 234
	largest	<sup>4</sup> 29 392	<sup>5</sup> 106 160	<sup>6</sup> 12 038
	median	2 956	2 051	1 097
Files parsed	total	3 235 852	3 234 178	507 612
	smallest	<sup>7</sup> 946	<sup>2</sup> 964	<sup>8</sup> 10
	largest	<sup>9</sup> 300 185	<sup>10</sup> 312 229	<sup>11</sup> 61 843
	median	11 171	10 685	2 198
Lines read	total	1 370 998 072	961 974 773	194 758 719
	smallest	<sup>12</sup> 328 118	<sup>13</sup> 55 971	<sup>8</sup> 212
	largest	<sup>9</sup> 194 875 095	<sup>5</sup> 160 422 521	<sup>6</sup> 18 194 261
	median	3 406 857	2 235 404	669 704
Runtime	parsing	13:25h	52:12h	29:09h
	analysis	1:35h	2:20h	1:00h
	total <sup>A</sup>	18:43h	57:20h	30:33h
	shortest <sup>A</sup>	<sup>7</sup> 10s	<sup>13</sup> 9s	<sup>15</sup> 9s
	longest <sup>A</sup>	<sup>9</sup> 2:44h	<sup>5</sup> 8:43h	<sup>6</sup> 5:56h
	average <sup>A</sup>	11:14min	34:24min	18:20min
	median <sup>A</sup>	2:15min	4:54min	3:43min
	total avg./rev. <sup>B</sup>	84ms	401ms	531ms
	median avg./rev. <sup>B</sup>	30ms	116ms	166ms

<sup>A</sup>) Including time spent cloning repositories and persisting results.<sup>B</sup>) Excluding time spent cloning repositories and persisting results.<sup>1</sup>javaassist <sup>2</sup>shadowsocks-windows <sup>3</sup>hain <sup>4</sup>cloudstack <sup>5</sup>mono<sup>6</sup>ember.js <sup>7</sup>commons-email <sup>8</sup>awesome-react <sup>9</sup>xtext <sup>10</sup>ravendb <sup>11</sup>babel<sup>12</sup>Android-Universal-Image-Loader <sup>13</sup>Fody <sup>15</sup>debug <sup>16</sup>d3

times when doing exploratory research but also enables the analysis of large samples without excessive wait times.

**Low-resolution historical data.** Analyzing all 1 341 802 revisions, capturing the detailed history of 300 projects using LISA took 4½ days. To put this into perspective, Tufano et al. report that in a recent study, the computation of metrics and code smells in 579 671 revisions of 200 projects required 8 weeks on a 28-core machine [20]. We argue that tools like LISA can enable new kinds of studies, which are currently not feasible because of the tremendous effort involved in creating high-resolution historical code measurements, and we demonstrate such a study in the following section.

**Isolation of programming ecosystems.** We demonstrate that for the computation of basic metrics (which are known to correlate with more advanced metrics [52], [53]), light-weight entity mappings and the principle of avoiding explicit modeling of structure are suitable. Formulating analyses in our study did not require language-specific instructions. However, the mapping of some entities may not be entirely obvious: for example, for JavaScript we mapped the class concept to apply to root level AST vertices, because there is no convention on how to represent classes in JavaScript. However, we find light-weight mappings to work well in practice. LISA comes with additional mappings for Lua and Python 3.

**Performance compared to other tools.** A fair one-on-one comparison to other tools is not feasible, as each tool has a different feature sets, restrictions and capabilities. In previous work [32], we compared the performance of a LISA prototype, which lacked most of the performance-enhancing techniques discussed in section II, to two existing analysis tools, namely inFusion [54] and SOFAS [55] for analyzing the AspectJ project. In that comparison, the prototype took 1:31min




Fig. 3. Analysis duration vs. LOC for Java (•), C# (+) and JavaScript (×) projects.

to analyze a single revision, outperforming SOFAS by a factor of 9.8 and inFusion by a factor of 4.3. The average time needed to analyze one revision fell below 2 seconds when analyzing more than 100 revisions and below 900ms when analyzing more than 1000 revisions, whereas using the other tools, each additional revision to be analyzed incurs the same cost.

