

# AN ACTOR-BASED APPROACH TO COORDINATING CROWD-SOURCED SERVICES

Ahmed Abdel Moamen and Nadeem Jamali

Department of Computer Science, University of Saskatchewan, Saskatoon, SK, Canada

moamen@agents.usask.ca, jamali@cs.usask.ca

## Abstract

As personal mobile devices become gain popularity, not only is it possible receive a variety of services virtually anywhere, the sensors on the same devices can also actively contribute real-time data from their contexts to be used in services. A busy professional could find a restaurant to go to for a quick lunch based on information available from smartphones of people already there having lunch, waiting to be seated, or even heading there. Also imagine a mid-21<sup>st</sup> century fine-grained democracy where citizens might authorize use of a police officer's firearm – in real time – based on a live video feed of the scene. Although the programming required for offering a new service of this sort can be significant if done from scratch, we argue that in many cases it does not. We identify core communication mechanisms which underly many crowd-sourced services, and present preliminary design of a middleware which implements them. Service designers may launch novel services over this middleware by simply plugging in small pieces of service specific code.

This paper identifies the coordination mechanisms required for these crowd-sourced services as types of multi-origin communication. We present details of how these core mechanisms can be implemented using Actors, and introduce high-level programming constructs for launching a new service. In addition, we present our design of a middleware for crowd-sourced services using multi-origin communication. Finally, we use examples to illustrate the implementation of services.

**Keywords:** Crowd-sourced services, coordination, multi-origin communication, Actors.

## 1. INTRODUCTION

As personal computational devices such as smart phones, Google glasses, etc., become ubiquitous, so do the sensors (and potentially actuators) on those devices. Not only can people receive services virtually anywhere, there is a way for potential services to rely on real-time snapshot of every place that a device might be located. Consider a restaurant recommendation service which samples data collected about experiences of clients at a number of restaurants in a neighborhood and ranks them according to the service experience of the clients currently there. The source of the data could be sensor feeds on clients' mobile devices, which could guess whether they are waiting in line, seated, enjoying their meals, paying or leaving. Consider, another service combining real-time routing information -- such as is collected for showing traffic information on Google Maps -- with wait times at hospital emergencies to recommend which one to go to when in need of urgent care. If we also consider user input explicitly or implicitly entered into the devices, another class of services can be offered, from real-time polling, to instant censuses, and even voting in elections. The generation growing up in the world of Twitter and Facebook may find it quaint that democracies hold elections only once every few years. Why not have much more frequent votes on much finer grained decisions effecting citizens every day? It is entirely conceivable that

in the not so distant future, a form of fine-grained democracy may emerge where citizens get to vote – in real time – on if and when a police officer's firearm is enabled, based on a live video feed of the scene

We are interested in an opportunity created by the similarity in the patterns of communication required for such services, which we refer to as multi-origin communication. Previously (Geng & Jamali, 2013) differentiated between single-origin and multi-origin types of multi-sender communication.<sup>1</sup> The single-origin type of multi-sender communication is initiated by a single party which solicits interest from other parties to join together in sending a particular message. An example of this would be a workplace petition. Using email, the option usually available is for one person to be the recognizable active sender of the petition, with the remaining people passively listed in the "cc" field. The alternative we are interested in is to allow all senders to be equally responsible for sending such a message, despite its single point of origin. In multi-origin (implicitly also multi-sender) communication, the expectation is that that there is no single party who must

<sup>1</sup> We have previously referred to multi-sender communication as many-to-many communication; however, the emphasis was on the multiple senders, because there is sufficient existing research on multiple recipients. For multi-origin multi-sender messages, we skip multi-sender because it is implicit.

take the lead. In other words, multiple parties may autonomously launch messages which could then be aggregated in order to create a group message.

It turns out that unlike single-origin multi-sender messages, multi-origin messages require a setup in advance. Consider a public square where a number of citizens spontaneously begin to gather to party or protest. In this context, the physical space of the square serves as part of a setup which allows mutual observation, an opportunity to join in or leave, to endorse, reject or refine the collective message or experience over time. The closest electronic equivalent of such a physical space would be services such as Twitter, which allow people to observe others' tweets in an aggregate form (which is quite natural in physical space, but requires filtering and counting mechanisms in electronic space), endorse them by adopting hashtags, improve upon the message, and so on. In general, for a crowd (or mass) - conceived communication to happen, there is a need for a mechanism to be in place to coordinate the generation of the message by soliciting messages, receiving them, and then aggregating them into a group message. The solicitation lays out the rules to be followed for selection of the potential senders, receiving their messages, and aggregating them. For example, imagine a multi-stage communication with the first solicitation being to invite nominations for topics to have the message on, followed by a vote to select the topic, and followed by a solicitation of messages, followed by a final vote to agree on an aggregate message. The communication could be one-time, periodic, or continual. There may or may not be a time-out for responding to the solicitation. All these aspects would be laid out in the original solicitation.

Our approach is to construct key coordination mechanisms required for this class of services requiring multi-origin communication, and then allow service designers to provide service specific code -- which uses the available mechanisms -- in order to launch their services. A new service could then be implemented by a service provider by simply providing the code for service-specific tasks, such as for solicitation of messages from mobile devices, the needed abstraction of sensor feeds on the devices, and aggregation of the messages to create a group message. These communications could be set up as continual real-time updates for a web or mobile app based service, serving user requests as they arise, or they could be launched each time a request comes in. The pieces of code provided by the service provider would simply be plugged into the common coordination mechanisms to create a new service.

The rest of paper is organized as follows: Section 2 presents related work. Section 3 describes how multi-origin communication can support crowd-sourced services. Section 4 presents our preliminary design of design of a middleware for crowd-sourced services using multi-origin communication. Section 5 uses two very different type of examples to illustrate how services could be

implemented using this approach. Finally, Section 6 concludes the paper.

## 2. RELATED WORK

There have been a number of projects -- both in academia and industry -- involving crowd-sourced services. The term crowd-sourced can refer to two types of services: participatory sensing services and crowdsensing services. Participatory sensing involves explicit participation by the human being in possession of the mobile device, whereas crowdsensing relies on sensor feeds automatically flowing from devices to servers.

We first present some representative examples of both these types of crowd-sourced services, and then discuss some existing frameworks for enabling such services.

### 2.1 Crowd-Sourced Services

Some of the best examples of participatory sensing services can be found in services aimed at assisting automobile drivers.

Waze (United States Patents, US Patent No. 8,271,057, 2009) is one of the largest community-orientated mobile travel applications with users volunteering information about their driving experience in real time, by reporting on congestions, delays, and gasoline prices. These reports then become the basis for information displayed on other drivers' maps (on their mobile devices), to help them make routing decisions.

Similarly, TrafficPulse (Li, Liang, Lee, & Byon, 2012) combines sensor data from mobile devices with real-time traveler reports from frequent travelers, and then offers this information to other drivers in an aggregated form.

Crowd-sourcing has also been found to be useful in efforts to coordinate rescue efforts following major disasters, such as the Haitian earthquake in 2010 (Zook, Graham, Shelton, & Gorman, 2010). Information aggregated from social media (e.g., blogs, emails, tweets, and facebook status updates) was used to overcome challenges created by both the inadequacy of maps and the change in landscape because of the devastation.

