Adaptive Redistribution and Replication to Improve the Responsiveness of Mobile Web Apps

Kijin An[1,*] and Eli Tilevich[2]

[1]*Quality Tool Lab, Samsung Research, Seoul, Republic of Korea*
[2]*Software Innovations Lab, Virginia Tech, United States*
E-mail: kijin.an@samsung.com; tilevich@cs.vt.edu
[*]Corresponding Author

Abstract

In a mobile web app, a browser-based client communicates with a cloud-based server across the network. An app is statically divided into client and server functionalities, so the resulting division remains fixed at runtime. However, if such static division mismatches the current network conditions and the device's processing capacities, app responsiveness and energy efficiency can deteriorate rapidly. To address this problem, we present Communicating Web Vessels (CWV), an adaptive redistribution and replication framework that improves the responsiveness of full-stack JavaScript mobile apps. Unlike standard computation offloading, in which client functionalities move to run on the server, CWV's redistribution is bidirectional. Without any preprocessing, CWV enables apps to move any functionality from the server to the client and vice versa at runtime, thus adapting to the ever-changing execution environment of the web. Having moved to the client, former server functionalities become regular local functions. To further improve performance, CWV can replicate server-side functionalities on the client and keep the replicas consistent. By monitoring the network, CWV determines if a redistribution or

a replication would improve app performance, and then analyzes, transforms, sandboxes, moves, or replicates functions and program state at runtime. An evaluation with third-party mobile web apps shows that CWV optimizes their performance for dissimilar network conditions and client devices. As compared to their original versions, CWV-powered web apps improve their performance (i.e., latency, energy consumption), particularly when executed over limited networks.[1]

Keywords: Mobile web apps, JavaScript, dynamic adaptation, program analysis & transformation, web frameworks, replication.

1 Introduction

Mobile web apps are fundamentally distributed: browser-based clients communicate with cloud-based servers over the available networks. Distribution assigns an app component to run either on the client or on the server. Some distribution strategies are predefined; for example, user interfaces must display on the client. Other distribution strategies aim at improving performance; for example, a powerful cloud-based server can execute some functionality faster than can a mobile device. Network communication significantly complicates the device/ cloud performance equation. For a client to execute a cloud-based functionality, it needs to pass parameters and receive results over the network. Transferring data across a network imposes latency and energy consumption costs. For low-latency, high-bandwidth networks, these costs are negligible. For limited networks, these costs can grow rapidly and unexpectedly. The overhead of network transfer can not only negate the performance benefits of remote cloud-based execution, but also strain the mobile device's energy budget. Operating over limited high-loss networks requires retransmission, which consumes additional battery power [40]. Hence, fixed distribution can hurt app responsiveness and energy efficiency.

Changing the locality of a software component can be non-trivial due to the differences in latency, concurrency, and failure modes between centralized and distributed executions [3, 4, 41]. Researchers and practitioners alike have thoroughly explored the task of rendering local components remote. *Cloud offloading* moves local functionalities to execute remotely in the cloud [10, 25, 35, 42]. Nevertheless, standard offloading is *unidirectional*:

---

[1]This article is a revised and extended version of our prior paper, published in the 21st International Conference on Web Engineering (ICWE 2021) [9].

it can only move a client functionality to run on a server. If mobile web apps are to flexibly adapt to the ever-changing execution environment of the web, client and server functionalities may need to adaptively switch places at runtime.

We address this problem by adaptively redistributing the client and server functionalities of already distributed applications to optimize their performance and energy efficiency. Our approach works with full-stack JavaScript apps, written entirely (i.e., client and server) in JavaScript. By dynamically instrumenting and monitoring app execution, our approach detects when network conditions deteriorate. In response, it moves the JavaScript code, program state, and SQL statements of a remote service to the client, so the service can be invoked as a regular local function. To prevent cross-site scripting (XSS) or SQL injection attacks, the moved code is sandboxed, creating a separate context with reduced privileges for safe execution in the mobile browser. Thus, the same functionality can be invoked locally or remotely as determined by the current execution environment. To the best of our knowledge, our approach is the first one to support *bidirectional dynamic redistribution of distributed mobile web apps*. Moreover, to take advantage of our approach, a mobile app needs not be written against any specific API or be pre-processed prior to execution.

Some functionalities cannot be moved across execution sites, but they can be replicated. For example, a server component could reflect the latest updates to a shared database, modified by different applications. This component cannot be moved away from the physical database engine, running on the server. However, this component can be replicated on the client, with the server changes synchronized with the client copy and vice versa. In addition, replicating rather than moving a functionality can offer better performance.

We called the reference implementation of our approach – Communicating Web Vessels (CWV) – due to its reminiscence of *communicating vessels*, a physical phenomenon of connected vessels with dissimilar volumes of liquid reaching an equilibrium. CWV balances mobile execution by adaptively redistributing functionalities between the server and the client, thus optimizing app performance for the current execution environment. Our contribution is three-fold:

1. A novel bidirectional redistribution approach that dynamically adapts distributed mobile apps for the current execution environment.
2. A reference implementation of our approach, CWV, that works with increasingly popular full-stack JavaScript mobile apps. Requiring no

pre-processing, CWV dynamically adapts apps by redistributing and/or replicating their JavaScript code, program state, and SQL statements at runtime.
3. A comprehensive evaluation with 23 remote services of 8 real-world apps. To assess the effectiveness of CWV's adaptations, we report on their impact on execution latency and energy consumption.

The rest of this paper is structured as follows. Section 2 motivates and explains our approach. Section 3 describes the reference implementation of our approach. Section 4 presents our evaluation results. Section 5 compares our approach to the related state of the art. Section 6 presents concluding remarks.

2 Approach

We first present a motivating example, then give an overview of CWV, and finally discuss our performance model.

2.1 Motivating Example

Consider *Bookworm*, an e-reader app for reading books on mobile devices. The app also provides text analysis features that report various statistical facts about the read books. The app is distributed: the client hosts the user interface; the server hosts a repository of available books and a collection of text processing routines. The current architecture of *Bookworm* is well-optimized for a typical deployment environment: a resource-constrained mobile device and a powerful server, connected to each other over a reliable network. For limited networks, the performance equation can change drastically. Hence, to exhibit the best performance for all combinations of client and server devices and network connections, the app would have to be distributed in a variety of versions. Even if developers were willing to expend a high programming effort to produce and maintain all these versions, network conditions can change rapidly while the app is in operation, necessitating a different client/server decomposition. Clearly, achieving optimal performance under these conditions would require dynamic adaptation.

Our framework, CWV, can adapt *Bookworm*, so its remote text processing routines could migrate to the client at runtime for execution. CWV monitors the network conditions, migrating server-side functions to the client and reverting the execution back to the server, as determined by the network conditions. The app can start executing with all the text processing routines

Figure 1   Conceptual view of communicating web vessels (CWV).

running on the server. Once the network connection deteriorates, a portion of these routines would be transferred over the network to the client, so they could execute locally. CWV's static and dynamic analyses determine the dependencies across server functions and their individual computational footprints. This information parameterizes CWV's performance model, which determines which part of server functionality needs to migrate to the client under the current network conditions.

2.2 Approach: Communicating Web Vessels

To optimize the performance of mobile web apps for the current network conditions, CWV continuously applies the two operations depicted in Figure 1:

1. $f_r = insource(/service\_r)$: The client requests that the server transfer the remote functionality($/service\_r$)'s partition f_r to the client.
2. $revert(f_r)$: The client stops locally invoking the insourced partition f_r, and starts remotely invoking its original server version $/service\_r$.

