SCIENCE CHINA
Information Sciences

• RESEARCH PAPER •

May 2018, Vol. 61 050103:1–050103:20
https://doi.org/10.1007/s11432-017-9400-y

Special Focus on Internetware: Research Progress and Current Trend

NavyDroid: an efficient tool of energy inefficiency problem diagnosis for Android applications

Yi LIU[1,2], Jue WANG[1,2], Chang XU[1,2*], Xiaoxing MA[1,2] & Jian LÜ[1,2]

[1]*State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing* 210023*, China;*
[2]*Department of Computer Science and Technology, Nanjing University, Nanjing* 210023*, China*

Received 29 November 2017/Revised 5 March 2018/Accepted 27 March 2018/Published online 20 April 2018

Abstract Energy inefficiency is an influential non-functional issue for smartphone applications, causing increased concerns from users. Locating these problems is labor-intensive, thus automated diagnosis tools are in demand. Some existing approaches detect energy inefficiency problems by exploring application states with the JPF framework, and get favorable results. However, the effects of these approaches are restricted because of their imprecise application execution models and incomplete energy inefficiency patterns. This paper introduces NavyDroid, an effective and efficient tool of energy inefficiency problem diagnosis for Android applications. We constructed a comprehensive application execution model in the form of a state machine, which accurately simulates the runtime behavior of Android applications. We designed a parallel algorithm to systematically explore an application's state space. Our approach supports more energy inefficiency patterns, and is able to detect complicated wake lock misuses. We implemented our approach as a prototype tool and applied it to real-world applications. We evaluated NavyDroid with 19 real-world Android applications, and NavyDroid located more energy inefficiency bugs in these applications than the existing work E-GreenDroid did. Also, NavyDroid reduced the analysis time with its parallel state exploration algorithm. The experimental results demonstrated the effectiveness and efficiency of our approach for detecting energy inefficiency bugs in Android applications.

Keywords energy inefficiency, smartphone application, wake lock

Citation Liu Y, Wang J, Xu C, et al. NavyDroid: an efficient tool of energy inefficiency problem diagnosis for Android applications. Sci China Inf Sci, 2018, 61(5): 050103, https://doi.org/10.1007/s11432-017-9400-y

1 Introduction

Nowadays, as smartphones become more and more popular, the number of Android applications is growing rapidly. The data show that there are lots of applications and cumulative downloads on the Google Play Store[1)]. However, many applications suffer from energy inefficiency problems. These applications employ energy-consuming operations, such as location sensing, to provide a good user experience. At the same time, Android developers are responsible for managing the power of the device. Therefore, if these energy-consuming operations are not performed properly, much energy will be wasted. In this case, the device battery can be exhausted in a few hours, and this will result in user frustration and complaints. As energy inefficiency problems become more common, users become concerned about this issue.

It is difficult for developers to diagnose energy inefficiency problems. These problems occur only at certain application states. In order to reproduce the problems, developers often need to explore a

* Corresponding author (email: changxu@nju.edu.cn)

1) Google Play. https://en.wikipedia.org/wiki/Google_Play, 2017.

variety of application states. Therefore, automatic detection tools help locate the problems and increase debugging efficiency.

In the past several years, researchers have proposed different tools to detect energy inefficiency problems automatically. Pathak et al. [1] conducted the first study in the area of detecting energy bugs for smartphone applications. Zhang et al. [2] proposed ADEL to detect energy leaks of network data. These approaches employ different analysis techniques. Among them, several pieces of work simulate the execution of applications upon a verification framework JPF, and are shown to be effective [3–5]. GreenDroid [4] detected the missing releasing of sensors and wake locks, and analyzed whether sensor data are effectively utilized. CyanDroid [3] systematically generated multidimensional sensor data to reproduce bugs that require specific sensory data to manifest. E-GreenDroid [5] optimized the simulation execution, and updated the library modeling of GreenDroid, including some new features of the Android system.

Generally, these energy inefficiency detection tools are composed of two parts, namely, the simulation part and the monitor part.

(1) Simulation. The application execution engine simulates the execution of Android applications, and explores application states. The two main processes of the simulation part are event sequence generation and state space exploration. Both processes are guided by the application execution model.

(2) Monitor. The monitor part is based on the simulation part, and operates during the simulation execution. It monitors suspicious operations of the application, and checks for operations that match energy inefficiency patterns. An investigation shows that many energy inefficiency problems are associated with two types of energy inefficiency patterns [4]:

• **Missing sensor listener or wake lock deactivation.** An application registers a sensor listener to fetch sensor data, and acquires a wake lock to keep the CPU awake for long background tasks. The sensor listener should be unregistered when the senor is no longer needed, and the wake lock should be released when background tasks are done. If sensor listeners or wake locks are not deactivated in time, the battery power can be exhausted rapidly[2)3)4)].

• **Sensor data underutilization.** The sensor itself consumes energy to retrieve sensor data. Thus, applications should utilize sensor data in an effective way. Applications using sensor data inefficiently can be viewed as a waste of energy.

Existing approaches that simulate the execution of an Android application follow the above two-part architecture. However, these approaches have some limitations in both parts. As for the simulation part, the application execution model is imprecise. As for the monitor part, the energy inefficiency patterns are not complete. Here we discuss these shortcomings in detail and propose our solution.

In the simulation part, the application execution model guides the simulation execution of applications. However, due to the complicated component lifecycle and runtime behavior of Android applications, the behavior guided by the model is different from the actual behavior. Concretely, the models in existing approaches are not accurate in the lifecycle of activities. They do not include paused and killed states of activities.

Meanwhile, existing approaches have poor performance in exploring applications' states. E-GreenDroid randomly generates event sequences for state exploration. However, in this way duplicated event sequences are generated, and the exploration is not efficient. There should be a better approach to exploring application states. Also, existing approaches to generating event sequences can cause an application to quit abnormally, resulting in false positives.

Also, existing work did not summerize the complete energy inefficiency patterns in the monitor part. For instance, the status of a wake lock is not as simple as acquired and released. With reference counts, wake locks have more complicated misuse patterns, e.g., multiple lock acquisitions. Existing approaches do not check for these complicated misuse patterns.

2) Android location strategies. https://developer.android.com/guide/topics/sensors/sensors_overview.html, 2017.
3) Android sensors usage. https://developer.android.com/guide/topics/sensors/sensors_overview.html, 2017.
4) Managing Android device awake state. https://developer.android.com/training/scheduling/wakelock.html, 2017.

Therefore, in order to address the above issues, we propose our approach in this paper. We define an application execution model that precisely simulates the lifecycle of Android components. The model is constructed in the form of a strengthened deterministic finite automaton (DFA). The automaton contains the paused state and the killed state of an activity, as well as their related state transitions. We improve the process of event sequence generation to ensure that an application terminates normally in simulation execution. We design a state exploration algorithm to systematically explore the state space of applications, and parallelize it to accelerate the exploration. Also, we summerize new misuse pattern detecting policies for the monitor part.

We implemented our approach as a prototype tool named NavyDroid for evaluation. We selected 19 real-world Android application in order to evaluate the effectiveness and efficiency of NavyDroid in energy inefficiency diagnosis. We compared NavyDroid with E-GreenDroid [5], a state-of-the-art energy inefficiency detection tool. We analyzed the test subjects with both tools, and compared their analysis reports. As a result, NavyDroid located more energy inefficiency bugs than E-GreenDroid did, and located all the energy inefficiency bugs that E-GreenDroid reported. Moreover, we evaluated the efficiency of our parallel state exploration algorithm. The results demonstrate that it takes average 40% less time to explore application states compared with the random exploration used in E-GreenDroid, while preserving the same effectiveness. We can conclude from the evaluation results that NavyDroid detects energy inefficiency problems more effectively and efficiently.

In summary, our work makes the following contributions:

- We design an algorithm to systematically explore application states. We make the algorithm more efficient by parallelizing its execution.
- We construct an accurate application execution model. The model supports more activity states, and is more consistent with the lifecycle of Android components.
- We enhance the monitor part to detect more wake lock misuse patterns. Our approach supports more energy inefficiency patterns such as multiple lock acquisitions.
- We implement our approach as a prototype tool named NavyDroid and evaluate it with real-world Android applications. NavyDroid is able to detect more energy inefficiency bugs in the test subjects, indicating its effectiveness. Also, NavyDroid takes average 40% less time for state exploration compared with E-GreenDroid.

The work presented in this paper is based on our previous work [6], and has significantly extended it. Previously, we used the same random event sequence generation as E-GreenDroid. However, this approach generates many duplicated event sequences, and shows poor efficiency. In this work, we have extended our approach with a state exploration algorithm, which can systematically explore an application's state space in analysis and eliminate repeated event sequences. Moreover, we parallelized the exploration algorithm to accelerate the execution. Evaluation results show that the parallel algorithm takes average 40% less time compared with the original random exploration.

