An Infrastructure Supporting Considerate Sensor Service Provisioning

Jia Zhang[1], Nimish Radia[2], Zhipeng Li[1], Norman Xin[1], Yuan Ren[1], Prateek Sachdeva[1], Puja Subramanyam[1], Sky Hu[1], Song Luan[1], Lydian Lee[1], Bo Xing[2], Du Li[2], Jordan Cao[3], Ted Selker[1], Bob Iannucci[1], Martin Griss[1], Anthony Rowe[4]

[1]Carnegie Mellon University – Silicon Valley, USA

[2]Ericsson, USA

[3]SAP, USA

[4]Carnegie Mellon University, USA

jia.zhang@sv.cmu.edu, nimish.radia@ericsson.com, bo.xing@ericsson.com, du.li@ericsson,com, jordan.cao@sap.com, ted.selker@sv.cmu.edu, bob@sv.cmu.edu, martin.griss@sv.cmu.edu, anthony.rowe@cmu.edu

Abstract—The Internet of Things (IoT) aims to integrate the digital world of the Internet with our encompassing physical world. However, existing IoT lacks to provide considerate services. This paper reports our on-going work developing a sensor service federation and provisioning infrastructure. A novel approach has been presented to build social sensor networks to record and study historical interaction patterns among sensors. A case study is reported to leverage in-memory database to monitor and manage real-time sensor service provisioning. A two-way publish/subscribe pattern on top of message bus is established to assure scalability of sensor service communication. Workflow provenance is carried by a dynamic virtual device concept introduced.

I. INTRODUCTION

Sensor-equipped connected devices (Internet of Things or IoT) have been helping people sense the physical world [1] and yielding unprecedented services [2]. However, the considerate attribute of sensor service provisioning has not been studied significantly. In the human world, we think of people as considerate when for example they introduce themselves and their intentions without interrupting an ongoing conversation [3]. The term *considerate* labeled to a software system usually attempts to reduce user disruptions with communication support [4]. To facilitate human interactions, the system might coordinate, prioritize or time shift communications. It might also add other annotative statements as well to introduce and orient the information receiver.

This paper reports our efforts of developing a service oriented computing-empowered technology layer to provide "considerate" services over Internet of Things. Our technical contributions can be summarized in five-fold. (1) We have developed a novel approach to build social sensor networks to record and study historical interaction patterns among sensors. (2) We have designed and developed a sensor application server, a service-oriented middleware that supports considerate sensor service discovery and provisioning. (3) We have leveraged in-memory database to monitor and manage real-time sensor service provisioning. (4) We have proposed a two-way publish/subscribe pattern to support scalability of sensor service communication. (5) We have developed a dynamic virtual device concept to carry workflow provenance management and analysis.

The remainder of the paper is organized as follows. In Section II, we describe a motivating example that will be used to discuss our work throughout the paper. In Sections III, IV, V, VI, we present our sensor service modeling, discovery, scalable collaboration, and analysis, respectively. In Section VII, we present a considerate service framework prototype. In Section VIII, we discuss related work. In Section VIV, we draw conclusions.

II. MOTIVATING EXAMPLE

The motivating example of our project is a physical therapy scenario illustrated in Fig. 1. A patient John visits a physical therapist Bill at his office. Bill prescribes a therapy plan (workflow), which advises John to move his legs up, to the horizontal level, for 20 times within 8 minutes. The same exercise is advised to be repeated twice a day and last for 3 weeks before a follow-up visit is scheduled. Considering John's medical records, every time when he does the exercise, the process has to be monitored (i.e., heart rate, blood pressure, and body temperature).

In this scenario, sensor data is used to measure John's medical condition at exercise. John may bind a mobile phone which carries a motion sensor. When he moves his leg, the motion sensor will sense the activity and reveal readings that can be used to study the frequency and strength of the movement. Based on John's medical profile, after 20 leg movements in a certain time period (i.e., 8 minutes), his body temperature and heart rate should move up to a certain degree. His blood pressure may also move up. At the same time, the room temperature and humidity in
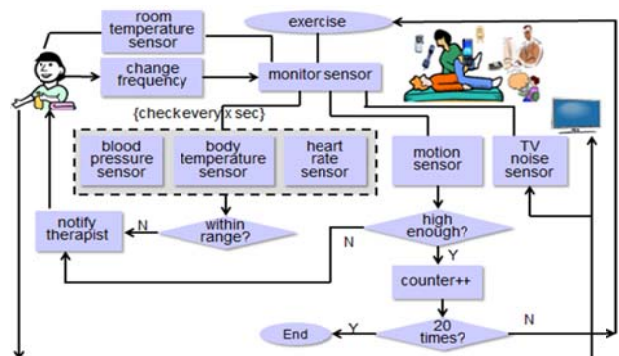

Fig. 1 Motivating example.

which John did the exercise are also factors that may have an impact on the measurement.