2.3 Reasoning About Responsiveness

Responsiveness is a subjective criteria: application is responsive if the user perceives the time taken to execute app functionalities as "short". For this reason, we define the responsiveness of a remote execution as the total execution time that elapses between the client invoking a remote functionality and the results presented to the user. We define the response time of a remote functionality f_r as $RT(f_r)$. The $RT(f_r)$ mainly depends on the "server speed" and "network speed" parameters. We simplify the responsiveness of f_r by means of the execution time f_r on the server $T_{server}(f_r)$ and the remaining remote execution overheads. The resulting Round Trip Time (RTT) is highly affected by the current network conditions. To estimate the network

conditions, CWV utilizes the RTT^{net} metrics, detailed in Section 3.5.1.

$$RT(f_r) = \begin{cases} T_{server}(f_r) + RTT^{net} & \text{remote exec.,} \\ T_{client}(f_r) & \text{local exec.} \end{cases} \qquad (1)$$

If f_r is executed locally, the responsiveness becomes the execution time f_r on the client $T_{client}(f_r)$.

3 Reference Implementation

To move a server-side functionality to the client at runtime, one has to migrate both the relevant source code and program state, which has to be captured and restored at the client. JavaScript has a powerful facility, the `eval` function, which executes a JavaScript program passed to it as a string argument. One could simply duplicate the entire server-side code and its state, passing them to a client-side `eval`. However, such a naïve approach would incur unacceptably high performance and security costs. Hence, our approach applies advanced program analysis and automated transformation techniques to minimize the amount of code to be transferred to and executed by the client (Sections 3.1 and 3.4). Furthermore, our approach establishes an efficient protocol for the transformed app to switch between different execution modes (Section 3.5), transferring the relevant code correctly and safely (Sections 3.6 and 3.7).

3.1 Execution Model for Client/Server Programs

CWV replicates a subset of the cloud server's state and code on clients. To determine which subset to replicate, it analyzes the runtime traces of the execution of cloud services. Consider a typical life-cycle of a cloud service: (1) the server initializes itself to the "init state" and (2) the server receives an HTTP request from the client and unmarshals the passed parameters, (3) the server passes the unmarshalled parameters to the the remote functionality f_r, which starts executing, and (4) upon completing its execution, the service marshals the response to return it over HTTP to the client. For brevity, we refer to the above steps (2)(3)(4) as the i^{th} execution or `exec` i.

$$f_r : \texttt{init}, \texttt{exec}\ i, \texttt{exec}\ i+1, \cdots \qquad (2)$$

To infer i^{th} execution corresponding to the remote functionality f_r, CWV not only extracts the code executed in step (3), but also captures the subset of the server's state required to initialize a replica to be executed in the client.

What complicates this procedure is that some cloud-based services are stateful, and as such change their state with every execution. Hence, it becomes problematic to detect their relevant subsets of the state to replicate. Even if such stateful services could be restarted anew for every execution, only some of their states would be restored, as they also often persist data in a database.

3.2 Isolating and Replicating States

CWV isolates state changes when analyzing the dynamic execution traces of server-based services. This *state isolation* ensures that the server's init state and the service's execution results remain fixed. To that end, after the `init` is executed, its state is checkpointed and subsequently restored before every repeated service execution as follows:

$$\text{init}, \text{save init}, \text{exec } i, \text{restore init}, \text{exec } i+1, \cdots \qquad (3)$$

To be specific, CWV identifies and replicates several different units of cloud-based services that typically include "a database", "files", and "program variables" [20].

Database Tables: CWV monitors dynamic traces of a Node.js cloud-based service with Jalangi [33], a dynamic JavaScript analysis framework. To identify database-related statements, CWV instruments all function invocations whose argument values are SQL commands. To that end, CWV modifies the INVOKEFUN(LOC, F, ARGS, VAL) callback API of Jalangi to be able to examine the arguments parameter, args. Then, it adds shadow executions into the identified SQL invocations loc to trace the changes to the database state. First, CWV appends shadow execution of the original SQL command with a SQL command to snapshot the entire dependent tables. Next, it adds transactional executions START TRANSACTION and ROLLBACK against the original SQL commands, to keep the database tables unchanged.

Files: In cloud-based services, files can be accessed both locally and remotely. To identify file accesses, CWV instruments all invocations whose arguments are file URLs. It then duplicates the identified files by copying or downloading.

Global variables: CWV adds get/set function instrumentation after the declarations of global variables to implement their save and restore operations. After the server has been initialized, CWV deeply copies all global variables and saves their states. The restore operation passes the saved states to each variable's set function.

3.3 Replicating Relevant Server Source Code and Program State

Server code comprises business logic and middleware libraries. The server-side business logic can store some states in the server program by including database access routines or other mechanisms. The portion that needs to be insourced is business logic only. In other words, business logic must be reliably separated from all middleware-related functionality. To that end, CWV identifies the entry and exit statements of the business logic portion, equivalently to unmarshaling/marshaling points for the steps of exec i, and then extracts all the code executed between these statements, converting that code to a new regular JavaScript function. All the dependent code of this new function is also extracted and transferred, thus producing a self-sufficient execution unit.

The specific steps are as follows. First, CWV normalizes the server code to facilitate the process of separating its business logic from middleware functionality. Then, CWV locates the statements that "unmarshal" the *client parameters* and "marshal" the *result* of executing the business logic. CWV automatically identifies these statements by capturing the client server HTTP traffic and instrumenting code at the server and at the client (Figure 2-(a)). To that end, CWV uses Jalangi [33], already used for isolating the server program's states. CWV modifies the built-in Jalangi's callback API calls to be able to detect the events that correspond to the "unmarshal/marshal" statements. By following these steps, CWV identifies the specific lines of code and variables that correspond to the entry and exit points of remote invocations, both at the server and the client.

The statements executed between these points comprise the server-side business logic and its dependent program states that may need to be moved to the client at runtime. To identify a subset of statements that satisfies a pair of entry/exit statements, CWV follows a strategy similar to that of other declarative program analysis frameworks that analyze programs by means of a Datalog engine [7, 36–38]. CWV encodes the declarative facts that specify the behavior of JavaScript statements of server program: (1) declarations of variables/functions, (2) their read/writes operations, and (3) control flow graphs. The dependency analysis query constructs a dependency graph between statements. Then, CWV solves constraints describing these points with the z3 engine [11] and then extracts them into a CWV-specific object that is movable between vessels (Figure 2-(b)).

