

# Hot Clones, Maintaining a Link Between Software Clones Across Repositories

Nikolaus E. Schwarz  
Software Composition Group  
University of Bern, Switzerland

Erwann Wernli  
Software Composition Group  
University of Bern, Switzerland

Adrian Kuhn  
Software Composition Group  
University of Bern, Switzerland

## ABSTRACT

Code duplication is common in current programming-practice: programmers search for snippets of code, incorporate them into their projects and then modify them to their needs. In today's practice, no automated scheme is in place to inform both parties of any distant changes of the code. As code snippets continues to evolve both on the side of the user and on the side of the author, both may wish to benefit from remote bug fixes or refinements—authors may be interested in the actual usage of their code snippets, and researchers could gather information on clone usage. We propose maintaining a link between software clones across repositories and outline how the links can be created and maintained.

## 1. INTRODUCTION

Since the rise of internet-scale code search engines, searching for reusable source code has quickly become a fundamental activity for developers (e.g. [4, 10]). Developers use search engines to find and reuse software. The promise of search-driven development is that developers will save time and resources by using search results. However, there are perils: the current practice of manually integrating code search results in a local code base leads to proliferation of untracked code clones. As an effect, in the authors' experience, bugs fixed in one clone typically do not traverse their new environment anymore, and the same holds true for extensions and code cleanups.

Even if they appear in the same project, software clones often cannot be eliminated [7]. But oversights in applying changes to clones consistently may introduce bugs into the system [3]. Therefore, tools have been proposed to maintain links between software clones [5, 3, 11], but they fail to link clones that are beyond project boundaries. Codebook by Begel et al. [1], is a social network in which people can befriend both other people and their work artifacts. Codebook is intended to maintain links between clones, it is however unclear how these links come into being. Begel et al. only

vaguely propose how that should be done: edges are to be added between a definition and its likely clones.

In this paper we propose a scheme to initially create and then maintain that link. A code search engine assists the developer by integrating its results into the source code. The IDE then remembers the origin of the code snippet and informs the repository that a clone was created, thus creating a link between original and copy. We will refer to clones creates in this way as hot clones. Whenever a hot clone changes, the linked clones' developers are informed and offered the option to update their instance. Also, whenever a method is inspected, its clones can be inspected too, providing valuable information. The connections between clone instances are thus proactive and bidirectional.

In the terminology of Koschke [8], we provide *compensative clone management*, i.e. we limit the negative impact of existing clones, but we also give developers benefits from the software duplication introduced by clones by providing developers with information on how their code is used and modified.

## 2. GOALS

The idea of hot clones scratches two itches. The first is to let developers benefit from code cloning, and the second is to provide researchers with more information on how cloning is used.

We believe that hot clones can ease backporting changes that occur in a linked clone. We believe that during development, hot clones will provide important feedback to developers. Contrasting ones own code with modified clones will give hints to bugs in related code, usage patterns, and plain examples of usage.

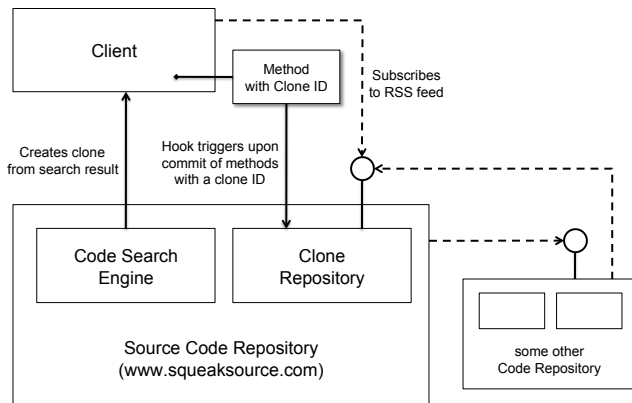
There has been research on the nature of clones within a single software project (e.g. [6]), but the evolution of code snippets copied from searches in software projects has not been studied. A prototypical implementation of hot clones would provide this opportunity. Being able to track the further evolution of code snippets after they are copied out of a search engine may give us great insight into the evolution of code, beyond the classes provided by Kapser and Godfrey [6]. If used by only a few developers, hot clones can provide insights from both a larger set of data than before, and from a wider range of uses. We plan to provide a prototype of hot clones.

## 3. SCOPE

In this section, we will discuss which tasks hot clones can assume and how this suits our goals. Software clones are

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IWSC '10, May 8 2010, Cape Town, South Africa  
Copyright 2010 ACM 978-1-60558-980-0/10/05 ...\$10.00.



**Figure 1: Proposed architecture of the “hot clone” prototype. It extends Monticello, a distributed version control system for Smalltalk, with two components: a code search engine and a clone repository.**

not typically exact copies of each other, but rather they start as exact copies and then evolve in different ways, for a number of reasons in accordance with the specific needs in their respective environments [7, 6].