CrowdHelp (Besaleva & Weaver, 2013) uses smartphones to collect direct feedback from mobile users about their medical condition, in combination with data coming from sensors in smartphones. This information is used to enable swift response to emergencies. For example, when CrowdHelp is used for emergency reporting, mobile users submit information relevant to an event (such as the number of injured people and their state) to a central server. This information is collected and sent to the nearest health care facility capable of treating the injured.

Among crowdsensing services, the real-time traffic information displayed on Google Maps is arguably the most widely used one. The service relies on location data voluntarily made available by users of Google's services,

which is then aggregated and then visualized on Google's Maps to show traffic flow. Since Google's acquisition of Waze in 2012, Waze's participatory sensing service has now been combined with Google's crowdsensing service for providing real-time traffic flow information.

Crowdsensing has also been used by Uga et al. (Uga, Nagaosa, & Kawashima, 2012) in an earthquake warning system, which uses data from accelerometers present in many modern mobile devices to detect seismic vibrations. Devices send reports of likely seismic activity to a server which then aggregates the reports received to send out warnings.

## 2.2 Mobile Crowd-Sourced Frameworks

Efforts to build frameworks for crowd-sourced applications have focused on narrow application areas, making it difficult to utilize them for a wider class of services. These frameworks are particularly lacking in reusable coordination mechanisms of the type we are proposing.

Medusa (Ra, Liu, La-Porta, & Govindan, 2012) is a programming framework for crowd-sourced applications. A task (such as video documentation or citizen journalism) is launched by a requester, and workers are solicited through Amazon's Mechanical Turk (AMT) service. These workers -- volunteering smartphone users -- then provide raw or processed data to be used as part of a social or technical experiment. Typically, an XML-based programming language, MedScript, is used to specify the required task as a series of several stages, from the initial recruitment of volunteer workers, to the workers' (say, for a video documentation task) recording videos on their smartphones, summarizing them, and then sending them back. The stages can involve actions selectable from a library of executables, which are downloaded to mobile devices from a cloud server. Medusa's limitation in terms of wider applicability

to a large class of crowd-source services lies in the limited types of activities that the tasks can involve, and the limited types of interactions that the parties can have.

AnonySense (Cornelius, et al., 2008) is another framework for collecting and processing sensor data, which pays particular attention to privacy concerns. AnonySense allows a requester to launch one of a selected group of applications with their parameters. The application then distributes sensing tasks across anonymous participating mobile devices (referred to as carriers), and finally aggregates the reports received from the carriers.

CDAS (Liu, et al., 2012) is an example of participatory crowd-sensing frameworks. In CDAS, the participants are part of a distributed crowd-sensed system. The CDAS system enables deployment of various crowd-sensing applications that require human involvement for simple verification tasks to deliver high accuracy services. Similar to CDAS, MOSDEN (Jayaraman, Perera, Georgakopoulos, & Zaslavsky, 2013) is a collaborative mobile sensing framework that operates on smartphones to capture and share sensed data between multiple distributed applications and users.

Mobile Edge Capture and Analysis middleware for social sensing applications (MECA) (Ye, Ganti, Dimaghani, Grueneberg, & Calo, 2012) is a middleware for efficient data collection from mobile device. It uses a multi-layer architecture to take advantage of similarities in the data required for different applications to lower the demand on the devices on which data is being collected. Although MECA takes an interesting approach in addressing the problem of growing demand of vertically integrated applications competing for limited resources on mobile devices, its focus is limited to a narrower class of applications, and does not address wider programmability challenges as we attempt to do in this work.

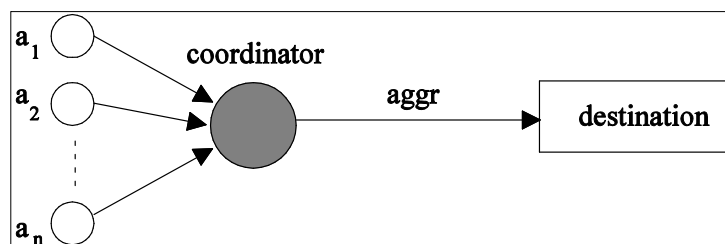


Figure 1. Multi-Origin Communication

## 3. SUPPORTING MULTI-ORIGIN COMMUNICATION

As illustrated by Figure 1, multi-origin communication involves a number of autonomous senders sending messages which are somehow aggregated into a group message. However, as previously explained in the introduction, this type of communication requires an advance setup for coordinating the communication.

We describe the implementation of this coordination setup as an Actor (Agha, 1986) program. Actors are autonomous concurrently executing primitive agents (i.e., active objects) which communicate using asynchronous messages.<sup>2</sup> We represent the different parties involved in a

<sup>2</sup> Actors are emerging as the model of choice for very large-scale applications such as Facebook chat service and Twitter have been written in actor languages (Agha G., 2014).

communication using actors, and define complex communications in terms of asynchronous actor messages.

The requester of a multi-origin communication makes a function call in order to launch the communication. The call passes parameters specifying the potential senders -- the constituency -- to be invited to participate in the communication, as well as the way in which the messages would be aggregated. As illustrated in Figure 2, invocation of the function results in the creation of a new coordinator

actor capable of coordinating the communication, which is next told to invite the constituency to participate. The coordinator then sends invitations to the members of the constituency (the senders) to send their message; when applicable, it also sends them parameters advising on how to construct their messages (such as by tapping into a set of sensors, or soliciting input from the user), how often to send them (once or periodically, how frequently), etc.

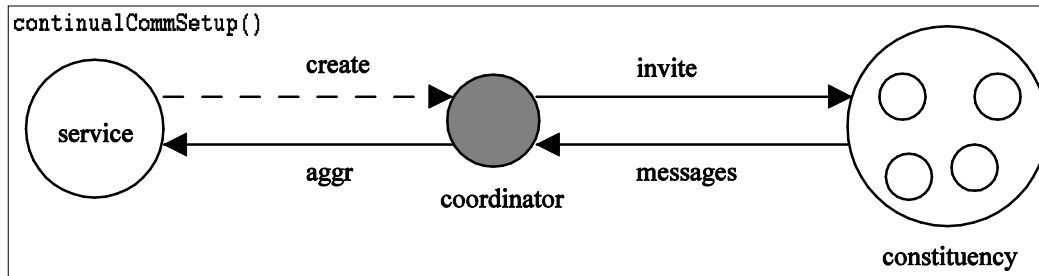


Figure 2. Service Setup

We assume that each sender is an actor with a method to receive these requests, and the capability to create the types of messages. Given that there are a relatively small number of sensors on mobile devices, the parameters could simply be specifying which sensors to be tapping into, with what frequency, and what periods to be averaging the feeds over, etc. However, coordinators for some services may be more interested in hearing about higher-level events -- such as a restaurant client sitting down at the table, finishing eating, paying the bill -- which would require more significant local processing to generate than simply receiving sensor feeds. This could be supported in various ways: by migrating an actor with the required behavior to the sender, by sending the code as a parameter to create an actor locally, or simply by frequently updating the sender-side application to include the functionality needed by every type of request.

As the senders send their messages, the messages are aggregated by the coordinator according to its own behavior, to generate group messages on behalf of the senders.