Some server-side program statements use third-party APIs, whose libraries and frameworks are deployed only at the server. CWV provides
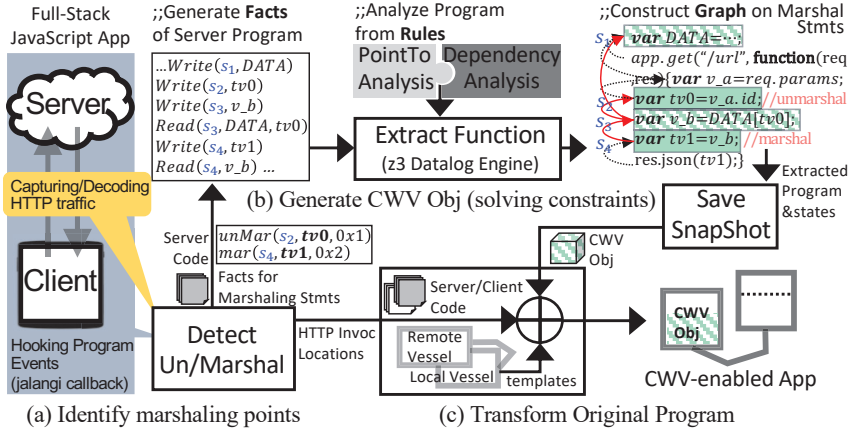
Full-Stack
JavaScript App

;;Generate **Facts**
of Server Program

;;Analyze Program
from **Rules**

;;Construct **Graph** on Marshal Stmts

Server

$...Write(s_1, DATA)$
$Write(s_2, tv0)$
$Write(s_3, v\_b)$
$Read(s_3, DATA, tv0)$
$Write(s_4, tv1)$
$Read(s_4, v\_b) ...$

PointTo Analysis   Dependency Analysis

Extract Function
(z3 Datalog Engine)

S_1 **var** $DATA=...$;
$app.get("/url",$ **function**(req
res→{**var** $v\_a=req.params$;
S_2 **var** $tv0=v\_a.id$; //unmarshal
S_3 **var** $v\_b=DATA[tv0]$;
S_4' **var** $tv1=v\_b$; //marshal
$res.json(tv1)$;}

Capturing/Decoding
HTTP traffic

(b) Generate CWV Obj (solving constraints)

Save SnapShot

Extracted Program &states

Server Code $\boxed{unMar(s_2, \boldsymbol{tv0}, 0x1)}$
$\boxed{mar(s_4, \boldsymbol{tv1}, 0x2)}$

CWV Obj

Client

Facts for Marshaling Stmts

Server/Client Code

Hooking Program Events
(jalangi callback)

Detect Un/Marshal

HTTP Invoc Locations

Remote Vessel

Local Vessel templates

CWV Obj

(a) Identify marshaling points

(c) Transform Original Program

CWV-enabled App

Figure 2   Automated program transformation for enabling CWV.

domain-specific handling of the statements that interact with relational databases. In particular, some statements interacting with a server-side relational database cannot be directly migrated to the client. As a specific example, consider the statement `mysql_server.query(SQL_STATEMENT)`, which queries the server-side MySQL database engine. Mobile clients can also use relational databases, but of a different type, a browser-hosted SQL engine. Hence, the database-related statement above should be replaced with `a_mobile_engine( SQL_STATEMENT)`. To identify such database-related statements, CWV instruments all function invocations whose arguments are SQL commands by using callback API of Jalangi. Despite the fragility of relying on the usage of SQL commands, our approach presents a practical solution for supporting domain-specific server-to-client migrations. Finally, CWV transforms the identified entry/exit points at the client and server sides to insert the CWV functionality with the local and remote vessels respectively that we explain in the next section (Figure 2-(c)).

3.4 Transforming Programs to Enable CWV

CWV enhances application source code to enable its transformation as follows.

3.4.1 Client enhancements

CWV transforms the identified HTTP invocation in the client program to be able to CWV's functionality as follow. The CWV-enabled client can operate

and switch between these two modes: *Original* and *Local*. In *Original* mode, the app operates the original remote execution and can switch to Local mode by means of *Insourcing*. The *Local* mode designates that the local version of the insourced remote functions is to be invoked and can revert to the original mode by means of *Reverting* (See Figure 3). To switch to a mode, the client invokes `fuzzMode(mode)` that simply fuzzes a certain parameter of the HTTP command that invokes the original remote service name. For instance, the client can dynamically fuzz a remote service `"/a_service"` (Original Request) into `"/a_service?CWVmode=Local"` (Local). And the app initiates the movement of the relevant remote server code and execution states $rcwv$ to the client by fuzzing the original invocation into `"/a_service?CWVmode=Insourcing"` (Insourcing Request).

Insourcing CWV moves a set of received server statements into a client's container, referred to as the *local vessel*. Initially, the local vessel is empty. When the client device determines to switch from the *Original* mode into the *Local* mode, the app issues the Insourcing Request and then invokes the `moveToLocalVessel(rcwv)` call, only then adding received server code and state to the local vessel. The client and server share all the referenced names for global entries added to the local vessels. To that end, CWV also adds a special-purpose global object for the client, $lcwv$. This object is used for storing functions and other JavaScript objects received from the server.[2] Finally, the app fuzzes the HTTP command into Local `"CWVmode=Local"` to change the current mode. After that, invoking the `rebalance()` function compares the local replica's execution time with that of its original remote version.

Reverting If the local execution stops being advantageous, the app with *Local* mode reverts to *Original* mode and clears the local vessel with `clearLocalVessel()`, overriding the local vessel into the empty function again. And then, the app switches the mode by fuzzing HTTP command into the original mode.

3.4.2 Server enhancements

In a CWV-enabled app, the server part can operate in one of three modes to respond the client's requests: *Original*, *Insourcing*, and *Local*. With the detected entry/exit points of a remote functionality, CWV transforms it to be able to detect the mode switching queries and switch to the client-requested modes. The *Original* mode refers to the original unmodified execution, with

---

[2]The properties of $lcwv$ are the same as of the remote object $rcwv$

the exception for the profiling of the time taken to execute the program statements that implement business logic $T_{server}(f_r)$ of the Equation (1). The client uses resulting performance profiles to ascertain the current network conditions RTT^{net} from the measured response time $RT(f_r)$. And $T_{server}(f_r)$ will be used to determine a threshold when to switch modes.

In the *Insourcing* mode, the server responds to the client's special insourcing query by serializing the relevant portions of a given remote functionality into a JSON string. To that end, CWV calls $saveSnapshot(f_r)$, whose invocation creates a snapshot of the remote functionality f_r. CWV adds to the server part a special-purpose global object, $rcwv$, which represents a *remote vessel*. This object's properties contain the extracted functions, $rcwv.main, rcwv.\ ftns[0], \cdots, rcwv.ftns[k]$ and their corresponding saved states for global variables $rcwv.gvars[0], \cdots, rcwv.gvars[l]$. To migrate f_r with database dependent statements, CWV takes a snapshot of database's table in terms of SQL commands to enable restoration in the client $rcwv.sql[0], \cdots, rcwv.sql[m]$. To implement $saveSnapshot(f_r)$, CWV instruments (1) the declarations of global variables and (2) *Call Expressions* of embedded SQL statements extracted by the constraints solving phrase. Finally, in the *Local* mode, the server executes no business logic, but responds to periodic pings from the client. Based on the roundtrip time of these pings, the client monitors the network conditions to detect if the *Local* mode execution no longer provides any performance advantages and then switches the app to the *Original* mode.