To cope with the changes that search-driven development introduces between clone instances, the connections between clone instances should be aware of the semantics of their differences. To retain an active connection between instances of a hot clone, it is paramount that the clones are not only aware of the lexical but also the semantic changes between their instances [9]. For example, a renamed method call in two different instances may be due to adherence to coding conventions, but may also be due to a deliberate change in one of the instances. So when owners of clones are informed about an incoming change, the semantics of the change should be taken into account. We propose capturing each change in a Changebox [2] annotated with the intent of a change.

Whenever an incoming change is presented to the developer, the semantic change history of all involved clones has to be taken into account. To consider an example, if the code search engine initially renamed all variable names from camel case to underscore (in order to fit the search result to the local naming convention) when the clone was created, this adaption should be applied to any incoming change as well.

Search results need to be adapted in order to fit into the target code. Integrating code search results into a local code base is, by its nature, unanticipated code reuse and will thus naturally require adaptations to the external search results such that they fit into the local code base. Search-driven development refers to this as *suitability* of search results [10].

#### 4. TECHNICAL FEASIBILITY

In this section, we discuss how hot clones can be recorded upon creation and how they should be presented to the user. A key challenge is tracking code snippets while they evolve to become evermore distinct. The versioning system can be informed of the original cloning and henceforth try to keep track of the clone automatically, even if it is not marked in the source code.

Figure 1 illustrates a sketch of the proposed prototype.

We plan to extend MONTICELLO, a distributed version control system for Smalltalk code, with two components: a code search engine and a clone repository.

The *code search engine* provides access to the full content of the <http://www.squeaksource.com> installation of Monticello, some 7.5 GB of Smalltalk source code<sup>1</sup>. On the client side, the search engine allows the developers to create semantically transformed clones that suit the local code base. Each method that belongs to the clone is tagged with the unique identifier of the clone.

The *clone repository* adds a commit hook to the MONTICELLO version control system that is triggered whenever someone commits a new version of a method that has (or previously had) a clone id. The hook adds the client to the list of linked clones (if not already present) and informs all linked clones of the update through an RSS feed. Clients that are subscribed to the RSS feed can present their developers with an option to update their instances of the clone in question. Also, the Pharo IDE can be extended to show all clones of a method from the context menu of that method.

#### Acknowledgments.

We gratefully acknowledge the financial support of the Swiss National Science Foundation for the project “Bringing Models Closer to Code” (SNF Project No. 200020-121594, Oct. 2008 - Sept. 2010). We thank Edouard Tavinor for his ideas and corrections.

#### 5. REFERENCES

- [1] A. Begel and R. DeLine. Codebook: Social networking over code. In *ICSE Companion*, pages 263–266, 2009.
- [2] M. Denker, T. Girba, A. Lienhard, O. Nierstrasz, L. Renggli, and P. Zumkehr. Encapsulating and exploiting change with changeboxes. In *ICDL '07*, pages 25–49, New York, NY, USA, 2007. ACM.
- [3] E. D. Ekoko and M. P. Robillard. Clonetracker: tool support for code clone management. In *ICSE '08*, pages 843–846, New York, NY, USA, 2008. ACM.
- [4] R. Hoffmann, J. Fogarty, and D. S. Weld. Assieme: finding and leveraging implicit references in a web search interface for programmers. In *UIST '07*, pages 13–22, New York, NY, USA, 2007. ACM.
- [5] D. Hou, P. Jablonski, and F. Jacob. Cnp: Towards an environment for the proactive management of copy-and-paste programming. In *2009 IEEE 17th ICPC*, pages 238–242. IEEE, May 2009.
- [6] C. Kapser and M. W. Godfrey. “cloning considered harmful” considered harmful. *WCRE '06*, 0:19–28, 2006.
- [7] M. Kim, V. Sazawal, D. Notkin, and G. Murphy. An empirical study of code clone genealogies. *SIGSOFT Softw. Eng. Notes*, 30(5):187–196, September 2005.
- [8] R. Koschke. Identifying and removing software clones. In *Software Evolution*, chapter 2, pages 15–36. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [9] G. Muller, Y. Padiouleau, J. L. Lawall, and R. R. Hansen. Semantic patches considered helpful. *SIGOPS Oper. Syst. Rev.*, 40(3):90–92, 2006.
- [10] S. P. Reiss. Semantics-based code search. In *ICSE Companion*, volume 0, pages 243–253, Los Alamitos, CA, USA, 2009. IEEE Computer Society.
- [11] M. Toomim, A. Begel, and S. L. Graham. Managing duplicated code with linked editing. In *VLHCC '04*, pages 173–180, Washington, DC, USA, 2004. IEEE.

<sup>1</sup>From a tweet by Lukas Renggli, project lead of Squeak-source: “<http://www.squeaksource.com> hosts 7.5 GB of Monticello versions,” posted on Jan 7, 2010, <http://twitter.com/renggli/status/7473119028>