



iPerfDetector: Characterizing and detecting performance anti-patterns in iOS applications

Sara Seif Afjehei¹ · Tse-Hsun (Peter) Chen¹  · Nikolaos Tsantalis²

Published online: 17 April 2019

© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

Performance issues in mobile applications (i.e., apps) often have a direct impact on the user experience. However, due to limited testing resources and fast-paced software development cycles, many performance issues remain undiscovered when the apps are released. As found by a prior study, these performance issues are one of the most common complaints that app users have. Unfortunately, there is a limited support to help developers avoid or detect performance issues in mobile apps. In this paper, we conduct an empirical study on performance issues in iOS apps written in Swift language. To the best of our knowledge, this is the first study on performance issues of apps on the iOS platform. We manually studied 225 performance issues that are collected from four open source iOS apps. We found that most performance issues in iOS apps are related to inefficient UI design, memory issues, and inefficient thread handling. We also manually uncovered four performance anti-patterns that recurred in the studied issue reports. To help developers avoid these performance anti-patterns in the code, we implemented a static analysis tool called iPerfDetector. We evaluated iPerfDetector on eight open source and three commercial apps. iPerfDetector successfully detected 34 performance anti-pattern instances in the studied apps, where 31 of them are already confirmed and accepted by developers as potential performance issues. Our case study on the performance impact of the anti-patterns shows that fixing the anti-pattern may improve the performance (i.e., response time, GPU, or CPU) of the workload by up to 80%.

Keywords iOS apps · Performance · Anti-pattern detection · Empirical study · Refactoring

Communicated by: David Lo, Meiyappan Nagappan, Fabio Palomba and Sebastian Panichella

✉ Tse-Hsun (Peter) Chen
peterc@encs.concordia.ca

Sara Seif Afjehei
seifaf@encs.concordia.ca

Nikolaos Tsantalis
tsantalis@encs.concordia.ca

¹ Department of Software Engineering and Computer Science, Software Performance, Analysis, and Reliability (SPEAR) Lab, Concordia University, Montreal, Canada

² Department of Software Engineering and Computer Science, Concordia University, Montreal, Canada

1 Introduction

The number of mobile phone applications (i.e., apps) has increased significantly in recent years. As of 2018, there are almost six million apps on the two most popular app stores: Google Play Store and Apple's App Store (Statista 2018). Such large number of apps makes the app stores highly competitive, since there exist many apps with similar functionalities. Hence, in addition to functional requirements, one important aspect that may affect users' perception of an app is the app's performance (Khalid et al. 2015; Liu et al. 2014).

Performance issues in mobile apps usually have a direct impact on the users. Prior studies (Khalid et al. 2015; Hu et al. 2018) found that performance issues are one of the most common complaints that app users have. However, existing research often focuses on studying performance issues in the system applications (e.g., enterprise applications or web servers) (Chen et al. 2014; Grechanik et al. 2012; Jin et al. 2012a; Nistor et al. 2015; Nistor et al. 2013b). Such performance issues may not be applicable to mobile apps, since mobile apps have different characteristics compared to system applications (Syer et al. 2013). For example, mobile phones have limited resources (e.g., battery and network), and mobile apps are usually UI-driven applications. A recent study by Liu et al. (2014) aimed to characterize performance issues in Android apps. However, the performance issues in Android may not be generalizable to iOS apps due to the differences in the two platforms and the used programming languages (Hu et al. 2018). Moreover, most prior research on mobile app APIs or code smells focused only on the Android platform (Martin et al. 2017); even though iOS is the second largest mobile platform in the world.¹

Therefore, to provide a better understanding of performance issues in mobile apps, we focus our study on iOS. We conducted an empirical study on 225 performance issue reports in four open source iOS apps written in the Swift programming language (i.e., Firefox, WordPress, Wire, and Chart) to study common types of performance issues. We also manually uncovered and documented four types of performance anti-patterns that recurred in the studied issue reports. Finally, we implemented a static analysis tool, called iPerfDetector, to detect the uncovered anti-patterns in Swift files. We collaborated with our industrial partner, and evaluated iPerfDetector on three commercial and eight open source apps. In total, iPerfDetector detected 34 performance anti-pattern instances, where 32 of them we manually verified as true positives. We reported these 32 anti-pattern instances to developers, and 31 of them are confirmed and accepted by developers. Our case study on the performance impact (i.e., in terms of response time, GPU, and CPU usage) of the anti-pattern shows that fixing the anti-pattern may improve the performance of the workload (e.g., scrolling a table in the app) by 5.8% to 80%. iPerfDetector received positive feedback from both commercial and open source developers, and it is now used by our industrial partner to ensure the performance of their iOS apps.

To the best of our knowledge, this is the first study on performance issues in iOS apps. In summary, we answer the following research questions:

- **RQ1: What are the common causes of performance problems in iOS apps?** We discuss our process of selecting and manually studying performance issue reports in open source iOS apps. We found three common types of performance issues in the four open-source apps that we studied: inefficient UI design, memory issues, and inefficient thread handling.

¹<https://www.statista.com/statistics/276623/number-of-apps-available-in-leading-app-stores/>

- **RQ2: What are the common iOS performance anti-patterns?** We document recurrent iOS performance anti-patterns that we found during our manual study. We found that accessing data model on the UI thread, updating UI controls on background threads, applying UI blurring effect inefficiently, applying UI transparency effect inefficiently, and object references causing retain cycles are the most common anti-patterns. We also provide a detailed discussion on the above-mentioned iOS performance anti-patterns.
- **RQ3: How many iOS performance anti-patterns can we detect and verify in both commercial and open source apps?** We implement a static analysis tool, iPerfDetector, to detect instances of the uncovered anti-patterns. iPerfDetector detected 34 anti-pattern instances in eight open source and three commercial apps, where 31 of them are already confirmed and accepted by developers as potential performance issues.
- **RQ4: What is the performance impact of the iOS anti-patterns?** We conduct a case study on the anti-patterns introduced in RQ2, to measure their performance impact. Our case study shows that fixing the anti-patterns may improve the performance by 5.8% to 67% in terms of GPU usage, 34.15% in terms of CPU usage, and by over 80% in response time.

Paper Organization Section 2 briefly introduces the background of iOS app development. Section 3 discusses the motivation, approach, and results of our research questions. Section 4 discusses threats to validity of our study. Section 5 surveys related work. Finally, Section 6 concludes the paper.

2 Background

iOS is a closed-source operating system developed by Apple for mobile devices such as iPhone and iPad. To date, iOS is one of the most widely used mobile operating systems in the world, and iOS apps generate much higher revenues compared to Android apps (Statista 2018). To provide better usability and functionalities in the app, developers usually rely on interacting with the APIs provided by the iOS Software Development Kit (SDK). Initially, iOS apps were all implemented in the Objective-C programming language. In 2014, Apple introduced the Swift programming language as an alternative for iOS app development. Compared to Objective-C, Swift offers a cleaner syntax that makes the apps easier to maintain (Apple 2018a). Since then, Swift has become more popular among developers compared to its predecessor (RedMonk 2018). Hence, in this paper, we focus our study on the iOS apps that are written in Swift.

iOS apps are often developed using the Model-View-Controller (MVC) design pattern. The MVC pattern separates the app's data access and business logic from the visual presentation. The layers in MVC help abstract the underlying device differences, such as screen sizes, and simplifies app development. Below, we briefly discuss the background of iOS development based on three major software layers in an app: model, view, and controller.

Model Layer Model layer is a set of UI independent objects, responsible for managing the persisted objects in the database. Developers often use the Core Data persistent framework that is provided by iOS SDK to manage object persistence in order to access the model layer (Guide 2017). By using Core Data, developers can directly access or modify object states in a database through API calls in the app's model layer, which reduces the complexity of the code (Apple 2018d).

View Layer View layer is responsible for displaying the user interface (UI) components and for users to interact with the app (e.g., by tapping on buttons). Developers can call subclasses of the *UIKit* class (e.g., *UILabel* or *UIButton* class) to design and organize UI components. Developers may also call APIs in the iOS SDK to customize UI components or add different UI effects (e.g., transparency or blurring). Since mobile apps are UI-driven, by default, every function call will be executed on the main thread (i.e., UI thread) (Apple 2018b). If the main thread is blocked or delayed due to heavy computation, users may experience poor UI responsiveness.

Controller Layer The controller layer implements the business of an app. For example, how often should the UI be updated/refreshed, or what should be stored in the model layer when users tap a button. In other words, the controller layer translates actions in the view layer to the corresponding actions in the model layer and vice versa (Apple 2018e). To ensure that the UI component is responsive when processing the business logic, iOS provides multi-threaded APIs. Heavy operations (e.g., complex calculation) or long-running processes (e.g., access the persisted objects in the model layer) can be executed on a background thread, instead of on the main thread.

3 iOS Performance Anti-Patterns and Their Impact

In this section, we answer the research questions we discussed in Section 1.

Figure 1 shows an overview of our study. We first manually study iOS apps issue reports to find the common root causes of performance issues (RQ1). Second, we extract and characterize iOS performance anti-patterns (RQ2). Third, we detect performance anti-patterns in commercial and open-source apps, and verify our findings with developers (RQ3). Finally, we study the performance impact of anti-patterns (RQ4). Below, we discuss each RQ in detail using the following template:

Motivation: We discuss the motivation of each RQ.

Approach: We describe the approach we take to answer the research question.

Results: We discuss our experiment results and answer the research question.

RQ1: What are the common causes of performance problems in iOS apps?

Motivation

In order to understand the common causes of performance problems in iOS apps, we manually study issue reports in open-source iOS apps. Understanding the common performance problems may help developers avoid the problems and inspire future research.

