Adaptive Deployment and Configuration for Mobile Augmented Reality in the Cloudlet

Tim Verbelen[a], Pieter Simoens[a,b], Filip De Turck[a], Bart Dhoedt[a]

[a]*Ghent University – IBBT, Department of Information Technology,*
Gaston Crommenlaan 8 bus 201, 9050 Gent, Belgium
[b]*Ghent University – IBBT, Ghent University College, Department INWE,*
Valentin Vaerwyckweg 1, 9000 Gent, Belgium

Abstract

Despite recent advances in mobile hardware, most mobile devices still fall short to execute complex multimedia applications with real-time requirements such as augmented reality (AR). Because offloading the application to the cloud is not always a solution due to the high and often unpredictable WAN latencies, the concept of cloudlets has been introduced: nearby infrastructure offering virtual machines for remote execution.

In this paper we present a cloudlet platform, providing two important contributions. First, the platform allows cloudlets to be formed in a dynamic way, including (fixed) virtualized infrastructure co-located with the wireless access point, as well as any device in the LAN network supporting the platform. The approach can also be extended towards the cloud, facilitating distribution of applications over three tiers (i.e. the device, the cloudlet and the cloud). Second, instead of moving a complete virtual machine to the cloudlet, we propose a more fine-grained approach, by managing and deploying applications on the component level. Application components are declared by the developer, together with their real-time constraints and configuration parameters, which are configured and distributed among the cloudlet by the platform at runtime depending on the application context, in order to meet the constraints and to optimize the user experience.

An OSGi-based prototype implementation on the Android platform is highlighted and evaluated using a mobile AR use case, showing the need for a component-based approach for the cloudlet.

Keywords: Cloudlet, Mobile Computing, Augmented Reality, OSGi

1. Introduction

As smartphones and tablets are gaining more and more capabilities (in terms of CPU power, network connectivity and sensors), they are becoming the preferred device for daily tasks such as browsing the web, e-mail, etc. Although high-end smartphones are equipped with dual- or quadcore processors, their processing capabilities remain an order of magnitude lower than their desktop counterparts, due to the limited battery capacity, form factor and passive cooling. Therefore, these mobile devices still fall short to execute future resource intensive multimedia applications, that involve real-time processing of video or audio, such as object recognition or mobile augmented reality [1].

To enable such resource-intensive applications, the concept of cyber foraging was introduced by Satyanarayanan et al. [2], where parts of the application are offloaded to nearby resources – called surrogates – connected via a wireless network. Such surrogates are recently defined as "cloudlets" [3], virtualized infrastructure co-located with the wireless access point, which allows mobile devices to offload applications to nearby resources in the LAN network.

Today many such cyber foraging systems exist, offloading applications either on a VM [4][3], software component [5][6][7] or a method level [8][9]. Although a lot of research has been done on profiling and partitioning applications and decision algorithms to decide when to offload, two important issues still remain.

First, almost all systems focus on the scenario where one mobile device offloads to one surrogate, called the cloudlet. However, in a realistic scenario multiple devices will be connected to the wireless network, some of them requiring resources to offload to (i.e. mobile devices or tablets) and other devices offering resources (i.e. servers or laptops). In this paper we present a cloudlet management platform, that extends the cloudlet to all devices in the network. This allows to take decisions optimizing behavior for all devices in the network, taking into account shared resources such as the network bandwidth, instead of optimizing on each device separately.

In addition to cloudlets consisting solely of fixed infrastructure (e.g. a corporate cloudlet), also ad hoc cloudlets can be formed with all devices available in the LAN network (e.g resources discovered in the home network or within a railway carriage), as depicted in Figure 1. The platform can also easily be extended towards the cloud, allowing application components to be
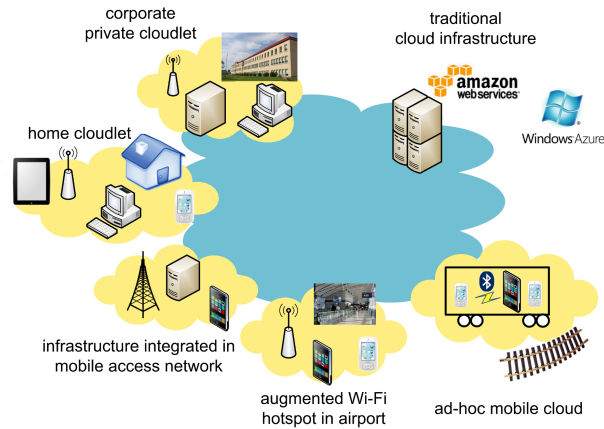
Figure 1: Static cloudlets can be provided by a corporation in a corporate cloudlet, or by a service provider in the mobile network. Ad hoc cloudlets can be formed in the home network or within a railway carriage.

distributed in a 3-tier fashion: locally on the mobile device, nearby on the cloudlet infrastructure or on a distant datacentre.

A second issue is that current cyber foraging systems aim to optimize a performance goal from a user perspective, such as energy usage, processing time or throughput. However, from a developer perspective, a more important goal is to achieve good performance on a wide range of devices, in which case "good" performance is an application specific metric. Such a metric can often be formulated as a number of constraints (e.g. maximum allowed execution time for a specific method call), especially in applications with real-time requirements.

To address this issue, our cloudlet platform manages applications on the component level, where the application developer declares application components with their offered functionality, configurable parameters and performance constraints. Application components are managed and configured at runtime, and can be distributed among resources in the cloudlet depending on the current context, in order to achieve the desired performance. In order to ease development for such constraint aware applications, we provide a programming model based on annotations, minimizing the burden put on the application developer.

We present a detailed architecture of the cloudlet platform. To show the need and effectiveness of the platform, we focus on an augmented reality use

3

case. An OSGi based prototype implementation is presented and evaluated on the Android mobile platform.