We can now compare the performance of LISA directly to the original prototype and by extension, to the tools used in the original study [32]. Analyzing a single revision of AspectJ, the prototype spent 1:31min while LISA spent only 37s. Of the total duration, the prototype spent 31s on parsing and building the code graph and 60s on the analysis, while LISA spent 22s and 15s respectively. When analyzing thousands of AspectJ revisions, the prototype spent 650ms on average per revision, while LISA spent only 34ms. Of this, the original prototype spent more than 500ms on parsing and graph building, and around 80ms on the analysis. LISA on the other hand spent 28ms on parsing and 5.9ms on the analysis. The resulting average of 34ms per revision is near the median average across all Java projects, as shown at the bottom of in table II.

The parsing speed improvement can be attributed both to the filtered parsing, enabled by the light-weight mappings (section II-D), and to the asynchronous multi-revision parsing algorithm (section II-A). The original prototype stored the entire parse trees as provided by the parser and could only parse files for one revision at a time. The analysis speed improvement follows naturally from the filtered parsing, as the signals need to travel much shorter distances within the graph. This demonstrates how the chosen meta-model (or rather, lack thereof) has a significant impact on the overall performance.

### C. Sampling Revisions in SE Research

An often recorded threat in existing code studies is the limited number of releases analyzed. While our tool can perform certain analyses for every commit, we hypothesize that there exists an optimal sampling interval which captures the evolution of a project with sufficient accuracy, such that slower, yet more feature-rich tools can be applied more effectively on a limited sample of revisions. However, to do so, we first need to identify this sampling interval.

The major releases of `facebook-csharp-sdk` are shown, again, in fig. 4, but with the fine-grained data obtained upon analyzing all 1611 commits using LISA overlaid. Even from the naked eye, it's clear that analyzing only major releases does not give a good approximation of the *overall*






Fig. 4. Method Count and Class Count over the entire history of the facebook-csharp-sdk project.

evolution of the project. Measuring the error between the assumed evolution based on sampling only major releases and the actual evolution in terms of the difference of the Area under the Curve (AUC) reveals a discrepancy upward of 30%.

To generalize this problem and to find an appropriate sampling interval, we apply the following procedure to each project. For 5 metrics from our data set (shown in fig. 5), we create two series: the first contains all, while the second contains  $s$  equally spaced commits from the entire history of the project, using an initial value of  $s = 4$ . We calculate the AUC of the high-resolution series  $AUC_{real}$  as well as the sparse series  $AUC_{sampled}$  and calculate the absolute difference  $AUC_{error} = |AUC_{real} - AUC_{sampled}|$ . We quantify the mismatch as  $e = AUC_{error}/AUC_{real}$ . We then keep incrementing  $s$  by 1 until we find the  $s$  for which  $e \leq 0.05$ . The sampling interval is therefore given by  $i = c/s$  where  $c$  is the number of commits. We repeat the experiment for targets  $e \leq 0.01$  and  $e \leq 0.1$  to see how sensitive the error is to different sampling intervals. The results for  $e \leq 0.05$  are visualized in fig. 5. We share the following insights:

- The appropriate sampling interval  $i$  varies for different programming languages. At  $e \leq 0.05$  the median  $i$  for Java projects is near 600, while for C#, it is close to 400. For JavaScript, even lower  $i$  are necessary.
- For Java and C#,  $i$  is stable across all metrics, while for JavaScript,  $i$  fluctuates depending on the metric.
- $i$  can vary significantly across projects. Sampling every 5000th commit can still be sufficient in some Java and C# projects, while other projects require  $i \leq 50$ .
- With  $e \leq 0.01$ , the median  $i$  drops significantly, to around 80 for Java, 40 for C# and below 5 for JavaScript. With  $e \leq 0.10$  it reaches  $\sim 800$  for Java,  $\sim 500$  for C# and  $\sim 200$  for JavaScript.

Software evolution studies need to carefully consider the number of revisions analyzed in order to maintain validity. The decision depends on the language, but also on the acceptable error margin for a particular study. A safe recommendation, at least given a large enough number of projects in a study, is to analyze every 250th commit for Java, every 150th commit for C# and every 25th for JavaScript projects. Finally, even when using sampling, projects with thousands of commits still require a large number of commits to be analyzed for an accurate description of their evolution.