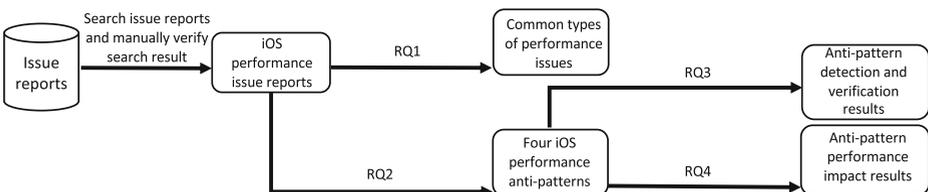


Fig. 1 An overview of our study

Table 1 An overview of the studied apps

System	Category	Num. of stars on GitHub	LOC	Num. of issue reports	Num. of perf bug reports
Firefox	Utilities	8K	39K	1.1K	78
Wordpress	Productivity	2.3K	85K	4.6K	82
Wire	Business	2.4K	61K	308	32
Charts	Chart library	18K	31K	2.9K	33

Approach

We conducted our manual study on open source iOS apps. We selected 21 popular open source Swift iOS apps according to Medium (2016) and the top 20 trending open source Swift repositories on GitHub (2018). Then, we applied four selection criteria on the apps. First, the selected apps should be related to iOS mobile app development, because Swift is a general purpose language that can be used for tasks such as implementing web servers. Second, a candidate app should use an issue report system, since we need to study the issue reports to uncover the types of performance issues. Third, the candidate app should contain issue reports that are related to performance issues. Finally, the app should be actively maintained. Since mobile APIs are constantly evolving (McDonnell et al. 2013), we want to ensure that new problems related to API usages are included in our study. We ended up with four open source iOS apps that met our criteria. Table 1 shows an overview of the studied apps. Mozilla Firefox client app (or simply Firefox) is a free and open-source web browser written in Swift.² We used the master branch of the Firefox iOS browser, which is written in Swift 4.2. WordPress is an open-source content management system (CMS).³ Wire is a cross-platform and encrypted instant messaging client.⁴ Charts is a library for drawing various types of charts for iOS apps.⁵ In general, the studied apps cover different categories, are popular (with 2.3K to 18K stars on GitHub), and are large in size (up to 39K LOC).

We used the keywords: “slow”, “performance”, “thread” and “memory” to search for performance-related issue reports. These keywords are commonly used in prior studies or are known as common performance problems (Liu et al. 2014; Jin et al. 2012b; Smith and Williams 2001; Chen et al. 2016b). In total, we found 960 issue reports in the studied apps that contain the keywords. We manually go through all 960 issue reports to identify the relevant issue reports. We first remove all the issue reports that do not have a fix (e.g., marked as *unresolved* or won’t fix). Then, we remove the false positives that are caused by the keywords we used. For example, by using the keyword “thread”, we found that in some cases, developers may be referring to “issue thread”. This filtering process was done by the first author of the paper and verified by the second author. After the manual filtering process, we ended up with 225 performance issue reports for our manual study.

Table 1 shows the number of performance issue reports that we studied in each app. To categorize the studied performance issues based on their effect, we manually studied all of the 225 issue reports and all associated information (e.g., pull requests, code changes, and developers’ discussions). We determined the categories of the anti-patterns based on their root causes. Our categorization process is as follows. We started our manual study with no

²<https://github.com/mozilla-mobile/firefox-ios>

³<https://github.com/wordpress-mobile/WordPress-iOS>

⁴<https://github.com/wireapp/wire-ios>

⁵<https://github.com/danielgindi/Charts>

specific category in mind and we took a note on the root cause of the problem. After the first author manually studied the issue reports, she created the categories based on the root causes that we found, and the second author verified the categorization.

Results

In general, we found three common types of performance issues in the studied iOS apps: inefficient UI design, memory issues, and inefficient thread handling. Note that since one issue may be assigned to more than one type (e.g., an issue report may be related to two different types of performance issues), the accumulated percentage may not be exactly 100%. There is also around 15% of the issues that do not belong to these three types. We found that such issues are often related to app-specific performance optimization, such as using more efficient data structures, reducing the number of log lines that are printed (Chowdhury et al. 2018), or optimizing database queries.

Below, we summarize the issues that we found in each type. To encourage future research on this topic, we also release our manually annotated data online (Afjehei and Chen 2018).

Inefficient UI Design We found that 20% of the issue reports tried to address inefficient UI design. One common UI-related performance issue that we studied happens when users scroll through a table. For example, developers may be using computationally intensive text styling in table cells (Wire bug 1751), dynamically calculating the height of each cell in a table instead of using a fixed height (Wire bug 5672), or displaying emojis in table cells (Wire bug 267). Such heavy UI-related operations in table cells may make scrolling a table sluggish. Another common problem that we saw is related to applying heavy UI effects. For example, developers may be applying UI blurring effect repetitively in a loop (Firefox bug 1221118). Since the main thread is responsible for UI rendering, applying heavy UI effects in a repeating fashion on the main thread will make the app temporarily unresponsive (i.e., the main thread is blocked to compute the effects). In summary, many performance issues that we saw are related to inefficient usage of iOS's UI-related APIs.

RQ2 further discusses the performance anti-patterns that are related to computing UI effects.

Memory Issues We found that 34% of the studied issue reports are related to memory issues. Furthermore, about 44% of the memory related issues address memory leaks. In particular, retain cycles are the most common cause of memory leaks that we found. Retain cycle happens when two or more objects keep references to each other (i.e., a reference cycle). This may prevent the objects from being garbage collected even if they cannot be accessed elsewhere from the heap. Other types of memory problems include memory leaks (e.g., in Firefox bug 1278006, UI objects are not released even if they are not visible to the user anymore) and out of memory errors (e.g., WordPress bug 7892).

Inefficient Thread Handling We found that threading-related issues are the most common causes of performance issues in our studied issue reports (36%). As mentioned in Section 2, there is a main thread (i.e., UI thread) and background threads in iOS. We found that most issues that we studied are related to inefficient uses of iOS's threading APIs. For example, developers may forget to execute heavy computation (e.g., data accesses) in a background thread, which may result in causing the UI to be temperately unresponsive (e.g., as discussed in WordPress 578 and Firefox 532).

We found that there are three main types of performance issues: inefficient UI design (20%), memory issues (34%), and inefficient thread handling (36%). We also found that most problems we studied are related to how developers use iOS-specific APIs (e.g., UI effect or threading).

RQ2: What are the common iOS performance anti-patterns?

Motivation

During our manual study of the iOS performance issue reports, we found several recurring code patterns that are common causes of the studied iOS performance issues. Hence, in this RQ, we document the recurring code patterns that we uncovered into four iOS performance anti-patterns.

Approach

In our study, we tried to derive the performance anti-patterns that are specifically related to iOS development and occur several times. Hence, we excluded other non-iOS and application-specific performance problems in the studied issue reports (e.g., displaying emojis, adding too many logs, tuning database queries, switching to a faster data structure etc). To derive the anti-patterns, we manually analyzed the fixes in the issue reports and took notes regarding code changes (e.g., what did developers change to fix the problem). Then, we went through the note and manually summarized and extracted the code changes as the anti-pattern.

Results

Table 2 summarizes the manually-uncovered iOS performance anti-patterns and the number of anti-pattern instances we found. Below, we discuss each anti-pattern in detail. To provide

Table 2 Description of the manually-uncovered iOS performance anti-patterns and their corresponding category

Category	Anti-pattern	Description	Abbr.	No of issues
Inefficient thread handling	Accessing data models on the UI thread	Developers access the model layer (i.e., access data) in the UI thread, which may cause bad user experiences due to unresponsive UI.	DM	17
	Updating UIs on background threads	UI updates (e.g. update label text) that are executed on a background thread, which may lead to app crashes.	UIBG	30
Inefficient UI design	Applying UI blurring effect inefficiently	Adding computationally expensive blurring effect to complex UI components (e.g., table with multiple cells).	BLEF	4
	Applying UI transparency inefficiently	Adding computationally expensive transparency effect to complex UI components (e.g., table with multiple cells).	TREF	4
Memory issues	Retain cycle	The condition when two objects keep a reference to each other and are retained, which creates memory leaks (Apple 2012).	RETAIN	22

more detailed information and breakdown of each studied iOS performance anti-pattern, we discuss each pattern using the following template:

Description Description of when and how the anti-pattern would take place.

Example An example of the anti-pattern from the real-world studied apps.

Examples of Developer Awareness We summarize some of developers' discussion that we found in the issue reports. We also searched on developer forums (e.g., Stack Overflow) and documents (e.g., Apple official documents) and summarize our findings on how the anti-pattern may affect other developers.

Possible Solutions We discuss possible solutions to resolve the anti-pattern.

Accessing Data Model on the UI Thread

Description Mobile apps that implement the MVC design pattern usually require a model layer to manage user information or application-specific data. In iOS app development, developers often use Core Data for data access (Apple 2018g). Core Data is an object persistent framework provided by Apple that allows developers to manage the model layer and to interact with persisted objects (Guide 2017). However, such data accesses can be time-consuming and computationally intensive. Hence, if Core Data is used inside the UI thread (i.e., the main thread), the UI may be temperately unresponsive until the data access is finished.

Example As an example, there is a discussion on fixing this performance anti-pattern on WordPress.⁶

The following function, *save*, is persisting the changed data to the model layer using Core Data (line 4).

```
(void) save {
    NSError *error;
    // saving changes to the model layer by calling
    // managedObjectContext, which is the API for using Core
    // Data
    if (![self managedObjectContext] save:&error) {
        DDLogError(@"Unresolved Core Data Save error %@ %@",
            error, [error userInfo]);
        exit(-1);
    }
}
...
(void) remove { // will be executed on the main thread
    ...
    [self save]; // calling data access through Core Data
}
```

⁶<https://github.com/wordpress-mobile/WordPress-iOS/pull/578>

However, this *save* function is called inside another function named *remove* (line 12), which is executed on the main thread. As a result, the UI will be temperately unresponsive due to having heavy data accesses on the main thread (i.e., because of a chain of function calls that result in executing the *remove* function).