The remainder of this paper is structured as follows. In the next section we present a mobile augmented reality use case that illustrates the need for a component-based approach. This use case has been fully implemented to evaluated the cloudlet platform. Section 3 gives an overview of related work in the field of application offloading and cloudlets. In Section 4 we propose a general architecture for a cloudlet framework, and in Section 5 we discuss in more detail the implementation of a prototype based on OSGi, and the integration with the Android OS. To ease component-based application development supporting the platform, an annotation-based programming model and corresponding developer tools are presented in Section 6. In Section 7 we show the need and effectiveness of our approach for the mobile augmented reality use case, while Section 8 concludes this paper and discusses future work.

2. Use case: Mobile Augmented Reality

As a use case, we present a mobile augmented reality application featuring markerless tracking as described by Klein et al. [10], combined with an object recognition algorithm presented in [11]. The application is shown in Figure 2. On the right a greyscale video frame is shown with the tracked feature points, from which the camera position is estimated. The left shows the resulting overlay with a 3D object, and a white border around the recognized book.

Figure 2: The augmented reality application tracks feature points in the video frames (right) to enable overlaying of 3D objects (left).

We have split up the augmented reality algorithms of the application

and redesigned them to arrive at the following components (also shown in Figure 3):

VideoSource The VideoSource fetches video frames from the camera hardware. These frames are analyzed by the Tracker, and rendered together with an augmented reality overlay by the Renderer.

Renderer Each camera frame is rendered on screen together with an overlay of 3D objects. These 3D objects are aligned according to the camera pose as estimated by the Tracker.

Tracker The Tracker analyses video frames and calculates the camera pose by matching a set of 2D image features to a known map of 3D feature points. The map of 3D points is generated and updated by the Mapper.

Mapper From time to time the Tracker sends a video frame to the Mapper for map generation and refinement. By matching 2D features in a sparse set of so-called keyframes, the Mapper can estimate their 3D location in the scene and generate a 3D map of feature points.

Relocalizer When no 2D image features are found in the video frame, the Relocalizer tries to relocate the camera position until tracking resumes.

Object Recognizer In the keyframes of the Mapper the Object Recognizer tries to locate known objects. When an object is found, its 3D location is notified to the Renderer that produces an overlay.

These components are not only very CPU intensive, some of them also have strict real-time constraints. The VideoSource and the Renderer have to be executed on the mobile device, as they access device specific hardware. In order to achieve an acceptable performance, the Tracker and the Relocalizer should be able to process frames within 30-50 ms, which translates into a frame rate of 20-30 frames per second. As the Mapper runs as a background task, constantly refining and expanding the 3D map, this component preferably runs on a device with considerable CPU resources. However the Mapper does not impose strict real time requirements. The Object Recognizer also has more relaxed requirements, as delays in the order of a second before recognizing an object are still tolerable to achieve an acceptable user experience.
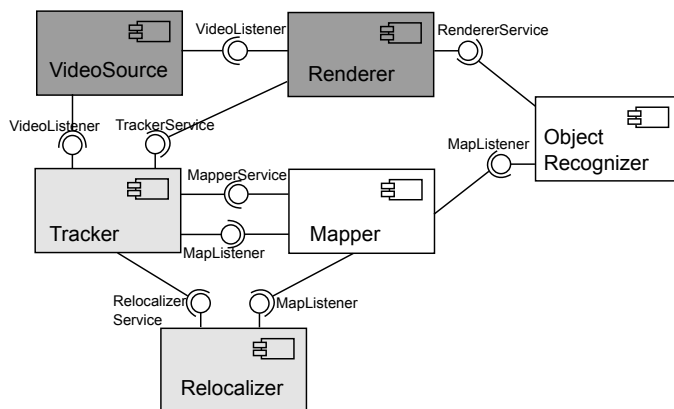
Figure 3: The component diagram of the AR application. The VideoSource and the Renderer (dark grey) are fixed on the mobile device. The Relocalizer and Tracker (grey) have real-time constraints ($\leq$ 50 ms). The Mapper and the ObjectRecognizer (white) do not have strict requirements.

The goal now is to run this application on a mobile device, while meeting all the required constraints, which have to be declared by the application developer. This use case shows that offloading the whole application to the cloudlet on a VM level is not sufficient, as sending frames to the cloudlet for processing and receiving rendered frames in return would require too much bandwidth, which is a scarce resource in current wireless access networks. Therefore we propose a cloudlet architecture that will manage the application on a component level, being able to distribute application components within the cloudlet or to other cloudlets.

In addition to component (re)distribution, QoS can also be assured by dynamically reconfiguring components, as shown in [12]. By adopting a component-based approach for the cloudlet, real-time constraints of applications can be achieved using combined dynamic reconfiguration and distribution. In our use case for example, the developer defined a parameter configuring the number of tracked features, which influences the frame processing rate and which can be adapted at runtime.

3. Related work

Offloading computation from mobile devices to remote surrogates, known as cyber foraging [13] has been subject for over a decade. In recent work,

Kristensen et al. [9] propose Scavenger, a cyber foraging system in Python which outsources Python methods, and the scheduler orchestrating this outsourcing process uses history-based profiles. MAUI [8] offloads methods of Microsoft .Net applications in order to save energy. At initialization time, MAUI measures the energy characteristics of the device, and at runtime the application and network characteristics are monitored. By solving an integer linear programming (ILP) problem a decision is made whether or not to offload a method. Chun et al. present CloneCloud [4], where virtualized clones of the mobile device are executed in the cloud. Different binaries of the application are generated in an offline profiling stage, with dedicated VM instructions added at migration points for selected methods. At runtime a clone VM is instantiated at the server side, and the application transparently switches between execution at the device or at the clone. Giurgiu et al. [5] and Verbelen et al. [6] use OSGi components offloadable units. Monitoring information is used to build a graph model of the software, from which graph cutting algorithms calculate the optimal deployment. Odesssa [7] is a cyber foraging system that focuses on data flow applications, trying to optimize the makespan and throughput of the application by exploiting offloading and data-parallelism.