3.5 Updating Modes and Cutoff Latency

The transition diagram in Figure 3 shows how an app can transition between different modes. CWV-enabled client always starts in the Original mode. An insourcing request issued in the Original mode can be either fulfilled (i.e., switching to the Local mode) or declined (i.e., continuing to execute remotely in the Original mode), with the latter incurring a large performance overhead. To avoid this overhead, the system determines the optimal time window for issuing "Insourcing Request" as soon as the app is automatically initialized with a couple of original executions. The procedure that determines the window is as follows. First, the client profiles both $RT(f_r)$ and $T_{server}(f_r)$ by means of multiple "Original Requests" during the initialization (Section 3.4.2). After that, the procedure invokes the "Insourcing Request" and extrapolates how much time it would take to execute the same business logic locally $T_{client}(f_r)$.
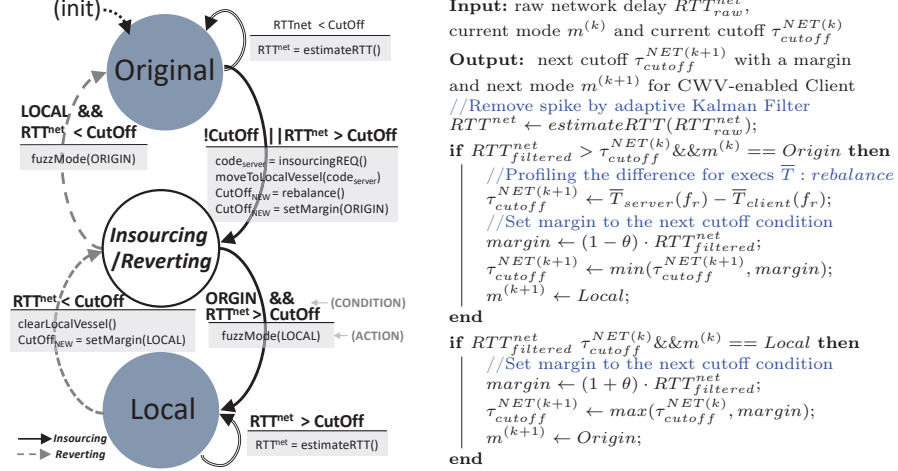
Figure 3 Transition diagram for CWV-enabled client (left). Algorithm for updating cutoffs and modes (right).

3.5.1 Estimating network delay

CWV-enabled mobile clients continuously monitor the underlying network conditions. The client collects the RTT_{raw}^{net} metric that represents raw network delay. Specifically, the client is continuously monitoring the RTT_{raw}^{net} by subtracting $T(f_r)$ from $RT(f_r)$, which are obtained from the server. Since the raw roundtrip is subject to sudden spikes [24], CWV filters out such temporary fluctuations by applying an adaptive filter [30], which calculates the covariance matrices and noise values for RTT_{raw}^{net} and then estimates the RTT^{net} metric in Equation (1).

3.5.2 Cutoff network latency

The resulting difference between the local and remote execution times is used as the threshold that determines when switching to the Local mode would become advantageous from the performance standpoint. In other words, the difference value is compared with the overhead of network communication, and when the latter starts exceeding the former, the app switches to the Local mode. We define this network condition as *cutoff network latency*, τ_{cutoff}^{NET}. Thus, a CWV-enabled app obtains this threshold as soon as it start executing, and then stays in the *Original* mode until reaching the *cutoff*. Then, it tries switching to the Local mode. Because this request is executed only upon reaching the *cutoff*, it is more likely to be fulfilled as offering better performance.

Since switching between modes incurs communication and processing costs, frequent switching in response to insignificant network changes should be prevented. To that end, the *margin* parameter expresses by how much the network conditions need to change and remain changed. The algorithm in Figure 3 explains how the *margin* and the current cutoff latency $\tau_{cutoff}^{NET(k)}$ determine the next cutoff latency $\tau_{cutoff}^{NET(k+1)}$. The margin parameter θ prevents switching in response to insignificant $\tau_{cutoff}^{NET(k)}$ changes. After switching to the Local mode, the app periodically pings the network to determine if the current conditions are advantageous for reverting to the Original remote mode. Figure 4 shows how applied filter removes the confusing noise. Compare the ground truth and CWV's switches, both the filter and the margin in CWV are important to ascertain the major trends in the changes of network delay.

3.5.3 Moving code before reaching a degraded network state

Notice that insourcing cannot be accomplished over a limited network. Hence, the procedure needs to be initiated when the network conditions start deteriorating, but before they have reached the point of becoming poor. Since the conditions of a typical mobile network can fluctuate, going up and down, the insourcing commences when the conditions degrade to a given threshold, at which it is still possible to transfer the required source code and state from the server to the client. After the insourcing, if the conditions deteriorate further, the execution switches into the local mode; however, if they improve, the insourcing is discarded, and the execution continues remotely.

Specifically, to detect the network degradation point, the CWV is monitoring successive increases in RTTs, which are larger than $RTT_{limited}^{NET}$, a configurable parameter used to identify if the network is becoming limited.[3]

In terms of the actual operations, in the original mode, the app can receive a transmitted code and state from the server at this network degradation point with Insourcing Request. However, this transmission is not applied to the local vessel by `moveToLocalVessel` until the cutoff point.

3.6 Synchronizing States

Some remote services can be invoked by means of HTTP POST, PUT, DELETE, which are all state-modifying operations. Invoking an insourced

---

[3]We set its default value to 4 secs, the ping command's default timeout.

Cutoff [ms]

$\tau(2)$

$\tau(0)$

"Margin"

$\tau(1)$

2 15

1 15

15

— Raw RTT
—+— Applying Filter
— — — Cutoff

| zero | +100 | +100 | Marginal Offsets ($< \pm\ \theta \cdot \tau(0)$) | -70 | -100 | -100 |

Delay Offsets [ms] (ground of the truth)

(a) Updates /bigtrip's cutoff τ on different network conditions ($\tau^{(0)}$=224ms)

Filter Margin — Mode Switching

Filter	Margin	Mode Switching					
✓	✓	**Origin**	↕	**Local**		↕	**Origin**
✓	−	Origin		↕ Loc ↕ Orig ↕ Loc ↕ Orig ↕ Loc ↕			Origin
−	−	Origin ↕↕ ↕↕ ↕↕ ··· ↕↕		↕ ↕ ↕↕		↕↕ ··· ↕↕ ↕↕ ··· ↕↕ ↕↕ Origin	

(b) Effects of Filter/Margin (θ=0.2) on mode switching for Fig.4 (a)'s scenario

Figure 4   Monitoring network conditions and adapting distribution.

stateful remote service locally modifies its state, which must be synchronized with its original remote version via some consistency protocol.

Mobile apps are operated in volatile environments, in which mobile devices become temporarily disconnected from the cloud server. To accommodate such volatility, CWV's synchronization is based on a weak consistency model. As an implementation strategy, we take advantage of a proven weak consistency solution, Conflict-Free Replicated Data Types (CRDT), which provide a predefined data structure, whose replicas eventually synchronize their states, as the replicas are being accessed and modified. In CRDTs, the concurrent state updates can diverge temporarily to eventually converge into the same state, as long as the replicas manage to exchange their individual modification histories [19].

Specifically, CWV wraps the replicated 'database' and 'global variables' of *cwv* objects into the 'CRDT-Table', and 'CRDT-JSON' of CRDT templates,[4] respectively. To keep track of changes and resolve conflicts, these CRDT-structures provide the API calls `getChanges` and `applyChanges`. By continuously applying/transmitting the reported changes, the device-based clients and the cloud-based server maintain their individual modification histories and exchange them, thus eventually converging to the same state. To that end, the cloud server periodically sends its state changes on *rcwv* to each client,

---

[4]https://github.com/automerge/automerge

while each client starts sending its state changes on $lcwv$ to the cloud server, as soon as this client reverts to executing remotely.

3.7 Sandboxing Insourced Code

Whenever code needs to be moved across hosts, the move can give rise to vulnerabilities unless special care is taken. The issue of insourcing JavaScript code from the server to the client is security sensitive. Server-side code has several privileges that cannot be provided by mobile browsers. In addition, as it is being transferred, the insourced code can be tempered with to inject attacks. Finally, the transferred segments of server-side database can be accessed by a malicious client-side actor. To mitigate these vulnerabilities, the insourced code is granted the least number of privileges required for it to carry out its functionality. To that end, we *sandbox* the insourced code. Specifically, CWV's sandboxing is applied to the entire local vessel. The insourced functionality has exactly one entry point through which it can be invoked. The sandbox guards the insourced execution from performing operations that require escalating privileges. Finally, because the insourced database data cannot be accessed directly, malicious parties would not be able to exfiltrate it.

