Robust Orchestration of Concurrent Application Workflows in Mobile Device Clouds

Parul Pandey, *Student Member, IEEE,* Hariharasudhan Viswanathan, *Student Member, IEEE,* and Dario Pompili, *Senior Member, IEEE*

Abstract

A hybrid mobile/fixed device cloud that harnesses sensing, computing, communication, and storage capabilities of mobile and fixed devices in the field *as well as* those of computing and storage servers in remote datacenters is envisioned. Mobile device clouds can be harnessed to enable innovative pervasive applications that rely on real-time, in-situ processing of sensor data collected in the field. To support concurrent mobile applications on the device cloud, a robust and secure distributed computing framework, called `Maestro`, is proposed. The key components of `Maestro` are (i) a task scheduling mechanism that employs controlled task replication in addition to task reallocation for robustness and (ii) `Dedup` for task deduplication among concurrent pervasive workflows. An architecture-based solution that relies on task categorization and authorized access to the categories of tasks is proposed for different levels of protection. Experimental evaluation through prototype testbed of Android- and Linux-based mobile devices as well as simulations is performed to demonstrate `Maestro`'s capabilities.

I. INTRODUCTION

The concept of *cyber foraging*—opportunistic discovery and exploitation of nearby compute and storage servers [2]—was conceived to augment the computing capabilities of mobile hand-held devices "in the field." Such augmentation would enable novel *compute-* and *data-intensive* mobile pervasive applications spanning across multiple domains, from education to infotainment,
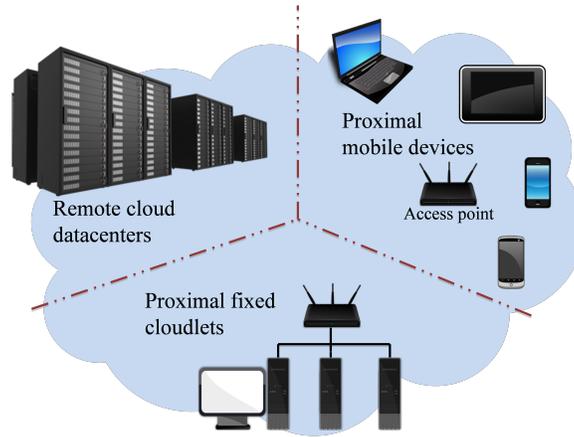
Fig. 1. Spectrum of computing resources in a mobile/fixed device cloud—mobile resources in the proximity, fixed (cloudlets) computing resources in the proximity usually tethered to Wi-Fi access points, and cloud resources.

from assisted living to ubiquitous healthcare. Such applications include (but are not limited to) object, face, pattern and speech recognition, natural language processing, and biomedical and kinematic signal processing for reality augmentation as well as for collaborative decision making. Research efforts towards realizing the vision of cyber foraging can be broadly classified into works in the fields of *mobile cloud computing*, *opportunistic computing*, and *mobile grid computing*. Prior work in mobile cloud computing has primarily focused on augmenting the computing capabilities of mobile devices in the field with *dedicated* and *trusted* computing resources, either situated remotely (in the *Cloud* [3]–[8]) or proximally (in *cloudlets* [9]–[11]). As cyber foraging was meant to convey a whole new pervasive-computing paradigm based on the principle of "living off the land" [2], works in the fields of Opportunistic Computing (OC) [12]–[15] and on mobile grid computing [16], [17] emerged due to the ubiquity and growing computing capabilities of mobile devices. These works have explored the feasibility of leveraging the computing and communication capabilities of other mobile devices *in the field* to enable innovative mobile applications.

Our Vision: We envision that the heterogeneous sensing, computing, communication, and storage capabilities of mobile and fixed devices in the field *as well as* those of computing and storage servers in remote datacenters can be *collectively* exploited to form a "loosely-coupled" mobile/fixed device cloud. In keeping with the broadest principles of cyber foraging, the device cloud's computing environment may be composed of (i) purely mobile resources in the proximity,

(ii) a mix of mobile and fixed resources in the proximity, or (iii) a mix of mobile and fixed resources in the proximity as well as in remote datacenters as shown in Fig. 1. We focus on the "extreme" scenario in which the device cloud is composed *purely* of proximal mobile devices as such scenario brings all the following concerns to the fore: robustness and security. One or more of these concerns do not arise in the other scenarios and, hence, solutions developed for the extreme case will easily extend to the other cases. The device cloud employs a role-based network architecture in which the devices may at any time play one or more of the following *logical roles*: (i) requester; (ii) Service (data or computing resources) Provider (SP), and (iii) broker. SPs notify the broker(s) of their capabilities and availability. The brokers are in charge of handling concurrent service requests as well as of orchestrating the execution of mobile pervasive applications on SPs.

Related Work and Motivation: In mobile cloud computing, researchers have primarily focused on augmenting the capabilities of mobile devices in the field by offloading expensive (compute and energy-intensive) tasks to dedicated wired-grid [18], [19] or cloud resources [5]–[7], [10] in a transparent manner. However, these approaches are not suitable for enabling data-intensive applications in real time due to prohibitive communication cost and response time, significant energy footprint, and the curse of extreme centralization. On the contrary, we explore the possibility of mobile devices offloading workload tasks to other devices in the proximity (mobile device clouds) so to enable innovative applications that rely on real-time, *in-situ* processing of sensor data. Research in the areas of mobile device clouds [16], [17] and OC [13]–[15] has explored the potential of code offloading to proximal devices by following two different approaches. Solutions for mobile grids advocate a structured and robust approach to workflow and resource management; whereas OC depends entirely on direct encounters and is highly unstructured with little or no performance guarantees. These works do not address the crucial research challenge of "real-time" concurrent applications management while also taking the following concerns into account: robustness, security, and privacy.

Management of concurrent workflows has been studied before in the context of wired-grid computing. In [20]–[22], different strategies for scheduling multiple workflows were investigated— namely, sequential, by interleaving tasks of the different workflows in a round-robin manner, and by merging the different workflows into one. Independently of the strategy, the workflow tasks are allocated using level-based [21], list-based [20], duplication-based [23], or clustering-

based [22] heuristics. In all these heuristics, all the tasks of a workflow are scheduled at the same time with the option of filling the "gaps" (schedule holes due to high communication cost between tasks) for efficient resource utilization.

However, none of the existing solutions can be adopted for workflow management mobile device clouds as they do not factor in task deduplication (for efficiency), reactive self-healing and proactive self-protection (for failure handling), security (from malicious resources), and privacy, which are *all* primary concerns in mobile cloud computing. Even though duplication-based scheduling provides some level of redundancy, it treats only fork tasks (ones with multiple successor tasks) as critical and does not protect other tasks that may be critical in the context of the application (or annotated by the developer or user as one). All the aforementioned shortcomings serve as a motivation for the clean-slate design of `Maestro`. To the best of our knowledge, ours is the first work to explore deduplication and scheduling of tasks belonging to concurrent real-time application workflows on mobile device clouds.

Contributions: We present `Maestro`, a framework for robust and secure mobile cloud computing where any mobile application is represented as a workflow, i.e., a Directed Acyclic Graph (DAG) composed of multiple parallelizable and sequential tasks along with dependencies. Often, multiple service requests are received *simultaneously* by brokers and, hence, tasks belonging to multiple workflows managed by the brokers have to be allocated and executed on the SPs in the mobile device cloud. `Maestro` employs controlled *replication of critical workflow tasks* and controlled *access to user data* (via multiple levels of authorization) to realize *secure* mobile computing. Controlled replication refers to the idea of replicating "critical" workflow tasks and only when needed (based on SPs' reliability). Not only task replication imparts *robustness* (against SP failures); it has the ability to handle *uncertainty* arising from device failures, denial of service, and intentional corruption of results. We categorize the workflow tasks according to the sensitivity of the data processed, and allow access only to authorized service providers.