We specifically introduce two types of such setups. The first -- one-off multi-origin communication -- is to solicit a group message from a number of senders with a termination condition and a timeout. This would be the type of communication used to serve one-time requests, such as to hold a census or an election, or to satisfy a one-off request to recommend a restaurant with a short waiting time. The second -- continual multi-origin communication -- is to solicit a continual feed of group messages from a number of senders. This would be useful for a service provided over the web or through a mobile application where site visitors or application users seek up-to-date information (say) on restaurant waiting times in a neighborhood. For some services, such as the one for restaurant recommendations, the choice of one or the other setup would depend on the frequency of requests, the number of potential senders of

messages, etc. For instance, it would not be useful to be maintaining up-to-date information about all restaurants when there are very few requests for recommendations; however, it would be wasteful to solicit one-off communications for frequent requests.

### 3.1 One-Off Multi-Origin Communication

In a one-off multi-origin communication, the coordinator actor expects at most one message from any sender. It collects messages until either a sufficient number of messages has been received (as can be tested using a termination function), or a timeout has been reached; it then proceeds to aggregate the messages, and sends the aggregate to the requester on behalf of the senders. An example of a multi-origin communication with timeout would be an electronic voting service, where the coordinator expects no more than one vote from each voter and there is a deadline by which all votes must be in.

Figure 3 illustrates the execution of a one-off multi-origin communication using an actor event diagram (Agha, 1986). In the figure,  $sender_1$  through  $sender_n$  are the prospective senders. There is a clock actor to which the requester sends a request to notify the coordinator when the timeout has been reached. We assume that the clock is local to the coordinator and has a way of notifying in a timely manner. The requester initiates the communication by calling the function `oneOffCommSetup(coordClass, constit, termCond, timeout)`, where `coordClass` is the desired behavior of the coordinator, `constit` is a list of senders, `termCond` is a function to test the termination condition indicating receipt of a sufficient number of messages, and `timeout` is a time when the coordinator would stop accepting messages from the senders.

Once the coordinator is created, it sends announcements to all senders, and begins collecting messages. The coordinator expects to receive the maximum of one message

from each sender. After the timeout is reached, the coordinator sends a message to the requester with an aggregate of all responses.

The coordinator actor's behavior can be defined by extending the `multicall` selective blocking broadcast operation defined in (Geng & Jamali, 2013) with support for timeouts, or directly using the following three methods:

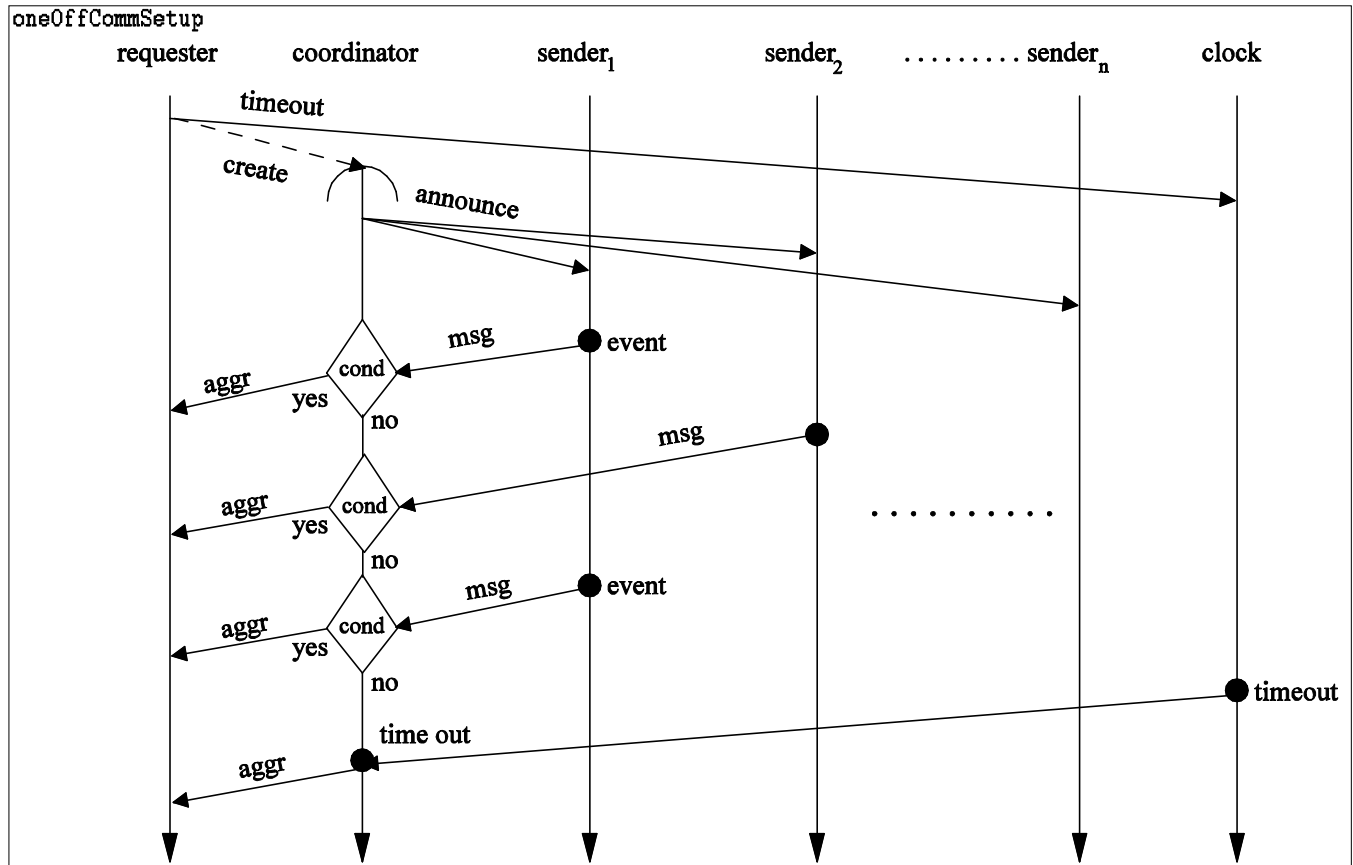


Figure 3. One-Off Multi-Origin Communication

- `announce(constit)`, used by the requester to instruct the coordinator to solicit messages from members of the constituency.
- `sendMessage(msg)`, used by the senders to send their messages to the coordinator.
- `timeout()`, used by the clock to tell the coordinator that the timeout has been reached.

```

void oneOffCommSetup(coordClass, constit,
    termCond, timeout, custom) {
    coordinator = createCoordActor(coordClass,
        termCond, custom);
    coordinator <- announce(constit);
    clock <- timeoutSetup(coordinator, timeout);
}
  
```

Figure 4. Pseudocode for `oneOffCommSetup`

A sender actor's communication behavior is defined by one method: `receiveAnnouncement(serviceParams)`. This is the method invoked when the solicitation is received from the coordinator, and it carries out the computations specified in `serviceParams` in order to create its message.

Figure 4 shows pseudocode for the `oneOffCommSetup` function. The `createCoordActor` function creates a new coordinator actor with the termination condition and application-specific customization initialized in its behavior, and returns the coordinator name. Once the coordinator has been created, a message is sent to the coordinator to broadcast an announcement to all senders. Another message is sent to the clock actor instructing it to notify the coordinator when the timeout is reached.