The rest of this paper is organized as follows. Section 2 introduces some background of Android programming and presents a motivating example. Section 3 elaborates on our approach to detecting energy inefficiency. Section 4 evaluates NavyDroid with real-world applications. Sections 5 and 6 discuss our work and some related work, and finally Section 7 concludes this paper.

2 Background and motivation

In this section, we introduce the basics of Android applications, and present a motivating example of our work.

2.1 Background

Application components are the essential parts of an Android application. There are four different types of application components[5)].

5) Android application fundamentals. https://developer.android.com/guide/components/fundamentals.html, 2017.

```
1  public class PlaybackActivity extends Activity {
2      private ImageButton mPlayPauseButton;
3      private PlaybackService mPlaybackService;
4      private ServiceConnection mConnection = new ServiceConnection() {
5          public void onServiceConnected(
6                  ComponentName name, IBinder binder) {
7              mPlaybackService = ((MyBinder) binder).getService();
8          }
9      };
10     public void onCreate(Bundle state) {
11         mPlayPauseButton.setOnClickListener(new OnClickListener() {
12             public void onClick(View v) {
13                 // toggle play/pause state
14                 mPlaybackService.playPause();
15             }
16         });
17         Intent i = new Intent(this, PlaybackService.class);
18         // start PlaybackService
19         startService(i);
20     }
21     public void onResume() {
22         Intent i = new Intent(this, PlaybackService.class);
23         // bind PlaybackService
24         bindService(i, mConnection, Context.BIND_ABOVE_CLIENT);
25     }
26 }
27
28 public class PlaybackService extends Service
29         implements MediaPlayer.OnPreparedListener {
30     private static String TAG = PlaybackService.class.getName();
31     private WakeLock mWakeLock;
32     private MediaPlayer mMediaPlayer;
33     public void onCreate() {
34         mWakeLock = getPowerManager().newWakeLock(
35                 PowerManager.PARTIAL_WAKE_LOCK, TAG);
36         mMediaPlayer = createAndSetupMediaPlayer();
37         mMediaPlayer.setOnPreparedListener(this);
38     }
39     public void onDestroy() {
40         // stop media when the service is destroyed
41         stop();
42     }
43     public void onPrepared() {
44         // start media when the media player get prepared
45         start();
46     }
47     public void playPause() {
48         if (mMediaPlayer.isPlaying())
49             pause();
50         else
51             start();
52     }
53     public void start() {
54         // play media and acquire wake lock
55         mWakeLock.acquire();
56         mMediaPlayer.start();
57     }
58     public void stop() {
59         // stop media and release wake lock
60         mMediaPlayer.stop();
61         if (mWakeLock.isHeld()) mWakeLock.release();
62     }
63     public void pause() {
64         // pause media and release wake lock
65         mMediaPlayer.pause();
66         if (mWakeLock.isHeld()) mWakeLock.release();
67     }
68     public class MyBinder extends Binder {
69         PlaybackService getService() {
70             return PlaybackService.this;
71         }
72     }
73     public IBinder onBind(Intent intent) {
74         return new MyBinder();
75     }
76 }
```

Figure 1 (Color online) Motivating example from the TomaHawk application (revision 543c3b9ab4).

Activity. Activities are the only components for graphical user interfaces (GUI). The GUI layouts corresponding to activities are declared in configuration files. The entry point of an application is usually an activity. During the execution, different running activities are organized in a back stack. The top activity of the back stack is called the activity at foreground.

Service. A service runs at the background to perform long-time tasks. An activity can start a service, or bind to a service and interact with it.

Broadcast receiver. A broadcast receiver is a component that responds to system-wide broadcast messages. A broadcast receiver can be another entry point into the application besides its own activities.

Content provider. A content provider manages a set of shared application data. Through the content provider, other components and applications can query or modify the data.

Each application component follows a prescribed lifecycle when it transits through different states[6]. As a component enters a new state, the Android framework invokes the corresponding event handlers. Scheduling these event handlers is essential to the simulation execution of an Android application.

2.2 Motivating example

We present two energy inefficiency bugs in TomaHawk[7] as the motivating example. Figure 1 shows a simplified version of the problematic code snippet. The `PlaybackActivity` and `PlaybackService` play the media specified by users. Users can switch the mode of media playing by clicking the play/pause button (lines 12–15, 48–51). `PlaybackService` is started by `PlaybackActivity` and runs at background (lines 17–19). The activity binds to the service for interactions (lines 22–24). In order to prevent the media playing from being interrupted, `PlaybackService` plays the media with a wake lock (lines 34–35). When the media playing starts, the wake lock will be acquired (lines 54–56); when the media playing pauses or stops, the wake lock will be released (lines 59–61, 64–66).

There are several operations on one wake lock in `PlaybackService`. When the media player gets prepared, the wake lock will also be acquired (lines 43–46). The service checks whether the wake lock is being held before releasing it (lines 61, 66), but acquires the wake lock directly without checking (line 55).

6) Android activity lifecycle. https://developer.android.com/guide/components/activities/activity-lifecycle.html, 2017.
7) Tomahawk. https://github.com/tomahawk-player/tomahawk-android, 2017.

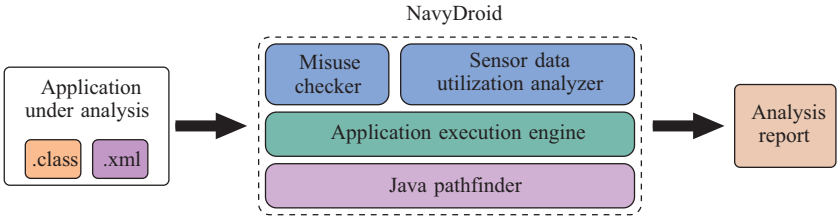

Figure 2 (Color online) Approach overview.

Therefore, the wake lock can be acquired more than once. If a user clicks the play/pause button twice after the media player gets prepared, the wake lock will be acquired twice but released only once. As wake locks are reference-counted by default, the wake lock calculates its reference count as the number of acquires minus the number of releases. The wake lock will keep being held because its reference count is larger than zero[8]. In this case, the working CPU will consume battery power without any user benefit. Existing tools, such as GreenDroid and E-GreenDroid, fail to detect this energy inefficiency problem, because of the incomplete wake lock misuse patterns. They cannot address this usage mode of wake locks.

Another energy inefficiency bug occurs when the activity and the service get killed by the Android system. When a user exits `PlaybackActivity`, both `onDestroy()` callbacks of the activity and the service will be invoked, and `PlaybackService` stops media playing and releases the wake lock (lines 40–41). However, the Android system may kill processes when the memory is insufficient, in which case the `onDestroy()` callbacks of killed activities and services will not be invoked. In this case, the wake lock will not be properly released, consuming the battery power. GreenDroid and E-GreenDroid, with imprecise application execution models, cannot simulate the runtime behavior of activities or services being killed. So they fail to detect this energy inefficiency problem.

The above example motivates us to propose an approach that can accurately simulate the execution of Android applications, and identify complex wake lock misuse patterns.

3 NavyDroid approach

In this section, we elaborate on our approach to detecting energy inefficiency problems in Android applications.

3.1 Overview

NavyDroid follows the two-part architecture described in Section 1. Its simulation part simulates the execution of Android applications, and explores application states. Its monitor part checks for energy inefficiency problems according to two types of misuse patterns, i.e., missing sensor listener or wake lock deactivation, and sensor data underutilization.

Figure 2 shows the high-level abstraction of NavyDroid. The simulation part is named application execution engine. The monitor part consists of two components, a misuse checker and a sensor data utilization analyzer. These two components check for energy bugs corresponding to the two energy inefficiency patterns, respectively.

NavyDroid takes as input an Android application's binary code and configuration files. It retrieves application components and GUI layouts from configuration files before executing the application. The application execution engine executes an application, and systematically explores its application states. The misuse checker monitors the operations on sensor listeners and wake locks during the execution. The sensor data utlization analyzer feeds sensor data to the application when related sensor listeners are registered. It then tracks where the sensor data propagate as the application executes, and analyzes

8) Powermanager.wakelock class. https://developer.android.com/reference/android/os/PowerManager.WakeLock.html, 2017.

how sensor data are utilized at different application states. At the end of the execution, NavyDroid compares sensor data utilization across explored application states to find underutilized states. It also reports misuses of sensor listeners and wake locks as output. We elaborate on these functional modules of NavyDroid in the following.

3.2 Supporting techniques

Java Path Finder (JPF) is a testing and verification framework for Java programs[9]. As shown in Figure 2, JPF is the base component and the supporting framework of NavyDroid. The core of JPF is a virtual machine for Java bytecode. In NavyDroid, the application execution engine executes an application in JPF's virtual machine. The implementation of NavyDroid utilizes JPF's facilities, including instruction listener and native peer.

• **Instruction listener.** Instruction listeners provide a way to inspect the execution of instructions[10]. NavyDroid keeps track of an Android application's execution by implementing instruction listeners. The sensor data utilization analyzer depends on this mechanism to trace the propagation of sensor data.

