

Specifying and Verifying the Correctness of Dynamic Software Updates

Christopher M. Hayden¹, Stephen Magill¹, Michael Hicks¹,
Nate Foster², and Jeffrey S. Foster¹

¹ Computer Science Department, University of Maryland, College Park
{hayden,smagill,mwh,jfoster}@cs.umd.edu

² Computer Science Department, Cornell University
jnfoster@cs.umd.edu

Abstract. Dynamic software updating (DSU) systems allow running programs to be patched on-the-fly to add features or fix bugs. While dynamic updates can be tricky to write, techniques for establishing their correctness have received little attention. In this paper, we present the first methodology for automatically verifying the correctness of dynamic updates. Programmers express the desired properties of an updated execution using *client-oriented specifications* (CO-specs), which can describe a wide range of client-visible behaviors. We verify CO-specs automatically by using off-the-shelf tools to analyze a *merged* program, which is a combination of the old and new versions of a program. We formalize the merging transformation and prove it correct. We have implemented a program merger for C, and applied it to updates for the Redis key-value store and several synthetic programs. Using Thor, a verification tool, we could verify many of the synthetic programs; using Otter, a symbolic executor, we could analyze every program, often in less than a minute. Both tools were able to detect faulty patches and incurred only a factor-of-four slowdown, on average, compared to single version programs.

1 Introduction

Dynamic software updating (DSU) systems allow programs to be patched on-the-fly, to add features or fix bugs without incurring downtime. DSU systems were originally developed for a few limited domains such as telecommunications networks, financial transaction processors, and the like, but are now becoming pervasive. Ksplice, recently acquired by Oracle, supports applying Linux kernel security patches dynamically [16]. The Erlang language, which provides built-in support for dynamic updates, is gaining in popularity for building server programs [2]. Many web applications employ DSU techniques to provide 24/7 service to a global audience—for these systems, there is no single time of day when taking down the service to perform upgrades is acceptable.

Given the increasing need for DSU, a natural question is: How can developers ensure a dynamically updated program will behave correctly? Today, developers need to reason manually about all the pieces of an updating program: the old

program version, the new program version, and code that transforms the state of the (old) running version into the form expected by the new version. Moreover, they need to repeat this reasoning process for each allowable “update point” during execution. In our experience this is a tricky proposition in which it is all too easy to make mistakes. Despite such difficulties, most DSU systems do not address the issue of correctness, or they focus exclusively on generic safety properties, such as type safety, that rule out obviously wrong behavior [7, 23–25] but are insufficient for establishing correctness [12].

This paper presents a methodology for verifying the correctness of dynamic updates. Rather than propose a new verification algorithm that accounts for the semantics of updating, we develop a novel program transformation that produces a program suitable for verification with off-the-shelf tools. Our transformation *merges* an old program and an update into a program that simulates running the program and applying the update at any allowable point. We formalize our transformation and prove that it is correct (Section 3).

We are particularly interested in using our transformation to prove execution properties from clients’ points of view, to show that a dynamic update does not disrupt active sessions. For example, suppose we wish to update a key-value store such as Redis [21] so that it uses a different internal data structure. To verify this update’s transformation code, we could prove that values inserted into the store by the client are still present after it is dynamically updated. We call such specifications *client-oriented specifications* (or *CO-specs* for short).

We have identified three categories of CO-specs that capture most properties of interest: *backward-compatible* CO-specs describe properties that are identical in the old and new versions; *post-update* CO-specs describe properties that hold after new features are added or bugs are fixed by an update; and *conformable* CO-specs describe properties that are identical in the old and new versions, modulo uniform changes to the external interface. CO-specs in these categories can often be mechanically constructed from CO-specs written for either the old or new program alone. Thus, if a programmer is inclined to verify each program version using CO-specs, there is little additional work to verify a dynamic update between the two. Nevertheless, some interesting and subtle properties lie outside these categories, so our framework also allows arbitrary properties to be expressed (Section 2).

We have implemented our merging transformation for C programs and used it in combination with two existing tools to verify properties of several dynamic updates (Section 4). We chose the symbolic executor Otter [22] and the verification tool Thor [17] as they represent two ends of the design space: symbolic execution is easy to use and scales reasonably well but is incomplete, while verification scales less well but provides greater assurance. We wrote two synthetic benchmarks, a key-value store and a multiset implementation, and designed dynamic patches for them based on realistic changes (e.g., one change was inspired by an update to the storage server Cassandra [5]). We also wrote dynamic patches for six releases of Redis [21], a popular, open-source key-value store. We used the Redis code as is, and wrote the state transformation code ourselves.

We checked all the benchmark programs with Otter and verified several properties of the synthetic updates using Thor. Both tools successfully uncovered bugs that were intentionally and unintentionally introduced in the state transformation code. The running time for verification of merged programs was roughly four times slower than single-version checking. This slowdown was due to the additional branching introduced by update points and the need to analyze the state transformer code. As tools become faster and more effective, our approach will scale with them. In summary, this paper makes three main contributions:

- It presents the first automated technique for verifying the behavioral correctness of dynamic updates.
- It proposes *client-oriented specifications* as a means to specify general update correctness properties.
- It shows the effectiveness of merging-based verification on practical examples, including Redis [21], a widely deployed server program.

2 Defining dynamic software update correctness

Before we can set out verifying DSU correctness, we have to decide what correctness is. In this section, we first review previously proposed notions of correctness and argue why they are insufficient for our purposes. Then we propose *client-oriented specifications* (CO-specs) as a means of specifying correctness properties, and argue that this notion overcomes limitations of prior notions. We also describe a simple refactoring that allows CO-specs to be used to verify client-server programs that communicate over a network.