Developers Awareness As mentioned in Apple’s iOS development document (Apple 2018g), any data processing should be avoided on the main thread. However, we still found that sometimes such anti-patterns may be difficult to find due to the complexity of the software design and interactions of function calls⁷ (Chen et al. 2014, 2016a). Namely, developers may not notice that the function that they call may eventually result in data access.

Possible Solutions One way to address this iOS performance anti-pattern is that developers can leverage APIs that are provided by Apple to execute the data access on a background thread asynchronously (Apple 2015). Another solution is to add a separate layer in the code to take care of *NSManagedObjectContext(s)* (i.e., call to Core Data). For example, developers may use the parent-child managed object contexts, where changes to the children objects will only be persisted when the parent object is persisted. Therefore, developers only need to make sure that calls to the parent objects are executed in background threads (i.e., developers do not need to make scatter code changes to make sure every child access is executed in the background thread). The first solution would require fewer changes to the code, but may introduce maintenance difficulties due to scattering code changes. On the other hand, the second solution may require significant code refactoring for the added abstraction layer. Developers may need to decide which solution is more suitable for the design of their app after considering the trade-offs, as indicated by WordPress developers.⁸

Updating UI Controls on Background Threads

Description We observed that sometimes developers may perform inappropriate performance optimization by moving heavy UI operations to a background thread. As an example, Firefox developers had a discussion regarding moving a UI rendering computation to a background thread in order to make the UI more responsive.⁹ However, such optimization may result in some problems. As part of the iOS architectural design, updating UI (e.g., resize or refresh) on a thread other than the main thread can be problematic. According to Apple’s development guideline, this anti-pattern may lead to problems such as missed UI updates, incorrect UI displays, or even crashes (Apple 2017a). Developers also discuss that such issues can cause delays in UI updates.¹⁰

Example Consider an example from WordPress.¹¹ In the code snippet below, there is a class function *resizeGalleryImageURL*, which resizes and scales the image (the implementation detail is omitted for better readability). The class function *resizeGalleryImageURL* is then used in another function named *formatContentString* (line 11).

⁷<https://github.com/wordpress-mobile/WordPress-iOS/pull/578>

⁸<https://github.com/wordpress-mobile/WordPress-iOS/pull/578>

⁹<https://github.com/mozilla-mobile/firefox-ios/pull/1215/commits/4ba952a58bf34ecfa555a8058f7c9a20c154997d>

¹⁰<https://stackoverflow.com/questions/28137380/updating-ui-from-background-thread-swift>

¹¹<https://github.com/wordpress-mobile/WordPress-iOS/pull/7864>

```

// resize an image
public class func resizeGalleryImageURL(_ string: String,
    isPrivateSite isPrivate: Bool) -> String {
    ...
}

...

public class func formatContentString(_ string: String,
    isPrivateSite isPrivate: Bool) -> String {
    ...
    // will result in resizing the image
    content = resizeGalleryImageURL(content, isPrivateSite:
        isPrivate)
    return content
}

```

The function *formatContentString* is called in the function *sanitizeCommentContent* (line 4 in the code snippet shown below), which is then called inside the *performBlock* of *managedObjectContext* (line 13). *managedObjectContext.performBlock* is a way of executing operations in an asynchronous fashion (i.e., in a background thread, when a private managed object context is defined). As a result, this caused WordPress to crash due to resizing an image in background threads.¹²

```

(NSString *)sanitizeCommentContent:(NSString *)string
    isPrivateSite:(BOOL)isPrivateSite{
    ...
    // calling formatContentString
    content = [RichContentFormatter formatContentString:content
        isPrivateSite:isPrivateSite]
    ...
}

...

[self.managedObjectContext performBlock:^(
    ...
    // May result in calling resizeGalleryImageURL in a
    background thread, which led to application crashes in
    WordPress
    remoteComment.content = [self
        sanitizeCommentContent:remoteComment.content
        isPrivateSite:isPrivateSite];
    ...
}];

```

¹²<https://github.com/wordpress-mobile/WordPress-iOS/pull/7864>

Examples of Developer Awareness In September 2017, Apple released a dynamic analysis tool called the Main Thread Checker that can detect UI API calls in background threads and provides warnings on the Xcode console (Apple 2017a). Even though such a tool is available, during the manual study, we still found that developers may not fully utilize the tool, and thus, resulted in app crashes. As an example, we counted the number of reports that were opened after the tool was released (by comparing the release date of the tool and the issue creation date), and found that about half of the anti-pattern instances in WordPress were reported by users after the tool was released. Namely, developers did not fully utilize the tool during development, so the anti-pattern instances remained unfixed when the app was released. One reason may be that the Main Thread Checker is only enabled when running apps with the Xcode debugger (Apple 2017a). Developers may not exercise the anti-pattern instances in the code when using the debugger, so the instances may not be detected. Compared to the dynamic approach provided by Main Thread Checker, iPerfDetector provides a lightweight static approach and does not introduce overhead during testing. Another reason may be that such warnings are difficult to notice due to hundreds or even thousands of other messages on the console (Jiang et al. 2008; Chen et al. 2017).

Possible Solutions All UI operations should be executed on the main thread. This can be done by dispatching the UI update operation to the main thread by calling the update operations in code blocks such as `DispatchQueue.main.sync{...}` (the code inside the code block would be executed in the main thread).

Applying UI Blurring Effect Inefficiently

Description As mentioned in Section 2, UI-related tasks (e.g., refresh or update) are processed on the main thread (i.e., UI thread). As a result, any time-intensive operation in the main thread may cause the UI to be temporarily non-responsive, which results in bad user experience. We found that blurring effect can be computationally expensive, especially when the blurring effect is applied in a repeated fashion (e.g., applying the effect on each cell of a table or in a loop). Applying such blurring effect inefficiently may result in noticeable delays in the UI.

Example Consider an example from Firefox¹³. In the code snippet below, we can see that an object instance of class `UIVisualEffectView` is instantiated (defines a blurring effect in line 4) and is referred by a variable named `backgroundEffect`. In the second function `init()` (line 13), where a cell is initialized, we can see that `backgroundEffect` is added (line 15) to each cell in a table (line 18) when initializing the table object. Hence, iOS needs to recalculate the blurring effect for each table cell. Applying the blurring effect repetitively is computationally expensive and can result in noticeable UI delays, especially when users are scrolling the table.

¹³https://bugzilla.mozilla.org/show_bug.cgi?id=1191058

```

// configuring the blur effect
lazy var backgroundEffect: UIVisualEffectView? = {
    ...
    let blur = UIBlurEffect(style: UIBlurEffectStyle.Light)
    let vib = UIVibrancyEffect(forBlurEffect: blur)
    let effect: UIVisualEffectView? =
        DeviceInfo.isBlurSupported() ?
            UIVisualEffectView(effect: blur) : nil
    ...
    return effect
}()

...

override init(frame: CGRect) {
    ...
    imageWrapper.addSubview(backgroundEffect)
    ...
    // adding blur effect to every cell in the table
    contentView.addSubview(imageWrapper)
    ...
}

```

Examples of Developer Awareness In the above-mentioned example from Firefox, developers complained that “the blur we render for each top site item is causing noticeable slowdowns in frame rate of the application”.¹⁴ We also found many other developers complained about the potential problems of applying blurring effect inefficiently during the manual study. In our study, we found that 33% of the UI-related performance issues reported in Firefox are related to applying UI blurring effect inefficiently.

Possible Solutions To fix this anti-pattern, developers sometimes choose to remove the blurring effect if the effect is not important to the UI design. For example, one Firefox developer in the above-mentioned issue report asked “just curious, how important is it to use blurring instead of a flat or alpha faded color?”.¹⁵ Another approach to fix the anti-pattern is to calculate the blurring effect once and reuse the calculated effect for every cell, instead of recomputing the effect for each cell. However, sometimes developers may apply different blurring parameters to each cell in a table. In addition, if the screen needs to be frequently refreshed, calculating the blurring effect for every UI frame may still be expensive. In such cases, removing the blurring effect may be a better performance optimization approach.

Applying UI Transparency Effect Inefficiently

Description Similar to blurring, we found that UI transparency effect may be computationally expensive if used inefficiently (e.g., applied repeatedly, which requires more computation). Such inefficient UI design may result in unresponsive UI and bad user experience.

¹⁴https://bugzilla.mozilla.org/show_bug.cgi?id=1191058

¹⁵https://bugzilla.mozilla.org/show_bug.cgi?id=1191058

Example Consider an example from Firefox. The UI transparency effect is configured and returned as the variable *faviconView* (an instance of the *UIImageView* class). Then, *faviconView* is used in the *init* function of a table cell (line 14), and the transparency effect is added to each cell in the table (i.e., by calling *addSubview(faviconView)*). When the table is refreshed (e.g., when scrolling), the transparency effect will be recomputed for each cell in the table and result in unresponsive UI.

```
lazy var faviconView: UIImageView = {
    ...
    // setting the transparency level
    faviconView.layer.borderColor = UIColor(white: 0, alpha:
        0.1).CGColor
    ...
    return faviconView
}()

...

override init(style: UITableViewCellStyle, reuseIdentifier:
    String?) {
    ...
    // adding transparency effect to each table cell
    contentView.addSubview(faviconView)
    ...
}
```

Examples of Developer Awareness In the studied issue reports, we found that some developers did not know the performance overhead of the transparency effect until users reported the problems. Namely, developers did not completely understand the performance profile of the app. For example, a WordPress developer was informed by a user that: “[d]rawing is improved when views that have solid backgrounds are set to opaque.”¹⁶ The manual observation also highlights the needs of developing research tools to help developers improve the performance of their apps.