• **Native peer.** A native peer class is like a mock class. It models a JVM's native class executed by JPF's virtual machine[11]. By native peers, we simulate the Android framework APIs in order to support the normal operation of applications. Also, we model some critical classes such as `LocationListener` and `WakeLock`. In this way, we generate some analysis-related data and import them into applications.

Since Android applications frequently interact with users, their executions are often triggered by user events. This event-driven feature separates program code into different event handlers, with implicit calling relationships. However, JPF is designed for conventional Java programs and is unable to analyze event-driven programs directly. Therefore, we generate user events, and guide JPF to schedule event handlers in the simulation part.

3.3 Application execution engine

As the simulation part of NavyDroid, the application execution engine simulates the execution of an Android application. It enables JPF to analyze event-driven Android applications. The engine first simulates user interactions by generating sequences of user events. This process is called event sequence generation. The engine explores application states by systematically generating event sequences. With event sequences generated, the engine then schedules the corresponding event handlers. The effect of evenet handler scheduling depends on the quality of the application execution model (AEM). In the following sections, we introduce these processes, i.e., event sequence generation, state exploration, and event handler scheduling.

3.3.1 *Event sequence generation*

The first process of the application execution engine is event sequence generation. NavyDroid generates user interactions and system messages to explore application states. We first define the concept of event sequence.

Definition 1 (event sequence)**.** An event sequence seq is an ordered list:

$$\text{seq} = [e_1, e_2, \ldots, e_n] \quad (e_i \in E, \quad i = 1, 2, \ldots, n),$$

where E denotes the set of events. An event sequece is a sequence of user and system events, which an application takes as input during its execution.

NavyDroid constructs the event set by static analysis. It retrieves GUI layouts of activites from the configuration files of an application, and extracts GUI widgets declared in the layout files. Each GUI widget (e.g., a button) receives a set of user actions (e.g., button clicks). We define the candidate event

9) Java pathfinder wiki page. https://babelfish.arc.nasa.gov/trac/jpf/wiki/intro/what_is_jpf, 2017.
10) Java pathfinder listeners wiki page. https://babelfish.arc.nasa.gov/trac/jpf/wiki/devel/listener, 2017.
11) Java pathfinder mji wiki page. https://babelfish.arc.nasa.gov/trac/jpf/wiki/devel/mji, 2017.

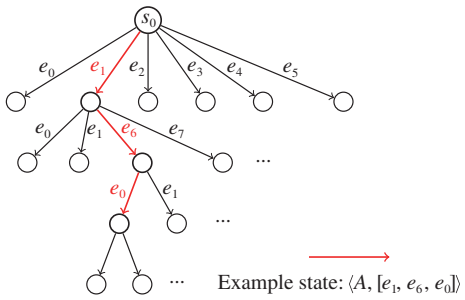

Figure 3 (Color online) State tree: the model of application states and event sequences.

set of an activity as the union set of these user actions. The candidate event set also contains some special events, such as physical keys (i.e., back, menu, and home) and the kill system event.

The application execution engine generates events in an iterating manner. An application first launches from its entry activity and enters the initial state. At this point, the first event can be generated and fed to the application. At runtime, the application execution engine watches the foreground activity. When the activity waits for user interactions, the engine generates an event selected from the candidate event set of this activity. This event generation process continues until all acitivities finish, or the number of generated events reaches an upper bound[12]. In the latter case, the engine generates additional back events to finish the running activities. When receiving a back event, the current activity is destroyed, and the previous activity comes to foreground. The application finishes when all the activities are destroyed. Therefore, by adding back events, the execution of the application terminates finally.

3.3.2 *State exploration*

In this section, we elaborate on our approach to systematically exploring application states in parallel. The application execution engine generates event sequences in order to explore application states. We first show that we can explore application states as long as we generate all possible event sequences. We define the correspondence between application states and event sequences as follows.

Definition 2. Given an application A (with initial state s_0) and an input event sequence $\text{seq} = [e_0, e_1, \ldots, e_n]$, the execution of the application is in an iterating manner. In the k-th iteration ($1 \leqslant k \leqslant n$), an event e_k is fed to the application, and the application transits from state s_{k-1} to state s_k. Each event e_k should be in the candidate event set of state s_{k-1}. Finally, application A reaches state s_n with event sequence seq, denoted as $s_n = \langle A, \text{seq}\rangle$.

As Definition 2 shows, we can explore application states as long as we generate all possible event sequences. In order to explore application states, we generate all possible event sequences, and execute the application according to the event sequences.

In order to facilitate understanding, we model all the event sequences and application states as a rooted tree called state tree, as shown in Figure 3. In this tree, a node represents an application state, and an edge represents an event. The root of the tree is the initial state of the application. All the leaving edges of an non-leaf node represent the candidate events of this state. For any state in the tree, its corresponding event sequence is the path from the root node to it. As mentioned before, as event sequences have an upper bound, the depth of the tree limited to a constant bound.

The application execution engine can generate event sequences randomly, or in a systematic manner. JPF can be used as a model checker to exhaustively explore all possible program states, like traversing the state tree. However, when we simulate the Android framework APIs in JPF, its model checker becomes fragile and often crashes. Therefore, E-GreenDroid generates event sequences randomly for state exploration. Although it is possible to reach any application state by randomly generating event sequences, many of these event sequences are duplicated. NavyDroid takes a different approach from

12) We restrict the length of event sequences in order to ensure the state exploration to finish in finite time. With a large enough bound, all representative states and event handlers can be explored, and this is sufficient for detecting energy problems.

the random mode. It explores application states systematically to avoid event sequence duplication and improve analysis efficiency.

The event sequences of an application cannot be generated statically, because the candidate event set of state s is unknown until an execution reaches s. Therefore, we design an algorithm for dynamic event sequence generation. The basic process of this algorithm is the breadth-first search (BFS) on the state tree, as is shown in Algorithm 1. We design the exploration algorithm in this BFS way for ease of parallelization. As a comparison, the depth-first search (DFS) algorithm depends on a stack, and it is error-prone for multiple processors to operate on a stack simultaneously.

Algorithm 1 State exploration

Input: Android application A.

```
1: Q ← ∅; {empty queue};
2: C_0 ← candidate event set of the initial state s_0;
3: for all e_0 ∈ C_0 do
4:     seq_0 ← [e_0];
5:     put seq_0 into Q;
6: end for
7: while Q is not empty do
8:     seq ← take an element from Q;
9:     execute the application with seq, and reach state s_k = ⟨A, seq⟩;
10:    if length(seq) < bound then
11:        C ← candidate event set of s_k;
12:        for all e_t ∈ C do
13:            seq′ ← seq + [e_t] {append e_t to the end of seq};
14:            put seq′ into Q;
15:        end for
16:    end if
17: end while
```

We maintain a queue Q (first in first out) of event sequences (line 1). First we put event sequences that contain exactly one initial event (an event which is accepted by the initial state) into Q (lines 3–6). In each iteration, we take an event sequence seq, execute the application according to this event sequence, and reach a new state s_k (lines 8–9). Then we append each event in the candidate event set of s_k to the end of seq to generate new event sequences (lines 12–15).

Although the lengths of event sequences are bounded, the total number of event sequences can still be large. Executing all the event sequences will cost a huge amount of time. Therefore, we parallelize our state exploration algorithm to accelerate the exploration. As our algorithm is a kind of breadth-first search and depends on a queue, we create multiple processors and perform putting/taking operations on the queue concurrently. There are $n+1$ processors, including one master processor and n worker processors. The master processor puts initial event sequences into Q, and starts all the worker processors. Each worker processor works the same as in Algorithm 1. It takes an event sequence seq from Q, executes it, and puts newly generated event sequences into Q. Finally, the master processor stops all the worker processors when Q becomes empty.

3.3.3 *Event handler scheduling*

Event handler scheduling is the process that invokes event handlers corresponding to the generated event sequences. The rules of scheduling is defined by AEM module. The AEM is the key to scheduling event handlers properly.

The application execution model specifies the rules of scheduling event handlers. Generally, an application execution model can be considered as an independent module. It takes an event as input, updates the application state, and determines which event handlers to invoke. The application's lifecycle state is stored in the AEM. Like event generation, the application execution engine also schedules event handlers in an iterating manner. Each time an event is fed to the application, the engine consults the AEM about the event handlers to schedule, and invokes a sequence of event handlers orderly.