2.1 Prior work on update correctness

Kramer and Magee [15] proposed that updates are correct if they are *observationally equivalent*—i.e., if the updated program preserves all observable behaviors of the old program. Bloom and Day [3] observed that, while intuitive, this is too restrictive: an update may fix bugs or add new features.

To address the limitations of strict observational equivalence, Gupta et al. [9] proposed *reachability*. This condition classifies an update as correct if, after the update is applied, the program eventually *reaches* some state of the new program. Reachability thus admits bugfixes, where the new state consists of the corrected code and data, as well as feature additions, where the new state is the old data plus the new code and any new data. Unfortunately, reachability is both too permissive and too restrictive, as shown by the following example. Version 1.1.2 of the `vsftpd` FTP server introduced a feature that limits the number of connections from a single host. If we update a running `vsftpd` server, we would expect it to preserve any active connections. But doing so violates reachability. If the number of connections from a particular host exceeds the limit and these connections remain open indefinitely, the server will never enter a reachable state of the new program. On the other hand, reachability would allow an update that

terminates all existing connections. This is almost certainly not what we want—if we were willing to drop existing connections we could just restart the server!

We believe that the flaw in all of these approaches is that they attempt to define correctness in a completely general way. We think it makes more sense for programmers to specify the behavior they expect as a collection of properties. Some properties will apply to multiple versions of the program while other properties will change as the program evolves. Because the goal of a dynamic update is to preserve active processing and state, the properties should express the expected continuity that a dynamic update is meant to provide to active clients. We therefore introduce *client-oriented specifications* (CO-specs) to specify update properties that satisfy these requirements.

2.2 Client-oriented specifications

We can think of a CO-spec as a kind of client program that opens connections, sends messages, and asserts that the output received is correct. CO-specs resemble tests, but certain elements of the test code are left abstract for generality (cf. Figure 1). For example, consider again reasoning about updates to a key-value store such as Redis. A CO-spec might model a client that inserts a key-value pair into the store and then looks up the key, checking that it maps to the correct value (even if a dynamic update has occurred in the meantime). We can make such a CO-spec general by leaving certain elements like the particular keys or values used unconstrained. Similarly, we can allow arbitrary actions to be interleaved between the insert and lookup. Such specifications capture essentially arbitrary client interactions with the server.

Our goal is to use our program transformation, defined in Section 3, to produce a *merged* program that we can verify using off-the-shelf tools. But existing tools only verify single programs in isolation, so we cannot literally write CO-specs as client programs that communicate with a server being updated. To verify a CO-spec in a real client-server program we replace the server’s main function the CO-spec and call the relevant server functions directly. In doing so, we are checking the server’s core functionality, but not its main loop or any networking code. For example, suppose our key-value store implements functions `get` and `set` to read and write mappings from the store, and the server’s main loop would normally dispatch to these functions. CO-specs would call the functions directly as shown in Figure 1. Here, `?` denotes a non-deterministically chosen (integer) value, and `assume` and `assert` have their standard semantics. If updates are permitted while executing either `get` or `set`, verifying Figure 1(b) will establish that the assertions at the end of the specification hold no matter when the update takes place.

In our experience writing CO-specs for updates, we have found that they often fall into one of the following categories:

- *Backward-compatible CO-specs* describe behaviors that are unaffected by an update. For the data structure-changing update to Redis mentioned earlier, the CO-spec in Figure 1(b) would check that existing mappings are preserved.

<pre> 1 int get(int k, int *v); 2 void set(int k, int v); 3 4 void arbitrary (int k1) { 5 int k2 = ?, v = ?; 6 if (k1 == k2 ?) 7 get(k2,&v); 8 else set (k2,v); 9 }</pre>	<pre> 10 void back_compat_spec() { 11 int k = ?, v_in = ?; 12 int v_out, found; 13 set(k, v_in); 14 while(?) arbitrary (k); 15 found = get(k,&v_out); 16 assert (found && 17 v_out == v_in); 18 }</pre>	<pre> 19 void post_update_spec() { 20 int k = ?; 21 int v_out, found; 22 while(?) arbitrary (?); 23 assume(is_updated); 24 delete (k); 25 found = get(k,&v_out); 26 assert (!found); 27 }</pre>
(a) interface, helper	(b) backward-compat. spec	(c) post-update spec

Fig. 1. Sample C specifications for key-value store.

- *Post-update CO-specs* describe behavior specific to the new program version. For example, suppose we added a `delete` feature to the key-value store. Then the CO-spec in Figure 1(c) verifies that, after the update, the feature is working properly. The CO-spec employs the flag `is_updated`, which is true after an update has taken place, to ensure that we are testing the new or changed functionality after the update. We discuss the semantics of this flag in Section 3.
- *Conformable CO-specs* describe updates that change interfaces, but preserve core functionality. For example, suppose we added namespaces to our key-value store, so that `get` and `set` take an additional namespace argument. The state transformation code would map existing entries to a default namespace. A conformable CO-spec could check that mappings inserted prior to the update are present in the default namespace afterward; in essence, the CO-spec would associate old-version calls with new-version calls at the default namespace. (Further details are given in our technical report [11].)

These categories encompass prior notions of correctness. Backward compatible specifications capture the spirit of Kramer and Magee’s condition, but apply to individual, not all, behaviors. The combination of backward-compatible and post-update specifications capture Bloom and Day’s notions of “future-only implementations” and “invisible extensions”—parts of a program whose semantics change but not in a way that affects existing clients [3]. The combination of backward-compatible and conformable specifications match ideas proposed by Ajmani et al. [1], who studied dynamic updates for distributed systems and proposed mechanisms to maintain continuity for clients of a particular version.

CO-specs can also be used to express the constraints intended by Gupta’s *reachability* while side-stepping the problem that reachability can leave behavior under-constrained. For example, for the `vsftpd` update mentioned above, the programmer can directly write a CO-spec that expresses what should happen to existing client connections, e.g., whether all, some, or none should be preserved. This does not fall into one of the categories above, demonstrating the utility of a full specification language over “one size fits all” notions of update correctness.