As a specific sandboxing mechanism, we take advantage of *iframe*, which has become a standard feature of modern browsers. An iframe creates a new nested browser context, separate from the global scope. Operating in a separate context precludes any shared state between the insourced code and the original client-based code. In addition, HTML5 supports the `sandbox` attribute to further restrict what iframes are allowed to execute.[5] It protects the client from the vulnerability related to client XSS. For instance, a sandboxed iframe is prohibited from accessing `window.localStorage[..]`.

4 Evaluation

Our evaluation seeks answers to the following questions:

- **RQ1: – Redistribution Adaptivity for different Devices:** How beneficial is CWV's redistribution for different mobile devices?
- **RQ2: – Redistribution Adaptivity for Networks:** How beneficial is CWV's redistribution and replication for different networks?

---

[5]https://developer.mozilla.org/en-US/docs/Web/HTML/Element/iframe

- **RQ3: – Energy Savings:** How does CWV's redistribution affect the energy consumption of mobile devices?
- **RQ4: – Overheads:** When integrated with mobile apps, what is the impact of CWV on their performance?

4.1 Device Choice Impact

4.1.1 Dataset

Our evaluation subjects are 23 remote services of 8 full-stack applications, 5 real-world full-stack mobile JavaScript applications, and 3 JavaScript distributed system benchmarks [42]. These subject apps use different middleware frameworks to implement their client/server (*tier-1/-2*) communication and database (*tier-3*), with these frameworks being most popular in the JavaScript ecosystem.

To that end, we searched the results based on combinations of keywords for popular server and client HTTP middleware frameworks, curated by the community. For server-side keywords, we used 'Express', 'Restify', etc., while for client-side keywords we used 'Ajax', 'Angular', etc. Table 1 summarizes their names and the number of source files; 4 subject applications contain database-dependent code. To answer **RQ1**, we tested how the introduced network delays affect different devices. At launch time for each device, CWV automatically calculates the *cutoff network latency* and applies it when scheduling mode switches to minimize the switching overhead. For example, CWV determined the cutoff network latency for the remote service "/hbone" as 26ms for device 1 (D1) in Table 1, having profiled the execution time at the server (T_{server}("/hbone")) and the client (T_{client}^{D1}("/hbone")) as 14ms and 40ms, respectively. Device 1 is a Qualcomm Snapdragon 616 (8 $\times$ 1.5 GHz), and Device 2 is an A8-iphone 6 (2 $\times$1.4 GHz); Device 1 outperforms Device 2. The server is an Intel desktop (i7-7700 4 $\times$ 3.6 GHz). We natively build the subject web apps (JavaScript, html, and CSS) for iOS and Android by using Apache Cordova, a cross-platform development framework. Table 1 demonstrates that the *cutoff latency* of Device 2 (τ_{cutoff}^{D2}) is always larger than that of Device 1 (τ_{cutoff}^{D1}).

4.2 Network latency impact

To answer **RQ2**, we set up a test-bed for evaluating network latency impact (See Figure 7-(a)). Even though, network latency can be changed by controlling RSSI levels, we change network conditions explicitly by means of an

Table 1 Subject remote services

Subject (# of Files)	Remote Services	τ_{cutoff}^{D1} (msec)	τ_{cutoff}^{D2} (msec)
Bookworm (*729 files*)	/ladypet	176*ms*	421*ms*
	/thedea	1120*ms*	2332*ms*
	/thered	158*ms*	424*ms*
	/thegift	97*ms*	120*ms*
	/bigtrip	146*ms*	224*ms*
	/offshore	619*ms*	1528*ms*
	/wallp	146*ms*	458*ms*
	/thecask	90*ms*	102*ms*
DonutShop (*4.9k files*)	/Donut	0.66*ms*	1.54*ms*
	/Donut:id	0.71*ms*	2.2*ms*
	/Empls	0.55*ms*	1.33*ms*
	/Empls:id	0.81*ms*	1.23*ms*
recipebook (*8k files*)	/recipe	0.7*ms*	1.66*ms*
	/recipe:id	0.68*ms*	1.1*ms*
	/ingts/:id	0.82*ms*	2.3*ms*
	/dirs/:id	0.75*ms*	2.1*ms*
pstgr-sql (*4k files*)	/user	1.33*ms*	2.71*ms*
	/user:id	1.72*ms*	2.92*ms*
chem-rules 2.8k files)	/hbone	26*ms*	59*ms*
	/molec	131*ms*	202*ms*
benchmark in [42] (*117 files*)			
str-fasta	/str-fasta	656*ms*	1424*ms*
fannk	/fannk	2576*ms*	4982*ms*
s-norm	/s-norm	1896*ms*	4873*ms*

application-level network emulator.[6] Then, we examine how CWV reacts by redistributing the running applications. In these experiments, the server and the mobile device are connected with a wireless router. We establish a high-speed wireless link between the router and the device (−55 dBm or better). By configuring the router to different delays, we simulate different network conditions in the increasing order of delay. Our test-bed has a minimum delay of about 100 ms for the simulator's zero delay. Therefore, our starting point is 100 ms, with the delays increased in the increments of 20 m, 50 ms, and 100 ms, based on the amount of *cutoff network latency* for each subject. For each increment, we measure the average delay in the execution of our subject applications (response time or responsiveness of a functionality), run in two
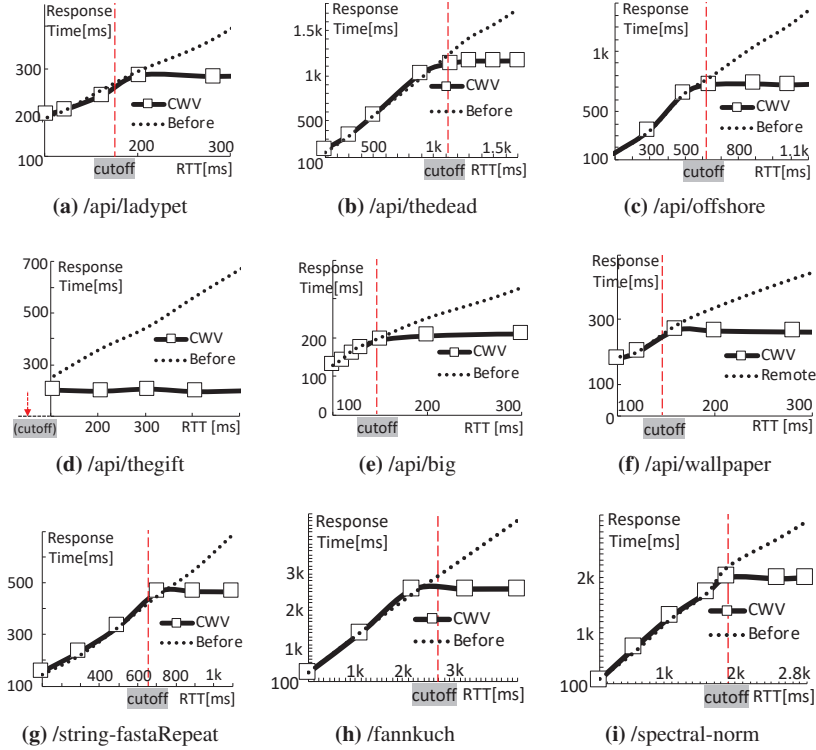
---

[6]https://github.com/h2non/toxy

Figure 5 Client's responsiveness comparison for different network latency. The cutoff equals to τ_{cutoff}^{D1} in Table 1.

configurations: (1) the original unmodified version (**Before**), (2) dynamically redistributed with CWV version (**CWV**). Figure 5 shows the performance results.