Based on sensor data retrieval, Bill may dynamically change the query to sensor data. For example, if John's heart rate goes up too fast, Bill may decide to conduct a closer monitoring, e.g., every 15 seconds instead of every 2 minutes. Such a change of request may lead to different sensor data gathering and analysis activities. Furthermore, such a process may require human helper on site to communicate with John and maintain contact with Bill. Thus, John may do the exercise either at Bill's office, at his home where his family may talk to him through a Skype meeting, or at a fitness center where a trainer could help him.

As shown in Fig. 1, a TV is part of the picture. Assume that John watches TV while he does the exercise. If John does not raise his leg high enough in the exercise, Bill will be notified and he will guide John. If the TV is too loud, its volume will be automatically turned down to facilitate effective communication between Bill and John.

This example shows that, to provide a considerate service, multiple sensors may need to be identified in real time, and their data should be dynamically federated to provide data analytics. Now that sensors can communicate with each other in sensor networks, how to leverage the latest advancement and establish an infrastructure over the sensor networks to enable and facilitate considerate service provisioning is the motivation of this project.

III. SENSOR SERVICE MODELING

In our previous work, we have introduced the concepts of virtual sensor and virtual device [5] to model sensor data federation. In this project, we extend our concepts to model workflow-oriented sensor services. Our study is based on our mapping between sensors to services. Since each sensor measures surrounding environment and provides readings, it can be viewed as a service provider. In other words, a sensor can be considered as a service.

A *virtual sensor* thus becomes a user-defined sensor service that reveals some information of a predefined environment. For example, a virtual sensor may be defined as a service to expose the average temperature of the room where John does exercises to help monitor his body condition. A virtual sensor may involve one or more physical sensors. A *virtual device* is a user-defined abstract container that comprises of one or more (virtual) sensors. It defines users' interested areas and maybe even a dynamic concept. For example, John may carry a mobile phone and a medical watch, each containing a collection of sensors. In this sense, John can be viewed as a virtual device.

In the service level of our study, we only care about virtual sensors that possess business logic (established by context). In other words, each virtual sensor provides a (business) sensor service to the outer world thus becomes an atomic unit of our study. A typical multi-step procedure is called a *workflow* that may invoke multiple virtual sensor services. Fig. 1 is an example of such a workflow involving a collection of sensor services.

As shown in the motivating example, every time when John conducts the therapy exercise (*workflow run*), the surrounding sensors will monitor and detect his body condition. Such data will help therapist Bill better understand John's reaction to the therapy and better personalize the treatment for John. Therefore, we associate each workflow with a virtual device in order to record workflow provenance (history) for later data analysis. The dependency between a workflow and sensor services is defined as follows.

Definition 1. A workflow W comprises a global schema G and a virtual device comprising a finite set of actions for each sensor service s_i of W. An action of service s_i is an expression *Update :- Condition* where *Condition* is a query over local schema L_{si}, and Update is a non-empty sequence of positive and negative relational literals over L_{si} such that each variable occurring in a negative literal also occurs in Condition.

Based on the association between a workflow and a virtual sensor, we model the relationship between a workflow and related sensor services. Leveraging social network theory and techniques, we aim to analyze usage history and behaviors of sensor services under a workflow context.

We model all sensors (Internet of Things) as social entities. The rationale is that sensors may be intelligently used together by people to better sense the world. We start from modeling the "use" relationship between a workflow and a sensor service. The workflow in the motivating example Fig. 1 monitors a set of sensors: a heart rate sensor, a blood pressure sensor, a body temperature sensor, and two room temperature sensors. Therefore, a social tie can be identified between them. A workflow-sensor network is thus established based on their inclusive relationships. Fig. 2(a) illustrates a segment of the network: an edge exists between a workflow and a sensor if the sensor's readings are leveraged in the workflow.

Such a network relationship can be formalized into a matrix R that describes the involvement relationships between workflows and sensors:

$$Q = [q_{ij}], 0 \le i \le m, 0 \le j \le n \text{, where:}$$

$q_{ij} = 1$ if workflow i retrieves sensor j.

As discussed earlier, all workflow runs are maintained in provenance. Therefore, such provenance data can be analyzed to extract the above relationships. Fig. 2(b) illustrates a part of a workflow-sensor network.

Based on the above relation Q, relation S can be retrieved:

$$S = Q^T \bullet Q = [s_{ij}], 0 \le i, j \le n \text{, where:}$$

$s_{ij} =$ number of workflows where both sensors i and j are retrieved; $s_{ii} =$ number of workflows where sensor i is retrieved.
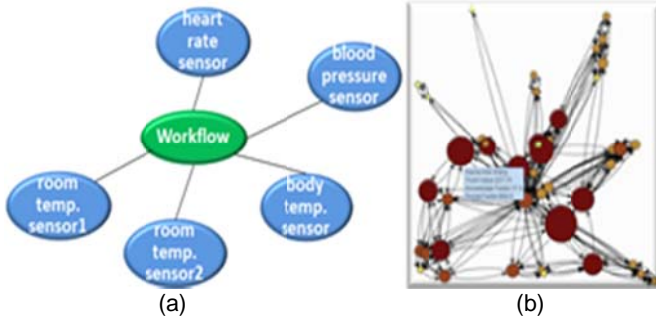
Fig. 2 A "social" relation between Internet of Things

Relation *S* represents a "social network" among sensors. Such a network of sensors forms a Social Internet of Things (SIoT). The semantic meanings of its comprising connections are: if two sensors are retrieved by the same workflow, there is a social tie between them. Note that rich context information may be carried by the edges in the network as labels. For example, an edge between two sensors may be labeled with the corresponding workflow run (i.e., query) that retrieves sensor readings among them. By digging out the details of a specific workflow run, we could understand when and under which circumstances the two sensors' readings were mashed up. The collaboration relationship maps to the association rules in social network analysis.