Note that, while concurrent workflows makes the real-time task-allocation problem complex, it also presents opportunities: there may be multiple duplicate service requests at a broker; similarly, there may be multiple duplicate tasks that are common across workflows. `Maestro` *deduplicates similar tasks across workflows* as it lends itself to minimization of duplication in services rendered. Task deduplication leads to efficient real-time, in-situ processing of simplified workflows (with fewer tasks than before) as well as to better utilization of resources. To address

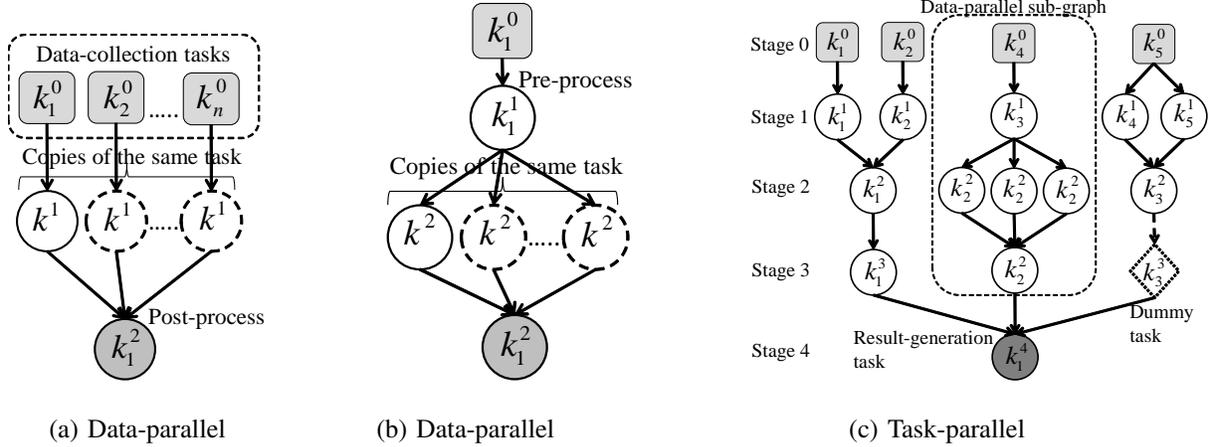(a) Data-parallel   (b) Data-parallel   (c) Task-parallel

Fig. 2. Workflows for *data-parallel* applications with n parallel tasks (a) when there are n separate data sources and (b) when data from a single source has to be divided into multiple chunks (pre-processing); (c) Workflow for a *task-parallel* application (with a data-parallel sub-graph in it)—here a "dummy task" is also represented in Stage 3 of the right branch.

the non-trivial research challenge of the identification of task duplicates and the creation of simplified workflows (at the brokers), we introduce Dedup, a sub-graph matching technique for task deduplication among DAGs. After deduplication, the tasks of the simplified workflows have to be scheduled for execution on the mobile/fixed cloud resources so to meet the user-specified deadline of the corresponding application workflow.

To sum up, *our contributions* are:

- We present a robust and secure mobile cloud computing framework, Maestro, for concurrent application workflow management on mobile device clouds;

- We impart robustness to Maestro via two core components, namely, 1) replication-based task-scheduling mechanism and 2) Dedup, for task deduplication among concurrent DAGs.

- We discuss the results from our experimental evaluation of Maestro in representative operating environments. We also present analysis of our replication-based task-scheduling mechanism for mobile device clouds through experiments on a prototype testbed of Android- and Linux-based devices.

Outline: In Sect. II, we present the architecture of mobile device clouds and our generalized workflow representation scheme; in Sect. III, we discuss the task scheduling mechanism for concurrent workflows and task deduplication algorithm to handle duplicate tasks in concurrent workflows; in Sect. IV, we discuss quantitative results that demonstrate the merits of our

contributions; finally, in Sect. V, we conclude with a note on future work.

II. MOBILE DEVICE CLOUD AND WORKFLOWS

In this section, firstly, we introduce a mobile device cloud, our envisioned heterogeneous computing environment; then, we discuss a generalized workflow representation scheme for depicting the data-processing chains in mobile applications.

A. Mobile Device Cloud

Our vision is to organize the heterogeneous sensing, computing, and communication capabilities of mobile devices in the proximity (as well as in remote datacenters) in order to form an *elastic resource pool*—a mobile device cloud (MDC). This cloud can then be leveraged to augment the capabilities of any mobile device in the network when required in order to enable novel mobile applications. The interested reader may refer to our work in [24] for details on the architecture of MDCs.

Logical roles: MDC is a hierarchical logical-role-based computing environment in which the devices may play one or more of the following logical roles: (i) *requester*, which place requests for application workloads that require additional data and/or computing resources from other devices, (ii) *service provider*, which can be a data provider (a sensing device), a resource provider (a computing device) or both; (iii) *broker*, which is in charge of handling requests and orchestrating the execution of applications on the MDC. This architecture enables easy management and does not suffer from the problem of extreme centralization (i.e., single point of failure) as these are only logical roles.

Service providers and broker: The spectrum of computing resources (SPs) in a MDC computing environment includes mobile (sensing and computing) resources in the proximity, fixed (dedicated cloudlets) computing resources in the proximity usually tethered to Wi-Fi access points or base stations, and fixed resources in remote datacenters (dedicated cloud resources). As mentioned earlier, we design solutions capable of handling the extreme case, i.e., a hybrid cloud composed *only* of proximal mobile computing resources. Another design decision that is faced with a spectrum of possibilities is the location of the broker. The role of a broker can be played by one of the mobile resources[1] (chosen based on centrality of location, battery level,

---

[1]Broker selection is out of the scope of this article.

and/or computing capabilities) or be one of the proximal fixed resources.

Service discovery: The role of broker is played by one of the proximal fixed resource tethered to the Wi-Fi access point or the base station so to ensure that all SPs are connected to a broker when they are in the network. Service discovery at the broker is achieved through service advertisements from the SPs. Service advertisements may include information about the sensor data, types of sensors (quality of data), amount of computing (in terms of unutilized CPU cycles [%]), memory ([Bytes]), and communication ([bps]) resources, the start and end times of the availability of those resources, and the available battery capacity ([Wh]) at each SP. The broker leverages this information to allocate workload to the SPs.

Uncertainty awareness: The broker extracts the following long-term statistics from its underlying resource pool: the average arrival (joining) rate of SPs ($\widetilde{W}$), the average SP availability duration ($\widetilde{T}$), and the average number of SPs associated with the broker at any point in time ($\widetilde{N}$) whose relationship is given by Little's law, i.e., $\widetilde{N} = \widetilde{W} \cdot \widetilde{T}$. These statistics help the brokers assess the *churn rate* of SPs, i.e., a measure of the number of SPs moving in to and out of their respective resource pools over a specific period of time. When the long-term statistics are not taken into account at the broker and when the durations advertised by the SPs are used to make workload allocation decisions, the mismatch between advertisements and ground reality will have an adverse effect on the performance of applications (particularly, in terms of response time). Uncertainty awareness at the broker enables design of robust workload scheduling algorithms.

B. Workflow Representation

Applications are composed of tasks whose order of execution is specified by workflows.