Another feature of CO-specs in these categories is that they can be mechanically constructed from CO-specs that are written for a single version. Thus, if a programmer was inclined to verify the correctness of each version of his program using CO-specs, the additional work to verify a dynamic update is not much greater. For details, see our technical report [11].

3 Verification via program merging

We verify CO-specs by *merging* an existing program version with its update, so that the semantics of the merged program is equivalent to the updating program. This section formalizes a semantics for dynamic updates to single-threaded programs, then defines the merging transformation and proves it correct with respect to the semantics. Many server programs for which dynamic updating is useful are single-threaded [13, 19, 12]. However, an important next step for this work would be to adapt it to support updates to multi-threaded (and distributed) programs.

3.1 Syntax

The top of Figure 2 defines the syntax of a simple programming language supporting dynamic updates. It is based on the Proteus dynamic update calculus [23], and closely models the semantics of common DSU systems, including Ginseng [19] (which is the foundation of our implementation), Ksplice [16], Jvolve [24], K42 [14], DLpop [13], Dynamic ML [25] and Bracha’s DSU system [4].

A *program* p is a mapping from function names g to functions $\lambda x.e$. A function body e is defined by a mostly standard core language with a few extensions for updating. Our language contains a construct `update`, which indicates a position where a dynamic update may take effect. To support writing specifications, the language includes an expression `?`, which represents a random integer, and expressions `assume v` , `assert v` , and `running p` , all of whose semantics are discussed below. Expressions are in administrative normal (A-normal) form [8] to keep the semantics simple—e.g., instead of $e_1 + e_2$, we write `let $x = e_1$ in let $y = e_2$ in $x + y$` . We write $e_1; e_2$ as shorthand for `let $x = e_1$ in e_2` , where x is fresh for e_2 .

3.2 Semantics

The semantics, given in the latter half of Figure 2, is written as a series of small-step rewriting rules between *configurations* of the form $\langle p; \sigma; e \rangle$, which contain the program p , its current heap σ , and the current expression e being evaluated. A heap is a partial function from locations l to values v , and a location l is either a (dynamically allocated) address a or a (static) global name g . Note that while the language does not include closures, global names g are values, and so the language does support C-style function pointers.¹

¹ Variables names x are values so that we can use a simple grammar to enforce A-normal form. The downside is that syntactically well-formed programs could pass

<p><i>Prog.</i> $p ::= p, (g, \lambda x.e) \mid \cdot$</p> <p><i>Exprs.</i> $e ::= v \mid v_1 \text{ op } v_2 \mid v_1(v_2) \mid ? \mid !v \mid \text{ref } v \mid$ $v_1 := v_2 \mid \text{if } v \ e_1 \ e_2 \mid \text{update} \mid$ $\text{let } x = e_1 \text{ in } e_2 \mid \text{assume } v \mid$ $\text{while } e_1 \text{ do } e_2 \mid \text{assert } v \mid$ $\text{running } p \mid \text{error}$</p> <p><i>Values</i> $v ::= x \mid l \mid i \mid (v_1, v_2) \mid ()$</p> <p><i>Locs.</i> $l ::= a \mid g$</p>	<p><i>Variables</i> x, y, z</p> <p><i>Globals</i> f, g</p> <p><i>Operators</i> op</p> <p><i>Integers</i> i, j</p> <p><i>Addresses</i> a</p> <p><i>Heaps</i> $\sigma \in \text{Locs} \rightarrow \text{Values}$</p> <p><i>Patch</i> $\pi ::= (p, e)$</p> <p><i>Labels</i> $\nu ::= \pi \mid \epsilon$</p>
<p>$\langle p; \sigma; v_1 \text{ op } v_2 \rangle \rightsquigarrow \langle p; \sigma; v' \rangle$</p> <p>$\langle p; \sigma; \text{ref } v \rangle \rightsquigarrow \langle p; \sigma[a \mapsto v]; a \rangle$</p> <p>$\langle p; \sigma; !l \rangle \rightsquigarrow \langle p; \sigma; v \rangle$</p> <p>$\langle p; \sigma; a := v \rangle \rightsquigarrow \langle p; \sigma[a \mapsto v]; v \rangle$</p> <p>$\langle p; \sigma; g := v \rangle \rightsquigarrow \langle p; \sigma[g \mapsto v]; v \rangle$</p> <p>$\langle p; \sigma; ? \rangle \rightsquigarrow \langle p; \sigma; i \rangle$</p> <p>$\langle p; \sigma; \text{let } x = v \text{ in } e \rangle \rightsquigarrow \langle p; \sigma; e[v/x] \rangle$</p> <p>$\langle p; \sigma; f(v) \rangle \rightsquigarrow \langle p; \sigma; e[v/x] \rangle$</p> <p>$\langle p; \sigma; \text{if } 0 \ e_1 \ e_2 \rangle \rightsquigarrow \langle p; \sigma; e_2 \rangle$</p> <p>$\langle p; \sigma; \text{if } v \ e_1 \ e_2 \rangle \rightsquigarrow \langle p; \sigma; e_1 \rangle$</p> <p>$\langle p; \sigma; \text{while } e_1 \text{ do } e_2 \rangle \rightsquigarrow \langle p; \sigma; \text{let } x = e_1 \text{ in}$ $\text{if } x \ (e_2; \text{while } e_1 \text{ do } e_2) \ 0 \rangle$</p> <p>$\langle p; \sigma; \text{update} \rangle \rightsquigarrow \langle p; \sigma; 0 \rangle$</p> <p>$\langle p; \sigma; \text{update} \rangle \rightsquigarrow^{\pi} \langle p_{\pi}; \sigma; (e_{\pi}; 1) \rangle$</p> <p>$\langle p; \sigma; \text{running } p \rangle \rightsquigarrow \langle p; \sigma; 1 \rangle$</p> <p>$\langle p; \sigma; \text{running } p' \rangle \rightsquigarrow \langle p; \sigma; 0 \rangle$</p> <p>$\langle p; \sigma; \text{assume } v \rangle \rightsquigarrow \langle p; \sigma; v \rangle$</p> <p>$\langle p; \sigma; \text{assert } v \rangle \rightsquigarrow \langle p; \sigma; v \rangle$</p> <p>$\langle p; \sigma; \text{assert } 0 \rangle \rightsquigarrow \langle p; \sigma; \text{error} \rangle$</p> <p>$\langle p; \sigma; \text{let } x = \text{error in } e \rangle \rightsquigarrow \langle p; \sigma; \text{error} \rangle$</p>	<p>$v' = \llbracket op \rrbracket(v_1, v_2)$</p> <p>$a \notin \text{dom}(\sigma)$</p> <p>$\sigma(l) = v$ and $l \notin \text{dom}(p)$</p> <p>$a \in \text{dom}(\sigma)$</p> <p>$g \notin \text{dom}(p)$</p> <p>for some i</p> <p>$p(f) = \lambda x.e$</p> <p>$v \neq 0$</p> <p>$x \notin \text{fv}(e_1, e_2)$</p> <p>$\pi = (p_{\pi}, e_{\pi})$</p> <p>$p' \neq p$</p> <p>$v \neq 0$</p> <p>$v \neq 0$</p>
$\frac{\langle p; \sigma; e_1 \rangle \rightsquigarrow \langle p'; \sigma'; e'_1 \rangle}{\langle p; \sigma; \text{let } x = e_1 \text{ in } e_2 \rangle \rightsquigarrow \langle p'; \sigma'; \text{let } x = e'_1 \text{ in } e_2 \rangle}$	