In addition, the social ties can be further differentiated into *tight* and *loose* collaboration relationships. Recall that our concept of *virtual sensor* implies that some formula is specified to federate readings from multiple homogeneous sensors, e.g., calculating an average room temperature from two sensors in a room. We call this relationship a tight collaboration tie, in the sense that their sensor readings are mashable, and the original sensor readings are not shown in the final values.

In contrast, the relationships among multiple sensors in a *virtual device* are called loose collaboration relationship. Sensors in a virtual device may not be the same type of sensors. Even they belong to the same type of sensor (e.g., the body temperature sensor and room temperature sensors in the motivating example), their readings are simply aggregated to show that there may be some correlation between them. For example, room temperature may have an impact on body temperature. Such loose collaboration relationship may be useful for users to conduct further data analysis.

After establishing the workflow-sensor networks, we calculated various metrics over them to comprehend the interaction patterns between sensors and workflows, as well as patterns among sensor usages. We adopted the tradition of calculation of *centrality and prestige* in social network analysis, including degree centrality (popularity), betweenness centrality, and clique. Through degree centrality analysis, we can identify highly used sensors based on the popularity of corresponding nodes. Through betweenness centrality analysis, we can identify hinge sensors. Through *clique*, we can identify collaborative sensors, i.e., association rules among services.

If a sensor is used in multiple workflows for multiple times, then popularity can be considered as a feature of annotation to be used for future purpose. For example, if a sensor is used heavily by many queries, then its readings may need to be pre-fetched, preprocessed, or redirected, etc. Using k-path betweenness algorithm, we can identify the key sensors, in the sense that they collaborate with other sensors in user queries (i.e., workflow runs).

Through these analyses, we can answer user queries regarding sensor query (usage) behaviors. For example, how are different sensors used together in workflows? And what types of workflows in which a sensor is usually used? For another example on sensor-sensor relationship, are there many sensors collaborate with each other in workflows, and how? And what are the key sensors in these collaborations? Such information will help our Sensor Service Discovery (SSD) service to help select proper sensors to provide considerate service provisioning, which will be discussed in the next section.

IV. SENSOR SERVICE DISCOVERY

Since our project aims to study considerate service provisioning, one core requirement is to understand personalization. For a specific service oriented to a specific user, a considerate service implies that we need to discover related sensors that can contribute for us to understand the context of the user, as well as the specific service query.

Based on a workflow instance assigned by therapist Bill to patient John, we need to dynamically construct its associated virtual device to identify physical sensors whose features satisfy the considerate requirements. In the motivating workflow, motion sensor, blood pressure sensor, body temperature sensor, heart rate sensor, and room temperature sensors at the user's location should be discovered. Along with this, users can also have various preferences, e.g., nearby and popular sensors will have higher priority.

Synchronization among sensors may also need to be considered. Consider the motivating example, the temperature sensor in which John does his exercise may be triggered to wake up to send readings if the motion sensor in the mobile phone bound on the patient's leg has sensed that the John has started an exercise. After the exercise is done, the temperature sensor does not have to send in readings. In this way, the temperature sensor does not have to send readings unless necessary – energy saving as another considerate service provisioning example.

Therefore, dynamic sensor service discovery is an integral part of our study. We have designed a service discovery engine that takes into account user preferences before identifying the sensors to be involved in a specific workflow run. Each sensor has a sensor type and carries values for a set of QoS attribute dimensions. Users may

configure their own preferences by selecting a subset of these QoS dimensions. Without losing generality, we consider the following five QoS dimensions: location, availability/reliability, popularity, accuracy, and usage.

Location: We care more about relative location instead of absolute GPS location. For example, if John does exercise in Building 23, room 120 - the sensors whose locations relative to the room within a threshold of distance are identified.

Availability/Reliability: Historical sensor data readings are examined to evaluate the availability and reliability of a sensor. Example measurements are: mean time between failures of the sensor, whether the sensor bandwidth exists some problem; or whether there is a need for real-time data or achieved data is sufficient.

Popularity: This dimension refers to the number of users who have used this sensor in the past, in other words, their workflows (virtual devices) invoke the sensor service.

Accuracy: The value of this dimension can be measured based on: whether the sensor readings are semantically reasonable; or whether the sensor readings are consistent to those from sensors close by.

Usage: This dimension aims to consider power consumption. For example, a patient asked to do the exercise 20 times may use the sensor data more often than a patient asked to do the exercise 10 times. This implies that a motion sensor used in the former case would consume more power than that used in the latter case. Therefore, if a sensor is required for a longer time, we aim to recommend a low power consumption sensor.