Structured workflows: Our generalized workflow is a *Directed Acyclic Graph (DAG)* composed of tasks (vertices) and dependencies (directed edges), as shown in Fig. 2. Tasks belong to one of the following three categories: 1) data-collection task, 2) computation task, or 3) result-generation task. These tasks are elementary and cannot be split further into *micro-tasks*, i.e., parallelization of elementary tasks does not yield any speed-up in execution time. The workflow is composed of multiple *stages* with a set of tasks to be performed at each stage. The resources that perform the tasks at Stage $i - 1$, where $i \geq 1$, serve as data sources for the tasks that have to be performed at Stage i. The data sources for tasks at Stage 1 are the sensors themselves (where Stage-0 tasks, i.e., data-collection tasks, are performed). There are no dependencies between the tasks

(a) Ubiquitous healthcare domain   (b) Distributed robotics domain   (c) Computer vision domain
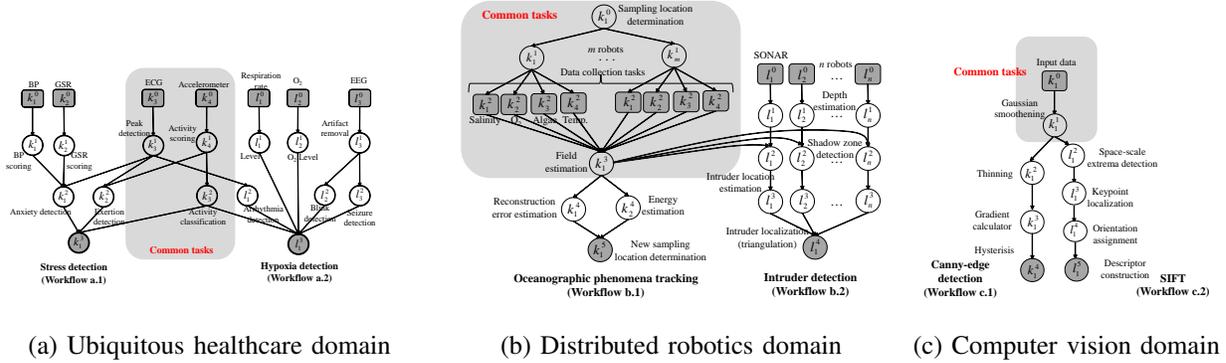
Fig. 3. Example workflows from two different domains, ubiquitous healthcare (a), distributed robotics (b), and computer vision domain (c). Common tasks across workflows in the same domain are highlighted. Workflows in (a) are purely task parallel, in (c) purely data parallel, while the ones in (b) are mixed.

at a particular stage and, hence, they can be performed in parallel. Also, without any loss in generality, we assume that there can be dependencies only between tasks of consecutive stages in the workflow. Whenever we have dependencies between tasks of non-consecutive stages, we introduce the notion of "dummy tasks," whose output equals the input and whose cost of operation (in terms of time and battery drain) is zero. Our *structured* workflow representation is rich in information. In addition to the regular information like task types and data dependencies, it also includes the following: task identifiers, task sizes, quality and quantity of inputs, the preferred interfaces to child and parent tasks, the implementation (when multiple exist), and the task criticality (either boolean or multiple degrees).

Concurrent workflows: Concurrent service requests are often received by each broker, i.e., multiple workflows have to be executed concurrently in the underlying pool. The aforementioned task-allocation problem, albeit complex, presents opportunities. There may be multiple similar service requests at a broker as well as multiple tasks that are common across different workflows. Deduplication of such common tasks leads to efficient real-time, *in-situ* processing of simplified workflows (with fewer tasks than before) as well as to better resource utilization.

Example applications: This representation is powerful as it captures both *data-* and *task-parallel* applications. Data-parallel applications are also referred to as "embarrassingly parallel" applications, in which an independent set of homogeneous tasks—working on disjoint sets of data—can be performed in parallel (preceded and succeeded by pre- and post-processing

tasks, respectively), as shown in Figs. 2(a) and 2(b). Task-parallel applications, on the other hand, have a set of sequential as well as parallel tasks with pre-determined dependencies and degree of parallelism. A task-parallel workflow may also have a data-parallel block built into it, as illustrated in Fig. 2(c). To understand the concept of workflows and their constituent computational tasks, consider the following example applications from different domains.

(a) Ubiquitous healthcare. Applications in this domain (also called data-driven, sensor-based healthcare) include stress detection, hypoxia (lack of oxygen) detection, alertness and cognitive performance assessment. Figure 3(a) depicts the task-parallel workflows for *stress detection* (Workflow a.1) and *hypoxia detection* (Workflow a.2), which use vital-sign data acquired from biomedical (e.g., Blood Pressure (BP), ElectroCardioGram (ECG), ElectroEncephaloGram (EEG)) as well as kinematic sensors (e.g., accelerometer, gyroscope) attached to a person. Dummy tasks are not shown for the sake of simplicity. The tasks in these two workflows belong to one of the following classes: data analysis, data manipulation, decision making; and they aid in determining the psychophysiological state of a person (knowledge) from raw sensor data. As mentioned earlier, applications belonging to the same domain may have similar component tasks, which can be deduplicated to achieve efficiency. For example, feature extraction from accelerometer outputs as well as ECG analysis (in Stage 1) are component tasks in the hypoxia-detection workflow (for context assessment, e.g., activity and arrhythmia detection, in Stage 2) as well as in the stress-detection workflow (for exertion detection in Stage 2), as shown in Fig. 3(a).

(b) Distributed robotics. Distributed decision-making applications in this domain include adaptive sampling, intruder detection, target tracking, data-driven path/trajectory planning, to name just a few. Figure 3(b) depicts workflows for *oceanographic phenomena tracking* (Workflow b.1) and coastal underwater *intruder detection* (Workflow b.2), which use all or a subset of the following data acquired using environmental and inertial navigation sensors as well as SONAR on autonomous underwater robots: temperature, salinity, pollutants, nutrients, position, and depth. Workflow b.1 depicts how field estimation is used to track the effect of oceanographic phenomena (e.g., temperature and salinity gradients, algae growth, nutrient concentration) on aquatic life. Workflow b.2 depicts intruder localization (using SONAR), which requires optimal positioning of the robots in order to avoid false positives due to severe transmission losses of the acoustic signal/waves traversing certain regions. Such regions of high transmission loss can again be determined from temperature, salinity, and depth field estimates. Here, field estimation

is a common task (between the two workflows), which can be deduplicated. Note that in this application the overall task-parallel workflow is composed of smaller data-parallel workflows.

(c) Computer vision. Many applications in this domain such as object recognition, face recognition, gesture recognition, and image stitching use two well-known algorithms shown in Fig. 3(c), *Canny edge detection* (Workflow c.1) and *Scale Invariant Feature Transform (SIFT)* (Workflow c.2). Workflow c.1 takes as input an image and uses a multi-stage algorithm to detect edges in the image. Workflow c.2 generates a large collection of feature vectors from an image, each of which is (i) invariant to image translation, scaling, and rotation, (ii) partially invariant to illumination changes, and (iii) robust to local geometric distortion. Both these workflows have a common task, Gaussian smoothening, which can be deduplicated.

III. MAESTRO

In this section, we present `Maestro`, a robust mobile cloud computing framework for concurrent workflow management on the MDC. Firstly, we present a concurrent workflow-scheduling mechanism designed for `Maestro`. Secondly, we discuss `Maestro`'s task-scheduling mechanism, which employs controlled task replication (for robustness) before scheduling the tasks for execution on the MDC's resources. At the end, we present `Dedup` (the sub-graph matching technique in `Maestro`) for task deduplication among DAGs.