Fig. 2. Syntax and semantics.

Most of the operational semantics rules are straightforward. We write $e[x/v]$ for the capture-avoiding substitution of x with v in e . We assume that the semantics of primitive operations op is defined by some mathematical function $\llbracket op \rrbracket$; e.g., $\llbracket + \rrbracket$ is the integer addition function. Loops are rewritten to conditionals, where in both cases a non-zero guard is treated as true and zero is treated as false. Addresses a for dynamically allocated memory must be allocated prior to assigning to them, whereas a global variable g is created when it is first assigned to. This semantics allows state transformation functions, described below, to define new global variables that are accessible to an updated program.

around unbound variables and store them in the heap. The ability to express such programs is immaterial to our modeling of DSU, and could be easily ruled out with a simple static type system.

The `update` command identifies a position in the program at which a dynamic update may take place. Semantically, `update` non-deterministically transitions either to 0, indicating that an update did not occur, or to 1 (eventually), indicating that a dynamic update was available and was applied.² In the case where an update occurs, the transition arrow is labeled with the patch π ; all other (unadorned) transitions implicitly have label ϵ . A patch π is a pair (p_π, e_π) consisting of the new program code (including unmodified functions) p_π and an expression e_π that transforms the current heap as necessary, e.g., to update an existing data structure or add a new one for compatibility with the new program p_π . In practice, e_π will be a call to a function defined in p_π . The transformer expression e_π is placed in redex position and is evaluated immediately; to avoid capture, non-global variables may not appear free in e_π . Notice that an update that changes function f has no effect on running instances of f since evaluation of their code began prior to the update taking place.

The placement of the `update` command has a strong influence on the semantics of updates. Placing `update` pervasively throughout the code essentially models asynchronous updates. Or, as prior work recommends [15, 1, 19, 12], we could insert `update` selectively, e.g., at the end of each request-handling function or within the request-handling loop, to make an update easier to reason about.

The constructs `running p`, `assume v`, and `assert v` allow us to write specifications. The expression `running p` returns 1 if p is the program currently running and 0 otherwise; i.e., we encode a program version as the program text itself. (In Figure 1(c) the expression `is.updated` is equivalent to `running p` where p is the new program version.) The expression `assert v` returns v if it is non-zero, and error otherwise, which by the rule for `let` propagates to the top level. Finally, the expression `assume v` returns v if v is non-zero, and otherwise is stuck.

3.3 Program merging transformation

We now present our program merging transformation, which takes an old program configuration $\langle p, \sigma, e \rangle$ and a patch π and yields a single *merged program* configuration, written $\langle p, \sigma, e \rangle \triangleright \pi$. We present the transformation formally and then prove that the merged program is equivalent to the original program with the patch applied dynamically. While we focus on merging a program with a single update, the merging strategy can be readily generalized to multiple updates (see our technical report [11] for details).

The definition of $\langle p, \sigma, e \rangle \triangleright \pi$ is given in Figure 3(e). It makes use of functions $\llbracket \cdot \rrbracket$ and $\{\! \{ \cdot \} \!\}$, defined in Figure 3(a)–(d). We present the interesting cases; the remaining cases are translated structurally in the natural way. For simplicity, the transformation assumes the updated program p_π does not delete any functions in p . Deletion of function f can be modeled by a new version of f with the same signature as the original and the body `assert(0)`.

² In practice, `update` would be implemented by having the run-time system check for an update and apply it if one is available [13].