Fig. 5. Maximum sampling interval accurately capturing the evolution within a 5% margin of error based on 100 projects for each of 3 languages.

#### D. Threats to Validity

**Threats to construct validity** concern the relationship between theory and observation. We ensure the correctness of analyses performed by LISA in a controlled environment by manually calculating metrics on exemplary multi-revision projects. This human judgement may be error-prone, but we tried to alleviate this issue by double-checking the results.

**Threats to external validity** concern the generalizability of our findings. One threat is the choice of projects used in our experimentation. To mitigate this issue we selected a large sample projects for each programming language and based our choice on their popularity according to multiple sources. Furthermore, all projects in our sample are open-source, and closed-source projects may generally undergo a different evolution. To extend the generality of our findings, we may replicate our study for additional projects and programming languages in future research.

#### IV. LIMITATIONS AND FUTURE WORK

Attempting to increase the performance of software analysis tools is not trivial and inescapably requires a trade-off between comprehensiveness of analyses and the speed of execution. Any sort of code analysis can be formulated in LISA, but considering it acts exclusively on source code (without access to a compiler), analyses which concern a single file (as opposed to cross-file analyses) are easiest to formulate. One could argue, that time used to reformulate analyses in LISA could be better spent using language-specific tools. However, other tools cannot easily take advantage of the zero-redundancy approach to multi-revision analysis of LISA. That said, the LISA graph can hold arbitrary data, as outlined in section II-A, so it would also be possible to, for example, integrate a Java class loader to load multiple pre-compiled versions of a program into LISA and perform much more complex analyses, such as coupling or control flow analyses, on the resulting multi-revision graph.

A technical limitation arises from how LISA identifies vertices in the graph. The identifier of a vertex always corresponds to the file and syntax tree location of a node. Hence, renaming a file constitutes the creation of a new file, whose code will be represented using new vertices, even if nothing inside the file has changed. Likewise, if the order of methods in a class is changed, the subtrees will be partially re-created. This problem is hard to solve in general [56], [57], even more so in a language-independent manner, because an alternative

identification scheme would be required, probably involving the actual names of entities within the file. Note however that this limitation *only* concerns the redundancy reduction in LISA. The actual analyses are robust with regard to these kinds of changes and their results are unaffected. If an analyst wishes to keep track of language-specific entities, she can simply formulate an analysis to connect all vertices belonging to the same program entity via additional edges.

It takes much more time to parse the source code than it takes to analyze it. It may be smart to find a way of persisting (and occasionally updating) the parsed graph such that when an analyst wants to run a computation, she does not need to re-parse everything. However, this has the downside that she must either decide on a rigid entity mapping or disable filtered parsing altogether, such that the entire syntax trees are represented in the graph (similarly to how BOA stores ASTs [29]), which in turn inflates memory requirements and computational load.

## V. RELATED WORK

In the past, code analysis frameworks had a strong focus on integrating different resources (like version control and bug trackers) as well as enabling replicable end-to-end analyses in the first place [28], [41], [55], [60]. More recently, efforts toward large-scale, high-performance analyses have yielded frameworks such as BOA [29], which utilizes a Hadoop cluster to permit the rapid execution of analyses on the metadata of over 8 million projects. BOA also provides access to parsed ASTs of numerous Java projects which can be analyzed using a visitor pattern. LISA, compared to these frameworks, is a lot more low-level: it is highly generic and provides algorithms to accelerate source code analysis, but does not target any particular infrastructure, language or type of analysis. Researchers can integrate it like any other JVM-compatible library to be used in a variety of scenarios.

Le et al. present a technique for patch verification across versioned control flow graphs (CFGs) [61] resembling LISA’s multi-revision code representation of ASTs, but it was designed for representing a few dozen revisions, not thousands. Similarly, TypeChef [62] uses a variability-aware parser to analyze multiple *configurations* of C programs (i.e., `#ifdefs`) in a shared graph. Both approaches have very different goals and are designed for problems on a much smaller scale compared to LISA.