Possible Solutions Similar to applying blurring effect inefficiently, the approach to fix this anti-pattern is to remove the UI transparency effect when the effect is not important to the UI design. As an example, Firefox developers even mentioned that they think avoiding partial transparencies, whenever the effect is not required, would be the best way to ensure the performance of the app.¹⁷ Other developers on Stack Overflow also discussed that the overhead to compute transparency is high and should be avoided if possible.¹⁸

¹⁶<https://github.com/wordpress-mobile/WordPress-iOS/issues/419>

¹⁷<https://github.com/mozilla-mobile/firefox-ios/pull/1215/commits/4ba952a58bf34ecfa555a8058f7c9a20c154997d>

¹⁸<https://stackoverflow.com/questions/9270723/what-is-the-better-way-to-set-an-uiview-backgroundcolor-to-transparent>

We found that the most common performance anti-patterns in iOS are: accessing data model on the UI thread, updating UI controls on background threads, applying UI blurring effect inefficiently, applying UI transparency effect inefficiently, and object references causing retain cycles.

RQ3: How many iOS performance anti-patterns can we detect and verify in both commercial and open source apps?

Motivation

We found that many of the anti-patterns that we studied are related to the complex calling relationship among the user-defined functions and iOS-specific APIs that may be difficult to find without tool assistance (Liu et al. 2014; Chen et al. 2016a). For example, developers may not notice that, after a chain of function calls, a function on the main thread will result in data accesses. Hence, to help developers improve the quality of their apps, we create a static analysis tool to detect the anti-patterns. In this RQ, we want to study whether our tool can detect the anti-patterns with a high precision.

Approach

We implemented a static analysis tool, called iPerfDetector, to detect instances of the uncovered anti-patterns. iPerfDetector analyzes the abstract syntax tree (AST) of the Swift code in iOS apps to detect the anti-patterns. iPerfDetector first uses an open source tool called SwiftAST (SwiftAST 2018) to parse the Swift code and generate the ASTs. Then, iPerfDetector applies different detection algorithms on the ASTs to detect the manually-uncovered iOS performance anti-patterns. Below we provide more details on how iPerfDetector detects each anti-pattern. Note that iPerfDetector focuses on detecting the types of performance anti-patterns that existing tools do not. Since prior tools (e.g., Facebook's Infer (2017) and Apple's Instruments (2018c)) focus on detecting memory-related issues, we implement iPerfDetector to detect inefficient thread handling and inefficient UI design.

Detecting Accessing Data Model on the UI Thread Developers use Core Data APIs by calling the class *NSManagedObjectContext* that is provided by the iOS SDK to interact with the persisted objects (Apple 2017b). Hence, every data access is made through this specific class. iPerfDetector first traverses all the classes in an app searching for objects that are instances of *NSManagedObjectContext*. Then, iPerfDetector applies points-to analysis to identify all the functions that contain data accesses (e.g., calling functions such as *save*, *fetch*, *insert*, or *delete* on the *NSManagedObjectContext* object). After this step, we obtain a list of functions that contain data access calls. Then, iPerfDetector constructs an interprocedural call graph, and traverses all the functions that are executed in the main thread. Since functions are executed in the main thread by default, iPerfDetector decides that a thread is executed in the main thread if the thread is *not executed* on background threads, such as not within in the code blocks of *DispatchQueue.global.async* and *performBackgroundTasks*. Finally, if any executed functions in the main thread may result in data access, an instance of this anti-pattern is detected by iPerfDetector.

Detecting Updating UI Controls on Background Threads To detect this anti-pattern, iPerfDetector traverses the entire source code and identifies all the instantiated UI objects (e.g., buttons, text fields, and labels). After this step, iPerfDetector obtains the locations of all the UI objects and their corresponding functions. Then, iPerfDetector constructs an

inter-procedural call graph. *iPerfDetector* traverses the call graph and looks for functions that are executed in background threads (e.g., called within the code block of *perform-BackgroundTask* or *DispatchQueue.global.async{...}*, which are the APIs provided by the iOS SDK for executing tasks in background threads). An anti-pattern instance is detected if we see that a UI-related object or function is called within the code blocks that will be executed in a background thread.

Detecting Applying UI Blurring Effect Inefficiently We found that there are two situations where applying blurring effects may result in bad user experience. The first situation is when an object instance of *VisualEffectView* (i.e., for setting visual effects) is added to a cell. A cell in iOS UI can be represented by any class that is extended from classes *TableViewCell* (a type of cell that repeats in a table vertically) or *CollectionViewCell* (a type of cell that repeats in a table horizontally). The second situation is when an object instance of *VisualEffectView* is added to another view (i.e., objects that are instances of *UIView*, which manages the contents on the screen) inside a loop. In these two situations, the blurring effect is computed multiple times (e.g., once for each cell) whenever the frame is updated (e.g., when scrolling a table).

To detect this anti-pattern, *iPerfDetector* first traverses all the source code files searching for object instances of *VisualEffectView* with the blurring effect turned on. *iPerfDetector* detects an instance of anti-pattern by using taint analysis. Namely, *iPerfDetector* checks the calls/usages of the object instance of *VisualEffectView*, and verifies whether the object instance will eventually be added as a subview to other object instances that represent cells in iOS UI (i.e., sub-classes of *TableViewCell* or *CollectionViewCell*). If the *VisualEffectView* object instance is used in a loop (e.g., *for* or *while*), an anti-pattern instance is also detected.

Detecting Applying UI Transparency Effect Inefficiently To detect this anti-pattern, *iPerfDetector* first traverses all the source code files searching for objects that are instances *UIView* and have transparency configured (i.e., the alpha value is larger than 0 and less than 1). Then, similar to detecting inefficient blurring effect, *iPerfDetector* uses taint analysis to verify where the partially transparent object instance is called. Namely, *iPerfDetector* tracks the usages and calls to object instance, and an anti-pattern instance is detected if the object instance is used in classes that are of the type *CollectionViewCell* or *TableViewCell*, and the object instance is added as a subview to the cell. An anti-pattern instance is also detected if the partially transparent object instance is used in a loop.

Furthermore, in order to evaluate *iPerfDetector*, we apply it on both commercial and open source apps, to find out how many anti-patterns *iPerfDetector* can detect. Then we report our findings to developers and ask them to verify our results.

We collaborated with a company named Seeb Smart Solutions,¹⁹ and used *iPerfDetector* to help improve the quality of their apps. Table 3 shows the applications that we used to evaluate *iPerfDetector*. We evaluated *iPerfDetector* on three apps from our industrial partner. App1 is an application for travel guides (similar to TripAdvisor to some extents). App2 is a

¹⁹<https://www.linkedin.com/company/seeb/>

Table 3 Apps that we use to evaluate iPerfDetector

App Name	Category	LOC	Availability	Version
Firefox	Utilities	145K	Open Source	8.0
WordPress	Productivity	84K	Open Source	8.5
Wire-iOS	Business	60K	Open Source	3.3
Charts	Library	31K	Open Source	3.0.3
Yep	Social Network	71K	Open Source	1.3
Sync	Library	8K	Open Source	3.2.3
TSWeChat	Social Network	8.5K	Open Source	1.0
SugarRecord	Library	2.5K	Open Source	1.0
App1	Travel	5 to 25K	Commercial	Latest
App2	Travel	15 to 50K	Commercial	Latest
App3	News	5 to 25K	Commercial	Latest

ride sharing application (e.g., similar to Uber) that is commonly used in Iran. Finally, App3 is a news app that informs users regarding technology news over the globe. In addition to commercial apps, we also applied iPerfDetector on the latest version of four additional open source iOS apps (i.e., eight open source apps in total). The categories of these apps range from social network to business and iOS libraries. These four apps are popular on Github (i.e., more than 2K stars) and are relatively large in size (more than 2.5K LOC). We did not select these four apps for our manual study in RQ1 due to a lack of performance-related issue reports. We applied iPerfDetector to the latest versions of the studied apps, and we focused our study on the source code files (i.e., we excluded the ones that were detected in the test files).

To verify the detection results, we took a two-step approach. First, we manually studied all of the detected anti-pattern instances to examine whether they are true positives or not. Then, we reported the detected anti-pattern instances to developers in order to receive their feedback.

Results

Table 4 shows the performance anti-pattern detection results. Note that the detected instances in WordPress, Wire-iOS, Firefox, and Charts are new issues that were not discovered/discussed in the manually studied issue reports. In total, iPerfDetector detected 34 anti-pattern instances in 11 studied apps. The execution time of running iPerfDetector is around three minutes or less per app, depending on the size of the app. We executed iPerfDetector on a desktop machine with 16 GB of memory and 2.6 GHz Intel-Core i5 CPU. Note that some anti-patterns may not be applicable in the studied iOS app (e.g., an app does not have data models or it is an iOS framework that does not have a UI), so we show the detection result for such anti-patterns as “N/A” in the table.