These cyber foraging systems address the situation where the application is distributed between one mobile device and the cloud. However, in a wireless network context, network bandwidth is a shared medium, and thus there is need for a global management entity that optimizes the bandwidth for all users. Also, it is often left undefined what should happen when no surrogate is available. From the application developer viewpoint, one would prefer the application to gracefully degrade when no remote resources are available, or conversely, treat the situation when parts are offloaded as a quality enhancement on top of a more basic application functionality. To achieve this, the offloading framework should also be able to reconfigure the application at runtime.

As surrogates for remote execution, Goyal et al. [14] propose the usage of virtual machine technology. Because the deployment of virtual machines in a cloud can lead to high WAN latencies, Su et al. present Slingshot [15], where the VMs are co-located with the wireless access point. Satyanarayanan et al. [3] proposed the concept a of cloudlet: a trusted, resource-rich computer or a cluster of computers well connected to the Internet and available for use by nearby mobile devices. Cloudlets offer their resources to mobile devices by dynamic VM synthesis, where small VM overlays are sent to the cloudlet

from which a complete VM is created running the mobile application.

An important challenge when using VM based offloading consists of the fact that either the complete application has to be executed on the remote VM using a thin client approach, or the developer himself is responsible for managing which parts should be executed locally on the mobile device and which on the remote VM. Also, the migration cost is rather high, as a complete VM has to be started and configured.

To tackle these issues, we propose a more fine grained approach for the cloudlet, where the problem is tackled at the component level through a middleware approach. A cloudlet platform configures and distributes the application components at runtime, alleviating the developer from the task to decide which components to offload and how to configure the application for the current context. To facilitate application development on our framework, we present an intuitive programming model based on code annotations. The framework also handles resource discovery, enabling the dynamic formation of cloudlets, and has a global overview of all resources, which makes global optimization possible.

4. Cloudlet architecture

We adopt a component-based cloudlet architecture as proposed in [16], which is shown in Figure 4. As mentioned above, a cloudlet consists of a collection of interconnected resources in each other's vicinity. This architecture has three interrelated management levels: the component level, the node level and the cloudlet level.

Components (units of deployment specified by their providing and required interfaces [17]) are managed by an *Execution Environment (EE)*, that can start and stop components, resolve component dependencies, expose provided interfaces, etc. To support distributed execution, dependencies can be resolved with other (remote) Execution Environments. In that case, proxies and stubs are generated and the components can communicate by remote procedure calls (RPCs). Components can also define performance constraints (e.g. the maximum execution time of a method), and expose configuration parameters to the EE. By monitoring the resource usage of each component, the EE can detect violations of the performance constraints and actions can be taken such as calculating a new deployment (i.e. offloading some resource intensive components) or adapting component configurations (i.e. lowering component quality).
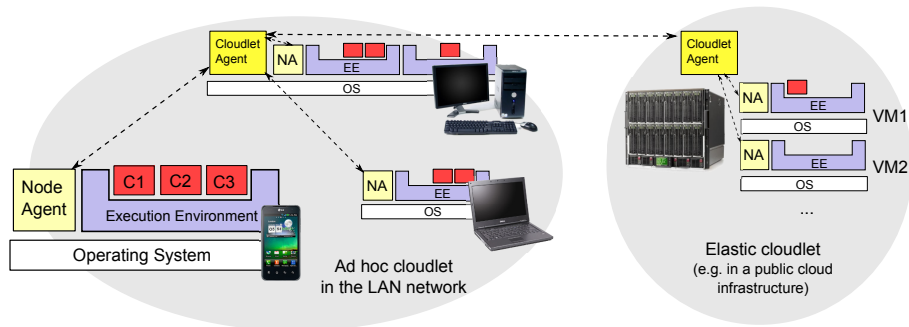
Figure 4: The application components are distributed among nodes in two cloudlets. An ad hoc cloudlet consisting of a mobile phone, a laptop and a desktop computer, and a distant elastic cloudlet in a public cloud infrastructure. All components are managed and monitored by an Execution Environment (EE). Different EEs on a node are managed by a Node Agent (NA), which are managed by the Cloudlet Agent (CA).

One or more Execution Environments run on top of an operating system (OS), which in turn can run on either virtualized or real hardware. The (possibly virtualized) hardware together with the installed OS is called a node, and is managed by a *Node Agent (NA)*. The Node Agent manages all the EEs running on the OS, and can also start or stop new Execution Environments, for example for sandboxing components. The NA also monitors the resource usage of the node as a whole, and has a view on the (maybe virtualized) hardware it runs on (e.g. the number of processing cores, processing speed, etc.).

Multiple nodes that are in the physical proximity of each other (i.e. low latency) form a cloudlet. The cloudlet is managed by a *Cloudlet Agent (CA)*, that communicates with all underlying Node Agents. Cloudlet Agents of different cloudlets can also communicate with each other, for example to migrate components between cloudlets. Within a cloudlet, the node with the most resources is chosen to host the Cloudlet Agent. The Cloudlet Agent has a global overview of all available resources, and is able to calculate a deployment, which should be globally optimal for all devices in the cloudlet.

There are two types of cloudlets, as shown in Figure 4: the ad hoc cloudlet and the elastic cloudlet. The ad hoc cloudlet consists of dynamically discovered nodes in the LAN network. These nodes run a Node Agent that can spawn Execution Environments to deploy components in. When nodes join or leave the cloudlet, the Cloudlet Agent will recalculate the deployments,

migrating components if needed. The elastic cloudlet runs on virtualized infrastructure, where nodes run in virtual machines. Here, the Cloudlet Agent can spawn new nodes when more resources are needed, or stop nodes when too much resources are allocated. This type of cloudlet comes close to the VM based cloudlet envisioned by Satyanarayanan [3], the main distinction being the additional platform components running in the VM (i.e. the Node Agent and the Execution Environment) managing the applications. An elastic cloudlet could run on virtualized hardware co-located with the wireless access point, or on a private cloud within the work environment.