Comparing our light-weight mappings to existing meta-models, the most significant difference is that the latter always imply a transformation of a source data structure (e.g., an AST) to a concrete instance of a meta-model. Meta-models such as FAMIX [63], KDM [64] or ASTM [65] as well as general-purpose models such as RSF [66] or GXL [67] model not only the kinds of nodes, but also their relationships and structure explicitly. The same is true for the cross-language model developed by Strein et al. [14] and  $M^3$  [68], which is specifically designed for use with Rascal Metaprogramming Language [69] and which also includes non-code concepts such as physical and logical source locations. Contrary to

this, the mappings in LISA only determine the *types of source nodes* which are loaded into the graph. The structure of the graph, however, will be identical to the original structure provided by the parser, minus any nodes which are not mapped. It is only in the context of a particular analysis that the mappings of individual nodes gain meaning. As such, our light-weight mappings can be considered a *view* onto an existing graph structure, rather than a meta-model. Furthermore, code models typically come with predefined entity types and relationships, whereas our light-weight mappings are formulated in the context of a particular analysis. Tichelaar et al., the makers of FAMIX, note that any code meta-model represents a trade-off between being too *coarse-grained* to be useful for a wide range of problems and being too *fine-grained* to remain sufficiently language-independent [63]. In this regard, LISA leaves it to the analyst to decide the level of coarseness.

## VI. CONCLUSION

We present several distinct redundancy removal techniques and demonstrate that in combination, they enable the comparatively rapid analysis of code contained 100 000s of commits. After formulating a particular analysis, the selection of projects, the creation of language mappings and the automated execution of analyses in our open-source tool, LISA, are straightforward and enable the quick extraction of fine-grained software evolution data from existing source code. We also present the idea of using light-weight mappings instead of traditional meta-models for static, structural analyses of code written in different languages. The light-weight mappings not only represent a simple, analysis-specific bridge between different languages, but they also play an important role in improving LISA’s performance, as they enable the filtering of unnecessary source data without sacrificing knowledge relevant to analyses. Our evaluation on a large sample of projects allowed us to observe that code in different programming languages evolves differently and that an accurate representation of their evolution is possible only with a large enough number of sampled commits.

LISA fills a unique niche in the landscape of software analysis tools, occupying the space between language-specific tooling used for the in-depth analysis of individual projects and releases, and traditional software repository mining, where code analysis is typically restricted to merely counting files and lines of code. The techniques discussed in this paper could be adapted for existing solutions individually, but LISA also offers clean and easy-to-implement interfaces for additional sources (i.e. version control systems), parsers and storage methods.

## ACKNOWLEDGEMENTS

We thank the reviewers for their valuable feedback. This research is partially supported by the Swiss National Science Foundation (Project №149450 – “Whiteboard”) and the Swiss Group for Original and Outside-the-box Software Engineering (CHOOSE).