Here we model a user preference space P for user u_i as $p_{u_i} = \{p_1, p_2, ..., p_k\}$, where $p_i \in Q$ and p_i denotes a dimension affecting u_i's decision. $P \subseteq Q$.

Definition 2. Sensor Dominance: A sensor service s_i is said to dominate another sensor service s_j on Q iff $\forall q_i \in Q$, $s_i, q_i \geq s_j, q_i$ and $\exists q_t \in Q, s_i, q_t \geq s_j, q_t$.

Definition 3. Sensor Skyline: A sensor service s_i is a skyline sensor iff there does not exist a sensor service $s_j \neq s_i$ dominating s_j.

Definition 4. User Preference-Aware Sensor Dominance (UPD): Given a user preference space P, a sensor service s_i is said to dominate another sensor service s_j on P iff $\forall p_j \in P, s_i, p_j \geq s_j, p_i$ and $\forall p_t \in P, s_i, p_t > s_j, q_t$.

Definition 5. User Preference-aware sensor Skyline (UPS): A sensor service s_i is a skyline sensor on P iff there does not exist a sensor service $s_j \neq s_i$ dominating s_i on P.

We thus decompose the process of sensor service discovery into the following two phases:

Local Optimization: Given a workflow $\mathcal{W}$, a user preference space P, and a candidate sensor service class S_i for a task $t_i \in \mathcal{W}$, compute the UPD(S_i,P)|t_i.

Global Optimization: Given UPD(S_i,P) for each task $t_i \in \mathcal{W}$, compute the top-k Virtual Devices for workflow $\mathcal{W}$.

As shown below, the local optimization phase aims to select the top k sensor services based on their QoS attributes. Each QoS dimension is associated with an index. We apply a multi-index based algorithm to compute the UPD of individual QoS dimensions. For a workflow w_i, the UPD will be calculated to partition sensors based on user preferences, such that sensors are sorted in a descending order of maximum value in that dimension.

Availability	Accuracy	Usage
S3(0.90,0.45,0.55)	S4(0.44,0.98,0.54)	S1(0.34,0.54,0.88)
S2(0.83,0.55,0.67)	S5(0.67,0.76,0.65)	S6(0.67,0.59,0.78)

The sensors in the UPD are then ranked based on the scores and the top-k sensors are identified. Thus we get the top-k available sensors for the workflow locally. The score of a sensor is calculated as:

$$Score(si) = \sum_{j=1}^{a} \frac{qjmax - qj}{qjmax - qjmin}$$

where qj_{max} and qj_{min} are the maximum and minimum values of the candidate sensor class on the dimension j and q_j is the value of the sensor i for dimension j.

For global optimization, we leverage "social" relationships among sensors (as described in Section III above) to select composite services. For each service cluster in the virtual device, we obtain top-k candidate services from the local optimization procedure. We then get all combinations of composite services. A social tie of each combination is calculated. The top one will be recommended to the user. The pseudo code is summarized as below.

Global Optimization Algorithm (top-k):
1. Virtual_Device ← ∅;
2. Heap = Root Node comprising services with lowest score;
3. initialize(ParentTable);
4. **while** Heap ≠ ∅ **do**
5.   x = top(**Heap**);
6.   **if** (∄y ∈ Virtual_Device, dominated(x,y)) **then**
7.   insert(x, Virtual_Device)
8.   **end if**
9.   Child_Nodes = Expand (x);
10. **for** $\forall n \in Child\_Nodes$ **do**
11.   Number_of_Parents (n)**--**;
12.   **if** Number_of_Parents (n) == 0 **then**
13.   insert(n,Heap);
14.   **end if**
15. **end for**
16. **end while**

As illustrated, top-k services are sorted based on their scores. All combinations of composite services are enumerated using a lattice, based on a min-Heap strategy that extracts the composite services with minimum scores and compares the social coalition against other virtual devices. A composite service is added to the final list of virtual devices only if it is not dominated by any other composite services in the list.

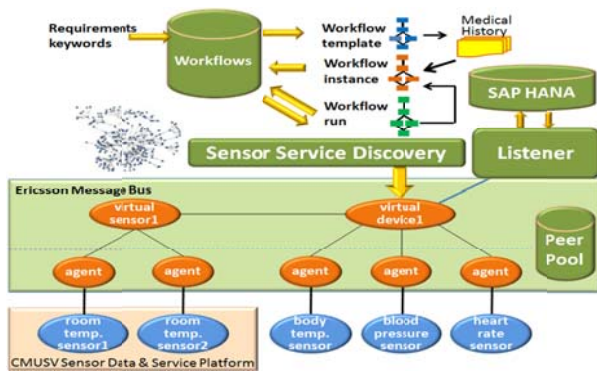The performance of the Global Optimization Algorithm

Fig. 3 Sensor application server.

is decided by two major factors:

1. Heap operations: It is affected by the heap size, which grows exponentially with the number of sensors. This behavior is observed because the size of the heap is bounded by the number of nodes on the middle level of the lattice. After every extraction the heap needs to be re-organized. With increasing number of members in the heap this operation becomes more costly.