By also deploying the cloudlet platform in the cloud, components can also be outsourced to the public cloud, which allows applications to be outsourced in a 3-tier fashion: hardware dependent components are deployed on the mobile device, latency constrained components are distributed on a nearby cloudlet and unconstrained components, exhibiting challenging CPU and/or data storage requirements, can be deployed on EEs in the public cloud.

5. Cloudlet platform design and implementation

The prototype implementation of the cloudlet platform is realized in OSGi [18], an industry standard module management system in Java. This technology was selected as the OSGi standard already specifies the software component concept, and the OSGi platform can run on multiple underlying platforms. In addition, a substantial amount of middleware building blocks (e.g. Configuration Admin, Declarative Services, ...) are standardized in the OSGi Compendium Specification [19]. In this section, we describe in detail the design and implementation of the Execution Environment and the Node and Cloudlet Agent, together with their integration in the Android platform. First the most important OSGi concepts are highlighted, on which the cloudlet platform is based.

5.1. OSGi

The OSGi core specification defines a service oriented module management system for Java, allowing to dynamically load and unload software modules – called bundles – at runtime. OSGi bundles can expose a service interface by registering an implementation of this interface with the OSGi service registry. When a bundle needs to call another service, the service registry is queried for available implementations. The portability of Java

10

enables the execution of the same code on different platforms and architectures, facilitating remote execution and code migration. Because OSGi was first designed for embedded devices, the incurred overhead is limited.

The OSGi module management system constitutes the core of the Execution Environment of the cloudlet platform. In addition to the core specification, we also use four compendium specifications: Declarative Services, the Configuration Admin Service, the Metatype Service and Remote Service Admin.

5.2. Execution Environment

An overview of the Execution Environment is shown in Figure 5. The EE bundle proxies the components to gather monitoring information (e.g. time spent executing a method) as well as to transparently forward method calls to remote instances through the Remote Service Admin. The latter component implements a remote procedure call protocol, such as R-OSGi [20]. Components come with three descriptors: the metatype descriptor exposes the configurable parameters, the service description declares the offered and required service interfaces and the SLA descriptor defines imposed constraints. The Declarative Services bundle ensures that component interfaces are registered with the OSGi runtime, and the component is bound to all its dependencies. The Metatype and Configuration Admin service allow the EE to discover the available configuration options and to set these configurations dynamically to meet the imposed constraints.

The Declarative Services specification presents a declarative model for publishing, finding and binding OSGi services. Instead of registering and looking up services programmatically in the source code, this is now done at runtime by the OSGi framework using an XML description that is bundled with the component, which describes the provided service interfaces, the service dependencies and how these dependencies should be bound to the component. For lifecycle management the developer can also define methods to be called when a component is activated or deactivated. This provides a simplified programming model to the developer, which has no more code dependencies to the OSGi core APIs, and does not have to handle with the complex dynamic service concept. An example of such a component description XML is shown in Listing 1.

```xml
<component name="ptam.tracker.Tracker">
    <implementation class="ptam.tracker.Tracker"/>
    <service>
```
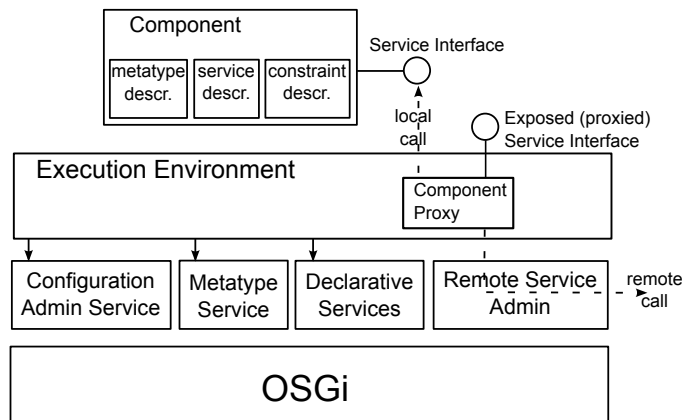
Figure 5: The Execution Environment proxies all component interfaces in order to monitor all service method calls, which are forwarded to either a local instance or a remote instance via R-OSGi. The Configuration Admin and Metatype service enable the EE to dynamically configure components, which are defined using the Declarative Services specification, in order to meet imposed constraints.

```
        <provide interface="ptam.tracker.api.TrackerService"/>
        <provide interface="ptam.video.api.VideoListener"/>
    ...
    </service>
    <reference name="mapper"
          interface="ptam.mapper.api.MapperService"
    ...
          bind="setMapper"
          unbind="unsetMapper"/>
    ...
</component>
```

Listing 1: Example snippet of a component description, defining the Tracker component. The Tracker component implements the TrackerService and VideoListener interfaces, and needs a reference to the MapperService, which is injected by the framework using the bind/unbind methods.

The ConfigurationAdmin (CA) decouples a component from its configuration parameters. The CA maintains a persistent repository of configuration data (a dictionary of key-value pairs) for all components. When a component is started, the CA provides the necessary parameters to correctly configure the component. When the configuration data is changed, components are updated accordingly, allowing at runtime reconfiguration.

12

Because the EE has to be able to adapt the component configuration at runtime using the ConfigurationAdmin, it should also know which configuration parameters are available and which values are allowed. This is where the Metatype Service comes in, which allows developers to describe attribute types in an XML format. An example metatype XML is shown in Listing 2. The metadata consists of an Object Class Definition (OCD), which defines a number of Attribute Definitions (ADs), and a Designate, which binds the OCD to a component.

```xml
<MetaData>
    <OCD description="PTAM Tracker  Configuration"
  name="PTAM  Tracker"
  id="ptam.tracker.TrackerOCD">
        <AD name="Fine  Points"
           id="tracker.finePoints"
           default="1000"
    min="200"
    max="1000">
        </AD>
    </OCD>
    <Designate pid="ptam.tracker.Tracker">
        <Object ocdref="ptam.tracker.TrackerOCD"/>
    </Designate>
</MetaData>
```

Listing 2: Example metatype information for the Tracker component that marks the number of tracked points as a configurable parameter.