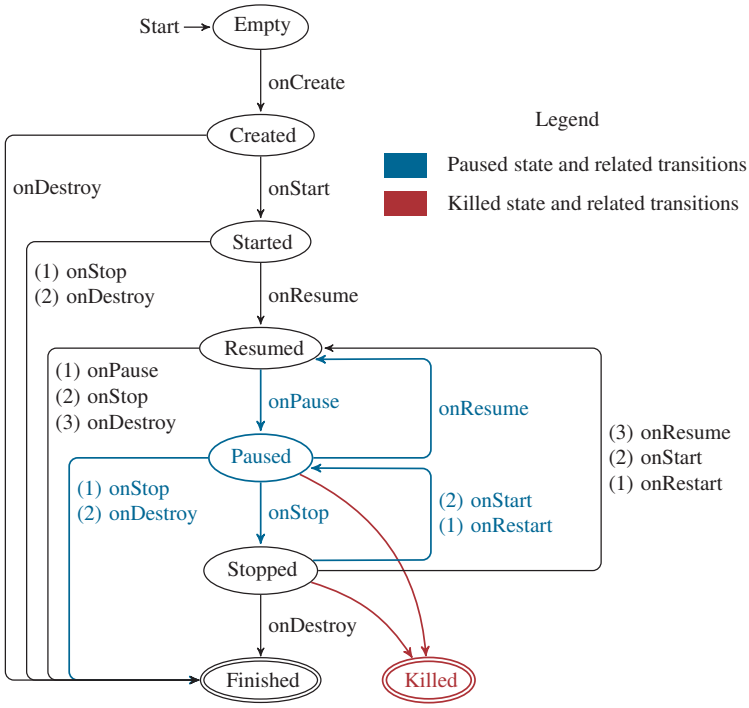

Figure 4 (Color online) NavyDroid AEM model (the activity part).

The application execution model has various forms in existing approaches. GreenDroid defines AEM in the form of temporal rules, denoted as $[\psi], [\phi] \Rightarrow \lambda$. This formula represents the expected event handers λ given the execution history ψ and the current situation ϕ [4]. E-GreenDroid translates temporal rules to an abstract state machine that is practical in programming [5]. However, the AEM of E-GreenDroid is not accurate enough in some lifecycle states of Android components. Therefore, we propose a new application execution model to simulate the runtime behavior of Android applications. We present the formal definition and the improvement of our AEM as follows.

The AEM is represented as a strengthened deterministic finite automaton (DFA). The inputs and inner states of the DFA represent user events, and an application's lifecycle states, respectively. It is also enhanced with additional outputs of event handlers.

Definition 3 (application execution model)**.** An application execution model AEM is a 7-tuple, $(S, E, H, \delta, f, s_0, F)$, consisting of

- a finite set of lifecycle states (S),
- a finite set of input events (E),
- a finite set of output event handler sequences (H),
- a transition function $(\delta : S \times E \to S)$,
- a scheduling function $(f : S \times E \to H)$,
- a start lifecycle state $(s_0 \in S)$,
- a set of final lifecycle states $(F \subseteq S)$.

Our AEM differs from a common DFA in a scheduling function f and an output set H. If an application receives user event e at lifecycle state s, a list of event handlers $h = f(s, e) \in H$ is expected to schedule. Figure 4 shows the transition diagram of our application execution model. In Figure 4, a node represents a lifecycle state $s \in S$, and an edge represents a state transition. Scheduled event handlers are marked on the edges, correponding to the state transitions. The states and transitions of AEM associate closely with the activity lifecycle. For example, when the AEM enters the resumed state, the `onResume()` callback is scheduled, and the activity comes to foreground; when the AEM enters the stopped state, the `onPause()` and `onStop()` callbacks are scheduled, and the activity switches to background.

Compared with E-GreenDroid's state machine model, our AEM is more accurate in modeling (1) an

Table 1 The correlation between activity state and the likelihood of the system's killing the process

Activity state	Likelihood of being killed
Created, started, resumed	Least
Paused	More
Stopped, destroyed	Most

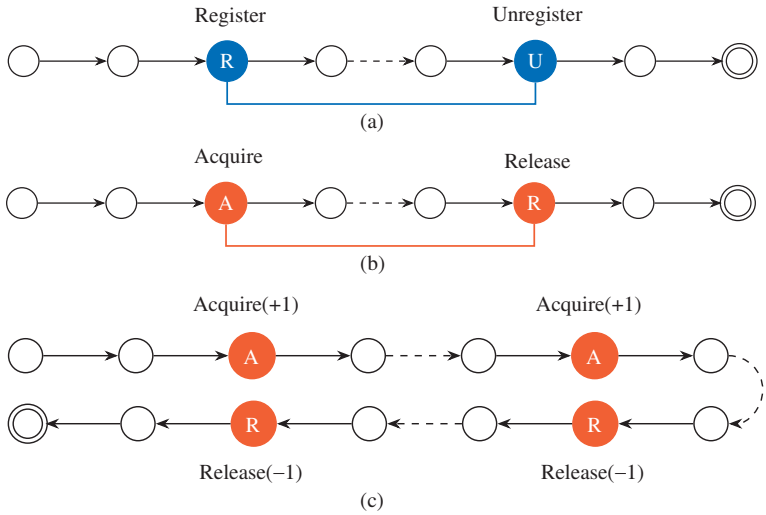

Figure 5 (Color online) Use modes of (a) sensor listeners, (b) wake locks (not reference counted), and (c) wake locks (reference counted).

activity that stays paused, and (2) an activity killed by the system.

An activity has an intermediate paused state between its resumed and stopped states. The paused state of an activity means that the activity is not at foreground, but still visible. An activity stays paused when a new, translucent activity (such as a dialog) keeps open. The Android system may kill an activity in the absence of memory. Table 1 illustrates that an activity is likely to be killed at the paused, stopped, and destroyed states.

In summary, we construct an application execution model that specifies the event handler scheduling rules accurately. With event sequence generation, state exploration and event handler scheduling, the application execution engine simulates the runtime behavior of an application, and reaches different application states.

3.4 Sensor and wake lock misuse check

As mentioned before, the misuse checker is one of the components of the monitor part. It corresponds to the first energy inefficiency pattern, i.e., missing sensor listener or wake lock deactivation. The misuse checker monitors the register/unregister of sensor listeners and acquire/release of wake locks during the simulation execution of an application.

Sensor listeners and wake locks are both interfaces that request system resources. A sensor listener requests corresponding sensors to work and fetches sensor data, and a wake lock requests the CPU to stay awake. Therefore, the misuse patterns of them are similar to some extent. In the following, we discuss our misuse patterns of sensor listeners and wake locks, and the detecting methods in NavyDroid.

Sensor listener misuse patterns. Figure 5(a) illustrates the usage mode of sensor listeners[13]. A sensor listener is registered for fetching sensor data. When the sensor data is no longer needed, the sensor listener should be unregistered. NavyDroid monitors and records all operations on sensor listeners. It checks execution paths for unregistered sensor listeners.

Wake lock misuse patterns. When without reference counts, wake locks has behavior similar to sensor listeners. However, the usage mode and misuse pattern of reference-counted wake locks are more

13) Liu Y P. Percom 2013 Presentation. http://sccpu2.cse.ust.hk/andrewust/files/PerCom_2013_Presentation.pdf, 2013.

Table 2 Common patterns of wake lock misuses and existing tools' capability

Misuse pattern	Number of issues	Related to energy waste?	E-GreenDroid solvable?	NavyDroid solvable?
Unnecessary wakeup	11	Yes	–	–
Wake lock leakage	10	Yes	Solvable	Solvable
Permature lock releasing	9	No	–	–
Multiple lock acquisition	8	Yes	–	Solvable
Inappropriate lock type	8	Yes	–	–
Problematic timeout setting	3	No	–	–
Inappropriate flags	2	Yes	–	–
Permission errors	2	No	–	–

complicated than sensor listeners. The reference counts of wake locks are like semaphores. An acquire operation increases the counter, while a release operation decreases the counter. A wake lock keeps held as long as its reference count is larger than zero. Therefore, in an execution path, the number of release operations of a wake lock should be no less than the number of acquire operations. Figure 5(b) and (c) illustrates the two different usage modes of wake locks.

An empirical study shows the eight common misuse patterns of wake locks [7]. As Table 2 shows, five out of these patterns relate to energy inefficiency. The wake lock leakage pattern corresponds to wake locks without reference counts, and the multiple lock acquisition pattern corresponds to wake locks with reference counts. Existing work, such as GreenDroid and E-GreenDroid, only detects the former pattern, while NavyDroid addresses these two patterns of wake lock misuses. We do not address all misuse patterns related to energy inefficiency, because some patterns lack detection criteria.

In summary, the misuse checker of NavyDroid monitors the operations of sensor listeners and wake locks to check for misuses. It supports a new misuse pattern of wake locks. In the analysis report, NavyDroid records the misused sensor listeners and wake locks with related application paths.

3.5 Sensor data utilization analysis

Sensor data utilization analyzer is another component of the monitor part. It corresponds to the second energy inefficiency pattern, i.e., sensor data underutilization. The sensor data utilization analyzer generates and propagates sensor data, and evaluates how sensor data is utilized at different application states.

As mentioned before, the sensor itself consumes energy to retrieve sensor data, thus applications should use sensor data efficiently. The utilization of sensor data can vary among application states. The analyzer reports application states with low utilization of sensor data. We do not analyze wake lock utilization, because wake locks only relate to CPU power management, and do not retrieve or use data like sensors.