2. Social tie comparisons: It grows exponentially with the number of sensors involved.

The high computational complexity of Global Optimization Algorithm makes it impractical to compute the virtual devices with a large number of sensors. A bottom-Up algorithm as shown in the above figure uses the same algorithm but iteratively thus offering significantly better scalability and a nearly optimal time complexity. Instead of creating combinations of all sensor types at one time, this algorithm iteratively combines these services thus providing a faster and more scalable solution.

V. Sensor Application Server

Based on our sensor service modeling and discovery techniques, we have developed a sensor-oriented infrastructure. We name it Sensor Application Server, analogous to the software-oriented application servers such as WebLogic and WebSphere. Particularly, our main goals are two-fold: (1) service collaboration provenance management, and (2) sensor service discovery and composition.

As explained in the earlier section, we have introduced a virtual device concept to carry contextual information (sensor data provenance) of a specific workflow execution (workflow run). Following this idea, we have differentiated three concepts: workflow template, workflow instance, and workflow run. A workflow template represents an abstract definition of stepwise activities; a workflow instance represents a personalized workflow template for a particular user; a workflow run represents an actual execution of a workflow instance.

As shown in Fig. 3, a workflow template is retrieved from a medical workflow repository representing a type of physical therapy procedure. Then Bill checks John's

medical record, in order to configure the workflow specific for John. Such a workflow instance is stored back to the database, as a customized workflow for John for a specific time period. Each time when John does the exercise, the corresponding workflow instance is triggered for execution as a workflow run. Provenance of each workflow run is also stored into database for Bill to analyze in later diagnosis.

As shown in Fig. 3, each workflow run will query the sensor service discovery (SSD) service to dynamically form a virtual device to support the execution of the workflow. Using the motivating example shown in Fig. 1, SSD locates two room temperatures in the room where John does his exercise. Therefore, Fig. 3 illustrates a formed virtual device comprising a virtual sensor representing the average room temperature based on two temperature sensors, a body temperature sensor, a heart rate sensor, and a blood pressure sensor. As shown in Fig. 3, a workflow run dynamically formed a virtual device comprising multiple sensors. In other words, these sensors collaborate with each other in the context of the workflow.

A. Message Bus

At Ericsson lab, we have developed a prototype of message bus, to enable message-based communication among sensors. As shown in Fig. 4, sensors are not isolated. Instead, they can "talk" to each other through message passing with the support of a centralized Messaging system, analogous to a post office system in the human life. Each message notifies an event happens at the corresponding sensor. As the following example shows, a sensor may declare that a corner read light is turned on by publishing the following message:

publish("turn on light", null, "color=red&location=corner");

A publish/subscribe pattern is adopted at the Message Bus. Sensors can publish messages onto the Message Hub; and sensors can subscribe to interested messages. The following example shows that a sensor is interested in receiving notifications when some corner red light is turned on.

subscribe("http://eus2.fuatara.com/callback",
  "turn on light",
  ["color=red", "location=corner", "eco-friendly=true"]);

As shown in Fig. 4, sensor peers thus form a loose coupling relationship. The Messaging Catalog acts as a filter, so that only allowed topics and metadata can be published and subscribed to devices. In addition, four forms of messaging are supported: unicast, multicast, broadcast, and publish-subscribe.

B. Message Bus-Empowered Service Collaboration

As shown in Fig. 3, all physical sensors are mapped into an agent in the Message Bus as its representative. It is analogous to an Entity Bean in a (software) application server to its persistent object. It is an agent's responsibility to monitor the health of the corresponding physical sensor.
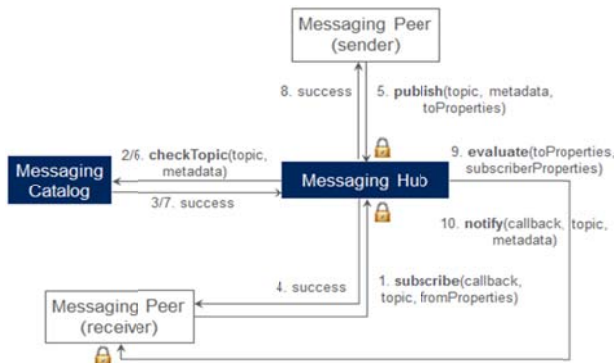
Fig. 4 Message Bus.

In other words, sensors interact with each other through their agents in the Message Bus.

Inspired by the EJB concepts that differentiate between entity beans and session beans, we differentiate between sensor-oriented agents and virtual entities, i.e., virtual sensors and virtual devices. As shown in Fig. 3, in the Message Bus, virtual entities are modeled as agents as well. However, they carry business logics.

Message Bus requires that the agents be deployed ahead of time. Therefore, for scalability, we have preconfigured and deployed a pool of anonymous peers (agents) when the Message Bus is deployed onto the Web server. At run time, peers will be extracted from the pool and dynamically assigned to act as various roles for workflow runs. Upon finished, the peers will be released back to the pool to be used by other workflow runs.