In our manual verification, we found only two false positives among all the 34 detected anti-pattern instances (i.e., a precision of 94%). The two false positives were both DM anti-pattern that were detected in Sync. We found that, in these two specific cases, Sync developers added some app-specific code to avoid data accesses in the main thread. The code checks whether the data access will be executed in the main thread, and if so, an error

Table 4 iOS performance anti-pattern detection result

App Name	DM		TREF		BLEF		UIBG	
	Detected	Verified	Detected	Verified	Detected	Verified	Detected	Verified
Wordpress	2	2	0	0	0	0	0	0
Wire-iOS	2	2	0	0	0	0	0	0
Firefox	N/A	N/A	1	1	1	1	0	0
Charts	N/A	N/A	0	0	0	0	0	0
Sync	18	16	N/A	N/A	N/A	N/A	0	0
Yep	N/A	N/A	3	3	0	0	0	0
TSWeChat	N/A	N/A	2	2	0	0	0	0
SugarRecord	1	1	N/A	N/A	N/A	N/A	0	0
App1	N/A	N/A	N/A	N/A	1	1	0	0
App2	N/A	N/A	0	0	1	1	0	0
App3	2	2	0	0	0	0	0	0

“N/A” means that the anti-pattern is not applicable to the studied iOS app (e.g., the app does not have data models or it is an iOS framework that does not have a UI). *Detected* shows the number of anti-pattern instances that are detected by iPerfDetector. *Verified* shows the number of anti-pattern instances that are manually verified to be true positives by the authors

will be thrown. Since iPerfDetector did not consider such app-specific code in the detection algorithm, these two anti-pattern instances were detected as false positives. We reported the remaining 32 detected anti-pattern instances (after removing the two above-mentioned false positives) to developers. We contacted the developers by creating issue reports on bug tracking systems (e.g., Bugzilla or Jira), or by sending emails to developers. Overall, we received positive responses for the 31 reported anti-pattern instances (we are still waiting for responses of one reported anti-pattern instance in SugarRecord). Developers acknowledged the reported problems and expressed interest in our detection tool.

For example, a developer from Sync agreed to the reported DM anti-pattern instances and said:

My recommendation is never do calls from the main thread. Another recommendation is: use `NSFetchedResultsController` when possible...²⁰

The developer suggests using the `NSFetchedResultsController` API or an API wrapper that he developed to avoid the DM anti-pattern. However, after manually checking the API and the wrapper, we found that the API is only able to handle one particular case of the DM anti-pattern: when the data retrieved from the database is used directly in a UI table. Hence, even though the developers are aware of the anti-pattern, the problems still exist in the app. This also shows that iPerfDetector is able to help developers detect undiscovered problems in iOS apps.

As another example, developers from Yep responded to our email and acknowledged the reported problem:

I think you're right about the anti-pattern issue, the code isn't perfect yet. This views are not in the main scenes, so the performance is tolerated [...]

²⁰<https://github.com/3lvis/Sync/issues/509>

Developers from TSWechat agreed that iPerfDetector is helping them detect problems in their app. They also mentioned how iPerfDetector helped their app detect the problems early in the development process:

Yeah, absolutely yes. I didn't do anything about the performance of this project so far. So we can improve the performance of this project [based on the anti-pattern instances] that you [detected].²¹

We found that in some cases, developers acknowledged that the reported anti-pattern instances are problematic, but the performance impact may not be clear. For example, a Firefox developer mentioned that:

...Changing this alpha might not make a measurable impact on CPU perf, maybe it saves some power demand on the GPU though...²²

Nevertheless, some developers think these anti-patterns are bad practices that one should always avoid. However, fixing the anti-patterns in a matured app can be a challenging task (e.g., may require refactoring the design of the app). For example, developers from WordPress mentioned that:

In my opinion it's not a matter of whether to do it or not, but more about deciding when to do it...Moving CoreData operations to a background thread could be one of those improvements as it would allow us not to lock the main thread... we have to evaluate how moving that operation to the background (ie: making it asynchronous) would impact the rest of the code.²³

We also got positive feedback from commercial apps' developers. All the reported anti-pattern instances were acknowledged as problematic by our industrial partner, and developers are in the process of fixing them. iPerfDetector is now integrated into the quality assurance process of our industrial partner to help them ensure the performance of their apps.

iPerfDetector detected a total of 34 anti-pattern instances in the studied apps, where we manually verified 32 of them as true positives. We reported the detected anti-pattern instances to developers and 31/32 are confirmed as potential performance problems (we are still waiting for developers' response for the remaining case).

RQ4: What is the performance impact of the iOS anti-patterns?

Motivation

To provide a better understanding of the iOS performance anti-patterns, in this RQ, we conduct a case study and see if these anti-patterns actually have a negative performance impact (i.e., in terms of GPU usage, CPU usage and response time).

Approach

We choose one detected anti-pattern instance from each of DM, BLEF, and TREF in the studied apps. We did not conduct the study on all the detected problems because generating a fix, in many cases, requires significant manual investigation and refactoring of the existing

²¹<https://github.com/hilen/TSWeChat/issues/42>

²²<https://github.com/mozilla-mobile/firefox-ios/issues/3961>

²³<https://github.com/wordpress-mobile/WordPress-iOS/pull/578>

code. However, for each anti-pattern, the root cause is the same across all detected instances. We choose TREF and BLEF from Firefox due to its popularity and size. We choose DM from one of the commercial apps to show that the studied anti-patterns exist and have a performance impact in both open source and commercial apps. We did not include UIBG in this study since the effect of UIBG is usually related to app crashes or missing UI updates (i.e., visual performance problems from the perspective of the users). We measured GPU usage when BLEF and TREF anti-patterns were happening, because as mentioned by developers in forums and Apple Documents, color blending parts of the screen, which are common causes of sluggish scrolling in tables, are highly associated with GPU usages (Medium 2015; Apple 2018f). For the DM instance, we measure the CPU usage and response time, since as stated in Apple Core Data related tutorials, processing data on the main thread specifically affects CPU (Apple 2018g).

We first measure the performance of the workload (e.g., scrolling a table) before fixing the anti-pattern. Then, we measure the performance improvement after we manually fix the anti-pattern. We used an iPhone 6 for the experiment in our case study and we used the Apple's Instrument to collect performance metrics (Apple 2018c). We repeat each execution 10 times, and take the average value to minimize the effect of fluctuation during performance measurement (Chen et al. 2014). Note that the performance impact of each anti-pattern instance may vary for different cases (e.g., some anti-pattern instances that are executed more frequently or are related to larger tables would have a larger performance impact) (Chen et al. 2016a). Hence, our goal in this case study is to find out whether the anti-pattern has an impact on the performance in a controlled environment.

Results

For both TREF and BLEF, we measured the GPU usage since the anti-patterns are related to processing UI effect. We apply the same experimental setting to these two anti-pattern instances. Table 5 shows the GPU usage before and after fixing the BLEF and TREF anti-patterns. The executed workload involves scrolling a table in Firefox. We set the table to contain 15 and 30 rows to study the performance impact when the size of a table increases. We monitor the GPU usage before and after removing the heavy UI effect. We wrote test cases to automate the scrolling and profile the GPU usage during test execution. We found that, when the table contains 15 rows, fixing the TREF anti-pattern instance improved the GPU usage by 5.8% (improved from 15.3% to 14.4%). On the other hand, fixing the BLEF anti-pattern instance improved the GPU usage by 72% (improved from 43.4% to 12.1%). When the table contains 30 rows, fixing the TREF anti-pattern instance improved the GPU usage by 29% (improved from 22.5% to 16%). For BLEF, we observed less improvement when fixing the anti-pattern instance in a larger table. The GPU usage was improved by 70% (improved from 50.2% to 15.3%) when BLEF is fixed.

Table 5 GPU usage (percentage) changes in median, average, and standard deviation, after resolving the BLEF and TREF anti-pattern instances

	15 rows				30 rows			
	Before		After		Before		After	
	Median	Avg±StDv	Median	Avg±StDv	Median	Avg±StDv	Median	Avg±StDv
BLEF	43.5%	43.4%±3.02%	12.5%	12.1%±4.04%	50%	50.2%±3.19%	16%	16.5%±1.71%
TREF	13%	15.3%±4.05%	13.5%	14.4%±3.74%	21.5%	22.5%±8.19%	18%	20.1%±4.2%

Table 6 CPU usage changes (percentage) in median, average, and standard deviation, and response time changes (percentage) after resolving the DM anti-pattern instance

	Before		After	
	Median	Avg±StDev	Median	Avg±StDev
CPU	27%	28.91%±6.42%	18.66%	19.03%±4.15%
Resp time		-	80%	

Due to the confidentiality agreement, we only show the percentage improvement for the response time

For DM, we monitored response time and CPU usage of App3, when one anti-pattern instance is happening before and after moving Core Data access to a background thread (the operation takes less than a second). Table 6 shows the CPU usage before and after fixing the DM anti-pattern. It also shows the response time improvement after fixing the anti-pattern. The executed workload involves fetching 90 objects from the database. Although we cannot report the actual response time due to the confidentiality agreement, the response time of the task that contained the anti-pattern instance improved by over 80%, after the anti-pattern was fixed. The average CPU usage improved by 34.15% (from 28.9% to 19.03%).

In summary, we found that fixing the anti-pattern instances has a non-negligible performance improvement. Moreover, our performance measurements show that the average values of the performance metrics are very close to the median values, which shows there is little fluctuation in our analysis (i.e., the results are stable). iPerfDetector can help developers detect the anti-pattern instances, and developers may decide when and how to fix the problem (e.g., problems that are commonly triggered by the users may need to be fixed earlier).

We found that fixing the anti-pattern instances may improve performance (i.e., in terms of response time, GPU, and CPU usage) by 5.8% to 67% in GPU usage, 34.15% in CPU usage, and by over 80% in response time of the exercised workload.

4 Threats to Validity

4.1 Internal Validity

In this paper, we conducted a manual analysis on the performance issue reports. Similar to prior studies (Liu et al. 2014; Jin et al. 2012b), we used keywords such as “performance” and “slow” to identify performance-related issue reports. We found that some issue reports may not be related to performance issues even though they contain the keywords. Therefore, to reduce the bias in the data, we manually went through all 960 matched issue reports to identify the relevant reports. Although we tried to use more general search terms, we may still miss some performance-related issue reports. However, we have a much larger number of performance issues compared to prior performance studies (e.g., we studied 225 issue reports after the manual filter process, where Liu et al. (2014) only studied 70 issue reports).