## REFERENCES

- [1] T. Mende and R. Koschke, "Revisiting the evaluation of defect prediction models," in *Proceedings of the 5th International Conference on Predictor Models in Software Engineering*, ser. PROMISE '09. ACM, 2009, pp. 7:1–7:10.
- [2] E. Kocaguneli, T. Menzies, and J. Keung, "On the value of ensemble effort estimation," *Software Engineering, IEEE Transactions on*, vol. 38, no. 6, pp. 1403–1416, 2012.
- [3] M. D'Ambros, M. Lanza, and R. Robbes, "Evaluating defect prediction approaches: a benchmark and an extensive comparison," *Empirical Software Engineering*, vol. 17, no. 4-5, 2012.
- [4] M. VanHilst, S. Huang, J. Mulcahy, W. Ballantyne, E. Suarez-Rivero, and D. Harwood, "Measuring effort in a corporate repository," in *IRI. IEEE Systems, Man, and Cybernetics Society*, 2011, pp. 246–252.
- [5] T. Mens and T. Tourwe, "A survey of software refactoring," *Software Engineering, IEEE Transactions on*, vol. 30, no. 2, pp. 126–139, 2004.
- [6] G. Szöke, C. Nagy, R. Ferenc, and T. Gyimóthy, "A case study of refactoring large-scale industrial systems to efficiently improve source code quality," in *Computational Science and Its Applications – ICCSA 2014*, ser. Lecture Notes in Computer Science. Springer, 2014, vol. 8583, pp. 524–540.
- [7] B. D. Bois, P. V. Gorp, A. Amsel, N. V. Eetvelde, H. Stenten, and S. Demeyer, "A discussion of refactoring in research and practice," *Tech. Rep.*, 2004.
- [8] M. Kim, J. Nam, J. Yeon, S. Choi, and S. Kim, "Remi: Defect prediction for efficient api testing," in *Proceedings of the IEEE/ACM International Conference on Automated Software Engineering*. ACM, 2010, p. To Apper.
- [9] N. Nagappan, T. Ball, and A. Zeller, "Mining metrics to predict component failures," in *Proceedings of the 28th International Conference on Software Engineering*, ser. ICSE '06. ACM, 2006, pp. 452–461.
- [10] J. Bosch, "From software product lines to software ecosystems," in *Proceedings of the 13th International Software Product Line Conference*, ser. SPLC '09, 2009, pp. 111–119.
- [11] T. Mens, M. Claes, P. Grosjean, and A. Serebrenik, "Studying evolving software ecosystems based on ecological models," in *Evolving Software Systems*. Springer Berlin Heidelberg, 2014, pp. 297–326.
- [12] G. Bavota, G. Canfora, M. Di Penta, R. Oliveto, and S. Panichella, "How the apache community upgrades dependencies: an evolutionary study," pp. 1–43.
- [13] G. Bavota, G. Canfora, M. Di Penta, R. Oliveto, and S. Panichella, "The evolution of project inter-dependencies in a software ecosystem: The case of apache," in *2013 IEEE International Conference on Software Maintenance*, 2013, pp. 280–289.
- [14] D. Strein, H. Kratz, and W. Lowe, "Cross-language program analysis and refactoring," in *Source Code Analysis and Manipulation, 2006. SCAM '06. Sixth IEEE International Workshop on*, Sept 2006, pp. 207–216.
- [15] T. Arbuckle, "Measuring multi-language software evolution: A case study," in *Proceedings of the 12th International Workshop on Principles of Software Evolution and the 7th Annual ERCIM Workshop on Software Evolution*, ser. IWPSE-EVOL '11, 2011, pp. 91–95.
- [16] T. Mens and S. Demeyer, Eds., *Software Evolution*. Springer-Verlag, 2008. [Online]. Available: <http://www.springer.com/computer/swe/book/978-3-540-76439-7>
- [17] T. Mens, M. Wermelinger, S. Ducasse, S. Demeyer, R. Hirschfeld, and M. Jazayeri, "Challenges in software evolution," in *Eighth International Workshop on Principles of Software Evolution (IWPSE'05)*, 2005, pp. 13–22.
- [18] R. Marinescu, "Detection strategies: metrics-based rules for detecting design flaws," in *Software Maintenance, 2004. Proceedings. 20th IEEE International Conference on*, 2004, pp. 350–359.
- [19] M. Munro, "Product metrics for automatic identification of "bad smell" design problems in java source-code," in *Software Metrics, 2005. 11th IEEE International Symposium*, 2005, pp. 15–15.
- [20] M. Tufano, F. Palomba, G. Bavota, R. Oliveto, M. Di Penta, A. De Lucia, and D. Poshyvanyk, "When and why your code starts to smell bad," in *Software Engineering (ICSE), 2015 IEEE/ACM 37th IEEE International Conference on*, vol. 1, May 2015, pp. 403–414.
- [21] B. Fluri, M. Wuersch, M. Pinzger, and H. Gall, "Change distilling: Tree differencing for fine-grained source code change extraction," *IEEE Trans. Softw. Eng.*, vol. 33, no. 11, pp. 725–743, 2007.
- [22] M. D'Ambros, H. C. Gall, M. Lanza, and M. Pinzger, "Analysing software repositories to understand software evolution," in *Software Evolution*, 2008, pp. 37–67.