A. Concurrent Workflows Scheduling

The brokers receive multiple workflow execution requests over a period of time from the service requesters. The tasks of these workflows have to be allocated to SPs in the MDC. While submitting a service request, the application can specify the *absolute deadline* (D [s]) within which the workflow execution has to be completed for it to be useful. There is also a notion of an acceptable probability of failure (P^{fail}) for each workflow. This probability can be a service-level guarantee advertised by the broker or negotiated a priori between brokers and service requesters. `Maestro`'s task-scheduling mechanism at the broker is in charge of determining (i) the set of workflow tasks that are ready to be allocated, (ii) the relative priority among the ready tasks[2], and (iii) the amount of replication and the appropriate SP(s) for each ready task. In `Maestro`,

---

[2]A ready task is one that does not have any unresolved dependencies, i.e., all its parent tasks have completed their execution.

tasks can be immediately allocated *as and when* they become ready *or* the ready tasks can be accumulated (over a waiting period, δ_b^{ready} [s]) and then allocated for a more efficient schedule in terms of *makespan* (i.e., total workflow execution time) and number of replicas (i.e., battery drain). This waiting period is again a tunable parameter; the larger the waiting period, the greater the chances of finding the most appropriate SPs (lower makespan and fewer replicas). However, δ_b^{ready} cannot be too large due to real-time constraints of the application.

Task Prioritization: Determining the relative priority among ready tasks from the same or different workflows requires incorporation of computation-time information and deadline requirements, as discussed in prior work on workflows management in computational grids [21]. Firstly, we determine the *level* of task, which is the length of the longest path from that task to an exit task. The length of a path in a DAG is the sum of the average computation time of that task and the average computation times of all the successor tasks along the path. The average communication times between successive tasks should also be taken into account if the Communication-to-Computation-costs Ratio (CCR) of the workflow DAG is high. The level Δ^k [s] of a task k is given by,

$$\Delta^k = \overline{\alpha^k} + \max_{c \in \mathcal{C}^k}\{\overline{\beta^{kc}} + \Delta^c\}, \tag{1}$$

where $\overline{\alpha^k}$ is the average computation time of a task k on the SPs in the MDC, $\mathcal{C}^k$ is the set of child tasks of k, and $\overline{\beta^{kc}}$ is the average communication time for data transfer between tasks k and c when executed on the SPs in the MDC. Once the level of each ready task is known, their *slack* S [s] (maximum allowable wait time before execution of that task) at any time t is determined, for task k, as,

$$S^k(t) = D^k - \Delta^k - t, \tag{2}$$

where D^k is the absolute deadline for the workflow which task k belongs to. The ready task k^* with the *smallest* slack has the *highest* priority, i.e., $k^* = \arg\min_k S^k(t)$.

After prioritization of the ready tasks according to this criterion, the most appropriate SP for allocation and the amount of replication (when necessary) are determined. Each SP in the MDC has a task queue. For a ready task k with the highest priority, the SP n that provides the earliest finish time $(t_n^{k,fin})$ is the most preferred. Finish times are considered due to heterogeneity in capabilities of service providers. In a homogeneous environment, start times are sufficient to

make allocation decisions. The t^{fin}'s are obtained as,

$$t_n^{k,fin} = t_n^{k,start} + \alpha_n^k, \tag{3}$$

where $t^{k,start}$ is the start time for task k on SP n. The t^{start} depends on the number and type of existing tasks in the queue. However, there is uncertainty associated with the availability of the SPs in a MDC for the required duration and this has to be taken into account in the scheduling mechanism.

Controlled Replication: An effective way to overcome the uncertainty (due to failures) is the reallocation of failed tasks (also called "healing"). However, healing is *not* suited for tasks with large computation times and for tasks that are critical for multiple workflows. Though healing provides robustness, it does increase the makespan as it waits for at least the task's computation time before making a decision (i.e., it is *reactive*). Conversely, we replicate *critical* tasks at multiple service providers (*proactively*) in order to ensure the completion of those tasks on time. Proactive task replication avoids unnecessary idle waiting times incurred in reactive failure handling (i.e., healing). Tasks that have to be replicated are allocated to the SP that provides the next earliest $t^{k,fin}$. Note that, as replicas have the same priority as the original, they are allocated together with the original before the other tasks that have lower priority.

All tasks in a workflow should not be replicated as it will increase the total number of tasks to be executed, in turn leading to massive queuing delays and large makespans. The application developer may explicitly *annotate* certain workflow tasks as *non critical*. All other tasks are treated as *blocking tasks*, i.e., the progress of the workflow depends on their completion. The decision to replicate a task k initially allocated to SP n is taken based on how the task-completion probability of n compares with the "required" success probability for that task derived from the pre-specified P^{fail}. The required success probability p^{succ} for each task in the set of incomplete tasks $\mathcal{K}$ of a workflow is obtained by solving,

$$(p^{succ})^{|\mathcal{K}|} = 1 - P^{fail}. \tag{4}$$

The task-completion probability $p_n^{k,succ}$ of a task k at SP n is,

$$p_n^{k,succ} = \Pr\{t + T^n > t_n^{k,start} + \alpha_n^k\}, \tag{5}$$

where t is the current time and T^n is the "actual" availability duration of SP n. Without any loss in generality, we assume that the distribution of SP availability duration is known while

determining $p_n^{k,succ}$. It is quite straightforward to obtain and maintain such statistics at the brokers. When $p_n^{k,succ} < p^{succ}$, a replica is allocated to the next best SP as mentioned before. Replicas of task k are created and allocated to SPs until the following condition is satisfied for the first time , i.e.,

$$1 - \prod_{n \in \mathcal{N}} (1 - p_n^{k,succ}) \geq p^{succ}, \tag{6}$$

where $\mathcal{N}$ is the set of SPs which the replicas are allocated to. As the tasks of a workflow are completed over time with probability one, the required success probability of remaining incomplete tasks decreases; this allows the scheduler to use some less reliable SPs, resulting in load balancing.

For "fork" tasks that are common across multiple workflows, the more stringent condition on the required probability of success is taken into account. To avoid uncontrolled replication, we use a maximum replication limit. This limit is different for different types of tasks. For example, the fork tasks, which are crucial for the success of multiple workflows, have more replicas than the other tasks. This difference in the level of protection is crucial to avoid blocking of MDC resource by tasks that are not so critical as the ones that follow them.

Levels of protection: Certain tasks in an application might work with sensitive (personal) user data that the user does not want to share. In such situations, these tasks can be given to only trusted service providers. To give different levels of protection to different tasks, we present a hierarchical approach similar to that of a social network. We assign tasks to different service providers based on level of trust, i.e., we determine what computing resources the different tasks are assigned to. We elaborate on our idea under the context of a ubiquitous health monitoring application. Data-analysis tasks are basic statistical methods that run over a tremendous amount of time-series data. The knowledge of the "data type" and context is inconsequential for the data-analysis tasks and, hence, the data can be anonymized so to not provide any private information (e.g., participant's identity and health status). Therefore, data-analysis tasks (*public* tasks) can be performed on *any* "volunteered resource" in proximity without any concerns over the level of trust of the computing resource. Conversely, data-manipulation tasks (e.g., artifact removal in biomedical signals) need to be aware of the data type and, hence, can be carried out *only* on "trusted resources". However, they do not need any contextual information or identity of the participant whom the data belongs to. Trusted resources include service providers belonging

to family members and friends (on social networks and real life) and this category of tasks is referred to as *protected* tasks. Differently from the other two, decision-making tasks require the participant's identity and contextual information to generate baseline information (e.g., health status of participants in a biomedical application). Therefore, these *private* tasks can *only* be performed on the participant's "personal mobile devices" (highest level of trust).