### 3.2 Continual Multi-Origin Communication

In a continual multi-origin communication, the coordinator expects multiple messages from each sender over time, and periodically aggregates them and sends updates to the communication's requester. When a new message arrives, the coordinator checks whether it warrants an update, or whether the interval for which it was to collect messages has passed. In either case, it forwards an aggregate

of messages received since the beginning of the interval to the requester. An example of continual communication would be that of a restaurant recommendation service available over the web, which attempts to offer up-to-day information to site visitors. The service could also be customized for individual visitors, based on their geographic locations, preferences, etc.

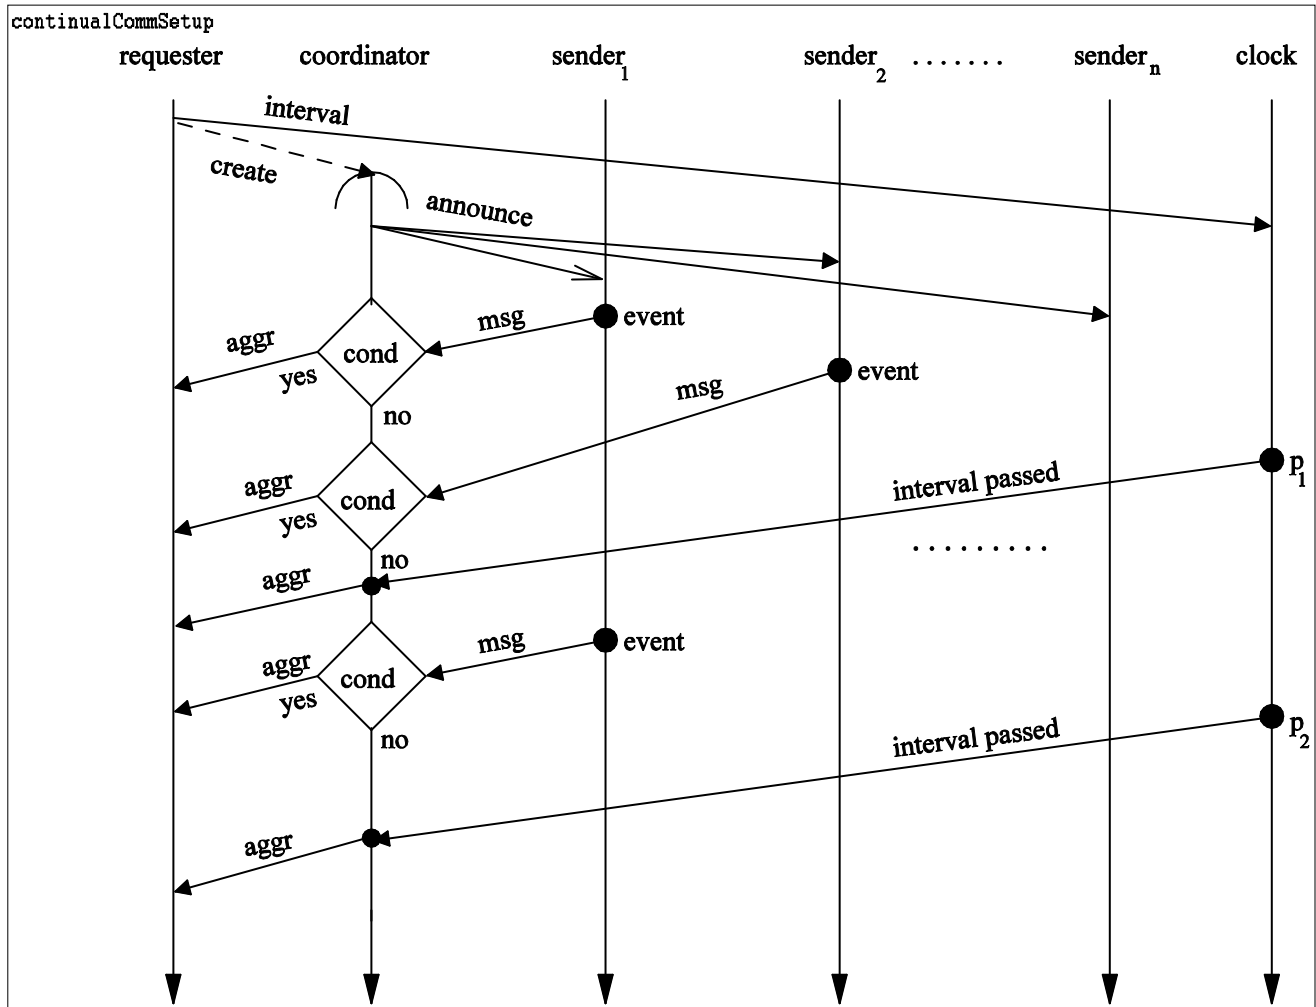


Figure 5. Continual Multi-Origin Communication

Figure 5 illustrates the execution of a continual multi-origin communication using an actor event diagram.  $sender_1$  through  $sender_n$  now send multiple messages over time, reporting local updates. Also, the clock actor periodically

A continual communication is initiated by the requester by calling the function `continualCommSetup(coordClassconstit, updateCond, interval)`, where `coordClass` is the desired behavior of the coordinator, `constit` is the list of prospective senders, `updateCond` specifies the condition in which the requester should be

(i.e., after every `interval` period of time) notifies the coordinator of the passage of an interval, at which time the coordinator computes a new aggregate.

immediately updated<sup>3</sup>, and `interval` specifies the intervals at which the coordinator would be notified by the clock.

Once the coordinator has been created, it broadcasts an announcement to all senders, and then waits to receive messages. Senders either send updates periodically or when

<sup>3</sup> This should also lead to resetting of the interval with the clock; this is not shown in the event diagram to avoid making it too crowded.

they observe an interesting event (such as a change in the level of activity in a restaurant, for example).

A coordinator actor's behavior is defined by the following methods:

- `announce(constit)`, used by the requester to instruct the coordinator to solicit messages from members of the constituency.
- `sendMessage(msg)`, used by the senders to send messages to the coordinator.
- `interval()`, used by the clock to inform the coordinator of the passage of each interval.

A sender actor's behavior is defined by one method: `receiveAnnouncement(serviceParams)`. This is the method invoked when the solicitation is received from the coordinator, and it carries out the computations specified in `serviceParams` required for creating its messages.

```
void continualCommSetup(coordClass, consti,
                        updateCond, interval, custom) {
    coordinator = createCoordActor(coordClass,
                                    updateCond, custom);
    coordinator <- announce (consti);
    clock <- intervalSetup(coordinator,
                           interval);
}
```

Figure 6. Pseudocode for *continualCommSetup*

Figure 6 shows the pseudocode for function `continualCommSetup`. The `createCoordActor` function creates a new coordinator actor with an update condition and application-specific customization initialized in its behavior, and returns the coordinator name. Once the coordinator has been created, a message is sent to the coordinator<sup>4</sup> to broadcast an announcement to all senders. Another message is sent to the clock actor instructing it to notify the coordinator every time the required interval has passed.

## 4. CROWD-SOURCED SERVICE MIDDLEWARE DESIGN

Our design of a crowd-sourced service (CSS) middleware builds on the mechanism for multi-origin communication described in the previous section. As illustrated in Figure 7, the sensing crowd becomes the constituency whose input is solicited. The service continually aggregates the feeds arriving from the crowd to create up-to-date custom views for various types of clients. For example, if the service were for recommending restaurants, one interface could be for prospective diners, another for the restaurant managers making real-time

staffing plans, yet another could be for a vehicular routing system interested in improving downtown traffic flow at lunch time.