- [23] H. C. Gall, M. Jazayeri, R. Klösch, and G. Trausmuth, "Software evolution observations based on product release history," in *1997 International Conference on Software Maintenance (ICSM '97), Proceedings*, 1997, p. 160.
- [24] T. Mens, "Introduction and roadmap: History and challenges of software evolution," in *Software Evolution*. Springer Berlin Heidelberg, 2008, pp. 1–11. [Online]. Available: [http://dx.doi.org/10.1007/978-3-540-76440-3\\_{\\_}1](http://dx.doi.org/10.1007/978-3-540-76440-3_{_}1)
- [25] J. Oosterman, W. Irwin, and N. Churcher, "Evojava: A tool for measuring evolving software," in *Proceedings of the Thirty-Fourth Australasian Computer Science Conference - Volume 113*, ser. ACSC '11. Australian Computer Society, Inc., 2011, pp. 117–126.
- [26] S. Kim, K. Pan, and E. E. J. Whitehead, Jr., "Memories of bug fixes," in *Proceedings of the 14th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, ser. SIGSOFT '06/FSE-14. ACM, 2006, pp. 35–45.
- [27] *Supporting and Accelerating Reproducible Research in Software Maintenance Using TraceLab Component Library*. IEEE Computer Society, 2013.
- [28] J. Bevan, E. J. Whitehead, Jr., S. Kim, and M. Godfrey, "Facilitating software evolution research with kenyon," pp. 177–186, 2005.
- [29] R. Dyer, "Bringing ultra-large-scale software repository mining to the masses with boa," Ph.D. dissertation, Ames, IA, USA, 2013, aA13610634.
- [30] W. Shang, Z. M. Jiang, B. Adams, and A. Hassan, "Mapreduce as a general framework to support research in mining software repositories (msr)," in *6th IEEE International Working Conference on Mining Software Repositories, 2009.*, 2009, pp. 21–30.
- [31] W. Shang, B. Adams, and A. E. Hassan, "An experience report on scaling tools for mining software repositories using mapreduce," in *Proceedings of the IEEE/ACM International Conference on Automated Software Engineering*. ACM, 2010, pp. 275–284.
- [32] C. V. Alexandru and H. C. Gall, "Rapid multi-purpose, multi-commit code analysis," in *Software Engineering (ICSE), 2015 IEEE/ACM 37th IEEE International Conference on*, vol. 2, May 2015, pp. 635–638.
- [33] L. Zhang, Y. Zou, and B. Xie, "A scalable crawler framework for FLOSS data," in *Proceedings of the 5th Asia-Pacific Symposium on Internetware*, ser. Internetware '13. ACM, 2013, pp. 8:1–8:7.
- [34] G. Ghezzi and H. Gall, "Replicating mining studies with sofas," in *Mining Software Repositories (MSR), 2013 10th IEEE Working Conference on*, 2013, pp. 363–372.
- [35] R. Dyer, H. Rajan, and T. N. Nguyen, "Declarative visitors to ease fine-grained source code mining with full history on billions of ast nodes," pp. 23–32, 2013. [Online]. Available: <http://doi.acm.org/10.1145/2637365.2517226>
- [36] D. Chen, G. Doumeings, and F. Vernadat, "Architectures for enterprise integration and interoperability: Past, present and future," *Comput. Ind.*, vol. 59, no. 7, pp. 647–659, Sep. 2008.
- [37] K. Kontogiannis, P. K. Linos, and K. Wong, "Comprehension and maintenance of large-scale multi-language software applications," in *22nd IEEE International Conference on Software Maintenance (ICSM 2006), 24-27 September 2006, Philadelphia, Pennsylvania, USA*, 2006, pp. 497–500.
- [38] L. Hochstein and Y. Jiao, "The cost of the build tax in scientific software," in *Proceedings of the 5th International Symposium on Empirical Software Engineering and Measurement, ESEM 2011, Banff, AB, Canada, September 22-23, 2011*, 2011, pp. 384–387.
- [39] S. McIntosh, M. Nagappan, B. Adams, A. Mockus, and A. E. Hassan, "A large-scale empirical study of the relationship between build technology and build maintenance," *Empirical Software Engineering*, vol. 20, no. 6, pp. 1587–1633, 2014.
- [40] S. McIntosh, B. Adams, and A. E. Hassan, "The evolution of java build systems," *Empirical Software Engineering*, vol. 17, no. 4, pp. 578–608, 2011.
- [41] H. Gall, B. Fluri, and M. Pinzger, "Change analysis with evolizer and changedistiller," *Software, IEEE*, vol. 26, no. 1, pp. 26–33, 2009.
- [42] S. Kim and M. D. Ernst, "Which warnings should i fix first?" in *Proceedings of the the 6th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering*. ACM, 2007, pp. 45–54.
- [43] E. Ligu, T. Chaikalis, and A. Chatzigeorgiou, "Buco reporter: Mining software and bug repositories," in *BCI (Local)*, ser. CEUR Workshop Proceedings. CEUR-WS.org, 2013, p. 121.