As described earlier, sensors communicate with each other through message passing. As shown in Fig. 3, our strategy is to establish a listener over the Message Bus. All sensor messages will be scanned and stored. As shown in Fig. 3, the Message Bus is treated as a standalone generic service. After a processing a message, the Messaging Hub forwards the message to our message listener in the format of a JSON object. We have built a message bus listener as a separate service. A log sending mechanism is added in the Message Bus. Every event in Message Hub will trigger a log message to be sent to the listener server.

By building a listener for message bus, the sensor application server knows what exactly happen in the system at runtime. In this way, we can use this information to determine which service/sensor need to be used in real-time conditions, and how to properly compose service to make the system considerate based on real-time needs. By representing sensors as agents with methods, we can model a certain service workflow by using methods as primitives. Sensors (corresponding sensor agents) publish their readings onto the Message Bus. Workflow runs subscribe to specific sensor readings (through service discovery) will receive the sensor readings. At the same time, workflow runs leverage HANA's analytical methods to monitoring the incoming streaming data at real time (will be discussed in the next section). If there is a change of plan (e.g., change the frequency of heart rate monitoring), the workflow run will

publish a request to the Message Bus.

We have designed a two-phase publish/subscribe pattern. When a workflow instance is established for a workflow run, it will construct a representative agent in the Message Bus. A two-way relationship between the workflow agent and the corresponding sensor agents. On one hand, the workflow run will subscribe to all identified sensor agents for their readings. On the other hand, the sensor agents will subscribe to the workflow agent to receive particular requests.

VI. INFRASTRUCTURE SUPPORTING CONSIDERATE SERVICE

We have designed a service oriented infrastructure supporting considerate service provisioning. As shown in Fig. 5, our infrastructure builds a layer on our Sensor Data & Service Platform (SDSP) and Message Bus. It comprises four horizontal layers and four vertical layers. For all sensors registered in our SDSP, the service modeling and publishing layer analyzes their interactions and establishes a sensor social network. The service composition and collaboration layer monitors the sensor interactions through a listener service, and handles dynamic service composition. The service visualization and analytics layer helps to visualize and analyze sensor relationships. The service provisioning layer interfaces with external sensors and deliver considerate services.

Vertically, the service discovery layer is in charge of dynamic sensor service discovery and integration. The considerate layer controls the quality of service provisioning focusing on the considerate property. The data architecture layer decides the data model and structure of service management. The governance layer monitors the health of the overall system.

One specific feature that we consider is the scalability. Our previous work adopted the traditional client-server model [5], where SDSP receives all sensor readings and forwards them to persistent storage (i.e., Amazon DynamoDB). Although straightforward, the SDSP server becomes the bottleneck. To support streaming sensor data and real-time data analytics, we decided to adopt the Message Bus in our infrastructure. As shown in Fig. 5, all
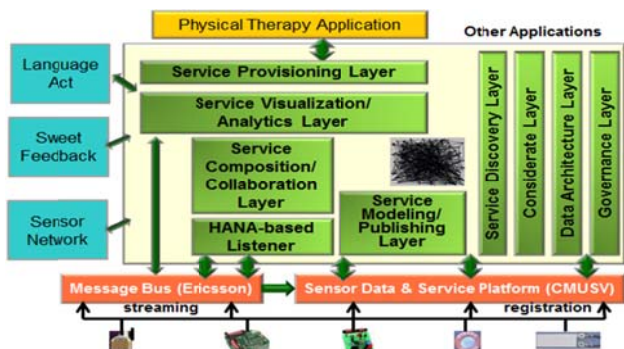

Fig. 5 Supporting considerate service provisioning.

registered sensors directly send notifications to the message bus.

To make the sensor network considerate, it is critical to seamlessly integrate human interaction. Without losing generality, we have incorporated language processing and language feedback to our social sensor network. John (or treating subject) interacts with the social sensor networks through voice, and the social sensor networks provide voice feedback according to the subject's words, the prescription for this subject, and the data collected from this subject. Such a feature becomes extremely useful when patient feels uncomfortable during a physical treatment, which indicates the necessity of the modification of the prescription.

When the exercise starts, the therapist may start a query, to ask to send to him the patient's average heart rate every 2 minutes. The service layer will then generate a workflow, caching the heart rate sensor's streaming data and send back average values every 2 minutes.

During the exercise (workflow run), the patient may feel uncomfortable and types "hard to breathe." The text information is passed to the language processing module to analyze the purpose and goal of the patient. It may be translated to language act "short of breath" to be sent to the therapist Bill, who may then change his query as to send the patient's heart rate every 15 seconds. Then the service will not cache heart rate sensor readings. Instead, all readings will be sent to the therapist directly for him to better monitor the patient's body situation.

VII. SCALABILITY STUDY

A. Sensor Service Analytics

The scalability and performance are the major concerns of our infrastructure. For analyzing sensor behaviors, all messages are stored. We have decided to adopt in-memory database to address the performance concern. Based on the collaboration between Carnegie Mellon University and SAP, we have established an HANA instance at Amazon Web Service (AWS), to study its feasibility of supporting message-based considerate service discovery.