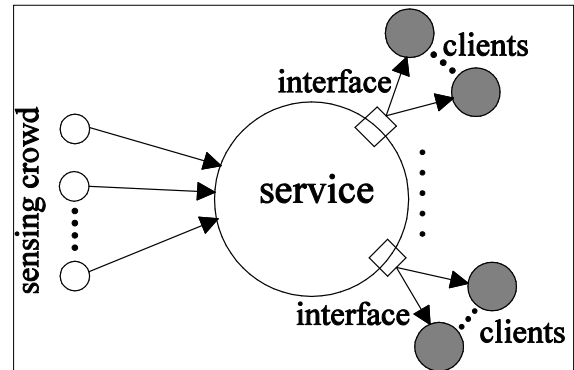


Figure 7. Crowd-Sourced Service

### 4.1 New Service Setup

Setting up of a new service can be requested by specifying the service. This could be done by either instantiating objects from a given class of services with parameter values, or by providing actual code. On receiving the request, the service platform uses the `continualCommSetup()` primitive to first create a custom service coordinator and then invite members of the identified constituency (i.e., the crowd) to begin sending their feeds to the coordinator. The decision to have the service platform (and not the service coordinator) invite the constituency helps support dynamically evolving crowds of relevance to a service, who could be identified based on their geographical location or other locality characteristics. The service coordinator periodically reports the messages received from the crowd to the service platform in aggregate form, to be then delivered to the service's clients through their custom interfaces.

#### 4.1.1 Contribution Requests

Each member of the crowd is represented by an actor executing on some device. The usual way of specifying the behavior of an actor is by defining specific methods which can be invoked on the actor as a result of incoming asynchronous messages. Each contributor actor has a method to receive parameterized requests from coordinators, and the capability to construct the requested types of messages. Given that there are a relatively small number of sensors on the types of mobile devices of interest, the parameters could simply be specifying which sensors to be tapping into, with what frequency, and what periods to be averaging the feeds over, etc. However, coordinators for some services may be more interested in hearing about higher-level events -- such as a restaurant client sitting down at the table, finishing eating, paying the bill -- which would require some amount of local processing in order to generate rather than simply forwarding raw sensor feeds. This could be supported in various ways: by migrating a custom-

<sup>4</sup>  $a \leftarrow m(p)$  means message  $m$  with parameters  $p$  is sent asynchronously to actor  $a$ .

designed actor with the required behavior to the sender, by sending the code as a parameter to create an actor locally on the contributor device, or simply by frequently updating a

contributor-side application to include the functionality needed by every new type of request.

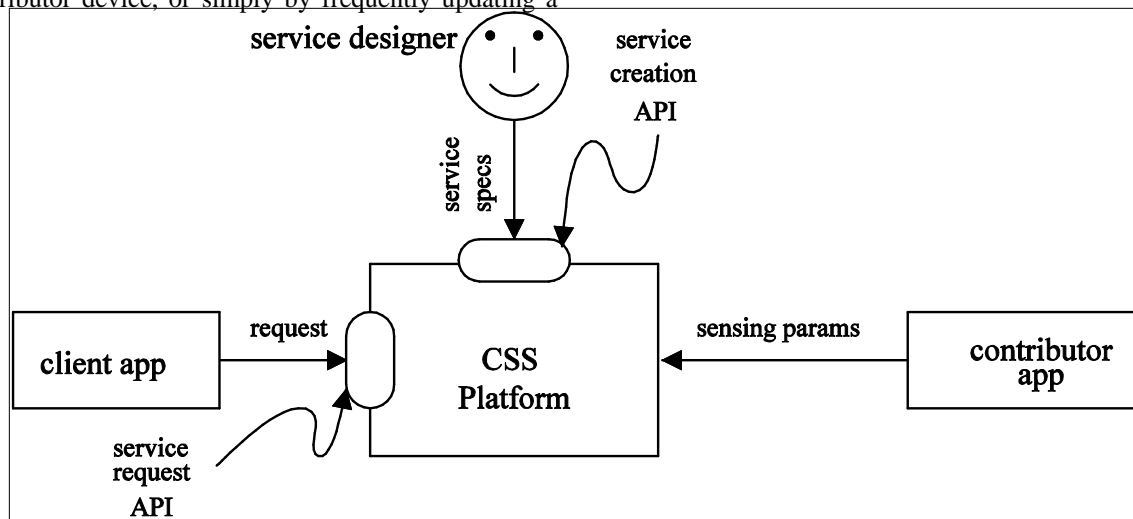


Figure 8. Service Platform and APIs

#### 4.1.2 CSS Platform APIs

Figure 8 illustrates the Crowd-Sourced Service (CSS) platform and its two main APIs. The first -- the Service Creation API -- is what a service designer uses to request the launching of a new crowd-sourced service. Service specifications are passed as parameters to specify the

constituency to be invited to participate in the communication, and the aggregation method to be used to aggregate the incoming feeds. The second -- the Service Request API -- is used by clients interested in using an existing service; each service may have multiple client interfaces delivering specific *views* of the service.

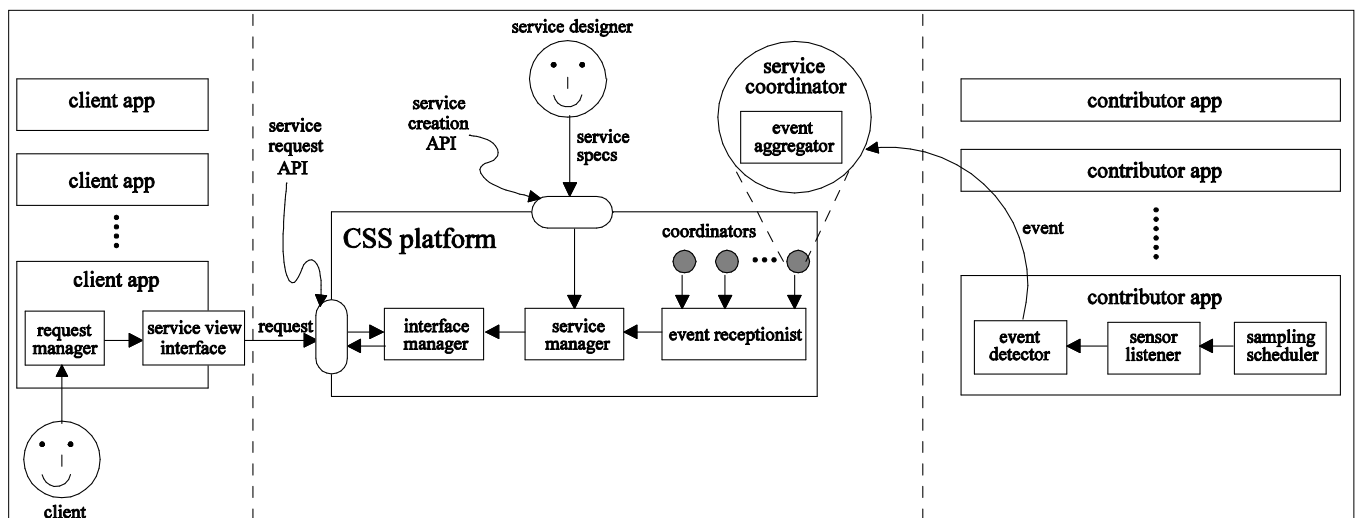


Figure 9. System Architecture

## 4.2 Distributed Runtime System

Figure 9 illustrates how the distributed run-time system for the middleware is organized with parts executing on the service platform, on devices of members of the constituency, as well as client devices. We discuss these three parts separately in the rest of this section.