The analysis of sensor data utilization can be divided into the following three phases.

Tainting. The analyzer generates sensor data object, each with a unique taint mark. The sensor data are generated from a data pool. NavyDroid feeds sensor data to the application when the application registers a sensor listener. When sensor data objects are constructed, NavyDroid assigns taint marks to these objects.

Propagation. The analyzer keeps track of the transformation of sensor data, and propagates taint marks via data flow. The propagation is at the level of bytecode instruction. The result of an instruction is tainted if one of the operands of this instruction is tainted. NavyDroid follows a collection of propagation rules [4]. The tainting and propagating phase both leverages JPF's object attribution functionality.

Evaluation. The analyzer evaluates the utilization of sensor data at executed application states. We calculate the data utilization coefficient (DUC) of each application state as [8]

$$\mathrm{DUC}(d,s)=\frac{\mathrm{usage}(s,d)}{\max_{s'\in S,d'\in D}\mathrm{usage}(s',d')}. \tag{1}$$

The DUC value of sensor data d at state s is defined as the ratio between the usage of d at state s and the maximum usage of any sensor data at any state. The usage of sensor data d at state s is defined

as [8]

$$\text{usage}(s,d) = \sum_{i \in \text{instr}(s,d)} \text{weight}(i,s) \times \text{rel}(i), \tag{2}$$

where $\text{instr}(s,d)$ denotes the set of bytecode instructions executed after d is fed to the application. Indicator function $\text{rel}(i)$ tests whether an instruction i uses any data with the same mark as d has. Function $\text{weight}(i,s)$ assigns a weight to instruction i to measure the benefits that i brings to the user.

After calculating DUC values of application states, we filter out application states with low DUC values. A low DUC value represents a low utilization of sensor data, and indicates an energy inefficiency bug. In this way, the analyzer identifies the inefficiency in sensor data utilization at certain application states. In the analysis report, NavyDroid records low sensor data utilization with related application paths and states.

4 Evaluation

In this section, we demonstrate the effectiveness and efficiency of our approach by experiments. NavyDroid is extended from E-GreenDroid, as E-GreenDroid is a state-of-the-art energy inefficiency detection tool. We replace the application execution model of E-GreenDroid with our strenthened DFA, in order to simulate the paused and killed states of activities accurately. First, we compare the effectiveness of NavyDroid and E-GreenDroid. On one hand, we evaluate whether NavyDroid can detect the equivalent energy inefficiency bugs that E-GreenDroid reports. On the other hand, we evaluate the enhanced abilities of NavyDroid as compared with E-GreenDroid. Second, we compare the efficiency of E-GreenDroid, sequential NavyDroid, and parallel NavyDroid. We aim to answer the following three research questions:

- **RQ1: ability equivalence.** Does NavyDroid hold the abilities as E-GreenDroid does; i.e., can NavyDroid conduct effective analysis on those applications with energy inefficiency bugs that E-GreenDroid can effectively analyze?
- **RQ2: ability enhancement.** Can NavyDroid detect energy inefficiency problems that E-GreenDroid is not able to detect?
- **RQ3: efficiency of state exploration.** With its (parallel) state exploration algorithm, does NavyDroid analyze applications more efficiently than E-GreenDroid?

4.1 Experimental setup

We selected different test subjects for the three research questions. For RQ1, we selected all the 13 Android applications that were used in the evaluation of E-GreenDroid [5]. The basic information of these 13 applications is shown in Table 3. E-GreenDroid reports energy inefficiency bugs in all of them. For RQ2, we selected six Android applications as test subjects. The basic information of these applications is shown in Table 4. They are all real-world applications with a relatively large scale (more than 1000 lines of code). For RQ3, we took all the applications in RQ1 and RQ2 as test subjects and evaluated the efficiency of analyzing them. We obtained the source codes of all these applications, as they are all open source projects.

Our experiments were conducted on a quad-core computer with Intel Core i7 CPU and 8GB RAM. The machine was installed with Ubuntu 16.04 LTS. We compiled each application on Android 5.0 for our experiments. For RQ1 and RQ2, we controlled E-GreenDroid to randomly generate 5000 event sequences, and NavyDroid to systematically explore application states. The lengths of event sequences were limited to six at most. This length is enough for E-GreenDroid and NavyDroid to explore considerable application states and detect energy inefficiency bugs. For RQ3, we ran parallel state exploration with four working threads.

We ran both E-GreenDroid and NavyDroid to diagnose all the test subjects, and examined the analysis reports to compare their abilities of locating energy inefficiency bugs. For RQ3, we evaluated and compared the efficiency of (1) E-GreenDroid exploring application states randomly; (2) NavyDroid exploring application states sequentially (with a single thread) and (3) NavyDroid exploring application states

Table 3 Information and analysis results of RQ1's test subjects

Application	Revision	Lines of code	E-GreenDroid results [a)]	NavyDroid results [a)]	Equivalent [b)]
AndTweet	V-0.2.4	8908	WLM	WLM	Yes
AAT	V-0.9-alpha	52800	SDU	SDU	Yes
BabbleSink	R-d12879a	1718	WLM	WLM	Yes
CWAC-Wakeful	R-d984b89	896	WLM	WLM	Yes
GPSLogger	R-15	659	SLM, SDU	SLM, SDU	Yes
GPSLogger-new	–	789	SLM, SDU	SLM, SDU	Yes
LocWriter2	V-0.1.1	1542	SDU	SDU	Yes
OmniDroid	R-863	12427	SDU	SDU	Yes
OsmDroid	R-750	18091	SDU	SDU	Yes
Recycle Locator	R-68	3241	SLM	SLM	Yes
RedBlackTree	R-0	483	WLM	WLM	Yes
Sofia Public Transport Nav.	R-114	1443	SDU	SDU	Yes
Ushahidi	R-9d0aa75	10186	SLM	SLM	Yes

a) We denote SLM as sensor listener misuse, WLM as wake lock misuse, and SDU as sensor data underutilization.
b) Whether E-GreenDroid and NavyDroid report the equivalent results.

Table 4 Information and analysis results of RQ2's test subjects

Application	Revision	Lines of code	E-GreenDroid results [a)]	NavyDroid results [a)]	Improved [b)]	Cause [c)]
AndroidRun	R-dbf6428	1649	SLM [d)]	SLM (1) [e)]	Yes	AEM
VLC	R-abe60f5	6839	NPD	WLM (1)	Yes	MLA
CSipSimple [f)]	R-152	13107	NPD	NPD (1)	No	–
TomaHawk	R-543c3b9	4601	NPD	WLM (2)	Yes	AEM, MLA
AndrOBD	R-d1a6ba0	24161	NPD	WLM (2)	Yes	AEM
MTPMS	R-935ceeb	1217	NPD	SLM (1)	Yes	MLA

a) We denote NPD as no problem detected, SLM as sensor listener misuse, and WLM as wake lock misuse. The numbers in parentheses indicate the number of bugs that NavyDroid detected for each application.
b) Explaining whether the results of NavyDroid are better than E-GreenDroid.
c) Explaining why NavyDroid reports better analyzing results than E-GreenDroid. AEM represents the improvement in application execution model, and MLA represents the ability of monitoring the misuse pattern of multiple lock acquisition.
d) The SLM bug that E-GreenDroid reports in AndroidRun is a false positive. We will discuss this in detail.
e) NavyDroid reports multiple energy inefficiency problems, but they are caused by one defect after manual analysis. We exclude the essentially equivalent reported bugs.
f) CSipSimple has six problematic revisions according to the empirical study. We compress these six revisions because the results are all NPDs.

parallely (with multiple threads). Also, we showed by examples that the systematic state exploration algorithm of NavyDroid improves efficiency by eliminating duplicated event sequences. In the following, we elaborate on our experimental results with respect to the three research questions.

4.2 RQ1: ability equivalence

In order to answer RQ1 about the ability equivalence of NavyDroid and E-GreenDroid, we ran both E-GreenDroid and NavyDroid to analyze the test subjects. There are three types of energy inefficiency bugs, namely, sensor listener misuse, wake lock misuse, and sensor data underutilization. NavyDroid and E-GreenDroid detected the same types of energy inefficiency bugs.

Table 3 lists the analysis results collected from E-GreenDroid and NavyDroid. E-GreenDroid and NavyDroid reported all misuses of sensor listeners and wake locks if they detect these two types of bugs. For sensor data underutilization bugs, they reported all suspicious application states with DUC values less than 1.0 for user's decision. In our experiments, we defined a DUC value that is less than 0.5 as severe underutilization. We considered a severe sensor data underutilization at an application state as an energy inefficiency bug.