As mentioned in Section 3, iPerfDetector does not detect retain cycles. Facebook’s Infer focuses on using static analysis to detect null pointer dereferences and memory leaks, which

makes it a great complement to our tool. Our tool is detecting performance anti-patterns that are not implemented in Infer's detection algorithm. Apple's Instruments, on the other hand, is a profiling tool that is able to show the performance information (e.g., object memory graphs and CPU usages of user events) when running an app. Therefore, developers may use Instruments to monitor the detailed activities of an app and try to diagnose the problems. However, compared to our approach, Instruments offers more flexibility in detecting performance problems, but it can not directly pinpoint developers to the root causes (i.e., anti-pattern in the code) and requires developers to dynamically exercise the problematic code in the app. Moreover, our empirical study shows the common performance problems that happen in iOS apps and compare our findings with prior studies in Android apps. Future studies may build upon our findings and provide better techniques to help improve the performance of iOS apps.

A number of projects may still use Objective-C code in some parts of the app (e.g., not yet migrated or in third party libraries). In such cases, our tool will not be able to detect the anti-patterns. Even though our tool can only detect anti-patterns in Swift code at the moment, the studied anti-patterns are related to how developers interact with the iOS SDK (i.e., the problem can occur whether an app is implemented in Swift or Objective-C). For example, developers may be calling iOS APIs to compute heavy UI effects in tables when developing apps in either language. Hence, the anti-patterns that we studied are applicable to both languages.

4.2 External Validity

We found that there is a limited number of mature and open source iOS apps that contain performance issue reports. Hence, we only conducted the manual study on four open source iOS apps that met the selection criteria. Although these apps are widely used and are large in scale, the performance issue reports that we studied may not be generalizable to other apps. However, we found that the manually-uncovered performance anti-patterns also exist in seven other iOS apps (three commercial and four open source apps), which means that these performance problems are not unique to the four manually-studied apps. We conducted a case study to measure the performance improvement when the anti-pattern instances are fixed. However, the performance improvement may vary in different cases. Various factors, such as number of cells in a table, number of iterations in loops, and how frequently users may trigger the anti-pattern instance in the code, may all affect the performance impact of the anti-pattern instance. Nevertheless, our result highlights the potential impact of the performance anti-patterns. Moreover, our tool can help detect the anti-pattern instances, and developers may decide when or how they want to fix the problem (e.g., the problems that are frequently triggered by users may need to be fixed earlier).

The performance problems that we found are mostly related to data access and UI rendering. The reason may be that we focus the study on iOS apps, so most of other problems (e.g., network-related issues) that we found are related to testing and exception handling. For example, we found some developers discussed how to test various features in an app when the network connection is slow.^{24,25} In most cases, app developers do not have control of external systems (e.g., the network speed and the servers that they are

²⁴<https://github.com/wordpress-mobile/WordPress-iOS/issues/9993>

²⁵<https://github.com/wordpress-mobile/WordPress-iOS/pull/10111>

communicating with). We did find some problems that are implicitly related to network, where developers are storing the received data in the database inefficiently. In such cases, we classify the problem as the data model (DM) anti-patterns. Future research should consider studying performance problems that exist in servers that are handling requests from iOS apps.

4.3 Construct Validity

We use static analysis to detect the iOS performance anti-pattern in the studied apps. Since we do not have the ground truth of the anti-pattern instances in the studied apps, it is impossible for us to compute a recall value. However, iPerfDetector was able to detect the anti-pattern instances that we manually studied in the issue reports, as well as the new anti-pattern instances that we detected in the latest version of the app.

During the manual study, we found that UIBG (i.e., doing UI-related operations in background threads) is a common anti-pattern (nearly 18%) among all the studied issue reports. However, iPerfDetector did not detect any such anti-pattern instance in the latest versions of the apps. We conjecture that the Main Thread Checker may have helped developers avoid such problems during app development. As we mentioned in Section 3, the Main Thread Checker is a dynamic analysis tool released by Apple that helps developers detect UIBG when running the app in debug mode (Apple 2017a). To verify the assumption, we conduct a small experiment on the studied issue reports in WordPress. Interestingly, we found that almost half of UIBG anti-patterns were reported by users after *the Main Thread Checker was released* for a while. Namely, developers did not fully utilize the Main Thread Checker to detect these anti-pattern instances during development, so these problems were reported by the users after the apps were released. The reason that iPerfDetector did not detect any UIBG may be that the studied apps have been used by many users for a long period of time, so most UIBG instances were already reported and fixed. Future studies are needed to study how developers interact with different bug detection tools to provide better development supports.

Finally, our tool currently analyzes the source code of an app, so if the source code is not available, we cannot detect the anti-patterns. Hence, if an app is using external libraries, which the source code is not available, we will not be able to detect the problems.

5 Related Work

In this section, we discuss the related work along two dimensions: 1) detecting performance anti-patterns in software applications, 2) code smells in mobile apps, and 3) empirical studies on Swift-based applications.

5.1 Detecting Performance Anti-patterns in Software Applications

Much research effort has been devoted to understanding and detecting performance anti-patterns in system applications (e.g., enterprise or database systems). According to Fowler et al. (1999), a code smell is any characteristic in the code (e.g., code design or pattern) that may indicate a deeper problem. Zaman et al. (2012) empirically studied the differences between performance and non-performance issues, and suggested that performance issues have their unique characteristics that would require specialized detection approaches. Nistor et al. (2013a) further supported the finding, where they found that, unlike functional

Table 7 Comparison between our study and prior studies in terms of the studied performance anti-patterns

Prior study	Anti-pattern category			Platform	
	UI	Memory	Multi-threading		Other
Our study	Applying UI blurring effect inefficiently, Applying UI transparency inefficiently	Retain cycle	Accessing data models on the UI thread, Updating UIs on background threads	—	iOS
Code smells in iOS apps: How do they compare to Android? (Habehi et al. 2017)	—	Ignoring low-memory warning	Blocking the main thread	Heavy enter-background tasks, download abuse	iOS
Study and refactoring of android asynchronous programming (Lin et al. 2015)	—	—	AsyncTask (used for encapsulating short running tasks) to IntentService (good choice for long-running tasks)	—	Android
Performance-based guidelines for energy efficient mobile applications (Cruz and Abreu 2017)	ViewHolder, ObsoleteLayoutParam, DrawAllocation, ObsoleteLayoutParam, overdraw	Recycle, DrawAllocation	—	Wakelock, UnusedResources	Android
On death, taxes, and sleep disorder bugs in smartphones (Jindal et al. 2013)	—	—	—	No-Sleep Bugs, Under-sleep bugs, over-sleepbugs	Android
Bootstrapping energy debugging on smartphones: a first look at energy bugs in mobile devices (Pathak et al. 2011)	—	—	—	OS processes, configuration changes, loop bug, immortality bug	Android

Table 7 (continued)

Prior study	Anti-pattern category		Memory	Multi-threading	Other	Platform
	UI					
Characterizing and detecting performance bugs for smartphone applications (Liu et al. 2014)	Wasted computation for invisible GUI		–	Lengthy operations in main threads	Frequently invoked heavy-weight callbacks	Android
An empirical study of the performance impacts of android code smells (Hecht et al. 2016)	–		–	–	Member ignoring method (mim), internal getter/setter (IGS), HashMap Usage (HMU)	Android
Saving energy on mobile devices by Refactoring (Gottschalk et al. 2014)	Backlight		–	–	Data transfer, binding resources too early, statement change, cloud computing	Android
Lightweight detection of android-specific code smells: the adocor project (Palomba et al. 2017)	–		Leaking inner class, leaking thread, no low memory resolver	Blocking the main thread	Data transmission without compression, debuggable release, durable wakelock, inefficient data format and parser, inefficient data structure, inefficient SQL query, Internal getter and setter, member ignoring method, public data, rigid alarm manager, slow loop, unclosed closable	Android
A study and toolkit for asynchronous programming in C# (Okur et al. 2014)	–		–	Fire & Forget methods, unnecessary async/await methods, using long-running operations under Async methods, unnecessarily capturing context	–	Windows phone

issues, most performance issues are uncovered through code reasoning with tool supports. Jin et al. (2012c) empirically studied 109 performance issues in system applications (e.g., MySQL) and derived rules for detecting performance anti-patterns. Nistor et al. (2013b, 2015) proposed both dynamic and static approaches to detect inefficient memory access in loops. Chen et al. (2014, 2016a, b) proposed a series of approaches that help developers detect database access performance anti-patterns in enterprise applications.

5.2 Code Smells in Mobile Applications

Most other prior studies focus on studying code smells in Android apps. Table 7 further clarifies the differences and similarities between our work and prior work on mobile performance code smells (i.e., anti-patterns). Most anti-patterns discussed in the prior studies are different than the ones we studied. We classify the anti-patterns based on the three main categories that we manually uncovered when studying the issue reports. We find that although some anti-patterns may be similar to the ones that we studied (e.g., blocking the main thread), prior studies are mostly all in the Android platform and our study revealed some unique problems that were not studied before. Below, we discuss each related work in detail.

To the best of our knowledge, there is only one prior study on iOS code smells. (Habchi et al. 2017) conducted a study on the Object-Oriented code smells in iOS apps and compare with that of Android apps. They found that iOS apps are less prone to Object-Oriented code smells. They also documented several iOS code smells (e.g., abuse of the singleton design pattern, ignoring low memory warning, and downloading the same files multiple times) based on reading Apple's iOS documentation, developer blogs, and Stack Overflow posts. However, most of the code smells that they found are different from the ones that we studied in this paper, and they did not provide an evaluation of the detection results. Also, different than their study, we studied bug reports from popular open source projects to find the real-world root-causes of anti-patterns.