To be able to distribute components among multiple resources in the network, we use R-OSGi as provider for our Remote Service Admin, which enables remote service binding and remote service calls.

In order to monitor all application components, the Execution Environment bundle creates a proxy for each service interface provided by a component. Using the OSGi hooks API from the core specification, the original service is hidden for other components, that can only bind to the proxied one. This way, all service calls pass through the proxy, which is able to monitor all service calls. The proxy forwards the call to an actual component interface, either locally or remote via R-OSGi. Using the Metatype service, the EE can discover the configurable parameters for each component, and modify these using the Configuration Admin. The configuration is adapted in order to meet constraints that are declared by the developer in the constraints descriptor as shown in Listing 3.

13

```
<sla>
  <service name="ptam.video.api.VideoListener">
    <method name="processFrame">
      <constraint type="maxTime">50</constraint>
    </method>
  </service>
</sla>
```

Listing 3: Example constraint description for the Tracker component. A video frame should be processed within 50 ms.

Here a constraint is defined that a video frame should be processed by the Tracker component within 50 ms. The Execution Environment implements a feedback loop that will constantly monitor the time to process a video frame. When the processing time exceeds 50 ms, the Tracker configuration parameters are changed in order to lower the processing time. The EE also notifies the Node Agent, which will ask the Cloudlet Agent whether also components need to be offloaded to enhance the application quality.

In order to create an application that exploits the cloudlet platform, the developer needs to design the application as a number of OSGi bundles. For each bundle he has to define the component descriptor that describes the provided and required services of the component. Optionally, the developer can also declare the configurable parameters in a metatype descriptor, and required constraints in a constraint descriptor. Because this can be a tedious and error prone task, we propose an annotation based programming model in Section 6 to alleviate development.

5.3. Node and Cloudlet Agent

The Node and Cloudlet agents are also implemented as OSGi bundles, running on an OSGi runtime in a separate process, as shown in Figure 6. When an EE is started, it registers with the Node Agent running on that node via R-OSGi on localhost.

When first started on a device, both the Node Agent and the Cloudlet Agent bundle are active. Other devices are discovered using JSLP, a Java implementation of the SLP discovery protocol [21]. When another CA is discovered in the LAN network, the CA on one of the two devices (the weakest of the two) is stopped, and the remaining CA becomes the master of the cloudlet. If the connection to the CA is lost, the Cloudlet Agent bundle is started again locally. This way, a cloudlet can be formed in an ad hoc manner, without the need for virtualized hardware. When the discovered

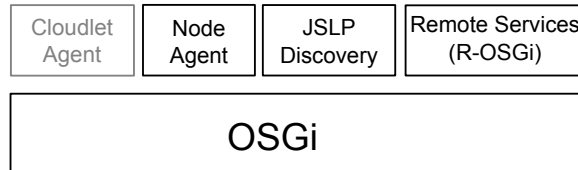| Cloudlet Agent	Node Agent	JSLP Discovery	Remote Services (R-OSGi)
OSGi			

Figure 6: The Node Agent and Cloudlet Agent bundles run on an OSGi runtime in a separate process. The Cloudlet Agent bundle is only present when the current node is also the master of the cloudlet.

Cloudlet Agent is not deployed nearby, both CAs communicate as peers, allowing for inter-cloudlet component offloading.

5.4. OSGi and Android integration

To deploy our cloudlet platform on the Android platform, the platform needs to strictly comply with Androids application model. Android applications are written in Java and compiled to Dalvik bytecode (which run on the Dalvik Virtual Machine), and are composed of different components: Activities, Services, Content Providers and Broadcast Receivers. An Activity provides the basic interaction logic with the user, containing a user interface and offering basic computing capabilities. An Android Service is a component that runs in the background, mainly used for long-running background processes (e.g. playing music) without blocking the user interface. Content Providers are used for managing a shared set of application data, and Broadcast Receivers are small components that respond to system-wide broadcast announcements (e.g. announcement when running low on power).

The OSGi runtime with the Node and Cloud Agent are embedded in a separate Cloudlet Android application that starts an Android Service at boot time, which allows the Node Agent to run in the background. The Execution Environment is embedded within an Android application as depicted in Figure 7. Within the Android application, components are identified that should be offloadable or that should be monitored and configured by the EE. The Android application calls these components by looking up their services with the OSGi runtime, and can also register services to enable callbacks from the components. Because the offloadable components can be executed both on the Android device as on a remote server, the code is compiled in both the Dalvik bytecode format and the regular Java bytecode format.
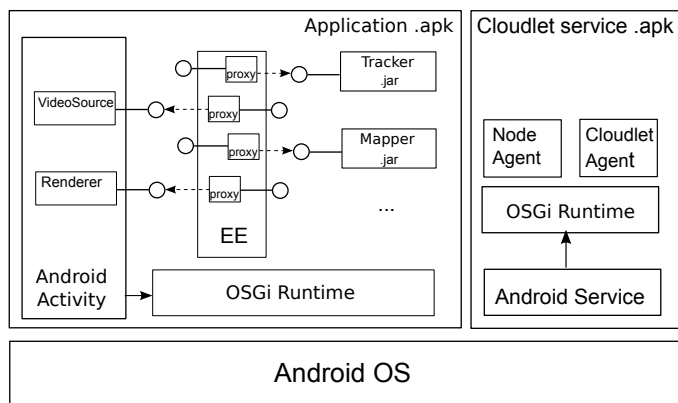
15

Figure 7: On the Android platform a separate application hosts the Node Agent and the Cloudlet Agent bundles, that run in a background process using an Android Service. Applications having offloadable and/or configurable components deploy these on an OSGi runtime hosting the EE, which proxies the provided interfaces. Components requiring the Android internals (e.g. drawing to screen) are embedded in the Android Activity and also provide their interfaces through the EE.