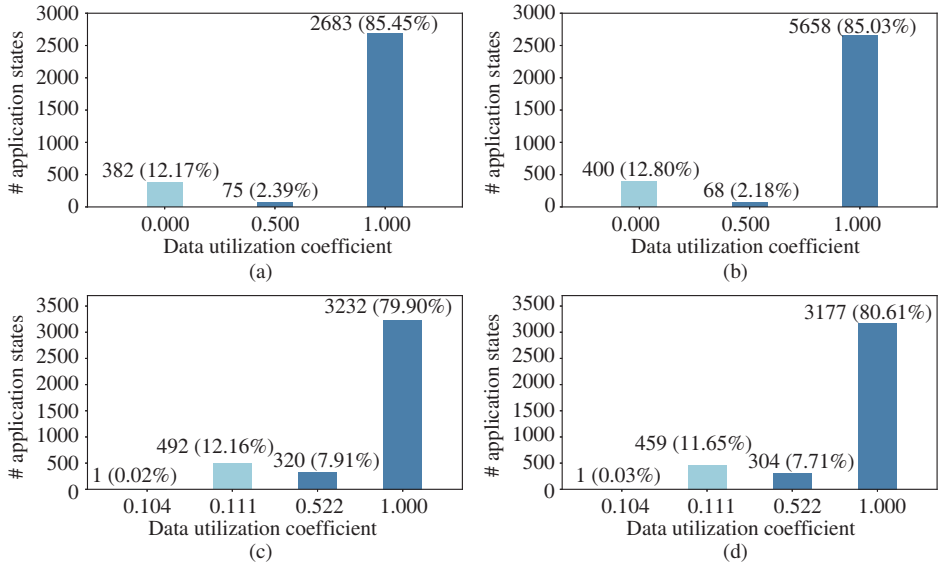

Figure 6 (Color online) DUC value distribution bar charts. DUC values of (a) GPSLogger analyzed by E-GreenDroid, (b) GPSLogger analyzed by NavyDroid, (c) OsmDroid analyzed by E-GreenDroid, (d) OsmDroid analyzed by NavyDroid.

We compared the detailed information in the analysis reports of E-GreenDroid and NavyDroid to further demonstrate the equivalence of energy inefficiency bugs. For energy inefficiency bugs of types sensor listener misuse and wake lock misuse, we compared the execution traces of them. Only two bugs with the same type and the equivalent execution traces are considered as equivalent[14]. For energy inefficiency bugs of type sensor data underutilization, we considered all application states with DUC values less than 1.0, and compared the distributions of DUC values. As is shown in Figure 6(a)–(d), the DUC value distribution for GPSLogger and OsmDroid in the report of E-GreenDroid and NavyDroid were essentially the same. In summary, we can answer RQ1 that NavyDroid locates all the energy inefficiency bugs reported by E-GreenDroid and holds the abilities of E-GreenDroid.

4.3 RQ2: ability enhancement

In order to answer RQ2 about the ability enhancement of NavyDroid as compared with E-GreenDroid, we analyzed several new test subjects with E-GreenDroid and NavyDroid. We selected six real-world Android applications as test subjects. Table 4 lists the basic information and analysis results of these applications. We collected and compared the analysis reports of E-GreenDroid and NavyDroid in the same way as RQ1. E-GreenDroid either detected no energy inefficiency bugs in an application, or reported false positives. Whereas, NavyDroid reported seven energy inefficiency bugs in six applications. We discuss some test subjects and the detected problems as follows.

AndroidRun. AndroidRun is a running and biking aiding application[15]. It tracks the location and calculates the speed of users. NavyDroid reported a sensor listener misuse in AndroidRun. The main activity registers a location listener (a kind of sensor listeners) in its `onCreate()` callback to collect GPS data. The unregister operation of this location listener is in the `onDestroy()` callback. However, if the activity gets killed by the Android system when it switches to the background, the `onDestroy()` callback will never be invoked, leaving the location listener still registered. E-GreenDroid also reported this sensor listener misuse. However, in E-GreenDroid's monitor part, after all the events in a event sequence are fed into the application and the corresponding event handlers are scheduled, the misuse checker immediately checks for sensor listener misuses. By this time, activities have not finished their lifecycles yet. Therefore, we considered the reported sensor listener misuse as a false positive. NavyDroid avoids this false positive,

14) The execution traces in analysis reports are not exactly the same because NavyDroid represents application states in an alternative way. We define two execution traces as equivalent when their event handlers and components are the same.

15) AndroidRun. https://sourceforge.net/projects/androidrun/, 2017.

but it can still report the misuse of this location listener, because it properly simulates the killed state in the application execution model.

VLC. VLC is an open source media player and framework[16]. NavyDroid reported an unbalanced acquires and releases of a wake lock in VLC. The `VideoPlayerActivity` plays user-specified videos. It acquires a wake lock to keep the screen on while playing the video. When a user plays or pauses video playing, the wake lock gets acquired or released. However, the wake lock can also be acquired when this activity is resumed and finishes loading the media file. If a user navigates to this activity and clicks the play/pause button twice, the wake lock will be acquired twice but released only once. According to the multiple lock acquisition pattern, a wake lock misuse happens, and NavyDroid detected this misuse. E-GreenDroid failed to detect this wake lock misuse because it does not consider the number of acquires and releases of wake locks.

CSipSimple. CSipSimple is a communication application, which supports Session Initiation Protocol (SIP) connection over the Internet[17]. A recent empirical study [7] reports wake lock misuses in CSipSimple. However, both E-GreenDroid and NavyDroid detected no energy inefficiency problems in it. CSipSimple maintains a `SipService` at background to register accounts. When there is successful registeration of an account, the service acquires a wake lock to keep the screen on. The number of acquire operations is not restricted, as the registeration of accounts can be repeated. Nevertheless, the service sets the wake lock to be non reference-counted, thus the wake lock will not keep held after one release operation. So we considered that CSipSimple has no misuses about this wake lock, and NavyDroid reported a reasonable analysis result.

TomaHawk. TomaHawk is a music player which supports multi-source media. NavyDroid reported two wake lock misuses in TomaHawk. As described in the motivating example in Subsection 2.2, the two bugs occur in different cases. If a user clicks the play/pause button twice after the media player gets prepared, the wake lock is acquired twice but released once, and remains being held as its reference count is larger than zero. E-GreenDroid fails to detect this wake lock misuse because it does not monitor the reference counts of wake locks. Another bug occurs when the Android system kills the activity and the service associated with media playing, in which case the `onDestroy()` callback is not invoked, causing wake lock leakage. E-GreenDroid fails to detect this wake lock misuse because it does not simulate the killed state of an activity.

4.4 RQ3: efficiency of state exploration

In order to evaluate the efficiency of the parallel exploration algorithm, we took all the applications in RQ1 and RQ2 as test subjects. We analyzed these applications with (1) the random event sequence generation of E-GreenDroid, (2) the sequential state exploration of NavyDroid, and (3) the parallel state exploration of NavyDroid. For all three cases, we compared the average execution time of applications. As discussed in experimental setup, we ran parallel exploration with four working threads. Also, we restricted the sequential algorithm to work on one CPU core. We controlled E-GreenDroid to generate 5000 event sequences for each test subject. For NavyDroid, because some of the test subjects have too many states to explore, we restricted the state exploration to generate at most 5000 event sequences.

We recorded the number of executed event sequences, and the execution time for the three cases. Then we calculated the average time of each test subject. Table 5 and Figure 7 presents the analysis results of the three cases. Although the execution time of NavyDroid is larger than that of E-GreenDroid by about 50%, the parallelization of the state exploration algorithm reduces execution time by about 60%. Overall, NavyDroid saves about 40% execution time compared with E-GreenDroid.

Moreover, NavyDroid improves efficiency by generating more unrepeated event sequences. The exploration algorithm ensures that all the generated event sequences are unrepeated. However, if generated randomly, the event sequences will contain repeated ones. We selected AndTweet and OmniDroid as representatives in order to show the efficiency of removing duplicated event sequences. Our state exploration

16) VLC. https://github.com/mstorsjo/vlc-android, 2017.
17) Csipsimple. https://github.com/r3gis3r/CSipSimple, 2017.