Hecht et al. (2016) conducted an empirical study on the impacts of three Android code smells (i.e., private getter/setter, member ignoring method, and using hashmap instead of array map). They found that fixing some code smells may reduce the number of delayed frames, improve the UI drawing performance, or improve energy efficiency in some cases. However, some of the performance problems may no longer exist in newer versions of Android. We study a broader range of anti-patterns and also provide a tool to detect those anti-patterns. Reimann et al. (2014) compiled a list of common code smells and refactoring solutions in Android apps. Palomba et al. (2017) implemented a tool to detect Android-related code smells that are presented in the study by Reimann et al. (2014). All of the anti-patterns they discussed are different from the ones we found in our study. Mannan et al. (2016) conducted a literature survey to identify common code smells in mobile apps. They found that there was no prior study on code smells in iOS apps. They further compared the types and distributions of code smells in mobile and desktop applications. Prior studies Gottschalk et al. (2014) and Morales et al. (2017) also studied the relationship between the above-mentioned code smells and energy consumption in Android apps. Jindal et al. (2013) study wakeLocks and the scope of the code paths which should be protected by wakeLocks and look for cases of where it is over-used, which increases battery consumption, or under-usage which affects apps correctness. Pathak et al. (2012) developed an energy profiler named eprof, to help developers find instances of energy bugs caused by wakeLocks. In another study, Pathak et al. (2011) also provided the taxonomy of energy

bugs (ebugs) in a variety of categories, such as hardware ebugs, Software ebugs, and ebugs caused by external conditions. Under software energy bugs category, they focused on bugs caused by OS processes and configuration changes, and applications and framework ebugs, such as unnecessary loops (i.e., app repeatedly trying to connect to the remote server) and apps respawning in spite of being killed by users. These two studies also focus on Android platform.

The energy code smells introduced by these studies, are only validated in Android platform, and the results cannot be generalized to iOS platform. Moreover, we believe that our performance case study may help identify which types of code smell or anti-pattern may consume more energy (e.g., base on the CPU or GPU usage), without needing to deploy complex hardware environment.

Morales et al. (2017) proposed a search-based approach to provide refactoring suggestion to developers. They found that fixing the code smells may reduce the number of delayed frames and improve the UI drawing performance. A study by Cruz and Abreu (2017) also focuses on performance related UI intensive anti-patterns, such as ViewHolder, Obsolete-LayoutParam, and also memory issues such as Recycle and DrawAllocation, all in Android mobile apps. They selected the most common performance related suggestions given by lint, and published a benchmarking suite for energy consumption in Android.

Lin et al. (2015) implemented a refactoring tool for fixing an async programming bad practice in Android apps. The anti-patterns include improper uses of primary construct, AsyncTask, which can lead to memory leaks, results getting lost, and waste of energy.

A study by Okur et al. (2014) focuses on uses and misuses of asynchronous constructs in Windows Phone applications. They also developed a set of tools for refactoring 2 of the common async programming bad practices. The anti-patterns studied in this paper are specific to async constructs in Windows Phone.

The study by Liu et al. (2014) is the closest to the problems that we studied in this paper. They manually studied 70 performance issue reports in Android apps and developed a static analysis tool to detect two manually-uncovered performance anti-patterns. In this paper, we focus our study on another popular mobile app platform, iOS. iOS apps were rarely studied in prior research, even though they have different characteristics compared to Android apps (Hu et al. 2018). Different from the study by Liu et al. (2014), we conducted our study on 225 performance issue reports. We found that many performance anti-patterns in the studied iOS apps are related to inefficient UI design (e.g., using inefficient UI effects such as blurring) and thread handling. We also provided a case study to illustrate the performance impact of the studied anti-patterns. To the best of our knowledge, we are the first study on performance issues in iOS apps. Future studies are needed to further help mobile developers with not only Android, but also iOS app development.

5.3 Empirical Studies on Swift-based Applications

Cassee et al. (2018) conducted a study on most frequently asked questions about Swift programming language on Stack Overflow. They found that 14.6% of the questions are about data storage, 12.8% are related to UI actions, and around 9% are related to multi-threading issues. Their results also echo with the importance of the anti-patterns we studied. They also surveyed iOS developers and discovered that they found Swift easy to adopt. Rebouças et al. (2016) studied open-source repositories to find out to what extent Swift developers apply best practices of error handling provided by guidelines and tutorials.

6 Conclusion

In this paper, we conducted an empirical study on 225 performance issue reports in four open-source iOS applications (i.e., apps) that are written in Swift. Our manual study on the issue reports found that inefficient UI design, memory issues, and inefficient thread handling are the most common types of performance issues. In particular, we found that most problems we studied are related to how developers use iOS-specific APIs (e.g., UI effect or threading). Hence, we manually derive four performance anti-patterns that are related to iOS API usage. We documented these four anti-patterns and implemented a static analysis tool, called iPerfDetector, to detect these patterns. We evaluated iPerfDetector on 11 iOS apps (eight open source and three commercial apps). iPerfDetector was able to detect a total of 34 anti-pattern instances (where we manually verified two as false positives). We reported 32 of the detected anti-pattern instances and 31 of them are accepted by developers as potential performance problems. iPerfDetector is now used by our industrial partner to ensure the quality of their apps. To the best of our knowledge, this is the first study on iOS performance issues. Our study highlights some common problems that may be unique to iOS development. However, the current research community often only focuses on Android app development. Future studies should investigate problems in iOS development and provide further support to developers.

References

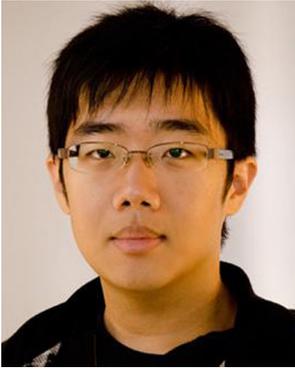
- Afjehei SS, Chen T-HP (2018) iOS performance issue reports. https://docs.google.com/spreadsheets/d/1hX8IBcYIVv6x4nWfWczT5oLV0b1SMQK03q6_xC9B8eQ/edit#gid=1079185239. Accessed 2018-07-23
- Apple (2012) Practical memory management. <https://developer.apple.com/library/archive/documentation/Cocoa/Conceptual/MemoryMgmt/Articles/mmPractical.html>. Accessed 2018-06-15
- Apple (2015) More core data and concurrency. <https://cocoacasts.com/more-core-data-and-concurrency>. Accessed 2018-06-15
- Apple (2017) Apple threading programming guide. https://developer.apple.com/documentation/code_diagnostics/main_thread_checker. Accessed 2018-06-15
- Apple (2017) Nsmanagedobject programming guide. <https://developer.apple.com/documentation/coredata/nsmanagedobjectcontext>. Accessed 2018-06-15
- Apple (2018a) Apple guide on swift programming language. <https://developer.apple.com/swift/>. Accessed 2018-06-15
- Apple (2018b) Apple UIKit. <https://developer.apple.com/documentation/uikit>. Accessed 2018-06-15
- Apple (2018c) Instruments. <https://help.apple.com/instruments/mac/current/#/dev7b09c84f5/>. Accessed 2018-07-23
- Apple (2018d) Making core data your model layer. https://developer.apple.com/documentation/coredata/making_core_data_your_model_layer. Accessed 2018-06-15
- Apple (2018e) Model-view-controller. <https://developer.apple.com/library/archive/documentation/General/Conceptual/DevPedia-CocoaCore/MVC.html>. Accessed 2018-06-15
- Apple (2018f) Monitoring your app's graphics activity. https://developer.apple.com/documentation/metal/tools_profiling_and_debugging/gpu_activity_monitors/monitoring_your_app_s_graphics_activity. Accessed 2018-07-23
- Apple (2018g) Using core data in the background. https://developer.apple.com/documentation/coredata/using_core_data_in_the_background. Accessed 2018-06-15
- Cassee N, Pinto G, Castor F, Serebrenik A (2018) How swift developers handle errors. In: 15Th international conference on mining software repositories (MSR 2018)
- Chen T-H, Weiyi S, Jiang ZM, Hassan AE, Nasser M, Flora P (2014) Detecting performance anti-patterns for applications developed using object-relational mapping. In: Proceedings of the 36th international conference on software engineering (ICSE), pp 1001–1012