$\llbracket p', (g, \lambda y. e) \rrbracket^{p, \pi} \triangleq$ $\llbracket p' \rrbracket^{p, \pi}, (g, \lambda y. \llbracket e \rrbracket^{p, \pi}),$ $(g_{ptr}, \lambda y. \text{let } z = \text{isupd}() \text{ in if } z \text{ } g'(y) \text{ } g(y))$ $\llbracket \cdot \rrbracket^{p, \pi} \triangleq (\cdot, (\text{isupd}, \lambda y. \text{let } z = !uflag \text{ in } z > 0))$ <p style="text-align: center;">(a) Old version programs</p>	$\{\!\{ p', (g, \lambda y. e) \}\!\}^p \triangleq$ $\{\!\{ p' \}\!\}^p, (g', \lambda y. \{\!\{ e \}\!\}^p)$ $\{\!\{ \cdot \}\!\}^p \triangleq \cdot$ <p style="text-align: center;">(b) New version programs</p>
$\llbracket g \rrbracket^{p, \pi} \triangleq$ $\begin{cases} g_{ptr} & \text{if } p(g) = \lambda x. e \\ g & \text{otherwise} \end{cases}$ $\llbracket \text{running } p'' \rrbracket^{p, (p_\pi, e_\pi)} \triangleq$ $\begin{cases} \text{let } z = \text{isupd}() \text{ in } z = 0 & \text{if } p = p'' \\ \text{isupd}() & \text{if } p_\pi = p'' \\ 0 & \text{otherwise} \end{cases}$ $\llbracket \text{update} \rrbracket^{p, (p_\pi, e_\pi)} \triangleq$ $\text{let } z = \text{isupd}() \text{ in}$ $\text{if } z = 0 \text{ (} uflag := ?;$ $\text{let } z = \text{isupd}() \text{ in if } z \text{ (}\{\!\{ e \}\!\}^{p, \pi}; 1) \text{ 0)}$ <p style="text-align: center;">(c) Old version expressions</p>	$\{\!\{ g \}\!\}^p \triangleq$ $\begin{cases} g' & \text{if } p(g) = \lambda x. e \\ g & \text{otherwise} \end{cases}$ $\{\!\{ \text{running } p'' \}\!\}^p \triangleq$ $\begin{cases} 1 & \text{if } p = p'' \\ 0 & \text{otherwise} \end{cases}$ $\{\!\{ \text{update} \}\!\}^p \triangleq 0$ <p style="text-align: center;">(d) New version expressions</p>
$\langle p; \sigma; e \rangle \triangleright \pi \triangleq \langle \bar{p}, \bar{\sigma}[uflag \mapsto i], \bar{e} \rangle$ $\text{where } (p_\pi, e_\pi) = \pi \quad \bar{p} = \{\!\{ p_\pi \}\!\}^{p, \pi}, \llbracket p \rrbracket^{p, \pi} \quad \bar{e} = \llbracket e \rrbracket^{p, \pi}$ $i \leq 0 \quad \bar{\sigma} = \{ l \mapsto \llbracket v \rrbracket^{p, \pi} \mid \sigma(l) = v \}$ <p style="text-align: center;">(e) Merging a configuration and a patch</p>	

Fig. 3. Merging transformation (partial).

The merging transformation renames each new-version function from g to g' , and changes all new-version code to call g' instead of g (the first rewrite rules in Figure 3(b) and (d), respectively). For each old-version function g , it generates a new function g_{ptr} whose body conditionally calls the old or new version of g , depending on whether an update has occurred (Figure 3(a)). The transformation introduces a global variable $uflag$ (Figure 3(e)) and a function $isupd$ to keep track of whether the update has taken place (bottom of Figure 3(a)). All calls to g in the old version are rewritten to call g_{ptr} instead (top of Figure 3(c)).

The transformation rewrites occurrences of **update** in old-version code into expressions that check whether $uflag$ is positive (bottom of Figure 3(c)). If it is, then the update has already taken place, so there is nothing to do. Otherwise, the transformation sets $uflag$ to $?$, which simulates a non-deterministic choice of whether to apply the update. If $uflag$ now has a positive value, the update path was chosen, so the transformation executes the developer-provided state transformation e , which must also be transformed according to $\{\!\{ \cdot \}\!\}^p$ to properly reference functions in the new program. While this transformation results in multiple occurrences of the expression e , in practice e is a call to a state

transformation function defined in the new version and so does not significantly increase code size.

Version tests running p are translated into calls to *isupd* in the old version, and to appropriate constants in the new code (since we know the update has occurred if new code is running).

3.4 Equivalence

We can now prove that an update to an old-program configuration is correct if and only if the result of merging that configuration and the update is correct. This result lets us use stock verification tools to check properties of dynamic updates using the merged program, which simulates updating, instead of having to develop new tools or extend existing ones.

We say that a program and a sequence of updates are *correct* if evaluation never reaches error (i.e., if there are no assertion failures). More formally:

Definition 1 (Correctness) *A configuration $\langle p; \sigma; e \rangle$ and an update π are correct, written $\models \langle p; \sigma; e \rangle, \pi$, if and only if for all p', σ', e' it is the case that $\langle p; \sigma; e \rangle \xrightarrow{\pi}^* \langle p'; \sigma'; e' \rangle$ implies e' is not error.*

The expression e at startup could be a call to an entry-point function (i.e., *main*). A correct program need not apply π , though no other update may occur. When no update is permitted we write $\models \langle p; \sigma; e \rangle$.

Theorem 1 (Equivalence) *For all p, σ, e, π such that $\text{dom}(p_\pi) \supseteq \text{dom}(p)$ we have that $\models \langle p; \sigma; e \rangle, \pi$ if and only if $\models (\langle p, \sigma, e \rangle \triangleright \pi)$.*

The proof is by bisimulation and is detailed in our technical report [11].

Observe that type errors result in stuck programs, e.g., `!1` does not reduce, while the above theorem speaks only about reductions to *error*. We have chosen not to consider type safety in the formal system to keep things simple; adding types, we could appeal to standard techniques [23–25, 7]. Our implementation catches type errors that could arise due to a dynamic update by transforming them into assertion violations. In particular, we rename functions and global variables whose type has changed prior to merging, essentially modeling the change as a deletion of one variable and the addition of another. Deleted functions are modeled as mentioned above, and deleted global variables are essentially assigned the error expression. Thus, any old code that accesses a stale definition post-update (including one with a changed type) fails with an assertion violation.