- [44] C. V. C. de Magalhães, F. Q. B. da Silva, and R. E. S. Santos, "Investigations about replication of empirical studies in software engineering: Preliminary findings from a mapping study," in *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering*, ser. EASE '14, 2014, pp. 37:1–37:10.
- [45] A. Hindle, D. M. German, and R. Holt, "What do large commits tell us?: A taxonomical study of large commits," in *Proceedings of the 2008 International Working Conference on Mining Software Repositories*, ser. MSR '08, 2008, pp. 99–108.
- [46] K. Agrawal, S. Amreen, and A. Mockus, "Commit quality in five high performance computing projects," in *Proceedings of the 2015 International Workshop on Software Engineering for High Performance Computing in Science*, ser. SE4HPCS '15, 2015, pp. 24–29.
- [47] P. Stutz, A. Bernstein, and W. Cohen, "Signal/collect: Graph algorithms for the (semantic) web," in *Proceedings of the 9th International Semantic Web Conference on The Semantic Web - Volume Part I*, ser. ISWC'10. Springer-Verlag, 2010, pp. 764–780.
- [48] T. McCabe, "A complexity measure," *Software Engineering, IEEE Transactions on*, vol. SE-2, no. 4, pp. 308–320, Dec 1976.
- [49] "We analyzed 30,000 github projects – here are the top 100 libraries in java, js and ruby," <http://blog.takipi.com/we-analyzed-30000-github-projects-here-are-the-top-100-libraries-in-java-js-and-ruby/>, accessed: 2016-03-20.
- [50] "Gathering statistics on JavaScript projects since 2015," <http://stats.js.org/>, accessed: 2016-03-20.
- [51] M. Lanza, R. Marinescu, and S. Ducasse, *Object-Oriented Metrics in Practice*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2005.
- [52] R. Subramanyam and M. S. Krishnan, "Empirical analysis of ck metrics for object-oriented design complexity: Implications for software defects," *IEEE Trans. Softw. Eng.*, vol. 29, no. 4, pp. 297–310, Apr. 2003. [Online]. Available: <http://dx.doi.org/10.1109/TSE.2003.1191795>
- [53] V. R. Basili, L. C. Briand, and W. L. Melo, "A validation of object-oriented design metrics as quality indicators," *IEEE Trans. Softw. Eng.*, vol. 22, no. 10, pp. 751–761, Oct. 1996. [Online]. Available: <http://dx.doi.org/10.1109/32.544352>
- [54] "infusion by intooitus s.r.l." <http://www.intooitus.com/products/infusion>, accessed: 2014-03-30.
- [55] G. Ghezzi and H. Gall, "Sofas: A lightweight architecture for software analysis as a service," in *Software Architecture (WICSA), 2011 9th Working IEEE/IFIP Conference on*, June 2011, pp. 93–102.
- [56] M. Kim and D. Notkin, "Program element matching for multiversion program analyses," in *Proceedings of the 2006 International Workshop on Mining Software Repositories*, ser. MSR '06. New York, NY, USA: ACM, 2006, pp. 58–64. [Online]. Available: <http://doi.acm.org/10.1145/1137983.1137999>
- [57] W. Yang, S. Horwitz, and T. Reps, "A program integration algorithm that accommodates semantics-preserving transformations," *ACM Trans. Softw. Eng. Methodol.*, vol. 1, no. 3, pp. 310–354, Jul. 1992. [Online]. Available: <http://doi.acm.org/10.1145/131736.131756>
- [58] M. Fischer, M. Pinzger, and H. Gall, "Populating a release history database from version control and bug tracking systems," in *Software Maintenance, 2003. ICSM 2003. Proceedings. International Conference on*, 2003, pp. 23–32.