We use SAP HANA as a scalable solution to the volumes of memory, which is an in-memory data platform that can be deployed on the cloud, in contrast to row-based storage in traditional relational database. The platform is used for providing real-time analytics. Different from other database engines, SAP HANA provides column-based storage which allows us to easily monitor data. As explained in Section V, sensor communications are represented by topic-based messages. Each topic is represented by a <key, value> pair. The sensor messages can be stored in SAP HANA and can be used for real-time analytics. Another reason why we studied SAP HANA is that it is well suited for enormous data sets from different sources. Since we have to store messages passed around, the data can be stored easily using SAP HANA. SAP HANA provides the feasibility of getting answers to complex questions at high speed from the enormous data sets.


Fig. 6 Scalability study.

We store raw sensor readings using the schema of <time stamp, device ID, sensor type, sensor reading>, and also aggregated the readings in a 1-minute time window. Since we have to store messages passed around, the data can be stored easily using SAP HANA. HANA provides the feasibility of getting answers to complex questions at a high speed from the enormous data sets.

The platform is used for providing real-time analytics. We run K-means algorithm from HANA's Predictive Analysis Library to monitor sensor health. It detects anomaly in sensor reading's interval, and flags outliers as sensor readings having an abnormal interval from the previous reading.

B. Experimental Study

We have conducted simulations to evaluate the scalability of our infrastructure. We simulated different scales of a sensor network, which comprising from 1 to 30,000 physical sensors. On each of the simulated sensor network, we simulated different numbers of concurrent users: from 1 to 50 users. As shown in Fig. 6, when the scale of a sensor network is moderate (comprising less than 1,000 physical sensors), our system scales well even when the number of concurrent users go up to 50. To better illustrate the experimental results, we keep only significant data in Fig. 6.

Note that the aforementioned experiments were conducted based on the SAP HANA database. In our earlier implementation, all sensor readings were stored in NoSQL DynamoDB database. We repeated the same set of experiments. As shown in Fig. 6, the scalability of the system relying on HANA significantly surpasses that on DynamoDB. This symptom is resulted from the specific pattern of frequent query over certain types of sensors, e.g., a user is working on a virtual sensor involving multiple temperature sensors. Thus, the specific column-based query used by HANA yields better results.

VIII. RELATED WORK

Mohamed and Al-Jaroodi [6] surveyed the requirements of service-oriented middleware for wireless sensor network. Internet-scale Resource-Intensive Sensor Network Services (IrisNet) [7] offers a set of generic functionalities and APIs to allow user to query and process distributed collections of high-bit-rate sensors. On top of IrisNet, Chen et al. [8]

proposed *X-Tree Programming*, a database-centric approach to programming over a large-scale of Internet-connected sensing devices. Microsoft SenseWeb [9] provides a Web 2.0 platform for users to upload and access sensor data streams from shared sensors on the Internet. SensorBase [10] built a centralized data storage and management platform that allows users to publish and share ("slog") sensor network data using a blog-like approach. Global Sensor Networks (GSN) [11] adopts a scalable P2P model in favor of integrating heterogeneous sensor network technologies. In this work, we explore how SOA can be leveraged to provide considerate sensor services.

Some researchers focus on building tools to support sensor data manipulation. For example, the Desthino (Distributed Embedded Things Online) project identifies a practical set of software tools to help users collect and store sensor data from heterogeneous distributed sensors [12].

Sensor Observation Service (SOS) is a standard Web service specification, aiming to standardize the way of requesting, filtering, and retrieving sensors and sensor data to enhance sensor interoperability [13]. Some researchers explore how the Semantic Web can be integrated with a Sensor Web, such as SemSOS [14] and Semantic Sensor Web [15]. Some researchers, such as Liu et al. [16], study the scalability of sensor networks. A concept of virtual sensor is introduced in GSN [11] to abstract sensor data as temporal streams of relational data, and to represent derived views or a combination of sensor data from different sources. SenseBox [17] introduced an autonomous computing unit encapsulating environment and REST APIs. In our earlier work, we have developed an SOA-based Web 2.0 platform that allows users to federate heterogeneous sensor data sources [5]. In this project, we have extended our earlier work and build an infrastructure to support considerate sensor service discovery and composition.

Our another previous work leverages social network techniques to facilitate workflow reuse, by analyzing usage history and behaviors of software components [18]. In this project, we have extended our technique to facilitate sensor federation.

IX. CONCLUSIONS

This paper has reported our efforts of exploring how to develop a service oriented computing-empowered infrastructure to provide "considerate" services. By mapping sensors to services and modeling sensor mashup as social networks, we leverage past sensor usage history to provide more considerate services.

In our future work, we plan to refine our prototyping system and provide our middleware to the community usage. We also plan to conduct performance study between using various databases as backend including in-memory database, NoSql database, and relational database.