4 Experiments

To evaluate our approach, we have implemented the merging transformation for C programs, with the additional work to handle C being largely routine. We merged several programs and dynamic updates and then checked the merged programs against a range of CO-specs. We analyzed the merged programs using

two different tools: the symbolic executor Otter, developed by Ma et al. [22], and the verification tool Thor, developed by Magill et al. [18]. The tools represent a tradeoff: Otter is easier to use and more scalable but provides incomplete assurance, while Thor can guarantee correctness but is less scalable and requires more manual effort. Overall, both tools proved useful. Otter successfully checked all the COs-specs we tried, generally in less than one minute. Thor was able to fully verify several updates, though running times were longer. Both tools found bugs in updates, including mistakes we introduced inadvertently. On average, verification of merged code took four times longer than verification of a single version. Since our approach is independent of the verification tool used, its performance and effectiveness will improve as advances are made in verification technology.

4.1 Programs

We ran Otter and Thor on updates to three target programs. The first two are small, synthetic examples: a multiset server, which maintains a multiset of integer values, and a key-value store. For each program, we also developed a number of updates inspired by common program changes such as memory and performance optimizations and semantic changes observed in real-world systems such as Cassandra [5]. The third program we considered is Redis [21], a widely used open-source key-value server. At roughly 12k lines of C code, Redis is significantly larger than our synthetic examples, and is currently not tractable for Thor. We developed six dynamic patches for Redis that update between each pair of consecutive versions from 1.3.6 through 1.3.12, and we also wrote a set of CO-specs that describe basic correctness properties of the updates.

As we mention in Section 2, we join each CO-spec with the server code and have the main function invoke the CO-spec after it initializes server data structures. The new-version source code includes the state transformation code, which is identified by a distinguished function name recognized by the merger.

Synthetic Examples. Figure 4 lists the synthetic benchmarks we constructed for our multiset and key-value store programs. Each grouping of rows shows a dynamic update and a list of CO-specs we wrote for that update. The multiset program has routines to add and delete elements and to test membership. The updates both change to a set semantics, where duplicate elements are disallowed. The first (correct) state transformer removes all duplicates from a linked list that maintains the current multiset. The second update has a broken state transformer that fails to remove duplicates.

The key-value store program also implements its store with a linked list. The updates are inspired by code changes we have seen in practice and include a bug fix (bindings could not be overwritten), a feature addition (adding namespaces), and an optimization (removing overwritten bindings), where for this last update the state transformer was broken at first.

The properties span all the categories of CO-specs that we outlined in Section 2. Backward compatible specs, such as add-mem, check core functionality

Program – <i>change</i> CO-specs	Thor time (s)			Otter time (s)		
	old	new	mrg	old	new	mrg
Multiset – <i>disallow duplicates</i> (correct)						
mem-mem ^b	90.11	121.27	1003.22	6.29	9.72	49.37
add-mem ^b	64.17	89.71	537.01	3.26	10.48	50.84
add-add-del-set ^g			–			4.04
Multiset – <i>disallow duplicates</i> (broken)						
mem-mem ^b	25.33	57.78	133.68	6.28	9.77	42.5
add-mem ^b	15.68	33.50	80.07	3.25	9.94	33.53
add-add-del-set-fails ^g			122.71			5.49
Key-value store – <i>bug fix</i>						
put-get ^b	27.01	26.13	41.62	3.28	2.54	18.42
new-def-shadows ^g			–			4.19
new-def-shadows-bc-fails ^b	38.97	41.52	117.56	3.88	2.06	19.03
Key-value store – <i>added namespaces</i>						
new-def-shadows-post ^p		–	–		1.02	2.99
put-get ^p		–	–		18.32	228.69
new-def-shadows-conf ^c	–	–	–	1.19	1.93	7.53
put-get-conf ^c	–	–	–	4.23	7.09	61.41
Key-value store – <i>optimization</i> (broken)						
put-get-back ^b	42.133	–	–	2.08	11.01	56.44
new-def-shadows-back ^b	15.344	–	–	2.14	11.33	56.03
Key-value store – <i>optimization</i> (correct)						
put-get-back ^b	41.87	–	–	2.07	10.87	69.31
new-def-shadows-back ^b	15.72	–	–	2.14	10.96	68.95

b – backward compatible *p* – post update *c* – conformable *g* – general

A dash indicates that the example could not be verified.

Fig. 4. Synthetic examples.

that does not change between versions (**add** actually adds elements, **delete** removes elements, etc.). Post-update and general CO-specs are used to check that functionality *does* change, but only in expected ways. For example, `new-def-shadows` in the *bug-fix* update checks that, following the update, new key-value bindings properly overwrite old bindings (which was not true in the old version).

We wrote specifications to be as general as possible. For example, `add-mem`, on the second line of the table in Figure 4, checks that after an element is added, it is reported as present after an arbitrary sequence of function calls that does not include `delete()`. The code for our synthetic examples and their associated CO-specs is available on-line.³

Redis. Figure 5 lists the updates and CO-specs for Redis. Four of the six updates required writing state transformers, often just to initialize added fields but sometimes to perform more complex transformation, e.g., the update to 1.3.9 required some reorganization of data structures storing the main database.