Figure 7 also shows how the use case components are embedded in the Android application. The Renderer and VideoSource component reside within the Android activity, as they need to render to screen, or need access to the Android internals to fetch video frames from the hardware. They expose their interfaces to the other components through OSGi and the EE. The other components such as the Tracker and the Mapper are packaged within the .apk as separate jars, which allows these components to be offloaded at runtime, and are deployed on the OSGi runtime.

The resulting application acts as a regular Android application, and no additional SDK is needed. When the application starts, the EE is created which registers with the Node Agent running on the device in the background. However, a huge burden is put on the application developer when he wants to develop such a component-based application. Necessary steps to take include: identifying the components, implementing these separately as OSGi-bundles together with the accompanying descriptors, embedding an OSGI runtime together with the EE, and building and packaging the components as an Android application. To simplify this process, a lightweight programming model is presented in the following section.

6. Programming model

To ease the development of cloudlet enabled Android applications, we present a straightforward programming model based on Java annotations. These annotations are analyzed in a preprocessing step in the build process, from which the required XML descriptions are generated. In the case of an Android application, the build system will also generate the code necessary to embed the OSGi runtime and the cloudlet platform, and create the desired components.

Listing 4 shows the annotated source code of the Tracker class, from which the Tracker component is generated, as well as the descriptors as described in Section 5.

```java
package ptam.tracker;

@Component
public class Tracker implements TrackerService, VideoListener {

  @Property(min=200, max=1000)
  private int finePoints = 1000;

  @Reference(policy=dynamic)
  private MapperService mapper;

  @Constraint(maxTime=50)
  public void processFrame(byte[] data) {
    // process frame
    ...
  }

  ...

}
```

Listing 4: Example annotated Java source file from which a component is generated, together with the required XML descriptors.

Classes annotated by `@Component` are handled as component implementation classes, providing all implemented interfaces. A reference to another service is created by declaring a variable of the desired type and annotate it with the `@Reference` annotation. This will generate a set and unset method to inject a reference to the required service at runtime, and create the reference attribute in the XML description. Configurable properties are annotated by `@Property`, from which the metatype description is generated. To declare

constraints on service methods, they can be annotated with the `@Constraint` annotation, describing the type constraint and the imposed threshold.
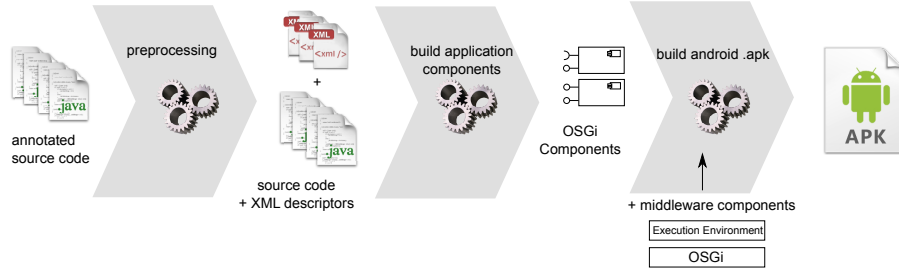

Figure 8: Overview of the steps done by the build system. First, the annotated source code is preprocessed and the required XML descriptors for the components are generated. Second, the component jars are built from the source code and the generated XML files. Finally the application components are bundled with the middleware components to create the resulting Android .apk which can be deployed on the Android OS.

The resulting build process is shown in Figure 8. First the source files are preprocessed and XML descriptors are generated from the detected annotations. In the second step, the source code of the components is compiled and packaged together with the XML descriptors. Finally the build system combines the application components with the middleware components and packages the resulting Android .apk.

7. Evaluation

To show the need for a cloudlet platform presented in this paper, we implemented the augmented reality use case presented in Section 2 using the proposed component model. The AR code was executed on an ad hoc cloudlet consisting of a mobile device and a laptop connected via WiFi. The laptop is equiped with an Intel Core 2 Duo CPU clocked at 2.26GHz. As mobile device we use a HTC Desire, with a single core Qualcomm 1 GHz Scorpion CPU, and an LG Optimus 2x powered by a dual core Nvidia Tegra 2 CPU, also clocked at 1GHz. Both devices capture camera frames at a 800x480 resolution, which is the same size as their screen resolution. The results show both the need for dynamically configurable components in the case of the Tracker, and the benefits of offloading resource intensive components such as the Mapper and Object Recognizer.

7.1. The need for dynamic configuration

The performance of the Tracker component, which has to process all video frames, is critical to the user experience, as too few frames processed per second deteriorates the application quality. However, in this case offloading is not an option, as not enough bandwidth is available to send and receive the video frames in a timely manner. The Tracker exposes one configuration parameter representing the maximum number of feature points that can be tracked. A high number of points will increase the tracking accuracy, making the Tracker more resilient to fast camera movements or motion blur, at the cost of more processing power.