### 4.2.1 Service Platform Side

The service designer uses the service creation API to create and launch a new crowd-sourced service. A set of parameters stating service specifications is passed through the API. These specifications identify the contributors to be invited to participate in the service, the aggregation method to be used, as well as a description of the feeds solicited

from the contributors in terms of specific events of interest, such as arrival at a restaurant, being seated at the table, etc.

To launch a new service, the service manager (see server in Figure 9) creates a new service coordinator to coordinate the communication between the contributors and the CSS platform, which is capable of coordinating the communication between the contributors and the CSS platform. Next, it sends invitations to the contributors to send their events -- when one is detected -- to the coordinator. It also sends them parameters advising on how to detect events, construct their messages, and how often to send them (once or periodically, how frequently, etc.).

Contributor events received by a service coordinator are handled by its event aggregator, which in turn reports the events in aggregate form to the CSS platform's event receptionist. The aggregated events are then passed on to the service manager, which processes them to update the service's state, which is forwarded to the service interface manager to deliver appropriate views requested by clients through custom interfaces.

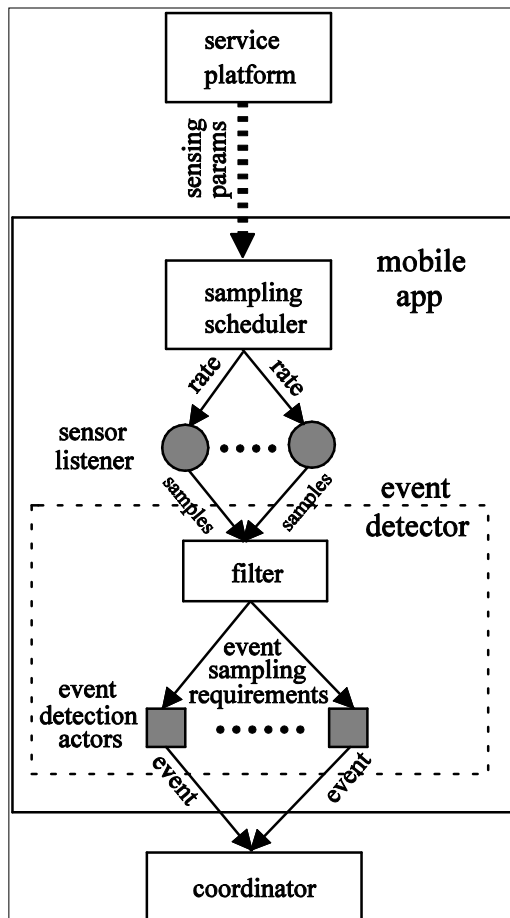


Figure 10. Contributor Side

#### 4.2.2 Contributor Side

To launch a service, the platform's service manager sends invitations to contributors to participate in the service.

It also sends them parameters advising on how to detect events and construct their messages (i.e., sensing parameters). Event detection is carried out by dedicated *event detection actors*, who generate event feeds using relevant sensor feeds, which are then sent to the service coordinator.

As shown in Figure 10, an optimizing *sampling scheduler* schedules the sampling of each sensor based on the sensing requirements received from the service coordinator for each service being served at the time. The scheduler attempts to optimize the sampling rate of each sensor exploiting opportunities for different services to share sensor feeds when possible. This could also be helped by setting granularity restrictions on the sampling rates for improving performance and conserving power.

The sensor listener is responsible for sampling sensor data according to the sampling rate received from the sampling scheduler. However, because sensor feeds are for all services, there is a *filter* to extract the required sub-feeds to be sent to the event detection actors. Each event detection actor uses all the sensor feeds it requires in order to detect events and generate its event feed to the service coordinator.

#### 4.2.3 Client Side

A service can have various types of clients subscribed to different views of the service's state, each provided by a custom interface. When a client requests subscription to a particular type of view, the request manager inside the client app constructs a custom view subscription request. This request is passed on to the service view interface, which is transmitted through the service request API of the CSS platform (see Figure 9). The platform adds the client to a list of subscribers to that view of the service, and begins sending it all updates.

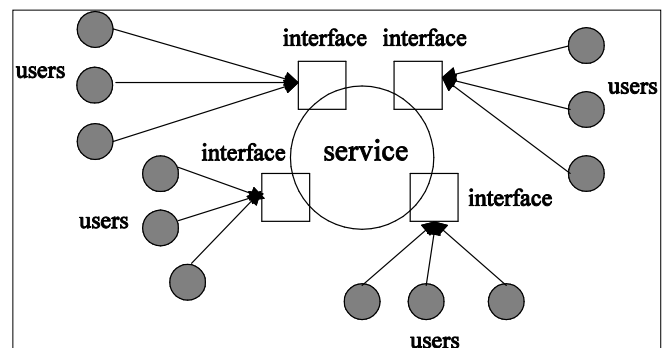


Figure 11. Design of Crowd-Sourced Services

## 5. USAGE EXAMPLES

Figure 11 illustrates how a crowd-sourced service could be designed using the multi-origin communication primitives we have described in the previous section. The service would solicit and receive one-off or continual multi-

origin communications from a target constituency. These communications would enable the service to track the state of an activity, and various types of users would be able to request relevant views of the state through custom interfaces.

This section presents two case studies to illustrate the use of the two multi-origin communication primitives we have designed. Both examples use the continual multi-origin communication primitive; one-off versions of the examples can be easily adapted from these solutions using the one-off communication primitive.

### 5.1 A Restaurant Recommendation Service

Figure 13 shows the type of restaurant recommendation service introduced earlier in the paper), where devices of people visiting restaurants in a neighborhood automatically send real-time updates about the service they are receiving to a service provider, who then aggregates this information for people searching for restaurants. We assume that information required for generating these real-time updates can be gathered automatically by a personal device (such as a smartphone) by tapping into various sensors to determine when some arrives at a restaurant, when they are waiting to be seated, when they sit down, when they are served, when they finish eating, and when they leave. The information could be coarser or finer grained depending on the device, usage habits, quality of behavior detecting software, etc. An aggregation of these updates could then be aggregated by the service provider to rank restaurants according to criteria such as the amount of wait time before being seated, the length of time taken dining (shorter or longer, as preferred), the total amount of time that the user could expect to travel to the restaurant, dine, and be back. The ranking could also include information about the number of people being sent to various restaurants by the service itself.

```
void start() {
    * choose restaurants to track; assign them IDs;
    place them in restIDList with coordinates *

    for each restID in restIDList {
        * collect names of devices in or near restaurant ID *
        continualCommSetup(restCoordClass,
            deviceNameList, sigChange, null, restID);
    }
}

rankedRestList getView(location, rankParams) {
    return rank(filter(restIDList, location),
        rankParams);
}

void update(stateUpdate, restID) {
    * update global state with restID's new state *
}
```

Figure 12. Methods Defining Behavior of Restaurant Service Actor

This service can be launched by creating and launching of a service actor, which in turn makes a number of calls to set up continual multi-origin messages, one for each restaurant, each geographical area, etc., depending on the degree of distribution required or desired. The `start` method in Figure 12 shows how this could be done if a separate coordination were needed for each restaurant. The restaurants of interest are chosen, assigned unique IDs, and placed in a `restIDList`. Then for each `restID`, mobile devices in and near the restaurant are identified, say by tracking automatic check-ins. Finally, a call is made to set up a multi-origin communication primitive for each restaurant, with the nearby devices identified as the constituency.