³ <http://www.cs.umd.edu/projects/PL/dsu/data/dsumerge-examples.tar.gz>

We found that across these updates, there were four different kinds of behavioral changes, each of which suggested a certain strategy for developing CO-specs; we employed CO-specs in each of the classes described in Section 2:

- *Unmodified behavior*: We adapted two CO-specs from our synthetic key-value store example (Figure 4), *put-get* and *new-def-shadows*, to check correct behavior of Redis’ SET and GET operations over string values. As these CO-specs concern behavior that all versions of Redis should exhibit, we applied them as backward compatible CO-specs.
- *New operations*: The HASHINCRBY operation, which adds to the numeric value stored for a hash key, first appeared in version 1.3.8. We check the operation’s correctness using a post-update CO-spec, *hashincrby*. The HASHINCRBY operation is supported by all later versions, and so we also developed a backward compatible *hashincrby* CO-spec for subsequent updates.
- *Modified semantics*: Before Redis version 1.3.8, a set whose last element was removed would remain in the database. We use the backward compatible CO-spec *empty-set-exists* to check this property against the patch to 1.3.7. Then for the patch to 1.3.8, which causes the server to remove a set when it becomes empty, we use a general CO-spec *empty-set-notexists* to ensure that sets are removed if they become empty after the update. Subsequent versions preserve this behavior, which we specify using a backward compatible CO-spec.
- *Conformable changes*: Redis’s ZINTER operation, which computes the intersection of two sorted sets, was renamed to ZINTERSTORE in version 1.3.12. We use a conformable CO-spec, *zinter*, to specify correct behavior regardless of when an update occurs.

To make symbolic execution tractable for Redis, we had to bound the non-determinism in our CO-specs, e.g., by limiting “arbitrary behavior” to a single operation, non-deterministically chosen from a subset of commands that relate to the specified property (rather than from the full set of Redis operations).

4.2 Effectiveness

In most cases, checking CO-specs validated the correctness of our dynamic patches. In some cases the checking found bugs. For example, in the state transformer for the multiset-to-set update, we inadvertently introduced a possible null pointer dereference when freeing duplicates. Verification with Thor discovered this problem. For Redis, we experimented with omitting state transformation code or using code with a simple mistake in it. In all cases, checking our specifications with Otter uncovered the mistakes.

Figures 4 and 5 show the running times for each of the update/CO-spec/tool combinations, listed under the **mrg** heading. As a baseline, we also list the running times for the backward-compatible specifications on both individual program versions, and for post-update specifications on the new version—this lets us compare the relative slowdown incurred by reasoning about updates.

	Specification	Otter time (s)			Specification	Otter time (s)			
		old	new	mrg		old	new	mrg	
→1.3.7	put-get ^b	9.76	9.52	24.99	→1.3.10	put-get ^b	9.22	10.05	27.37
	new-def-shadows ^b	2.19	2.19	3.97		new-def-shadows ^b	2.70	2.69	4.79
	empty-set-exists ^b	9.95	9.92	29.15		hashincrby ^b	14.86	15.26	46.74
→1.3.8	put-get ^b	9.20	9.58	28.53	→1.3.11*	empty-set-notexists ^b	11.14	11.36	35.01
	new-def-shadows ^b	2.17	2.27	4.14		put-get ^b	9.85	10.04	50.73
	hashincrby ^p		3.02	14.81		new-def-shadows ^b	2.69	2.77	6.30
→1.3.9*	empty-set-notexists ^g			27.58	→1.3.12*	hashincrby ^b	15.19	15.51	77.80
	put-get ^b	9.14	9.31	48.08		empty-set-notexists ^b	11.33	11.57	72.40
	new-def-shadows ^b	2.27	2.66	5.46		put-get ^b	10.32	9.72	49.23
→1.3.10*	hashincrby ^b	14.23	14.83	77.14	→1.3.11*	new-def-shadows ^b	2.85	2.92	6.27
	empty-set-notexists ^b	10.56	11.13	62.88		hashincrby ^b	15.20	14.79	77.27
						empty-set-notexists ^b	11.58	11.67	72.16
					zinter ^c	60.30	59.73	294.05	

b – backward compatible *p* – post update
c – conformable *g* – general * – xform

Fig. 5. Otter checking times for Redis

Otter. We performed experiments with Otter on a machine with a dual-core Pentium-D 3.6GHz processor and 2GB of memory. The running times range from seconds to a few minutes, depending on the complexity of the specification and the program. For example, the CO-specs for the multiset-to-set example were expensive to symbolically execute because each set insertion checks for duplicates, which induces many branches when symbolic values are involved.

We also see that, across the synthetic examples and Redis, it takes four times longer to analyze merged programs versus individual versions on average, and 6.4 times longer in the worst case. We investigated the source of the slowdown, and found it was due to the extra time required to model update points and state transformers, which is fundamental to verifying updating programs, rather than an artifact of our merging strategy. In particular, Otter runs on the merged versions, so it must explore additional program paths to model each possible update timing; on average, CO-specs reached 3.7 update points during execution and, loosely speaking, each update point could induce another full exploration through the set of non-updating program paths. State transformation is also executed following updates, so the expense of symbolically executing the transformer is multiplied by the number of times an update point is reached. Nevertheless, despite this slowdown, total checking time was rarely an impediment to checking useful properties.

Thor. We ran Thor on a 2.8GHz Intel Core 2 Duo with 4GB of memory. The average slowdown was 3.9 times, and ranged from 1.5 times to 8.3 times. Much of the slowdown derived from per-update-point analysis of the state transformation function; tools that compute procedure summaries or otherwise support modular verification would likely do better. Thor could not verify all our examples, owing to complex state transformation code and CO-specs that specify very precise properties. For example, for the multiset-to-set example, Thor was able to prove

that the state transformer preserves list membership (used to verify *mem-mem*), but not that it leaves at most one copy of any element in the list (needed for *add-add-del-set*).

The CO-specs we considered lie at the boundary of what is possible for current verification technology. To verify all our examples requires a robust treatment of pointer manipulation, integer arithmetic, and reasoning about collections. We are not aware of any tools that currently offer such a combination. However, we hope that the demonstrated utility of such specifications will help inspire further research in this area.