B. Task Deduplication

To handle duplicate tasks in different workflows arriving at the broker, we present a task deduplication and workflow consolidation mechanism. The brokers group service requests before proceeding with task deduplication. The duration (time window) for which a broker waits (δ_b^{wait} [s]) before deduplication is a tunable parameter. For a given rate of service request arrivals, the larger the window, the greater the chances of finding task duplicates. However, the windows cannot be too large as the workflow requests have to be serviced real time. This "pause-aggregate-service" strategy eliminates the unrealistic assumption of *strictly simultaneous* workflow arrivals at the broker.

Dedup—at the broker—parses the workflow descriptions to identify task duplicates and to create simplified workflows (with fewer tasks than before). Dedup looks for matching sub-graphs (connected group of tasks) between a pair of DAGs. Trivially, every single vertex in a DAG (workflow) is a sub-graph. Dedup starts with the comparison of Stage-0 tasks in the two workflows, as shown in Algorithm 1. Two tasks are considered to be "similar" when the following attributes match: task identifier (i.e., type), number and types of inputs (i.e., set of parent tasks, $\mathcal{P}$), and inputs' sizes (quantity of inputs), as shown in Algorithm 2. When tasks in two DAGs are similar, their corresponding sets of child tasks ($\mathcal{C}$s) are recursively checked for similarity. This recursive step is aimed at growing the size (i.e., number of tasks) of the matched sub-graph. In the recursive procedure, when the tasks under comparison, say task k of workflow 1 and l of workflow 2, cease to be similar, a link is created from k's parent to l. Also, l is added to the children set of Parent(k). The tasks belonging to the duplicate subgraph in workflow 2 are discarded. Note that, while checking for similarity, tasks that have been visited and have tested positive for similarity are marked, so that they need not be checked again. The worst-case time complexity of Dedup is $\mathcal{O}(|\mathcal{V}_1| \cdot |\mathcal{V}_2|)$, where $\mathcal{V}_1$ and $\mathcal{V}_2$ are the sets of vertices in the two input DAGs.

Algorithm 1: `Dedup`

Input: $\mathcal{K}$ and $\mathcal{L}$ are initially set to stage-0 tasks of the concurrent workflows

Output: Simplified workflows

  for every $k \in \mathcal{K}$ **do**

    if (visited(k) == `true`) **then**

      continue

    end if

    for every $l \in \mathcal{L}$ **do**

      if (visited(l) == `true`) **then**

        continue

      end if

      if (`checkSimilarity`(l,k) == `true`) **then**

        `Dedup`($\mathcal{C}^k$,$\mathcal{C}^l$)

      else

        Parent(l) = Parent(k)

        addChild(Parent(k),l)

      end if

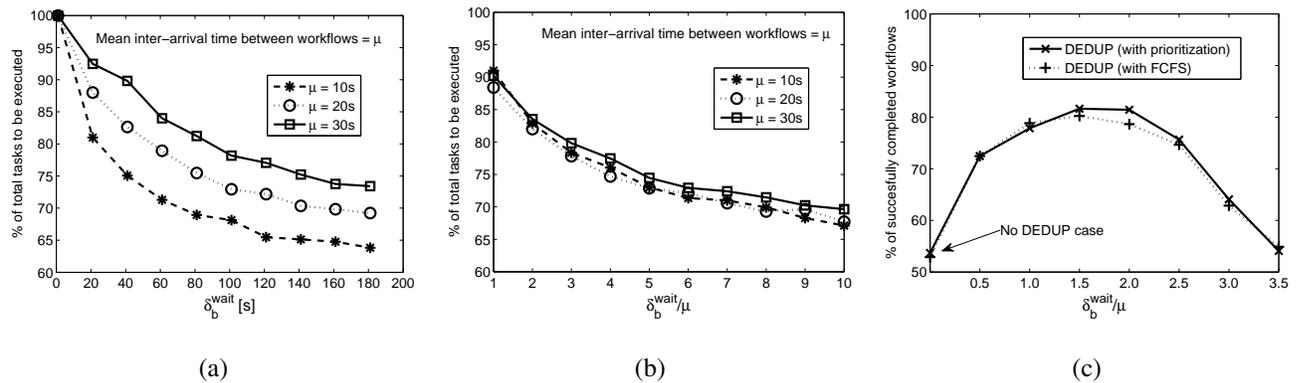   **end for**

  end for

Fig. 4. Decrease in percentage of the total number of tasks to be executed (while using `Dedup`) with increase in δ_b^{wait} (a) when δ_b^{wait} is not adapted to the arrival rate of workflows; (b) when δ_b^{wait} is adapted to the arrival rate of workflows (proportionally); (c) Behavior of `Dedup` in terms of percentage of successfully completed workflows with increase in δ_b^{wait}.

In the resulting workflow, k's immediate predecessor task will be "fork" point from which other deduplicated workflow branches out. In Fig. 3(a), the vertices corresponding to Peak detection

---

Algorithm 2: `checkSimilarity`

Input: Tasks k and l

Output: `true` or `false`

  if ($k.taskID == l.taskID$) AND ($k.output == l.output$) **then**

    if `checkSimilarity`($\mathcal{P}^k$,$\mathcal{P}^l$) **then**

      visited(k) = visited(l) = `true`

      visited($\mathcal{P}^k$) = visited($\mathcal{P}^l$) = `true`

      return `true`

    else

      return `false`

    end if

  else

    return `false`

  end if

---

in ECG signal and to Activity scoring of accelerometer output become fork points. Note that the time complexity of `Dedup` (or the total number of comparisons) is not altered by changing the order of comparison of different DAGs. Similarly, `Dedup` results in the same set of simplified workflows irrespective of the order of comparison of the different DAGs. It is not necessary that deduplication is done strictly before task allocation; deduplication is also achieved on the fly when tasks are already in execution. Under such circumstances, execution of duplicates is piggybacked, thus deduplicating at run time. Also, results of certain repetitive tasks are cached locally at the SPs so to deduplicate services. *Service forwarding*, either at run time or through caching, relies on inferential analysis from historical request traces.

IV. PERFORMANCE EVALUATION

We developed a simulator in Java[TM] to evaluate empirically the performance gains provided by the different components of `Maestro`. Simulations also allow us to evaluate at scale. In the following, firstly, we present details about our experiment methodology, specifically, the workflows, workflow traces, the service providers, and the SP dynamics. Then, we discuss specific simulation scenarios as well as provide the results that (i) demonstrate the benefits of `Dedup`, (ii) highlight the price of using different protection levels in mobile computing, and

(iii) illustrate the merits of replication-based failure handling.

Workflows: The workflows we used, inspired by the applications discussed earlier, are all task parallel with data-parallel sub-graphs built into some of them. Even though currently available example workflows (from the biomedical and robotics domains) can be used for preliminary evaluation of `Maestro`, evaluation at scales necessitates the creation of arbitrary workflows that are similar to the ones available in literature. Hence, we developed a workflow (DAG) generator to generate *arbitrary* workflows for large simulations. The synthetic workflows used in our simulations vary in terms of number of stages, tasks per stage, types and sizes of tasks at each stage and dependencies, and are representative of a wide range of task-parallel applications that `Maestro` can support.

Workflow Traces: At different instants, the service requesters (which may also be data providers) submit workflow requests to brokers while also specifying a deadline and a probability of success. The utility of the result from the workflow is assumed to be zero after the requester-specified deadline. Traces that capture workflow request arrivals over time in a mobile computing environment are not available in the literature. Workflow arrival traces in cloud and grid-computing environments cannot be adopted directly either. This is because the workflows, their deadline requirements, and arrival statistics are not representative of the applications or of the dynamics envisioned in `Maestro`. Hence, we developed a workflow-trace generator that can create multiple workflow arrival traces varying in terms of number of workflows and inter-arrival time between workflows as well as request-specific deadline and probability of success.