Across all experimental subjects, the CWV-enabled configuration consistently outperforms the original version, once the network latency surpasses the *cutoff network latency* mark. Once the network delay reaches the *cutoff network*, the difference in performance starts increasing by a large margin, as accessing any remote functionality becomes prohibitively expensive. Before reaching the *cutoff network* mark, the majority of CWV-enabled apps and their original version exhibit comparable performance since two versions are operated in remote execution. When operating over a high-speed network, CWV-enabled apps remain in the original mode due to the remote execution's performance advantages. Some subjects consistently exhibit better performance when executed locally. These subjects with their relatively
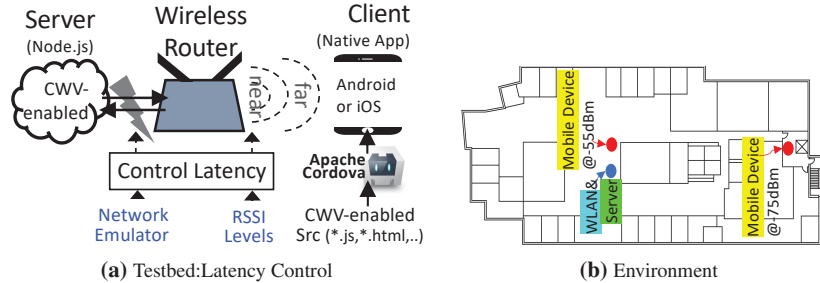
(a) Testbed:Latency Control

(b) Environment

Figure 6   Testbed and testing environments.

low utilization of server resources are better off not making any remote invocations, as the overhead of network delays is not offset by the server's superior processing capacity.

4.3 Energy Consumption

Next, we evaluate how much energy is consumed by a mobile device executing CWV-enabled and original versions of the same subjects. To that end, we use Qualcomm's Trepn Profiler [1] to profile the energy consumption of an Android device running the Snapdragon chipset. Trepn is a self-metering profiler that dynamically generates power models, based on the information collected from the device, thus requiring no pre-training of power models [23].

We executed each subject 100 times and collected the profiled results for battery power (mW). Figure 7 shows the obtained samples of the battery power measurements over time.

To test the consumed energy under a real-world's low speed network environment, we placed the Android client device far from the wireless router, so the signal strength level(RSSI) was -75 dBm. Figure 7 shows that CWV always uses more power than the original version despite shortening the execution time. Remote execution consumes no device power, even if it takes much longer for the client to receive the results. By removing the need to communicate with the server over low-speed high-latency networks, our approach shortens the overall execution time. As the total consumed energy is the product of the power and the elapsed time, our approach consumes less energy **241.8J** (2492 mW·97.05 sec) as compared to **266.3J** (1753 mW·151.92 sec) in the original remote version. Compared to the original version, our approach improves energy efficiency by as much as **10.13%** for poor network conditions.

Figure 7   Testbed and consumed energy.

This result is not unexpected, as a large RTT causes longer idle periods between TCP windows [13]. Even tough, the device switches into the low power mode during the idle states, the longer total execution time still causes larger overall energy consumption.

4.4 Communication Overhead

To insource server execution, CWV serializes relevant code and state to transfer and reproduce at the client. To evaluate the resulting communicating overhead (**RQ3**), we compared the amount of network traffic during the regular remote execution for unmodified version (Tr_{reg}) vs. the additional traffic resulting from CWV insourcing server execution (Tr_{cwv}).

Among our subjects, the Bookworm app exhibits the largest of Tr_{orig}, as this app's remote services need to transfer not only the book content but also the statistical information extracted from that content. Whereas, the med-chem app shows the largest of Tr_{cwv}, as CWV needs to replicate about 10K server-side DB entries. However, the transmitting overhead is occurred only once at initialization as these services are stateless. The resulting over-all overhead ratio $\frac{Tr_{cwv}}{Tr_{reg}}$ turned out to be **2.4** on average for our subjects (Table 2).

Among our subjects, the Bookworm app exhibits the largest of Tr_{orig}, as this app's remote services need to transfer not only the book content but also the statistical information extracted from that content. Whereas, the med-chem app shows the largest of Tr_{cwv}, as CWV needs to replicate all server-side DB entries. However, the transmitting overhead is occurred only once at initialization as these services are stateless. The resulting overall overhead ratio $\frac{Tr_{cwv}}{Tr_{reg}}$ turned out to be **2.4** on average for our subjects (Figure 8). To quantify the benefits of CWV's insourcing transferring only the necessary code and state, we also measured the overhead of the naïve approach, which

Table 2   Overhead comparisons of subject apps (KB)

Subjects	Tr_{reg}	Tr_{cwv}	Naïve
string-fst	0.38	1.3	374.0
fannkuch	0.37	0.99	375.0
spectral-n	0.39	1.3	375.2
recipe(4)$^{\Psi}$	0.45	1.1	**13k**
Bookw(8)	**42.0**	15.8	2.4k
donut-s(4)	0.46	1.1	7.5k
med-c(2)	21.9	**395**	7.5k
Average	**22.5**	**54.4**	**4154**
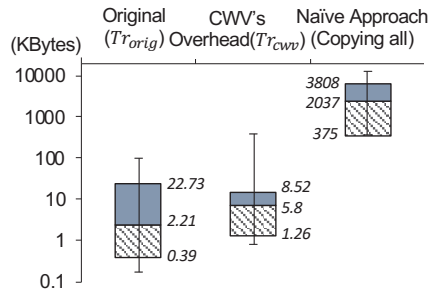
Figure 8   CWV's overhead.

transfers the entire server code and state to the client. The performance overhead of transferring everything is about two orders of magnitude slower than CWV, an unacceptable slowdown for any practical purposes (Figure 8).

4.5 Threats to Validity

Our evaluation is subject to both internal and external threats to validity that we discuss in turn next.

4.5.1 Internal threats

We chose to perform all code analysis and transformation at runtime as an implementation choice, even though some of these tasks could have been performed offline. In other words, we could have transformed the client code statically by adding to it the remote functions that might need to be executed locally. In fact, such static code transformation would allow us to take advantage of the advanced optimization capability of modern JavaScript engines that remove the overhead of invoking constructors. Nevertheless, we chose to transform the code at runtime for maximum flexibility at the cost

of performance. One can further optimize our implementation by replacing dynamic code migration with a separate post-compilation phase. Hence, our evaluation numbers are reflective of the slowest possible implementation strategy and do not unfairly characterize the efficiency of our approach.

4.5.2 External threats

To measure how much energy is consumed by Android devices, we utilize the Snap dragon profiler, which follows a model-based procedure for estimating energy consumption. If we were to use a power monitor instead, our energy measurements could have differed. Nevertheless, when evaluating energy consumption, our focus is the difference between the local and remote execution of certain functionalities rather than their raw energy consumption numbers. Hence, the obtained energy measurements are sufficient to answer our evaluation questions.

When evaluating the impact of a device choice on the mode switching points, we used two different mobile devices: Android and iOS. One could argue that the actual execution environment of JavaScript mobile clients is a mobile browser, whose execution is affected mainly by the underlying device's hardware components rather than the mobile platform. Indeed, as we have observed, the actual cutoff points are heavily affected by the device's CPU speed, with the differences stemming from the device's mobile platform being quite modest.