5 Related work

This paper presents the first approach for automatically verifying the correctness of dynamic software updates. As mentioned in the introduction, prior automated analyses focus on safety properties like type safety [23], rather than correctness. As described in Section 2, our notion of client-oriented specifications captures and extends prior notions of update correctness.

Our verification methodology generalizes our prior work [10, 12] on systematically *testing* dynamic software updates. Given tests that pass for both the old and new versions, the tool tests every possible updating execution. This approach only supported backward-compatible properties and does not extend to general properties (e.g., with non-deterministically chosen operations or values).

The merging transformation proposed in this paper was inspired by KISS [20], a tool that transforms multi-threaded programs into single-threaded programs that fix the timing of context switches. This allows them to be analyzed by non-thread-aware tools, just as our merging transformation makes dynamic patches palatable to analysis tools that are not DSU-aware.

An alternative technique for verifying dynamic updates, explored by Charlton et al. [6], uses a Hoare logic to prove that programs and updates satisfy their specifications, expressed as pre/post-conditions. We find CO-specs preferable to pre/post-conditions because they require less manual effort to verify, and because they naturally express rich properties that span multiple server commands.

6 Summary

We have presented the first system for automatically verifying dynamic-software-update (DSU) correctness. We introduced *client-oriented specifications* as a way to specify update correctness and identified three common, easy-to-construct classes of DSU CO-specs. To permit verification using non-DSU-aware tools, we developed a technique where the old and new versions are *merged* into a single program and proved that it correctly models dynamic updates. We implemented merging for C and found that it enabled the analysis tool, Thor, to fully verify several CO-specs for small updates, and the symbolic executor, Otter, to check and find errors in dynamic patches to Redis, a widely-used server program.

Acknowledgements. We thank Elnatan Reisner, Matthew Parkinson, Nishant Sinha, and the anonymous reviewers for helpful comments on drafts of this paper. This research was supported by the partnership between UMIACS and the Laboratory for Telecommunications Sciences, by ONR grant N00014-09-1-0652, and NSF grants CCF-0910530, CCF-0915978 and CNS-1111698. Any opinions, findings, and recommendations are those of the authors and do not necessarily reflect the views of the ONR or NSF.

References

1. S. Ajmani, B. Liskov, and L. Shriru. Modular software upgrades for distributed systems. In *ECOOP*, July 2006.
2. J. Armstrong, R. Virding, C. Wikstrom, and M. Williams. *Concurrent programming in ERLANG (2nd ed.)*. Prentice Hall International Ltd., 1996.
3. T. Bloom and M. Day. Reconfiguration and module replacement in Argus: theory and practice. *Software Engineering Journal*, 8(2):102–108, March 1993.
4. G. Bracha. Objects as software services. <http://bracha.org/objectsAsSoftwareServices.pdf>, Aug. 2006.
5. Cassandra API overview. <http://wiki.apache.org/cassandra/API>.
6. N. Charlton, B. Horsfall, and B. Reus. Formal reasoning about runtime code update. In *HOTSWUP*, 2011.
7. D. Duggan. Type-based hot swapping of running modules. In *ICFP*, 2001.
8. C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *PLDI*, 1993.
9. D. Gupta, P. Jalote, and G. Barua. A formal framework for on-line software version change. *IEEE TSE*, 22(2), 1996.
10. C. M. Hayden, E. A. Hardisty, M. Hicks, and J. S. Foster. Efficient Systematic Testing for Dynamically Updatable Software. In *HOTSWUP*, 2009.
11. C. M. Hayden, S. Magill, M. Hicks, N. Foster, and J. S. Foster. Specifying and verifying the correctness of dynamic software updates (extended version). Technical Report CS-TR-4997, Dept. of Computer Science, University of Maryland, 2011.
12. C. M. Hayden, E. K. Smith, E. A. Hardisty, M. Hicks, and J. S. Foster. Evaluating dynamic software update safety using systematic testing, Mar. 2011.
13. M. Hicks and S. Nettles. Dynamic software updating. *ACM TOPLAS*, 27(6), 2005.
14. The K42 Project. <http://www.research.ibm.com/K42/>.
15. J. Kramer and J. Magee. The evolving philosophers problem: Dynamic change management. *IEEE TSE*, 16(11), 1990.
16. Never reboot Linux for Linux security updates : Ksplice. <http://www.ksplice.com>.
17. S. Magill, M.-H. Tsai, P. Lee, and Y.-K. Tsay. THOR: A tool for reasoning about shape and arithmetic. In *CAV*, LNCS 5123, pages 428–432. Springer, 2008.
18. S. Magill, M.-H. Tsai, P. Lee, and Y.-K. Tsay. Automatic numeric abstractions for heap-manipulating programs. In *POPL*, 2010.
19. I. Neamtiu, M. Hicks, G. Stoye, and M. Oriol. Practical dynamic software updating for C. In *PLDI*, 2006.
20. S. Qadeer and D. Wu. KISS: Leap it simple and sequential. In *PLDI*, 2004.
21. The Redis project. <http://code.google.com/p/redis/>.
22. E. Reisner, C. Song, K.-K. Ma, J. S. Foster, and A. Porter. Using symbolic evaluation to understand behavior in configurable software systems. In *ICSE*, 2010.

23. G. Stoye, M. Hicks, G. Bierman, P. Sewell, and I. Neamtiu. *Mutatis Mutandis: Safe and flexible dynamic software updating*. *ACM TOPLAS*, 29(4), 2007.
24. S. Subramanian, M. Hicks, and K. S. McKinley. Dynamic software updates for Java: A VM-centric approach. In *PLDI*, 2009.
25. C. Walton. *Abstract Machines for Dynamic Computation*. PhD thesis, University of Edinburgh, 2001. ECS-LFCS-01-425.