Service Providers: In our simulations we use a heterogeneous pool of SPs. The factors that contribute to heterogeneity are processing speed or capability (in terms of number of instructions per second), communication capability (in terms of bps), rate of battery drain for computation (in terms of mAh per instruction), and finally the duration of availability. We use the mean availability duration (the duration for which the SP is in the MDC) and the mean away duration (the duration for which the SP is not in the MDC) as well as their respective distributions to control the dynamics in the mobile computing environment. We choose these durations carefully to maintain an average number of SPs in the MDC as the three variables are related by Little's theorem.

Benefits of Task Deduplication: To demonstrate the benefits of task deduplication, we performed two experiments to ascertain the following two factors with and without `Dedup`:

TABLE I

CHARACTERISTICS OF THE HETEROGENEOUS MOBILE COMPUTING DEVICES IN OUR TESTBED.

Devices	Samsung Galaxy Tab	ZTE Avid N9120	Huawei M931	Toshiba Satellite	Raspberry Pi
Type of devices	Tablet	Smartphone	Smartphone	Laptop	Netbook
No. of devices	2	3	1	1	1
CPU	1GHz Dual-core ARM	1.2GHz Dual-core	1.5GHz Dual-core	2.13 GHz i3 Intel	700 MHz ARM
OS	Android v4.0	Android v4.0	Android v4.0	Windows 7	Windows 7
RAM [GB]	1	0.512	1	4	0.512
Battery [mAh]/[V]	7,000/4	1,730/5	1,650/10.8	4,200/10.8	2,200/5

TABLE II

AVERAGE EXECUTION TIMES OF TASKS OF ROBOTIC APPLICATIONS ON OUR TESTBED.

Computing Task	Samsung Galaxy Tab	Raspberry Pi	Toshiba Satellite Laptop
Location determination [s]	143	1100	28.6
Field estimation [s]	28.5	273	5.7
Error estimation [s]	0.3	3.3	0.06
Energy estimation [s]	0.3	3.3	0.06

(i) the reduction in percentage of total number of tasks to be executed in the MDC and (ii) the percentage of successfully completed workflows among all the requests submitted.

Experiment 1: As δ_b^{wait} is a tunable parameter, we observed performance in terms of reduction in number of tasks by varying it. We studied the behavior of Dedup when δ_b^{wait} is adapted to the arrival rate of workflows and when it is not. We created three distinct workflow traces, each with 100 requests, with mean inter-workflow-arrival durations of $\mu = 10, 20,$ and 30 s, respectively. The number of distinct workflows in each trace was set to 40, and the number of SPs to 10.

Observations: Figures 4(a) and 4(b) show that the percentage of total number of tasks to be executed in the MDC decreases by 25% when δ_b^{wait} is five times the inter-arrival duration μ of the workflows. This decrease will be greater when the number of distinct workflows in the traces is reduced below the current value of 40. Figure 4(a) was obtained by varying δ_b^{wait} in increments of 20 agnostically to the workflow arrival rate. As a result, in comparison to the trace with $\mu = 20s$, the percentage of tasks to be executed is higher for the trace with $\mu = 30$ s while it is lower for the one with $\mu = 10$ s. This is because for the same δ_b^{wait}, the number of workflows considered together for deduplication decreases with increase in μ. Therefore, adaptation of δ_b^{wait} with respect to μ is key to achieve improved performance (see Fig. 4(b)).

Experiment 2: We observed the performance in terms of percentage of successful workflow completions by varying the waiting period δ_b^{wait}. We created a workflow trace with a total of 500 requests. The mean inter-workflow-arrival duration in the workflow trace was set to $\mu = 10$ s and the number of distinct workflows in the trace was set to 10. The deadline of each workflow request was chosen randomly between 40 and 80 s, and the number of SPs was set to 10.

Observations: Figures 4(c) shows that the percentage of successful workflow completions in the MDC increases to as much as 83% (compared to the baseline no `Dedup` case at 53%) when δ_b^{wait} is twice the inter-arrival duration μ of the workflows. This increase will be greater when the failed workflow tasks are discarded (which we did not do to study the worst case). Figure 4(c) clearly highlights the situation when `Dedup` may not be beneficial in `Maestro`. Even though an increase in δ_b^{wait} results in a decrease in the total number of tasks to be executed, as shown in Figs. 4(a) and 4(b), the decrease is only sub-linear. The accumulation of tasks over time may result in an overload for the underlying SP pool as is the case when $\delta_b^{wait}/\mu > 2.0$ in Fig. 4(c) where the gain drops until finally reaching the baseline at $\delta_b^{wait}/\mu = 3.5$. Also, it is important to note that the task prioritization in `Maestro` results in improved performance in comparison to a First-Come-First-Served (FCFS) scheduling policy. The difference in performance will widen further when the variance in deadlines is greater than what we used here.

Testbed: The focus of this subsection is geared towards presenting the performance of our solution for real-world workflows. We prepared a testbed for our experiments, which consisted of Android- and Linux-based mobile devices with heterogeneous capabilities (summarized in Table I). We consider biomedical, robotic, and computer vision workflows presented in Fig. 3 to show the performance of our proposed controlled replication approach in order to overcome
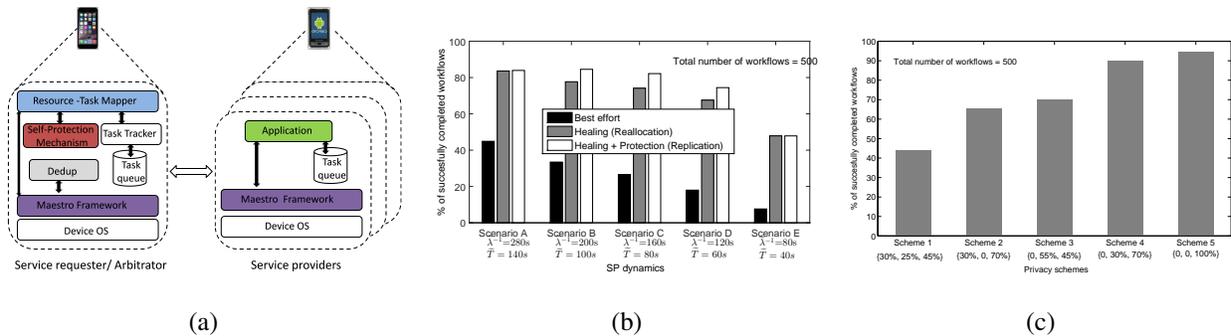
Fig. 5. (a) Architecture of `Maestro` including service requester and providers; (b) Percentage of successful workflow completions under different SP dynamics in the MDC. $\widetilde{\lambda}^{-1}$ [s] is the average inter-arrival duration and $\widetilde{T}$ [s] is the average availability duration of the SPs; (c) Percentage of successful workflow completed when different levels of protection schemes are employed.

uncertainty due to failure of SPs to finish allocated task. We profiled the time taken for various tasks in the workflow on all the devices in our testbed. The architecture of our testbed is shown in Fig. 5(a), which is based on our work in [24]. The service requester device contains the resource task mapper, which is responsible to allocate task to different service providers.

Input data set: For each of application given in Fig. 3 we obtained the execution time of each task by extensive offline profiling which involved running each task of an application with multiple input data. For applications in Fig. 3(a,b) we generated the input data artificially and for application in Fig. 3(c) we used input data as images from the Berkeley image segmentation and benchmark dataset [25].