For both mobile devices, the average processing time (and standard deviation) to track one video frame as a function of the number of tracked features is shown in Figure 9. For this case, we implemented a simple negative feedback loop that tries to keep the time needed for the Tracker to process one frame close to the imposed threshold of 50 ms by tuning the defined parameter. Of course, a more complex and more general adaptation strategy could be implemented. At runtime, the framework adapts the configuration parameter depending on the mobile device capabilities, resulting in a number of tracked features between 200 and 250 for the HTC Desire, and a value between 450 and 500 for the LG Optimus 2x.
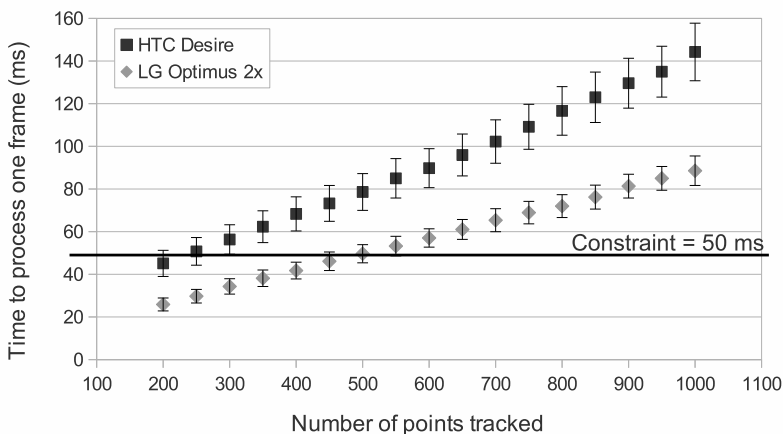

Figure 9: The time needed to track a frame as a function of the number of feature points. Depending on the device, the framework will choose a different parameter setting in order to meet the maximum processing time constraint.
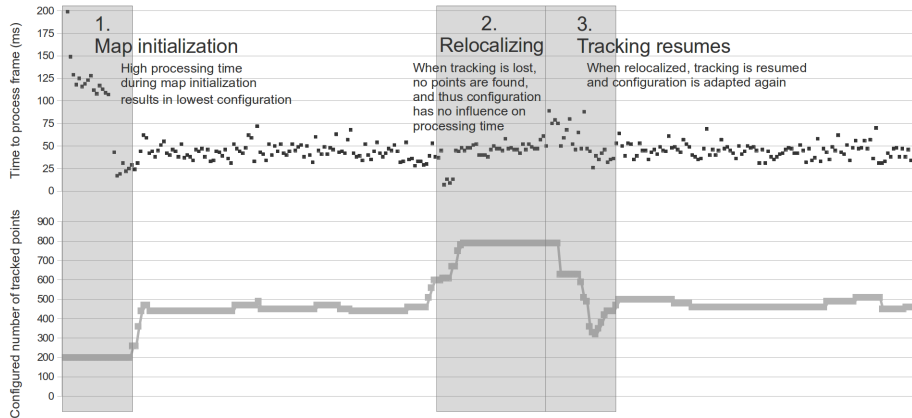
19

Figure 10: The framework dynamically adapts the configuration parameter in order to meet the constraint of 50 ms to process a frame.

The processing time per frame as well as the value of the configurable parameter (= maximum number of feature points tracked) while running the application on the LG Optimus 2x are shown in Figure 10.

The initialization of the map (1) results in high processing times, and thus the value of the parameter is reduced to the minimum. When tracking is started, the parameter value is increased to have the best tracking accuracy possible within the 50 ms imposed.

When the tracking is lost due to a too quick camera movement, no feature points are found, and the applications starts relocalizing (2). As no feature points are found, the processing time lowers the parameter value is increased. When tracking resumes (3), the processing time is too high due to the increased parameter, and the parameter value is adapted again to meet the imposed constraint.

7.2. The benefits of offloading

The goal of the cloudlet is of course to offload application components from the device to discovered resources nearby. In this experiment, an ad hoc cloudlet is formed between the laptop and the mobile devices. The Cloudlet Agent then selects resource intensive components suitable for offload, and migrates those to the laptop. In the case of the AR application presented in Section 2, the Mapper and Object Recognizer are selected for outsourcing.

When a new keyframe is added, all positions of the points in the map are optimized. This optimization problem scales with the total number of
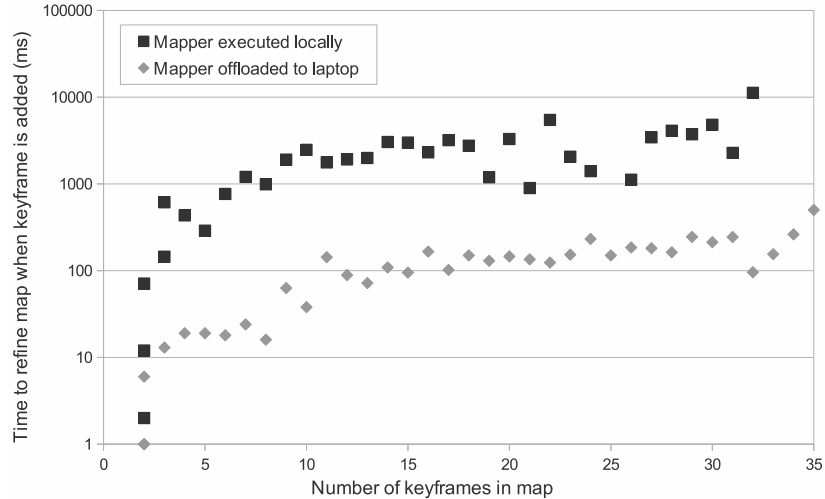
Figure 11: The time to refine the map when a new keyframe is added, as a function of the total number of keyframes in the map. Offloading the Mapper component results in a speedup factor of 10.

keyframes in the map. Figure 11 shows the time needed for this map refinement as a function of the keyframes in the map for one run, both executed locally on the LG Optimus 2x, and remote when the Mapper component is outsourced to the laptop. As the number of keyframes increases, the refinement process takes up to 10 seconds on the mobile device, while remote execution takes less then a second. Overall we witness a speed up of a factor 10 when outsourcing the the laptop.

For the Object Recognizer, larger differences are perceived. When performing an object search on a videoframe, this takes up to 30 seconds on the HTC Desire, and around 18 seconds on the LG Optimus 2x. When outsourcing the Object Recognizer to the laptop, the processing time is reduced to on average 1.5 seconds.

To decide automatically at runtime which components should be outsourced and which should be executed on the mobile device, partitioning algorithms could be used based on the monitoring information gathered by the ExecutionEnvironment [22].