Table 5 Average execution times of sequential and parallel state exploration algorithms [a)]

Application	Number of explored event sequences	Average execution time of E (s)	Average execution time of N1 (s) [b)]	Average execution time of N4 (s) [c)]
AAT	1456	1.284	2.048 (+60%)	0.816 (−60%) (−36%)
AndrOBD	1645	1.944	3.376 (+74%)	1.122 (−67%) (−42%)
AndroidRun	189	1.710	3.101 (+81%)	1.413 (−54%) (−17%)
AndTweet	1456	1.313	1.886 (+44%)	0.801 (−58%) (−39%)
BabbleSink	189	1.084	1.783 (+64%)	0.693 (−61%) (−36%)
CWAC-Wakeful	189	1.108	1.693 (+66%)	0.624 (−63%) (−39%)
GPSLogger	5000	1.310	1.954 (+49%)	0.778 (−60%) (−41%)
GPSLogger-new	5000	1.388	2.123 (+53%)	0.854 (−60%) (−38%)
LocWriter2	1456	1.244	1.992 (+60%)	0.786 (−61%) (−37%)
OmniDroid	5000	3.952	6.360 (+61%)	2.485 (−61%) (−37%)
OsmDroid	5000	1.926	3.093 (+61%)	1.331 (−57%) (−31%)
Recycle Locator	189	1.394	2.016 (+45%)	0.772 (−62%) (−45%)
RedBlackTree	5000	1.062	1.678 (+58%)	0.616 (−63%) (−42%)
Sofia Public Transport Nav.	432	1.778	3.421 (+92%)	1.461 (−57%) (−18%)
TomaHawk	1456	1.222	2.035 (+67%)	0.804 (−60%) (−34%)
MTPMS	189	1.078	1.720 (+60%)	0.683 (−60%) (−37%)
Ushahidi	432	1.596	2.468 (+55%)	1.012 (−59%) (−37%)
VLC	1277	1.326	2.107 (+59%)	0.892 (−58%) (−33%)

a) E denotes E-GreenDroid's random event sequence generation, N1 denotes NavyDroid's state exploration algorithm with one single thread, and N4 denotes NavyDroid's state exploration algorithm with four threads.

b) The percentages in parentheses indicate the rates of change relative to E-GreenDroid's average execution time.

c) The percentages in parentheses indicate the rates of change relative to the average execution time of single-threaded NavyDroid and to E-GreenDroid's average execution time, respectively.

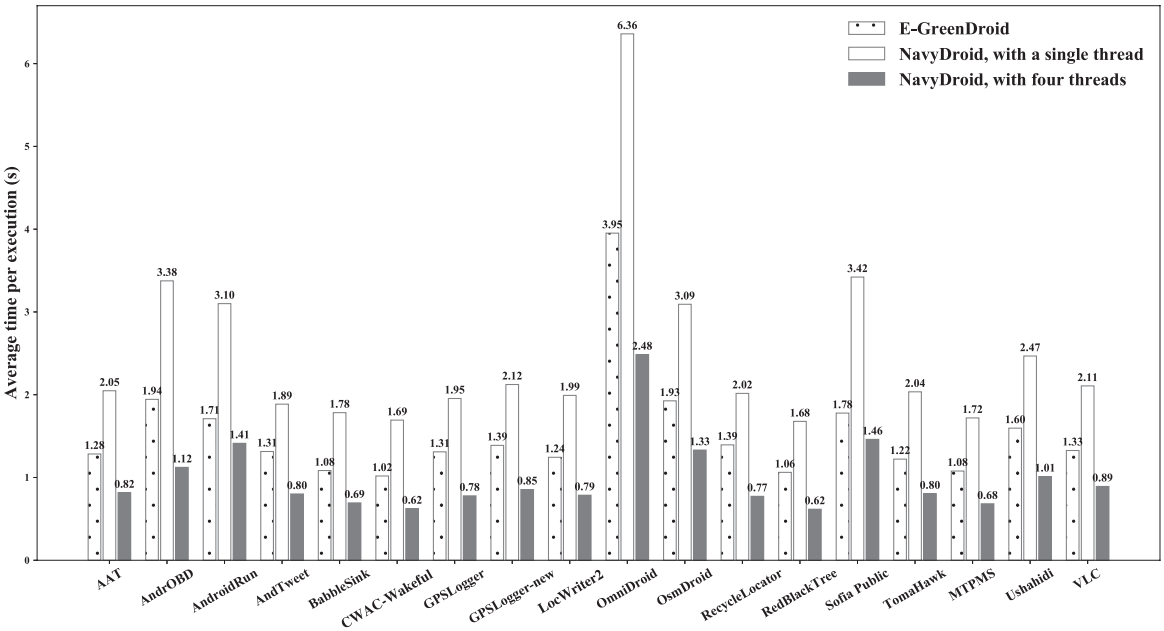

Figure 7 Average execution time of sequential and parallel state exploration.

algorithm generates 1456 and 13384 event sequences for AndTweet and OmniDroid, respectively.

Figure 8 shows the trend of unrepeated event sequence number as the total event sequence number increases. In Figure 8, the faster the curve rises, the higher the efficiency of state traversal is. Our state exploration algorithm reaches the optimal efficiency of state traversal (every generated event sequence

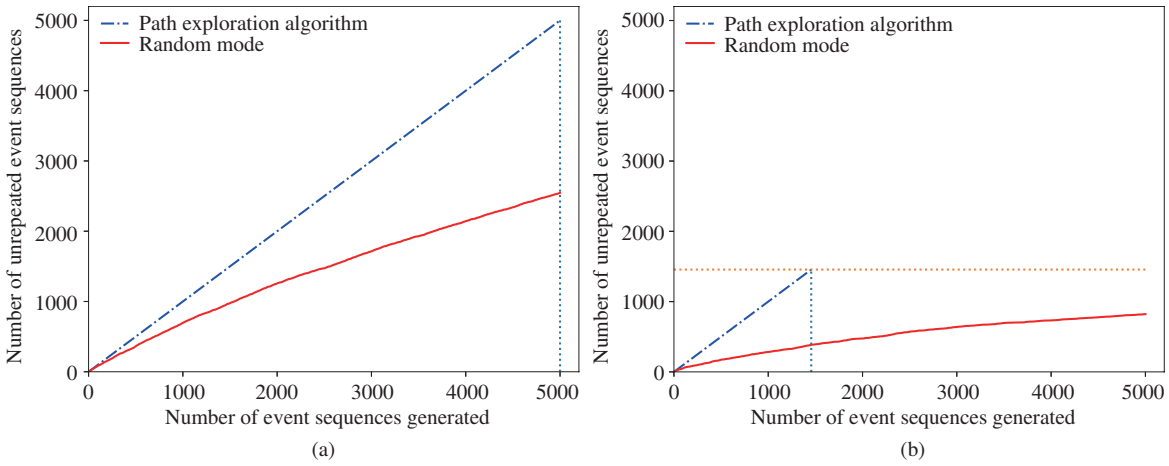

Figure 8 (Color online) The trend of unrepeated event sequence number. (a) OmniDroid; (b) AndTweet.

is unrepeated). As shown in Figure 8(a), for OmniDroid, in 5000 executions, the random mode only generates 2546 valid (unrepeated) event sequences. The efficiency is only 51% of the state exploration algorithm. As shown in Figure 8(b), for AndTweet, as the total number of unrepeated event sequences is 1456, the state exploration algorithm stops after generating 1456 event sequences. However, the random mode only generates 384 unrepeated event sequences in 1456 executions, and 822 valid event sequences in 5000 executions.

Therefore, it is not efficient to generate event sequences randomly due to lots of repeated event sequences. Also, as the number of randomly generated event sequences increases, the number of repeated ones grows and the efficiency decreases. Our state exploration algorithm exceeds random generation in efficiency by removing duplicated event sequences.

In summary, the state exploration algorithm of NavyDroid is efficient in two aspects. First, the state exploration algorithm generates more unrepeated event sequences than random generation within the same time. Second, by parallelizing the exploration algorithm, lots of time can be saved. If necessary, we can extend our parallel exploration algorithm to distributed systems and reduce execution time by running on multiple servers.

5 Discussion

Like E-GreenDroid, NavyDroid is also implemented on top of JPF. We choose JPF as the analysis framework because of its functionality of intercepting the execution of Java applications. One limitation of JPF is that it can only analyze traditional Java programs. In our approach, the application execution model guides the simulation execution of Android applications. It stores an application's lifecycle state, and determines the scheduling of event handlers. Our AEM may still be inaccurate when simulating the execution of Android applications in some cases. However, it is highly extensible and can be easily updated.

In this paper, we present NavyDroid as an energy inefficiency diagnosis tool. Furthermore, the application execution model can serve as a general dynamic analysis tool, since it generates event sequences for applications. As discussed before, the diagnosis tool consists of a simulation part and a monitor part. The application execution model is the key to the simulation part. It defines the scheduling policies of event handlers, and provides native framework libraries for applications. We can assemble the simulation part with different monitor parts, i.e., components with different monitor points and problem patterns. With different monitor policies, NavyDroid can detect more bugs in Android applications. Therefore, the application execution model and the simulation part can become a more general dynamic analysis tool for Android applications.

Currently, NavyDroid is able to detect two wake lock misuse patterns. NavyDroid cannot address the

unnecessary wakeup pattern, which is the most common pattern of wake lock misuse according to the empirical study [7]. We plan to extend NavyDroid to analyze this pattern in future. There are also some other misuse patterns of wake locks such as inappropriate lock type and inappropriate flags. However, these two patterns currently lack feasible diagnosis criteria, and thus are difficult to detect automatically.

The Android system hands over part of the energy management power to developers by exposing some resource management APIs to developers. Thus, misuse patterns are usually related to these resource management APIs. In this work, we select sensor listener and wake lock as two representative APIs, and derive three misuse patterns from them. We can probably find more misuse patterns from resource related APIs, so as to detect more bugs in applications. For instance, `MediaRecorder` is used to record audio and video, and `Camera` is used to take pictures, and they can potentially cause energy inefficiency [1].