Motivation for Controlled Replication: We utilize real-time distributed robotic applications, as in Fig. 3(b), to motivate the need for controlled replication. For such applications, e.g., detection of harmful chemicals in a field, quick execution of tasks given in the workflow is essential. Table II shows the set of tasks in the robotic application and execution times of these tasks. We see that location-determination and field-estimation tasks are critical because of their high execution times. If we fail to get results for these tasks from a SP due to network disconnection or lack of resources at that SP, the workflow may not meet the deadline. To avoid unnecessary idle waiting times incurred in reactive failure handling, a proactive approach like our proposed controlled replication will help meet the application deadline.

Benefits of Controlled Replication: `Maestro`'s task-scheduling mechanism employs proactive protection (selective controlled replication of large tasks) in addition to reactive healing

(reallocation of failed tasks) in order to provide robustness. To study the improvement in performance provided by healing and protection over the best-effort (baseline) case, we performed an experiment under different service provider dynamics. We set the *average number of active SPs* in the MDC to 30. We varied the SP dynamics in the MDC from highly volatile to highly stable by tuning the following parameters—*average inter-arrival duration* ($\widetilde{\lambda}^{-1}$ [s]) of SPs and *average availability duration* ($\widetilde{T}$ [s]) of SPs. The average number of SPs and the aforementioned parameters are related by Little's law. We created a workflow trace with a total of 500 requests. The mean inter-workflow-arrival duration in the trace was set to $\mu = 20$ s and the number of distinct workflows was set to 10. The trace has a mix of small (66%) and large (33%) workflows (differing in terms of task sizes and deadlines). The deadline of the small workflows was chosen randomly between 40 and 80 s, while the large workflows' deadline was picked randomly between 80 and 160 s.

Observations: Figure 5(b) shows five ordered scenarios, A through E, where A corresponds to a highly stable MDC and E corresponds to a highly volatile one. The performance of `Maestro` with only healing as well as with both healing and protection is always better than the baseline case. In Scenario B through D the use of protection in addition to healing prevents more workflows from failing than using only plain self-healing. This is because when using healing in isolation the time taken to recover from a failure of a task belonging to a "large" workflow is more than twice that the task execution time. However, selective replication of such critical tasks (which may easily jeopardize the workflow when they fail), as done in protection, prevents a greater percentage of workflows from failing. However, note that protection does not provide any additional gain over plain healing in Scenario A when the probability of SP failure during task execution is very low and in Scenario E when the probability is very high.

Price of Using Multiple Protection Levels: Even though `Maestro` provides multiple levels of protection to different tasks through controlled access by authorized service providers, there is a price to pay for it in terms of performance. Authorizing SPs to execute only certain types of tasks restricts the feasibility region of the solution to the problem the broker is trying to solve. In `Maestro`, the broker aims at scheduling tasks in the MDC in such a way that the percentage of successfully completed workflows be maximized. We performed an experiment to quantify the difference in performance when multiple levels of protection are employed by varying the percentage of private, protected, and public tasks in the workflows while keeping the percentage
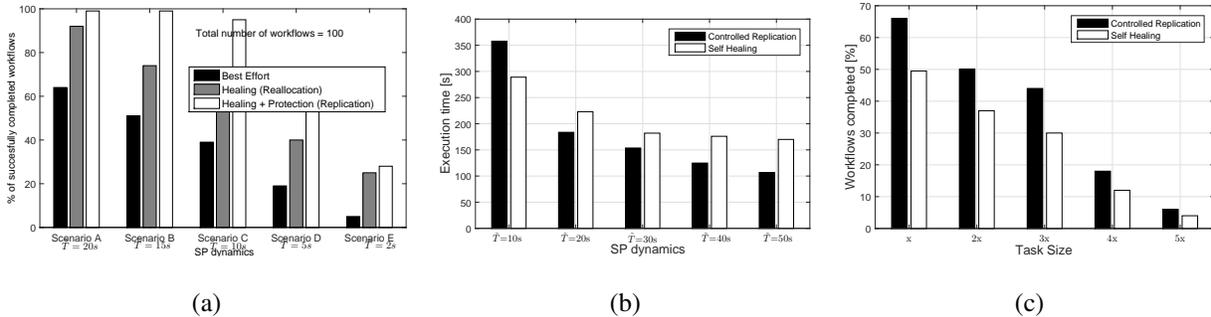
Fig. 6. (a) Comparison of controlled replication technique with best effort and healing approaches via experiments on our testbed; Comparison of performance of controlled replication and self-healing in terms of (b) execution time by varying the SP statistics in the MDC; (c) percentage of successful workflows by varying the task sizes.

of personal, trusted, and untrusted 3^{rd} party devices in the MDC fixed. These percentages are 1, 33, and 66%, respectively. We created a workflow trace with a total of 500 requests. The mean inter-workflow-arrival duration in the workflow trace was set to $\mu = 10$ s and the number of distinct workflows in the trace was set to 10. The deadline of each workflow request was chosen randomly between 40 and 80 s.

Observations: Figure 5(c) shows the five schemes with decreasing degree of levels of protection (or increasing number of public tasks in the workflows). We observed that, as the degree of protection is decreased, i.e., the percentage of public tasks increased, the performance in terms of percentage of workflow completions increased. Scheme 5 represents an unrestricted scenario where all tasks are public, while Scheme 1 is extremely restricted. Scheme 2 through 4 reflect what may be adopted in real-world deployments. The performance of Schemes 2 and 3 can be improved to match that of Scheme 4's by either relaxing the deadline requirements or by increasing the MDC size. A small relaxation in the deadline is a marginal cost to incur if multiple protection levels are desired.

Controlled Replication for Biomedical Applications: We focus on the stress-detection application, which receives input data from variety of smartphone sensors. To study the improvement in performance provided by self-healing and self-protection over the best-effort (baseline) case, we performed an experiment under different service-provider dynamics. The goal of our technique protection is controlled replication of collective tasks in the workflow. We studied the performance of our approach in case of real-time deadline constraints. From the concepts explained earlier, we determined the anxiety-detection task to be the most critical task of the

workflow. As a result, this task was replicated on a subset of SPs based on (5).

Observations: Figure 6(a) shows five ordered scenarios, A through E, where A corresponds to a highly stable MDC and E to a highly volatile one. A successful workflow is the one that completes all its tasks within the user-specified deadline. As seen earlier, the performance of `Maestro` with only healing as well as with both healing and protection is always better than the baseline case. As the mean duration of service time becomes higher, we observe both healing and protection with healing to give similar performance. We see that the performance of baseline technique with 63% successful workflows is much lower than protection with healing and only healing approach with above 85% successful workflows. In some cases, healing performs worse than combination of healing and protection as healing is a reactive approach and leads to increase in makespan due to failure of certain tasks. These experiments show the robustness of our technique for real-world, real-time workflows.

Controlled Replication for Computer-vision Applications: To study controlled replication in computer vision domain we implemented Workflow c1.a as shown in Fig 3(c). We consider two competing fault-tolerance mechanism, self-healing and controlled-replication. We consider one of the device in the testbed, (Samsung Galaxy Tab) to serve as the broker and it sends multiple execution requests of Workflow c.1 to two other devices (ZTE Avid, Huawei) forming the MDC. To implement self-protection, tasks given to device ZTE Avid, Huawei is also replicated on another devices from the testbed given in the Table I.