X. ACKNOWLEDGEMENT

VIII. REFERENCES

1. Nagpal, R., H. Shrobe, and J. Bachrach, "Organizing a Global Coordinate System from Local Information on an ad hoc Sensor Network", in Proceedings of *2nd International Conference on Information Processing in Sensor Networks (IPSN)*, 2003, Palo Alto, CA, USA, pp. 333-348.
2. Yick, J., B. Mukherjee, and D. Ghosal, "Wireless Sensor Network Survey", *Computer Networks*, Aug. 22, 2008, 52(12): pp. 2292-2330.
3. Rajan, R., J. Hsiao, and T. Selker, ""Roger that!" - The Value of Adding Social Feedback in Audio-mediated Communications", in Proceedings of *Interact*, Sep., 2013, Cape Town, South Africa, pp.
4. Selker, T., "Understanding Considerate Systems - UCS (pronounced: You See Us)", in Proceedings of *The 2010 International Symposium on Collaborative Technologies and Systems*, May 17-21, 2010, Chicago, IL, pp. 1-12.
5. Zhang, J., B. Iannucci, M. Hennessy, K. Gopal, S. Xiao, S. Kumar, D. Pfeffer, B. Aljedia, Y. Ren, M. Griss, S. Rosenberg, and A. Rowe, "Sensor Data as a Service - A Federated Platform for Mobile Data-Centric Service Development and Sharing", in Proceedings of *IEEE International Conference on Services Computing (SCC)*, Jun. 26-Jul. 2, 2013, Santa Clara, CA, USA, pp. 446-453.
6. Mohamed, N. and J. Al-Jarood, "A Survey on Service-Oriented Middleware for Wireless Sensor Networks", *Service Oriented Computing and Applications*, 2011, 5(2): pp. 71-85.
7. Gibbons, P.B., B. Karp, Y. Ke, S. Nath, and S. Seshan, "IrisNet: An Architecture for a World-Wide Sensor Web", *IEEE Pervasive Computing*, Oct.-Dec., 2003, 2(4): pp. 22-33.
8. Chen, S., P.B. Gibbons, and S. Nath, "Database-Centric Programming for Wide-area Sensor Systems", in Proceedings of *1st International Conference on Distributed Computing in Sensor Systems (DCOSS)*, Jun. 30-Jul. 1, 2005, Marina del Rey, CA, USA, pp. 89-108.
9. Grosky, W.I., A. Kansal, S. Nath, J. Liu, and F. Zhao, "SenseWeb: An Infrastructure for Shared Sensing", *IEEE MultiMedia*, Oct.-Dec., 2007, 14(4): pp. 8-13.
10. Chang, K., N. Yau, M. Hansen, and D. Estrin, "SensorBase.org-A Centralized Repository to Slog Sensor Network Data", in Proceedings of *International Conference on Distributed Computing in Sensor Network (DCOSS)/Euro-American Workshop on Middleware for Sensor Networks (EAWMS)*, 2006, San Francisco, CA, USA, pp.
11. Aberer, K., M. Hauswirth, and A. Salehi, "Infrastructure for Data Processing in Large-Scale Interconnected Sensor Networks", in Proceedings of *International Conference on Mobile Data Management*, May 7-11, 2007, Mannheim, Germany, pp. 198-205.
12. Santini, S. and D. Rauch, "Minos: A Generic Tool for Sensor Data Acquisition and Storage", in Proceedings of *19th IEEE International Conference on Scientific and Statistical Database Management*, 2008, pp.
13. OGC. *Sensor Observation Service (SOS)*. Open Geospatial Consortium 12/30/2012]; Available from: http://www.opengeospatial.org/standards/sos.
14. Henson, C.A., J.K. Pschorr, A.P. Sheth, and K. Thirunarayan, "SemSOS: Semantic Sensor Observation Service", in Proceedings of *2009 International Symposium on Collaborative Technologies and Systems (CTS)*, May 18-22, 2009, Baltimore, MD, USA, pp. 44-53.
15. Sheth, A., C. Henson, and S. Sahoo, "Semantic Sensor Web", *IEEE Internet Computing*, Jul./Aug., 2008: pp. 78-83.
16. Liu, Y., Y. He, M. Li, J. Wang, K. Liu, L. Mo, W. Dong, Z. Yang, M. Xi, J. Zhao, and X.-Y. Li, "Does Wireless Sensor Network Scale? A Measurement Study on GreenOrbs", in Proceedings of *IEEE International Conference on Computer Communications (INFOCOM)*, Apr. 10-15, 2011, pp. 873-881.
17. Bröring, A., A. Remke, and D. Lasnia, "SenseBox-A Generic Sensor Platform for the Web of Things", *Mobile and Ubiquitous Systems: Computing, Networking, and Services*, 2012, 104: pp. 186-196.
18. Zhang, J., W. Tan, J. Alexander, I. Foster, and R. Madduri, "Recommend-As-You-Go: A Novel Approach Supporting Services-Oriented Scientific Workflow Reuse", in Proceedings of *IEEE International Conference on Services Computing (SCC)*, Jul. 4-9, 2011, Washington DC, USA, pp. 48-55.