8. Conclusions and future work

In this paper, we present a cloudlet architecture that not only provides fixed infrastructure co-located with the WiFi access point, but also enables ad hoc discovery of devices in the vicinity to share resources among each other. Instead of providing infrastructure on a virtual machine level, a more fine-grained approach is presented, where the cloudlet framework manages applications on a component level. By adaptively configuring and outsourcing application components, the platform optimizes the application depending on the mobile device capabilities and the available resources in the cloudlet.

A prototype implementation of the framework based on the OSGi industry standard is presented, together with an annotation-based programming model which enables easy application development. We evaluated the cloudlet platform using an AR use case, showing the need for both adaptive configuration and component offloading to improve the application quality.

As future work, we plan to tackle new challenges with respect to deployment calculation and scheduling. In comparison with existing cyber foraging frameworks, a lot more decision options are generated by allowing different configurations of components. Moreover, instead of a fixed surrogate infrastructure, now multiple places for remote execution have to be considered in the cloudlet, that can dynamically join and leave. The hierarchical architecture also allows for decision taking for all users to reach a global optimum, instead of possibly conflicting local optima for each user individual.

9. Acknowledgement

References

[1] P. Jäppinen, R. Guarneri, L. M. Correia, An applications perspective into the future internet, Journal of Network and Computer Applicationsdoi:10.1016/j.jnca.2012.08.009.

[2] M. Satyanarayanan, Pervasive computing: Vision and challenges, IEEE Personal Communications 8 (2001) 10–17.

[3] M. Satyanarayanan, P. Bahl, R. Caceres, N. Davies, The case for vm-based cloudlets in mobile computing, IEEE Pervasive Computing 8 (4) (2009) 14–23.

[4] B.-G. Chun, S. Ihm, P. Maniatis, M. Naik, A. Patti, Clonecloud: elastic execution between mobile device and cloud, in: Proceedings of the sixth conference on Computer systems, EuroSys '11, ACM, 2011, pp. 301–314.

[5] I. Giurgiu, O. Riva, D. Juric, I. Krivulev, G. Alonso, Calling the cloud: enabling mobile phones as interfaces to cloud applications, in: Proceedings of the 10th ACM/IFIP/USENIX International Conference on Middleware, Middleware '09, Springer-Verlag, 2009, pp. 5:1–5:20.

[6] T. Verbelen, T. Stevens, P. Simoens, F. De Turck, B. Dhoedt, Dynamic deployment and quality adaptation for mobile augmented reality applications, Journal of Systems and Software 84 (11) (2011) 1871–1882.

[7] M.-R. Ra, A. Sheth, L. Mummert, P. Pillai, D. Wetherall, R. Govindan, Odessa: enabling interactive perception applications on mobile devices, in: Proceedings of the 9th international conference on Mobile systems, applications, and services, MobiSys '11, ACM, New York, NY, USA, 2011, pp. 43–56.

[8] E. Cuervo, A. Balasubramanian, D.-k. Cho, A. Wolman, S. Saroiu, R. Chandra, P. Bahl, Maui: making smartphones last longer with code offload, in: MobiSys '10: Proceedings of the 8th international conference on Mobile systems, applications, and services, ACM, 2010, pp. 49–62.

[9] M. D. Kristensen, N. O. Bouvin, Scheduling and development support in the scavenger cyber foraging system, Pervasive and Mobile Computing 6 (6) (2010) 677–692, special Issue PerCom 2010.

[10] G. Klein, D. Murray, Parallel tracking and mapping for small ar workspaces, in: Proceedings of the 2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality, ISMAR '07, IEEE Computer Society, 2007, pp. 1–10.

[11] D. G. Lowe, Distinctive image features from scale-invariant keypoints, Int. J. Comput. Vision 60 (2) (2004) 91–110.

[12] W. Li, Qos assurance for dynamic reconfiguration of component-based software systems, Software Engineering, IEEE Transactions on 38 (3) (2012) 658–676.

[13] R. Balan, J. Flinn, M. Satyanarayanan, S. Sinnamohideen, H.-I. Yang, The case for cyber foraging, in: EW 10: Proceedings of the 10th workshop on ACM SIGOPS European workshop, ACM, 2002, pp. 87–92.

[14] S. Goyal, J. Carter, A lightweight secure cyber foraging infrastructure for resource-constrained devices, in: WMCSA '04: Proceedings of the Sixth IEEE Workshop on Mobile Computing Systems and Applications, IEEE Computer Society, 2004, pp. 186–195.

[15] Y.-Y. Su, J. Flinn, Slingshot: deploying stateful services in wireless hotspots, in: MobiSys '05: Proceedings of the 3rd international conference on Mobile systems, applications, and services, ACM, 2005, pp. 79–92.

[16] T. Verbelen, P. Simoens, F. De Turck, B. Dhoedt, Cloudlets: bringing the cloud to the mobile user, in: Proceedings of the third ACM workshop on Mobile cloud computing and services, MCS '12, ACM, New York, NY, USA, 2012, pp. 29–36.

[17] C. Szyperski, Component Software: Beyond Object-Oriented Programming, 2nd Edition, Addison-Wesley Longman Publishing, 2002.

[18] The OSGi Alliance, OSGi Service Platform, Core Specification, Release 4, Version 4.2, aQute, 2009.

[19] The OSGi Alliance, OSGi Service Platform, Service Compendium, Release 4, Version 4.2, aQute, 2009.

[20] J. S. Rellermeyer, G. Alonso, T. Roscoe, R-OSGi: distributed applications through software modularization, in: Middleware '07: Proceedings of the International Conference on Middleware, Springer-Verlag, 2007, pp. 1–20.

[21] E. Guttman, C. Perkins, J. Veizades, M. Day, Service location protocol, version 2 (1999).

[22] T. Verbelen, T. Stevens, F. D. Turck, B. Dhoedt, Graph partitioning algorithms for optimizing software deployment in mobile cloud computing, Future Generation Computer Systemsdoi:10.1016/j.future.2012.07.003.