The performance of both controlled replication and self-protection over different SP dynamics in the MDC is shown in Fig. 6(b). We observe in Fig. 6(b) that for mean availability duration, $\tilde{T} = 10$s, self-healing performs better than self-protection, however, as the mean availability duration of the SPs increases controlled-replication performs better. However, this gain comes at the cost of employing higher the number of SPs than in the self-healing approach. We also calculate the percentage of successfully completed workflows within 200s for different task sizes in the workflow given in Fig. 6(c). The mean availability duration of SPs is considered to be, $\tilde{T} = 40$s. The task size here corresponds to the input data size k_1^0 in Workflow c.1. As discussed earlier that as the task size increases reactive approach such as self-healing does not perform well as it leads to wastage of resources and time in case of failure of execution of task at the SPs. We observe that via controlled replication we are able to execute higher number of workflows in the same amount of time. As the task sizes increase further we see that neither self-healing

or controlled replication perform well. To improve the performance of MDC the availability duration of SPs should be increased which will increase the number of completed workflows.

V. Conclusion

We presented our vision for harnessing the sensing, computing, and storage capabilities of mobile and fixed devices in the field *as well as* those of computing and storage servers in remote datacenters to form a mobile/fixed device cloud. We discussed `Maestro`, a framework for robust and secure pervasive mobile computing in a mobile/fixed device cloud, which can enable a wide range of mobile applications that rely on real-time, *in-situ* processing of data generated in the field. We focused on two key components of `Maestro`—(i) a task scheduling mechanism that employs controlled task replication in addition to task reallocation for robustness and (ii) `Dedup` for task deduplication among concurrent workflows. A flexible architecture to impart multiple levels of protection to tasks is also discussed. We are currently working towards addressing security issues that arise when malicious nodes provide incorrect results and not just denial-of-service attacks.

References

[1] H. Viswanathan, P. Pandey, and D. Pompili, "Maestro: Orchestrating Concurrent Application Workflows in Mobile Device Clouds," in *Proc. of Workshop on Distributed Adaptive Systems (DAS) at IEEE Intl. Conference on Autonomic Computing (ICAC)*, Wurzburg, Germany, July 2016.

[2] M. Satyanarayanan, "Pervasive Computing: Vision and Challenges," *IEEE Personal Communications*, vol. 8, no. 4, pp. 10–17, 2001.

[3] S. Deng, L. Huang, J. Taheri, and A. Y. Zomaya, "Computation Offloading for Service Workflow in Mobile Cloud Computing," *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 12, pp. 3317–3329, 2015.

[4] X. Chen, "Decentralized Computation Offloading Game for Mobile Cloud Computing," *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 4, pp. 974–983, 2015.

[5] B.-G. Chun, S. Ihm, P. Maniatis, M. Naik, and A. Patti, "CloneCloud: Elastic Execution between Mobile Device and Cloud," in *Proc. of The European Professional Society on Computer Systems (EuroSys)*, Salzburg, Austria, Apr. 2011.

[6] E. Cuervo, A. Balasubramanian, D.-k. Cho, A. Wolman, S. Saroiu, R. Chandra, and P. Bahl, "MAUI: Making Smartphones Last Longer with Code Offload," in *Proc. of Intl. Conference on Mobile Systems, Applications, and Services (MobiSys)*, San Francisco, CA, Jun. 2010.

[7] M. R. Ra, A. Sheth, L. Mummert, P. Pillai, D. Wetherall, and R. Govindan, "Odessa: Enabling Interactive Perception Applications on Mobile Devices," in *Proc. of Intl. Conference on Mobile Systems, Applications, and Services (MobiSys)*, Bethesda, MD, Jun. 2011.

[8] M. Othman, F. Xia, A. N. Khan *et al.*, "Context-Aware Mobile Cloud Computing and Its Challenges," *IEEE Cloud Computing*, vol. 2, no. 3, pp. 42–49, 2015.

[9] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The Case for VM-Based Cloudlets in Mobile Computing," *IEEE Pervasive Computing*, vol. 8, no. 4, Oct.-Dec. 2009.

[10] M. S. Gordon, D. A. Jamshidi, S. Mahlke, Z. M. Mao, and X. Chen, "COMET: Code Offload by Migrating Execution Transparently," in *Proc. of the USENIX Conf. on Operating Systems Design and Implementation (OSDI)*, Hollywood, CA, Oct. 2012.

[11] Y. Liu, M. J. Lee, and Y. Zheng, "Adaptive Multi-Resource Allocation for Cloudlet-Based Mobile Cloud Computing System," *IEEE Transactions on Mobile Computing*, vol. 15, no. 10, pp. 2398–2410, Oct 2016.

[12] A. Mtibaa, K. A. Harras, K. Habak, M. Ammar, and E. W. Zegura, "Towards Mobile Opportunistic Computing," in *Proc. of Intl. Conference on Cloud Computing*, 2015.

[13] M. Conti and M. Kumar, "Opportunities in Opportunistic Computing," *IEEE Computer*, vol. 43, no. 1, pp. 42–50, Jan. 2010.

[14] A. Passarella, M. Kumar, M. Conti, and E. Borgia, "Minimum-Delay Service Provisioning in Opportunistic Networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 22, no. 8, pp. 1267–1275, 2011.

[15] C. Shi, V. Lakafosis, M. H. Ammar, and E. W. Zegura, "Serendipity: Enabling Remote Computing among Intermittently Connected Mobile Devices," in *Proc. of Intl. Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Hilton Head Island, SC, June 2012.

[16] D. Chu and M. Humphrey, "Mobile OGSI.NET: Grid Computing on Mobile Devices," in *Proc. of IEEE/ACM Intl. Workshop on Grid Computing*, Pittsburgh, PA, Nov. 2004.

[17] L. dos S. Lima, A. T. A. Gomes, A. Ziviani, M. Endler, L. F. G. Soares, and B. Schulze, "Peer-to-peer Resource Discovery in Mobile Grids," in *Proc. of the Intl. Workshop on Middleware for Grid Computing (MGC)*, Grenoble, France, Nov. 2005.

[18] J. Hwang and P. Aravamudham, "Middleware Services for P2P Computing in Wireless Grid Networks," *IEEE Internet Computing*, vol. 8, no. 4, pp. 40–46, 2004.

[19] J. Roth and C. Unger, "Using Handheld Devices in Synchronous Collaborative Scenarios," *Ubiquitous Comp.*, vol. 5, no. 4, pp. 187–199, 2001.

[20] L. F. Bittencourt and E. R. Madeira, "Towards the Scheduling of Multiple Workflows on Computational Grids," *Journal of Grid Computing*, vol. 8, no. 3, pp. 419–441, 2010.

[21] G. L. Stavrinides and H. D. Karatza, "Scheduling Multiple Task Graphs in Heterogeneous Distributed Real-time Systems by Exploiting Schedule Holes with Bin Packing Techniques," *Simulation Modelling Practice and Theory*, vol. 19, no. 1, pp. 540 – 552, 2011.

[22] Y.-L. Tsai, K.-C. Huang, H.-Y. Chang, J. Ko, E. T. Wang, and C.-H. Hsu, "Scheduling Multiple Scientific and Engineering Workflows through Task Clustering and Best-Fit Allocation," in *Proc. of IEEE World Congress on Services (SERVICES)*, Honolulu, HI, Jun. 2012.

[23] G.-L. Park, B. Shirazi, and J. Marquis, "DFRN: A New Approach for Duplication based Scheduling for Distributed Memory Multiprocessor Systems," in *Proc. of Intl. Parallel Processing Symp.*, Geneva, Switzerland, Apr. 1997.

[24] H. Viswanathan, E. K. Lee, I. Rodero, and D. Pompili, "Uncertainty-aware Autonomic Resource Provisioning for Mobile Cloud Computing," *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 8, pp. 2362–2372, 2015.

[25] D. Martin, C. Fowlkes, D. Tal, and J. Malik, "A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics," in *Proc. of Intl. Conference on Computer Vision (ICCV)*, Vancouver, BC, July 2001.