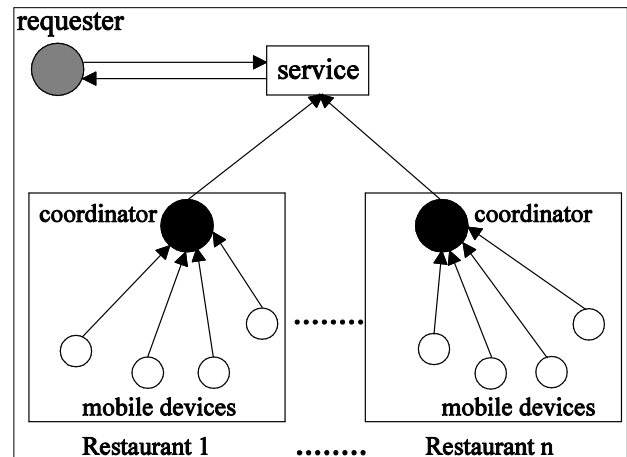


Figure 13. Restaurant Recommendation Service

Additional parameters specify the condition indicating significant change in the restaurant state warranting an update to the server, and null to indicate that there is no set interval at which updates must be made. Each of these calls creates a local restaurant coordinator which invites event updates from current diners' devices. The devices in turn have applications installed to tap into sensor feeds to recognize significant events, such as arriving at the restaurant, being seated at a table. If there are a number of similar services that the device's owner is interested in, then each would interpret the sensor feeds for the purposes of that service. As an event gets recognized by a device, it sends a message to its restaurant coordinator, invoking the coordinator's `sendMessage` method (Figure 14). `sendMessage` records the event in `eventList` and checks to see whether the event represents a significant change in the restaurant's state, and if so, sends an update message to the restaurant service -- known to the coordinator by its actor name `serviceName` -- to report the change. Invocation of `update` in the service updates the global state with the new information. In a real system, it would also make sense for both the restaurant coordinators as well as the global service to use aging functions to lower the relevance of obsolete information.

A user searching for restaurants would call the `getView` method on the server with `location` and `rankParams` as parameters, where `location` specifies the user's geographical coordinates, and `rankParams` specify the metrics by which to rank the restaurant (such as by the wait time). The server filters the restaurant list for relevance

```
void sendMessage(deviceName,event,restID)
{
  * record received event in eventList *
  if (sigChange(eventList))
    serviceName <- update(aggr(eventList),
                          restID);
}
```

according to the user's location, and then creates a ranking using `rankParams` to be returned to the client.

Figure 14. Methods Defining Behavior of a Regional Coordinator Actor

## 5.2 Twitter-like Messaging Service

Twitter serves a number of purposes, which include transmission of personal, organizational and news updates, networking, coordination of collective action, and sharing or propagation of opinions. Increasingly, it has also served as a source of information for journalists, opinion makers, politicians, etc. to acquire a sense of public sentiment. There are a handful of specific message formatting devices (particularly hashtags) which are created and subsequently adopted by contributors to indicate relationship with existing messages and conversations, and which enable some degree of analysis of sentiment. Here we discuss how to use the mechanisms we have presented in this paper to implement a service which allows users to both contribute their opinions, and obtain aggregate information helpful in assessing contributor sentiment.

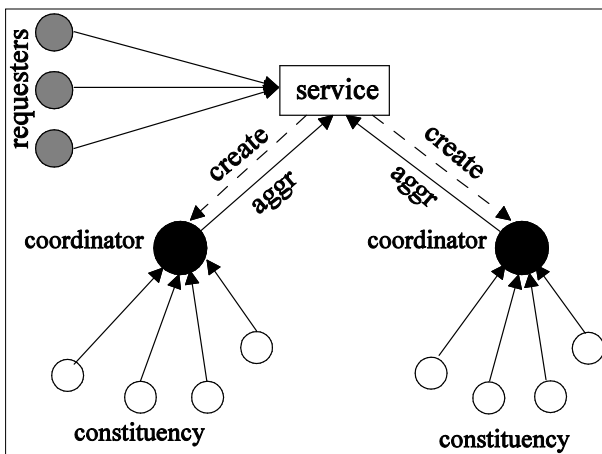


Figure 15. Twitter-like Messaging Service

Figure 15 shows how the service can be set up. The service is launched by the creation and launching of the messaging service actor, whose behavior is to receive

requests for creation of new discussions with identified constituencies. These requests are received in the form of `createDiscussion` message sends as shown in Figure 16.

When the service receives this message, it assigns a new ID -- `discussionID` -- to identify the discussion topic by, and calls the continual multi-origin communication setup primitive `continualCommSetup` with parameters specifying the discussion coordinator's behavior (`discCoordClass`), the `constit`, null for the update condition, `updateInterval` specifying the length of the intervals after which the service should receive updates from the coordinator, and finally `discussionID` to tell the coordinator its discussion topic ID. This call creates a dedicated discussion coordinator for that discussion, which in turn announces the discussion to the constituency. Once invited, members of the constituency are free to send messages to the discussion coordinator in the form of an asynchronous message invoking its `sendMessage` method (shown in Figure 17).

Figure 16. Methods Defining Behavior of Messaging

```
void createDiscussion(discussionTitle,
                    constit) {
  * assign unique ID to discussionTitle *
  continualCommSetup(discCoordClass,constit
                    ,null,updateInterval,discussionID);
}

void getView(userName,userID,discussionID,
            viewType,viewParams) {
  authenticate(userName,userID);
  userName <- view(filter(state,discussionID
                        ,userID,viewType,viewParams));
  * add username's record to the subscriber list *
}

rankedMessageList findMessages(userName,
                              discussionID,keywords) {
  * create ranked list of existing messages relevant to
  keywords *
  return * ranked message list *;
}

void update(votesUpdate,discussionID) {
  * update state with votesUpdate *
  for each entry e in subscriber list {
    e.userName <- view(filter(state,
                          e.discussionID,e.userID,
                          e.viewType,e.viewParams));
  }
}
```

Service Actor

`sendMessage` takes as parameter a list `voteList` of (message, weight) pairs, where message is either a new message drafted by the sender, or an existing message previously sent to the service (a ranked list of which can be obtained by calling the `findMessages` method of the messaging server), and weight indicates the proportional

weight that the sender intends that message to have of their vote. Each sender has exactly one vote for any discussion, which they are free to distribute between various messages under their discussion.

The service can have various types of users, subscribed to different views of the discussions' states provided by custom interfaces (see Figure 15). When a user requests subscription to a particular type of view -- `viewType` -- after authentication, it is sent the view (by having a `view` message sent to it), and is also added to a subscriber list to be sent future updates. The types of view may include a view for an analyst interested in tracking trends, or even a view for a message sender interested in staying up to date about a discussion to possibly revise their votes.