6 Related work

Our work relates to existing studies of several research topics, which includes energy inefficiency analysis and wake lock misuse detection. In this section, we discuss some representative pieces of work in recent years.

Energy efficiency analysis. In recent years, researchers have proposed various energy inefficiency diagnosis approaches for smartphone applications. Some pieces of work use static analysis techniques to detect energy inefficiencies. Pathak et al. [1] detected no-sleep energy bugs in Android applications with reaching-definition data-flow analysis. Their work is probably the first to explore and define the characteristics of no-sleep energy bugs. Many researchers proposed apporaches based on dynamic analysis techniques to avoid false positives in static analysis. Liu et al. [4] characterised energy inefficiency caused by sensor data underutilization, and proposed GreenDroid that searches application states for low sensor data utilization. GreenDroid detects energy inefficiency bugs with taint-based dynamic data-flow analysis. Zhang et al. [2] presented a framework named ADEL to detect energy leaks of network data. ADEL is similar to GreenDroid in evaluating the utilization of program data. CyanDroid systematically generates multi-dimensional sensor data to reproduce energy inefficiency bugs that require specific sensor data to manifest [3]. Wang et al. [5] updated and optimized the execution simulation process of GreenDroid.

Some approaches focus on optimizing and repairing energy inefficiency problems. Ma et al. [9] presented eDoctor, a practical tool that captures an application's time-varying behavior in order to identify the abnormal behavior of an application. It suggests repair solutions to users, and helps regular users troubleshoot energy inefficiency problems. Manotas et al. [10] presented a general framework that automatically selects the most energy-efficient library implementations for optimizing Java applications. Bouquet detects and bundles HTTP requests at runtime in order to reduce the energy consumption of HTTP requests [11]. Li et al. [12] proposed an approach to fixing sensor data underutilization automatically through instrumentation of applications. Banerjee et al. [13] developed EnergyPatch, a framework that detects, validates and repairs energy inefficiencies in Android applications. EnergyPatch uses a static analysis technique to detect potential energy bugs, and validates the presence of energy bugs with a dynamic analysis technique.

Energy consumption estimation. Energy consumption estimation profiles energy hotspots (i.e., components that consume the most energy) in an application. A common practice with this type of approaches is to execute the application with different test cases, while monitoring the energy consumption of the device. Pathak et al. [14] proposed Eprof that records invocations of system APIs, and estimates the energy consumption using a power model. Besides Eprof, several subsequent pieces of work proposed more fine-grained profiling techniques for estimating energy consumption. Both eLens [15] and vLens [16] estimate energy consumption for each line of source code in an application. Banerjee et al. [17] proposed a hardware-software hybrid approach to systematically generating test inputs that lead to energy bugs and energy hotspots. Energy consumption estimation approaches relate to energy inefficiency diagnosis approaches as they identify energy hotspots. However, energy hotspots are not equivalent to energy inefficiency bugs. The concept of energy inefficiency bugs emphasizes that the energy consumption is not

necessary and brings no user benefits.

Wake lock misuse detection. Some severe energy problems associate with misuses of wake locks. Pathak et al. [1] detected energy bugs caused by wake lock leakage using data-flow analysis. The approach proposed by Vekris et al. [18] also verifies wake lock misuses according to a collection of policies. GreenDroid and E-GreenDroid monitor and record the operations on wake locks during simulation executions [4, 5]. WLCleaner is another approach that repairs wake lock issues at runtime [19]. However, these pieces of work only address energy inefficiency bugs caused by wake lock leakage. Liu et al. [7] conducted an empirical study to understand common wake lock usage in practice, and summarized eight common patterns of wake lock misuses. Our work detects multiple wake lock misuses, including wake lock leakage, and unbalanced acquires and releases.

7 Conclusion

In this paper, we have proposed an approach to detecting energy inefficiency problems in Android applications. Our approach executes Android applications by generating event sequences and scheduling event handlers. The event handler scheduling is guided by an application execution model derived from Android specifications. We also designed a parallel algorithm to explore application states systematically. During the execution, it monitors the operations on sensor listeners and wake locks to detect misuses. It also evaluates sensor data utilization using a taint-based dynamic analysis.

Our approach exceeds existing approaches by a more accurate application execution model, and more energy inefficiency patterns. We implemented our approach as a prototype tool named NavyDroid, and evaluated it with real-world applications. The experimental results demonstrate its effectiveness and efficiency in detecting energy inefficiency problems.

In future, we plan to further extend this work to support more patterns of energy inefficiency problems. We also plan to re-implement our tool on top of the Android virtual machine. At present, dynamic analysis tools can only simulate the execution of Android applications and may be inconsistent with the actual execution of applications. The analysis can be more accurate by inspecting real application executions inside the Android virtual machine.

Acknowledgements This work was supported in part by National Key Research and Development Program of China (Grant No. 2017YFB1001801) and National Natural Science Foundation of China (Grant Nos. 61690204, 61472174). The authors would also like to thank the support of the Collaborative Innovation Center of Novel Software Technology and Industrialization, Jiangsu, China.

References

1 Pathak A, Jindal A, Hu Y C, et al. What is keeping my phone awake? Characterizing and detecting no-sleep energy bugs in smartphone apps. In: Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services, Low Wood Bay, 2012. 267–280

2 Zhang L D, Gordon M S, Dick R P, et al. Adel: an automatic detector of energy leaks for smartphone applications. In: Proceedings of the Eighth IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis, Tampere, 2012. 363–372

3 Li Q W, Xu C, Liu Y P, et al. CyanDroid: stable and effective energy inefficiency diagnosis for Android apps. Sci China Inf Sci, 2017, 60: 012104

4 Liu Y P, Xu C, Cheung S C, et al. GreenDroid: automated diagnosis of energy inefficiency for smartphone applications. IEEE Trans Software Eng, 2014, 40: 911–940

5 Wang J, Liu Y P, Xu C, et al. E-greendroid: effective energy inefficiency analysis for Android applications. In: Proceedings of the 8th Asia-Pacific Symposium on Internetware, Beijing, 2016. 71–80

6 Liu Y P, Wang J, Xu C, et al. NavyDroid: detecting energy inefficiency problems for smartphone applications. In: Proceedings of the 9th Asia-Pacific Symposium on Internetware, Shanghai, 2017

7 Liu Y P, Xu C, Cheung S C, et al. Understanding and detecting wake lock misuses for Android applications. In: Proceedings of the 24th ACM SIGSOFT International Symposium on the Foundations of Software Engineering, Seattle, 2016. 396–409

8 Liu Y P, Xu C, Cheung S C. Where has my battery gone? Finding sensor related energy black holes in smartphone applications. In: Proceedings of the IEEE International Conference on Pervasive Computing and Communications,

San Diego, 2013. 2–10
9 Ma X, Huang P, Jin X X, et al. Edoctor: automatically diagnosing abnormal battery drain issues on smartphones. In: Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation, Lombard, 2013. 57–70
10 Manotas I, Pollock L, Clause J. SEEDS: a software engineers energy-optimization decision support framework. In: Proceedings of the 36th International Conference on Software Engineering, Hyderabad, 2014. 503–514
11 Li D, Lyu Y J, Gui J P, et al. Automated energy optimization of HTTP requests for mobile applications. In: Proceedings of the IEEE/ACM 38th International Conference on Software Engineering, Austin, 2016. 249–260
12 Li Y C, Guo Y, Kong J J, et al. Fixing sensor-related energy bugs through automated sensing policy instrumentation. In: Proceedings of the IEEE/ACM International Symposium on Low Power Electronics and Design, Rome, 2015. 321–326
13 Banerjee A, Chong L K, Ballabriga C, et al. EnergyPatch: repairing resource leaks to improve energy-efficiency of Android apps. IEEE Trans Software Eng, 2017. doi: 10.1109/TSE.2017.2689012
14 Pathak A, Hu Y C, Zhang M. 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, Bern, 2012. 29–42
15 Hao S, Li D, Halfond W G J, et al. Estimating mobile application energy consumption using program analysis. In: Proceedings of the 2013 International Conference on Software Engineering, San Francisco, 2013. 92–101
16 Li D, Hao S, Halfond W G J, et al. Calculating source line level energy information for Android applications. In: Proceedings of the 2013 International Symposium on Software Testing and Analysis, Lugano, 2013. 78–89
17 Banerjee A, Chong L K, Chattopadhyay S, et al. Detecting energy bugs and hotspots in mobile apps. In: Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, Hong Kong, 2014. 588–598
18 Vekris P, Jhala R, Lerner S, et al. Towards verifying Android apps for the absence of no-sleep energy bugs. In: Proceedings of the 2012 USENIX Conference on Power-Aware Computing and Systems, Hollywood, 2012. 3
19 Wang X G, Li X F, Wen W. Wlcleaner: reducing energy waste caused by wakelock bugs at runtime. In: Proceedings of the 2014 IEEE 12th International Conference on Dependable, Autonomic and Secure Computing, Dalian, 2014. 429–434