- [59] S. S. Chawathe, A. Rajaraman, H. Garcia-Molina, and J. Widom, "Change detection in hierarchically structured information," in *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD '96, 1996, pp. 493–504.
- [60] T. Zimmermann, A. Zeller, P. Weissgerber, and S. Diehl, "Mining version histories to guide software changes," *Software Engineering, IEEE Transactions on*, vol. 31, no. 6, pp. 429–445, 2005.
- [61] W. Le and S. D. Pattison, "Patch verification via multiversion interprocedural control flow graphs," in *Proceedings of the 36th International Conference on Software Engineering*, ser. ICSE 2014. New York, NY, USA: ACM, 2014, pp. 1047–1058. [Online]. Available: <http://doi.acm.org/10.1145/2568225.2568304>
- [62] C. Kästner, P. G. Giarrusso, T. Rendel, S. Erdweg, K. Ostermann, and T. Berger, "Variability-aware parsing in the presence of lexical macros and conditional compilation," in *Proceedings of the 2011 ACM International Conference on Object Oriented Programming Systems Languages and Applications*, ser. OOPSLA '11. New York, NY, USA: ACM, 2011, pp. 805–824. [Online]. Available: <http://doi.acm.org/10.1145/2048066.2048128>
- [63] S. Tichelaar, S. Ducasse, S. Demeyer, and O. Nierstrasz, "A meta-model for language-independent refactoring," in *Principles of Software Evolution, 2000. Proceedings. International Symposium on*, 2000, pp. 154–164.
- [64] "OMG: KDM," <http://www.omg.org/spec/KDM/1.3/>, accessed: 2016-10-06.
- [65] "OMG: ASTM," <http://www.omg.org/spec/ASTM/1.0/>, accessed: 2016-10-06.
- [66] H. M. Kienle and H. A. Müller, "Rigi—an environment for software reverse engineering, exploration, visualization, and redocumentation," *Science of Computer Programming*, vol. 75, no. 4, pp. 247 – 263, 2010, experimental Software and Toolkits (EST 3): A special issue of the Workshop on Academic Software Development Tools and Techniques (WASDeTT 2008). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S016764230900149X>
- [67] A. Winter, B. Kullbach, and V. Riediger, "An overview of the gxl graph exchange language," in *Revised Lectures on Software Visualization, International Seminar*. London, UK, UK: Springer-Verlag, 2002, pp. 324–336. [Online]. Available: <http://dl.acm.org/citation.cfm?id=647382.724795>
- [68] A. Izmaylova, P. Klint, A. Shahi, and J. J. Vinju, "M3: an open model for measuring code artifacts," *CoRR*, vol. abs/1312.1188, 2013. [Online]. Available: <http://arxiv.org/abs/1312.1188>
- [69] M. Hills, P. Klint, and J. J. Vinju, *Program Analysis Scenarios in Rascal*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 10–30. [Online]. Available: [http://dx.doi.org/10.1007/978-3-642-34005-5\\_2](http://dx.doi.org/10.1007/978-3-642-34005-5_2)
- [70] M. Lanza, S. Ducasse, H. Gall, and M. Pinzger, "Codecrawler - an information visualization tool for program comprehension," in *Software Engineering, 2005. ICSE 2005. Proceedings. 27th International Conference on*, May 2005, pp. 672–673.
- [71] G. Rakić, Z. Budimac, and M. Savić, "Language independent framework for static code analysis," in *Proceedings of the 6th Balkan Conference in Informatics*, ser. BCI '13. New York, NY, USA: ACM, 2013, pp. 236–243. [Online]. Available: <http://doi.acm.org/10.1145/2490257.2490273>