On receiving a `sendMessage` message, the discussion coordinator first updates `recentUpdates` to reflect the new messages received, and then checks to see if it is time to aggregate received messages and report back to the service. If it is time,<sup>5</sup> it aggregates the updates and reports them to the server using an `update` message, which invokes the corresponding method in the server. The server's `update` method updates the state of the discussion, and then for every entry in the list of service subscribers, sends them the view that they are subscribed to.

```
void sendMessage(userName,userID,voteList) {
    authenticate(userName,userID);
    *record received votes in recentUpdates*
    if(currentTime >= lastAggregate + interval)
    {
        serviceName <- update(aggr(recentUpdates),
                                discussionID);
        lastAggregate += interval;
    }
}
```

Figure 17. Method Defining Behavior of a Discussion Coordinator Actor

The service maintains the current state for all discussions. In practice, the service itself could be distributed into a number of actors, each handling any number of discussions.

There are some noteworthy features of this approach. First, message contributors are authenticated, and the voting is fair in that each contributor has the same one vote in any discussion, which they may divide among the multiple messages they support. Second, the constituency for each discussion is explicitly specified. This would allow this approach to be used for holding credible votes. Third, the approach naturally aggregates by allowing contributors to vote for existing messages rather than having them send a fresh message each time.

<sup>5</sup> If messages are infrequent, a clock could be asked by the service to interrupt the coordinator at the end of each interval.

## 6. CONCLUSIONS

With the growing ubiquity of sensors – be it as part of special purpose sensor networks or as sensors on people's mobile devices – it is more possible than ever to offer innovative services based on what may be happening virtually anywhere that there are either people or critical infrastructure (with connected embedded sensors). People can be directed to the restaurant with the most available tables or the hospital with the shortest wait in the Emergency. People can also more actively participate in decision making such as in a mid-21<sup>st</sup> century fine-grained democracy. However, the barriers to offering such services continue to be prohibitive for most. Not only must these services be implemented, they would inevitably compete for resources on people's devices.

We have argued in this paper that many crowd-sourced services, including prominent social media services (if we consider their role of helping evolve collective messages), require similar communication mechanisms. We focus on one such mechanism -- multi-origin communication -- which allows a number of autonomous participants to contribute messages which can then be aggregated to create group messages on behalf of all. We introduce an approach to supporting crowd-sourced services using multi-origin communication, and present our design of an Actor-based middleware for crowd-sourced services as a platform for launching such services. Finally, we present two case studies to illustrate the use of the two multi-origin communication primitives we have designed. Both examples use the continual multi-origin communication primitive; one-off versions of the examples can be easily adapted from these solutions using the one-off communication primitive.

A prototype is currently in the process of being implemented over an Actor implementation also ported to the Android operating system to support crowdsensing. We are also looking at the patterns of communication in wireless sensor networks – which appear to broadly fit the criteria of multi-origin communication – to see if network routing approaches developed for WSN would also help optimize communication in our context.

In on-going work, we are simultaneously examining the possibility of further generalizing the class of services which can be supported with this approach, as well as simplifying the programmability of the most common types of services. We want to apply our approach for fine-grained resource coordination to refining the sensor sampling scheduler, and more generally to manage the resource demands that a larger number of services may place on resource-constrained mobile devices. Finally, we plan to experimentally evaluate the scalability of the approach.

## 7. ACKNOWLEDGMENT

Support from the Natural Science and Engineering Council (Canada), the Canadian Foundation for Innovation, and the Government of Saskatchewan is gratefully acknowledged.

## 8. REFERENCES

- AbdelMoamen, A., & Jamali, N. (2014). Coordinating Crowd-Sourced Services. Proceedings of the Third IEEE International Conference on Mobile Services. Alaska, USA: IEEE.
- Agha, G. (1986). *Actors: A Model of Concurrent Computation in Distributed Systems* (1st edition). Cambridge, MA, USA: MIT Press.
- Agha, G. (2014). Actors Programming for the Mobile Cloud. Proceedings of the 13th IEEE International Symposium on Parallel and Distributed Computing (ISPD), (pp. 3-9). Porquerolles Island, France.
- Besaleva, L., & Weaver, A. (2013). CrowdHelp: A crowdsourcing application for improving disaster management. Proceedings of the 2012 IEEE Conference on Global Humanitarian Technology (pp. 185-190). San Jose, California, USA: IEEE.
- Cornelius, C., Kapadia, A., Kotz, D., Peebles, D., Shin, M., & Triandopoulos, N. (2008). Anonymsense: Privacy-aware People-centric Sensing. Proceedings of the 2008 ACM International Conference on Mobile Systems, Applications, and Services (pp. 211-224). Breckenridge, CO, USA: ACM.
- Geng, H., & Jamali, N. (2013). Supporting Many-to-many Communication. Proceedings of the 2013 ACM SIGPLAN Workshop on Programming Based on Actors, Agents, and Decentralized Control (AGERE!@SPLASH) (pp. 81- 86). Indianapolis, Indiana, USA: ACM.
- Jayaraman, P., Perera, C., Georgakopoulos, D., & Zaslavsky, A. (2013). Efficient opportunistic sensing using mobile collaborative platform MOSDEN. Proceedings of the 2013 International Conference on Collaborative Computing: Networking, Applications and Worksharing, (pp. 77-86). Austin, Texas, United States.
- Levine, U., Shinar, A., Shabtai, E., & Shmuelelevitz, Y. (2009). United States Patents, US Patent No. 8,271,057.
- Li, R.-Y., Liang, S., Lee, D.-W., & Byon, Y.-J. (2012). TrafficPulse: A Mobile GISystem for Transportation. Proceedings of the 2012 ACM SIGSPATIAL International Workshop on Mobile Geographic Information Systems (pp. 9-16). Redondo Beach, California: ACM.
- Liu, X., Lu, M., Ooi, C., Shen, Y., Wu, S., & Zhang, M. (2012, June). CDAS: A Crowdsourcing Data Analytics System. The Journal of the VLDB Endowment (PVLDB), 5, 1040-1051.
- Ra, M.-R., Liu, B., La-Porta, T., & Govindan, R. (2012). Medusa: A Programming Framework for Crowd-sensing Applications. (pp. 337-350). Low Wood Bay, Lake District, UK: ACM.
- Uga, T., Nagaosa, T., & Kawashima, D. (2012). An emergency earthquake warning system using mobile terminals with a built-in accelerometer. Proceedings of the 2012 IEEE Conference on ITS Telecommunications, (pp. 837-842). Taipei, Taiwan.
- Ye, F., Ganti, R., Dimaghani, R., Grueneberg, K., & Calo, S. (2012). MECA: Mobile Edge Capture and Analysis Middleware for Social Sensing Applications. Proceedings of the 2012 International Conference Companion on World Wide Web (pp. 699-702). Lyon, France: ACM.
- Zook, M., Graham, M., Shelton, T., & Gorman, S. (2010). Volunteered Geographic Information and Crowdsourcing Disaster Relief: A Case Study of the Haitian Earthquake. World Medical and Health Policy, 2, 7-33.

## Authors



**Ahmed Abdel Moamen** is currently a third-year PhD candidate in the Department of Computer Science at University of Saskatchewan. He has received his B.S. (2007) and M.S. (2011) degrees in Computer Science at Faculty of Computers and Information, Cairo University, Egypt. His research interests are in the general areas of Concurrency, Distributed Systems, Mobile Computing and Crowd-sourcing.



**Nadeem Jamali** is an associate professor in the Department of Computer Science at the University of Saskatchewan. He received a Ph.D. from University of Illinois at Urbana-Champaign in 2004. His research interests lie in coordination, control and communication problems in parallel and distributed systems.