Similarly to the mobile devices running the client, the server machine running the remote functionality could also affect the resulting performance and the mode switching calculation. The more powerful is the server machine, the more advantageous it is to execute a functionality remotely, if at all possible. The cutoff network latency is always higher for powerful server machines, as the performance advantages of fast server execution can easily subsume the sluggishness of transmitting over slow networks.

4.6 Applicability and Limitations

Our approach's reference implementation works only with JavaScript source code and SQL database code. Nevertheless, some mobile web applications are multilingual (i.e., written in different programming languages and database query languages). Nevertheless, our key ideas and new technologies can be extended for multilingual distributed applications. As automatic language translation has been entrenched into modern software development [5] (e.g., transpiling TypeScript to JavaScript), integrating mature

language translators with our infrastructure is mainly an engineering issue. Furthermore, full-stack JavaScript applications are starting to dominate the development landscape of distributed web applications, as their monolingual nature lowers the development and maintenance burdens, requiring programming proficiency in only JavaScript, used across the development stack.

5 Related Work

5.1 Program Synthesis and Transformation

For a given remote service, CWV automatically identifies the relevant business functionality that satisfies the client's input and server's output constraints. In that respect, it can be seen as a variant of program synthesis [14–17, 21, 34], an active research area concerned with producing a program that satisfies a given set of input/output relationships. Recently, several techniques have been introduced that automatically integrate portions of a program's source in another program [10, 18, 35, 39]. CodeCarbonReply [35] and Scalpel [10] supporting this functionality for C/C++ programs. The programmer annotates the code regions to integrate, and a tool automatically adapts the receiving application's code to work seamlessly with the transferred functionality. In contrast, CWV is both fully automated and dynamic, integrating program code and state at runtime. CanDoR [6] fixes the bug in the centralized variant version with existing tools, then CanDoR applies the resulting fixes to the original distributed app by using program transformation.

5.2 Adaptive Middleware

Several middleware-based approach has been proposed to reduce the costs of invoking remote functionalities. APE [31] is an annotation based middleware service for continuously-running mobile (CRM) applications. APE defers remote invocations until some other applications switch the device's state to network activation. Similarly, to reduce the overhead of HTTP communication, HTTP requests in Android apps are automatically identified and then bundled into a single batched network transmission [28, 29]. The e-ADAM middleware [26] optimizes energy consumption by dynamically changing various aspects of data transmission (e.g., encoding and compression). APE [31] defers remote invocations until some other apps switch the device's state to network activation. DR-OSGi [27] enhances middleware mechanisms with resilience against network volatility. D-Goldilocks [8]

adapts distributed web apps to adjust their distribution granularity to improve both performance and invocation costs. CWV is yet another middleware, albeit tailored for the realities of adapting mobile apps by transforming their code at runtime.

5.3 Executing Code in a Mobile Browser

Ours is not the only approach that moves server-side components and data to the client. Meteor [32], a JavaScript framework, transparently replicates given parts of a server-side MongoDB database at the client, so these parts can be used for offline operations. Browserify [2] enables a browser to use modulesin the same way as regular Node.js modules at the server. WebAssembly [22] provides portable low-level bytecode to execute components written in a variety of programming languages in a browser. WebAssembly has been enhanced with formal type and memory safety guarantees [22, 43]. Servers and browsers execute code in dissimilar ways. For one, browser-based execution of JavaScript code is typically slower than server-based execution, as browsers must handle event-callback functions in a strict sequential order. RT.js [12] prioritizes the execution of browser-based real-time jobs within the browser's event queue, so they meet real-time timeliness constraints.

6 Future Work and Conclusions

Although CWV works well in the domain of full-stack JavaScript web applications, its applicability extends only to monolingual environments. To remove this applicability constraint would require extending CWV to work in environments, in which distributed parts are written in different languages. Before a vessel can execute an insourced functionality written in a different language, the functionality has to be adapted for the vessel's execution language and environment. Common approaches to such adaptation are transpilation and virtual execution, which we discuss in turn. Transpiler source-to-source translates source language statements into target language statements as they are being interpreted. This execution strategy is typically used with languages that lack their own execution environment. For example, TypeScript is transpiled into JavaScript, immediately interpreted by the available JavaScript interpreter. Universal Virtual Machines provide interoperability between languages via a shared runtime, thus obviating the necessity for source-to-source translating the insourced code in different languages. With this approach, CWV could generate a regular client vessel in a language different from that of

the client. Then the insourced functionalities could be invoked via a Universal Virtual Machine incorporated into a mobile browser.

This paper has presented Communicating Web Vessels (CWV), a dynamic adaptation approach that improves the responsiveness of mobile web apps under the ever-changing execution environment of the web. The CWV's reference implementation offers full automation and a low performance overhead. Through its powerful dynamic program analysis and transformation, CWV correctly and efficiently adapts web apps for dissimilar execution conditions by moving and replicating app functionalities across execution sites at runtime. The realities of the mobile web will continue to be defined by the fragmentation of the mobile device market and the necessity to operate mobile devices over a variety of dissimilar mobile networks. If mobile apps are to execute efficiently and reliably in the presence of these realities, developers would need to create advanced software adaptations. We hope that when creating these adaptations, developers would benefit from our experiences with designing, implementing, and evaluating CWV.

Acknowledgments

References

[1] Trepn profiler. https://developer.qualcomm.com/forum/qdn-forums/increase-app-performance/trepn-profiler/27700, 2018.

[2] Tim Ambler and Nicholas Cloud. Browserify. In *JavaScript Frameworks for Modern Web Dev*, pages 101–120. Springer, 2015.

[3] Kijin An. Facilitating the evolutionary modifications in distributed apps via automated refactoring. In *Web Engineering*, pages 548–553. Springer International Publishing, 2019.

[4] Kijin An. Enhancing web app execution with automated reengineering. In *Companion Proceedings of the Web Conference 2020*, pages 274–278, 2020.

[5] Kijin An, Na Meng, and Eli Tilevich. Automatic inference of translation rules for native cross-platform mobile applications. In *Proceedings of the 5th IEEE/ACM International Conference on Mobile Software Engineering and Systems(MobileSoft) 2018*, 2018.

[6] Kijin An and Eli Tilevich. Catch & release: An approach to debugging distributed full-stack JavaScript applications. In *Web Engineering*, pages 459–473, 2019.

[7] Kijin An and Eli Tilevich. Client insourcing: Bringing ops in-house for seamless re-engineering of Full-Stack JavaScript Applications. In *Proceedings of The Web Conference 2020*, pages 179–189, 2020.

[8] Kijin An and Eli Tilevich. D-goldilocks: Automatic redistribution of remote functionalities for performance and efficiency. In *Proceedings of the 27th IEEE International Conference on Software Analysis, Evolution and Reengineering(SANER)*, 2020.

[9] Kijin An and Eli Tilevich. Communicating web vessels: Improving the responsiveness of mobile web apps with adaptive redistribution. In *Proceedings of the 21st International Conference on Web Engineering (ICWE 2021)*, pages 388–403. Springer, 2021.

[10] Earl T. Barr, Mark Harman, Yue Jia, Alexandru Marginean, and Justyna Petke. Automated software transplantation. In *Proceedings of the 2015 International Symposium on Software Testing and Analysis*, ISSTA 2015, pages 257–269, 2015.

[11] Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340. Springer, 2008.

[12] Christian Dietrich, Stefan Naumann, Robin Thrift, and Daniel Lohmann. Rt.js: Practical real-time scheduling for web applications. In *2019 Real-Time Systems Symposium (RTSS)*.