- Chen T.-H., Shang W, Hassan AE, Nasser M, Flora P (2016a) Detecting problems in the database access code of large scale systems: an industrial experience report. In: Proceedings of the 38th international conference on software engineering companion, ICSE '16, pp 71–80
- Chen T-H, Shang W, Jiang ZM, Hassan AE, Nasser M, Flora P (2016b) Finding and evaluating the performance impact of redundant data access for applications that are developed using object-relational mapping frameworks. *IEEE Trans Softw Eng* 42(12):1148–1161
- Chen T-H, Syer MD, Shang W, Jiang ZM, Hassan AE, Nasser M, Flora P (2017) Analytics-driven load testing: an industrial experience report on load testing of large-scale systems. In: Proceedings of the 39th international conference on software engineering: software engineering in practice track, ICSE-SEIP '17, pp 243–252
- Chowdhury S, Di Nardo S, Hindle A, Jiang ZMJ (2018) An exploratory study on assessing the energy impact of logging on android applications. *Empir Softw Eng* 23(3):1422–1456
- Cruz L, Abreu R (2017) Performance-based guidelines for energy efficient mobile applications. In: 2017 IEEE/ACM 4th international conference on mobile software engineering and systems (MOBILESoft). IEEE, pp 46–57
- Facebook (2017) Facebook's infer static analysis tool. <http://fbinfer.com/>. Accessed 2018-06-15
- GitHub (2018) Trending swift projects on github. <https://github.com/trending/swift?since=monthly>. Accessed 2018-06-15
- Gottschalk M, Jelschen J, Winter A (2014) Saving energy on mobile devices by refactoring. In: *EnviroInfo*, pp 437–444
- Grechanik M, Fu C, Xie Q (2012) Automatically finding performance problems with feedback-directed learning software testing. In: Proceedings of the 34th international conference on software engineering (ICSE), pp 156–166
- Guide ACDP (2017) Core data programming guide. <https://developer.apple.com/documentation/coredata>. Accessed 2018-06-15
- Habchi S, Hecht G, Rouvoy R, Moha N (2017) Code smells in ios apps: how do they compare to android? In: Proceedings of the 4th international conference on mobile software engineering and systems. IEEE Press, pp 110–121
- Hecht G, Moha N, Rouvoy R (2016) An empirical study of the performance impacts of Android code smells. In: Proceedings of the international workshop on mobile software engineering and systems. ACM, pp 59–69
- Hu H, Bezemer C-P, Hassan AE (2018) Studying the consistency of star ratings and the complaints in 1 & 2-star user reviews for top free cross-platform android and iOS apps. *Empir Softw Eng*
- Jiang ZM, Hassan AE, Hamann G, Flora P (2008) Automatic identification of load testing problems. In: Proceedings of 24th international conference on software maintenance (ICSM), pp 307–316
- Jin G, Song L, Shi X, Scherpelz J, Lu S (2012a) Understanding and detecting real-world performance bugs. In: Proceedings of the 33rd ACM SIGPLAN conference on programming language design and implementation, PLDI '12
- Jin G, Song L, Shi X, Scherpelz J, Lu S (2012b) Understanding and detecting real-world performance bugs. In: Proceedings of the 33rd ACM SIGPLAN conference on programming language design and implementation, PLDI '12, pp 77–88
- Jin G, Song L, Shi X, Scherpelz J, Lu S (2012c) Understanding and detecting real-world performance bugs. *ACM SIGPLAN Not* 47(6):77–88
- Jindal A, Pathak A, Hu YC, Midkiff S (2013) On death, taxes, and sleep disorder bugs in smartphones. In: Proceedings of the workshop on power-aware computing and systems. ACM, p 1
- Khalid H, Shihab E, Nagappan M, Hassan AE (2015) What do mobile app users complain about? *IEEE Softw* 32(3):70–77
- Lin Y, Okur S, Dig D (2015) Study and refactoring of android asynchronous programming (t). In: 2015 30th IEEE/ACM international conference on automated software engineering (ASE). IEEE, pp 224–235
- Liu Y, Xu C, Cheung S-C (2014) Characterizing and detecting performance bugs for smartphone applications. In: Proceedings of the 36th international conference on software engineering. ACM, pp 1013–1024
- Mannan UA, Ahmed I, Almurshed RAM, Dig D, Jensen C (2016) Understanding code smells in android applications. In: 2016 IEEE/ACM international conference on mobile software engineering and systems (MOBILESoft). IEEE, pp 225–236
- Martin W, Sarro F, Jia Y, Zhang Y, Harman M (2017) A survey of app store analysis for software engineering. *IEEE Trans Softw Eng* 43(9):817–847
- Fowler M, Beck K, Brant J, Opdyke W, Roberts D (1999) Refactoring: improving the design of existing code. Addison-Wesley, Reading

- McDonnell T, Ray B, Kim M (2013) An empirical study of api stability and adoption in the android ecosystem. In: Proceedings of the 2013 IEEE international conference on software maintenance, ICSM '13, pp 70–79
- Medium (2015) Diagnose and solve performance problems with xcode instruments. <https://medium.com/@zhenya.peteliev/diagnose-and-solve-performance-problem-with-xcode-instruments-5c25c27f21d5>. Accessed 2018-11-05
- Medium (2016) 21 amazing open source iOS apps written in swift. <https://medium.mybridge.co/21-amazing-open-source-ios-apps-written-in-swift-5e835afee98e>. Accessed 2018-06-11
- Morales R, Saborido R, Khomh F, Chicano F, Antoniol G (2017) Earmo: an energy-aware refactoring approach for mobile apps. *IEEE Trans Softw Eng*, (1):1–1
- Nistor A, Jiang T, Tan L (2013a) Discovering, reporting, and fixing performance bugs. In: Proceedings of the 10th working conference on mining software repositories. IEEE Press, pp 237–246
- Nistor A, Song L, Marinov D, Lu S (2013b) Toddler: detecting performance problems via similar memory-access patterns. In: Proceedings of the 2013 international conference on software engineering, ICSE '13, pp 562–571
- Nistor A, Chang P-C, Radoi C, Lu S (2015) Caramel: detecting and fixing performance problems that have non-intrusive fixes. In: Proceedings of the 2015 international conference on software engineering, ICSE '15, pp 902–912
- Okur S, Hartveld DL, Dig D, Deursen AV (2014) A study and toolkit for asynchronous programming in C#. In: Proceedings of the 36th international conference on software engineering. ACM, pp 1117–1127
- Palomba F, Di Nucci D, Panichella A, Zaidman A, De Lucia A (2017) Lightweight detection of android-specific code smells: the adactor project. In: Proceedings of the 24th international conference on software analysis, evolution and reengineering, SANER '17, pp 487–491
- Pathak A, Hu YC, Zhang M (2011) Bootstrapping energy debugging on smartphones: a first look at energy bugs in mobile devices. In: Proceedings of the 10th ACM workshop on hot topics in networks. ACM, p 5
- Pathak A, Hu YC, Zhang M (2012) Where is the energy spent inside my app?: fine grained energy accounting on smartphones with eprof. In: Proceedings of the 7th ACM European conference on computer systems. ACM, pp 29–42
- Rebouças M, Pinto G, Ebert F, Torres W, Serebrenik A, Castor F (2016) An empirical study on the usage of the swift programming language. In: 2016 IEEE 23rd international conference on software analysis, evolution, and reengineering (SANER). IEEE, pp 634–638
- RedMonk (2018) The RedMonk programming language rankings: January 2018. <http://redmonk.com/sogrady/2018/03/07/language-rankings-1-18/>. Accessed 2018-06-11
- Reimann J, Brylski M, Aßmann U (2014) A tool-supported quality smell catalogue for android developers. In: Proceeding of the conference modellierung 2014 in the workshop modellbasierte und modellgetriebene softwaremodernisierung–MMSM
- Smith CU, Williams LG (2001) Software performance antipatterns; common performance problems and their solutions. In: Int CMG conference, pp 797–806
- Statista (2018) Number of apps available in leading app stores. <https://www.statista.com/statistics/276623/number-of-apps-available-in-leading-app-stores/>. Accessed 2018-07-23
- Statista (2018) Number of apps available in leading app stores as of 1st quarter 2018. <https://www.statista.com/statistics/276623/number-of-apps-available-in-leading-app-stores/>. Accessed 2018-06-15
- SwiftAST (2018) Swift-ast tool on github. <https://github.com/yanagiba/swift-ast>. Accessed 2018-6-30
- Syer MD, Nagappan M, Hassan AE, Adams B (2013) Revisiting prior empirical findings for mobile apps: an empirical case study on the 15 most popular open-source android apps. In: Proceedings of the 2013 conference of the center for advanced studies on collaborative research, CASCON '13, pp 283–297
- Zaman S, Adams B, Hassan AE (2012) A qualitative study on performance bugs. In: Proceedings of the 9th IEEE working conference on mining software repositories. IEEE Press, pp 199–208

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Sara Seif Afjehei is a master's student in Software Engineering at Concordia University. She works on different areas of software engineering, such as performance in mobile apps and mobile APIs design. She received her masters in MBA from Azad University of Qazvin, Iran, for research investigating software projects portfolio optimization. Contact her at s.seifaf@encs.concordia.ca.



Tse-Hsun (Peter) Chen is an Assistant Professor in the Department of Computer Science and Software Engineering at Concordia University, Montreal, Canada. He leads the Software Performance, Analysis, and Reliability (SPEAR) Lab, which focuses on conducting research on performance engineering, program analysis, log analysis, production debugging, and mining software repositories. He serves regularly as a program committee member of international conferences in the field of software engineering, such as ASE, ICSME, SANER, and ICPC, and he is a regular reviewer for software engineering journals such as JSS, EMSE, and TSE. Dr. Chen obtained his BSc from the University of British Columbia, and MSc and PhD from Queen's University. Besides his academic career, Dr. Chen also worked as a software performance engineer at BlackBerry for over four years. Early tools developed by Dr. Chen were integrated into industrial practice for ensuring the quality of large-scale enterprise systems. More information at: <http://petersehun.github.io/>.



Nikolaos Tsantalis is an associate professor in the Department of Computer Science and Software Engineering at Concordia University, Montreal, Canada, and holds a Concordia University Research Chair in Web Software Technologies. His research interests include software maintenance, software evolution, empirical software engineering, refactoring recommendation systems, refactoring mining, and software quality assurance. He has been awarded with two Most Influential Paper awards at SANER 2018 and SANER 2019, and two ACM SIGSOFT Distinguished Paper awards at FSE 2016 and ICSE 2017. He served as a program co-chair for SANER ERA track in 2015 and 2019, SANER Tool Demonstration track in 2017, and SCAM Engineering track in 2019. He serves regularly as a program committee member of international conferences in the field of software engineering, such as ASE, ICSME, SANER, ICPC, and SCAM, and he is a member of the IEEE TSE review board. He is a member of the IEEE and the ACM, and holds a license from the Association of Professional Engineers of Ontario.