[13] Ning Ding, Daniel Wagner, Xiaomeng Chen, Abhinav Pathak, Y Charlie Hu, and Andrew Rice. Characterizing and modeling the impact of wireless signal strength on smartphone battery drain. In *ACM SIGMETRICS Performance Evaluation Review*, pages 29–40, 2013.

[14] Kevin Ellis, Armando Solar-Lezama, and Josh Tenenbaum. Sampling for bayesian program learning. In *Advances in Neural Information Processing Systems*, pages 1297–1305, 2016.

[15] Yu Feng, Ruben Martins, Osbert Bastani, and Isil Dillig. Program synthesis using conflict-driven learning. In *ACM SIGPLAN Notices*, volume 53, pages 420–435. ACM, 2018.

[16] Yu Feng, Ruben Martins, Jacob Van Geffen, Isil Dillig, and Swarat Chaudhuri. Component-based synthesis of table consolidation and transformation tasks from examples. In *ACM SIGPLAN Notices*, volume 52, pages 422–436. ACM, 2017.

[17] John K Feser, Swarat Chaudhuri, and Isil Dillig. Synthesizing data structure transformations from input-output examples. In *ACM SIGPLAN Notices*, volume 50, pages 229–239. ACM, 2015.

[18] Matthew Fredrikson and Benjamin Livshits. Zø: An optimizing distributing zero-knowledge compiler. In *23rd USENIX Security Symposium (USENIX Security 14)*, pages 909–924, San Diego, CA, 2014.

[19] Victor BF Gomes, Martin Kleppmann, Dominic P Mulligan, and Alastair R Beresford. Verifying strong eventual consistency in distributed systems. *Proceedings of the ACM on Programming Languages*, 1(OOPSLA):1–28, 2017.

[20] Marco Guarnieri, Petar Tsankov, Tristan Buchs, Mohammad Torabi Dashti, and David Basin. Test execution checkpointing for web applications. In *Proceedings of the 26th ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 203–214, 2017.

[21] Sumit Gulwani, Oleksandr Polozov, Rishabh Singh, et al. Program synthesis. *Foundations and Trends in Programming Languages*, 4(1-2):1–119, 2017.

[22] Andreas Haas, Andreas Rossberg, Derek L Schuff, Ben L Titzer, Michael Holman, Dan Gohman, Luke Wagner, Alon Zakai, and JF Bastien. Bringing the web up to speed with webassembly. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 185–200, 2017.

[23] Mohammad Ashraful Hoque, Matti Siekkinen, Kashif Nizam Khan, Yu Xiao, and Sasu Tarkoma. Modeling, profiling, and debugging the energy consumption of mobile devices. *ACM Computing Surveys (CSUR)*, (3):39, 2016.

[24] Krister Jacobsson, Håkan Hjalmarsson, Niels Möller, and Karl Henrik Johansson. Estimation of rtt and bandwidth for congestion control applications in communication networks. In *IEEE CDC, Paradise Island, Bahamas*. IEEE, 2004.

[25] Young-Woo Kwon and Eli Tilevich. Power-efficient and fault-tolerant distributed mobile execution. ICDCS '13. IEEE, 2013.

[26] Young-Woo Kwon and Eli Tilevich. Configurable and adaptive middleware for energy-efficient distributed mobile computing. In *MobiCASE*, pages 106–115. IEEE, 2014.

[27] Young-Woo Kwon, Eli Tilevich, and Taweesup Apiwattanapong. DR-OSGi: Hardening distributed components with network volatility

resiliency. In Jean M. Bacon and Brian F. Cooper, editors, *Middleware 2009*, pages 373–392. Springer, 2009.

[28] Ding Li, Shuai Hao, Jiaping Gui, and William GJ Halfond. An empirical study of the energy consumption of Android applications. In *2014 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pages 121–130. IEEE, 2014.

[29] Ding Li, Yingjun Lyu, Jiaping Gui, and William GJ Halfond. Automated energy optimization of HTTP requests for mobile applications. In *Software Engineering (ICSE), 2016 IEEE/ACM 38th International Conference on*, pages 249–260. IEEE, 2016.

[30] Reiner Marchthaler and Sebastian Dingler. *Kalman-Filter*, volume 30. Springer, 2017.

[31] Nima Nikzad, Octav Chipara, and William G Griswold. APE: an annotation language and middleware for energy-efficient mobile application development. In *Proceedings of the 36th International Conference on Software Engineering*, pages 515–526. ACM, 2014.

[32] Josh Robinson, Aaron Gray, and David Titarenco. Getting started with meteor. In *Introducing Meteor*, pages 27–41. Springer, 2015.

[33] Koushik Sen, Swaroop Kalasapur, Tasneem Brutch, and Simon Gibbs. Jalangi: A selective record-replay and dynamic analysis framework for JavaScript. In *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering*, pages 488–498, 2013.

[34] Jiasi Shen and Martin C Rinard. Using active learning to synthesize models of applications that access databases. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 269–285. ACM, 2019.

[35] Stelios Sidiroglou-Douskos, Eric Lahtinen, Anthony Eden, Fan Long, and Martin Rinard. CodeCarbonCopy. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, pages 95–105, 2017.

[36] Chungha Sung, Markus Kusano, Nishant Sinha, and Chao Wang. Static DOM event dependency analysis for testing web applications. In *Proceedings of the 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 447–459, 2016.

[37] Chungha Sung, Markus Kusano, and Chao Wang. Modular verification of interrupt-driven software. In *2017 32nd IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 206–216, 2017.

[38] Chungha Sung, Shuvendu K. Lahiri, Constantin Enea, and Chao Wang. *Datalog-Based Scalable Semantic Diffing of Concurrent Programs*, page 656–666. Association for Computing Machinery, New York, NY, USA, 2018.

[39] Eli Tilevich and Yannis Smaragdakis. J-Orchestra: Automatic java application partitioning. In Boris Magnusson, editor, *ECOOP 2002 – Object-Oriented Programming*, pages 178–204, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.

[40] Vassilios Tsaoussidis, Hussein Badr, Xin Ge, and Kostas Pentikousis. Energy/throughput tradeoffs of tcp error control strategies. In *Proceedings ISCC 2000. Fifth IEEE Symposium on Computers and Communications*, pages 106–112. IEEE, 2000.

[41] Jim Waldo, Geoff Wyant, Ann Wollrath, and Sam Kendall. A note on distributed computing. In *International Workshop on Mobile Object Systems*, pages 49–64. Springer, 1996.

[42] Xudong Wang, Xuanzhe Liu, Ying Zhang, and Gang Huang. Migration and execution of JavaScript applications between mobile devices and cloud. In *Proceedings of the 3rd annual conference on Systems, programming, and applications: software for humanity*, pages 83–84, 2012.

[43] Conrad Watt. Mechanising and verifying the webassembly specification. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 53–65, 2018.

Biographies

Kijin An is a Research Staff Member at Samsung Research. His research interests are in Software Engineering and Web Engineering. He received his Ph.D. degree from Virginia Tech and his Master's degree from POSTECH. In the past, he conducted research at Robotics Research at KIST.

Eli Tilevich is a Professor in the Dept. of Computer Science at Virginia Tech, where he leads the Software Innovations lab. Tilevich's research interests lie on the Systems end of Software Engineering, with a particular emphasis on distributed systems, mobile/IoT applications, middleware, automated program transformation, energy efficiency, privacy & security, and CS education. He has published over 120 refereed research papers on these subjects. Tilevich has earned a B.A. *summa cum laude* in Computer Science/Math from Pace University, an M.S. in Information Systems from NYU, and a Ph.D. in Computer Science from Georgia Tech.