Provenance-aware Pervasive Computing in Clinical Applications

Yevgeniya Kovalchuk[1], Yuhui Chen[2], Simon Miles[2], Shao Fen Liang[3], Adel Taweel[2]
[1] Department of Primary Care and Public Health Sciences, School of Medicine,
[2]Department of Informatics, School of Natural and Mathematical Sciences,
[3]NIHR GSTFT/KCL Biomedical Research Centre,
King's College London, London, UK
{yevgeniya.kovalchuk, yuhui.chen, simon.miles, fennie.liang, adel.taweel}@kcl.ac.uk

Abstract—**Pervasive computing applications bring together heterogeneous network-connected devices, services and resources to enable context-aware information integration. The increasing adoption of pervasive computing technology in the healthcare domain offers a healthcare model that delivers high quality service with fewer resources. In this paper, we briefly review the existing pervasive healthcare solutions and propose a novel provenance-aware system design that can enhance the performance of such solutions by means of including provenance capture functionality. We argue that our system architecture can improve quality of clinical data, efficiency of its collection, and its integrating ability with other data sources. To demonstrate our system and explain its provenance capacity, we use a clinical research example in which patient's condition is closely monitored in order to assess the safety and efficacy of medications and treatments prescribed to him.**

Keywords— provenance; pervasive computing; PROV model; clinical research information system

I. INTRODUCTION

Pervasive computing has become an active area of research in the last two decades since its introduction by Weiser in the early 1990s [1]. The current post-desktop computing facility exploits an increasing availability of small, embedded pervasive devices such as mobile phones, tablets and other portable personal devices to interlink the cyber world with physical reality and provide people with a more natural way to interact with information and services [13-15].

The constantly increasing healthcare cost and lack of clinical professionals have stimulated the adoption of pervasive computing technologies in the healthcare domain [2-7]. While pervasive healthcare [8] is still in its infancy, pervasive computing technologies usher a promising future in developing clinical applications, where patients' condition must be monitored continuously. More specifically, modern low cost and low power mobile devices, portable wireless sensors and advanced communication technologies open the opportunity to construct a pervasive healthcare environment that surrounds an individual patient and enables gathering rich clinical data about him (e.g., patients' behavior, physiological parameters, social dynamics, etc.) [2, 13]. The information collected by the devices can be further transited to a centralized clinical information system and become a part of the patient's electronic healthcare record (EHR).

However, one of the main challenges for such a system in healthcare is its ability to collect quality data traceable to individual devices' context and usage and ensuring its validity for use, based on which critical decisions may be taken. Recent work in provenance research and the recent adoption of the provenance standard W3C PROV [17] provide a potential solution to address this challenge. The concept of provenance originates from the fine arts where it refers to the trusted, documented history of some work of art [22]. In computer systems, provenance is concerned with tracing the history of individual pieces of data.

In this paper, we discuss how to capitalize on the existing engineering solutions and previous research efforts in pervasive computing in order to support data collection, validation, and integration processes in clinical applications. We propose generic provenance-aware pervasive system architecture that allows recording not only clinical facts about patients at the point of care (the current state of the art), but also relevant provenance information. We are particularly interested in provenance documentation of the data collection processes carried out by pervasive computing technologies, so as data sources could be verified and collected data could be easily and accurately integrated across various devices and heterogeneous environments. This would provide an advantage of ensuring accurate and comparable data, or at least a mechanism to alert of any mismatches or ambiguities across data sets (e.g. when translating age into data of birth). This is especially important in clinical applications, for example when selecting patients for participation in clinical trials.

The paper is structured as follows. First, we discuss the advantages and challenges for pervasive computing in clinical applications that have motivated our work (Section II). We then present a state of the art synopsis on pervasive computing applications in the healthcare domain (Section III). Section IV outlines a clinical use case scenario, which we use to explain and demonstrate the proposed system design. Section V provides details on our solutions along with the discussion of the concept of provenance and its relevance and application to healthcare domain. With example graphs, we show how the provenance capacity of our solution can help to address the issues of traceability and validation when collecting patients' electronic data. Section VI concludes the paper and discusses our future work.

II. Advantages and Challenges For Pervasive Computing In Clinical Applications

Over the last decade, the pervasive adaptation of ICT has dramatically changed the healthcare landscape. The wide deployment of information system technology in healthcare centers has greatly improved the availability and accessibility of patients' electronic records. The growing volume of healthcare information can help to improve the quality of healthcare service, allowing more extensive clinical research, and supporting rapid decision making in treatment of a disease and epidemic control [9].

Clinical research includes patient-oriented research studies on human subjects. Several ICT application projects have emerged recently to facilitate acquisition and integration of medical information into clinical research [10, 11, 20, 21]. These projects aim to develop infrastructures for integrating distributed, and often heterogeneous, healthcare data sources for supporting clinical research. For instance, the Electronic Health Records for Clinical Research (EHR4CR) project [11, 21] proposes an integrated platform that provides controlled and regulated access to hospital healthcare and research information systems, such as Electronic Health Records (EHRs) and Clinical Research Information Systems (CRISs). These may be distributed throughout numerous hospitals in many countries. A clinical researcher may then query the clinical research data for a multi-site study, for example to determine the number of patients who meet a set of criteria that would make them eligible subjects for a clinical trial. Such systems bring both financial (e.g. by reducing costs) and nonfinancial (e.g. by reducing selection times) benefits to those setting up and conducting clinical trials.

Clinical research findings directly rely on complete and accurate data. Clinical research data capture is an extremely important part of a clinical research project. Traditionally, the process relies heavily on paper-based documentation which is time consuming and error prone. Over the last decades, paper case report forms (CRFs) at hospitals are increasingly replaced by computerized data capture, such as remote data capture (RDE) and electronic data capture (EDC) which are indeed early forms of pervasive computing applications. The adoption of these computerized data collection approaches were not initially motivated due to the lack of mobile hardware for data gathering [12]. With the advancement of mobile hardware and communication technologies however, pervasive computing becomes increasingly viable to replace the conventional data collection mechanisms in the healthcare domain.

At the same time, such issues as data traceability and validation still remain among the challenges that discourage computerized data collection in clinical applications [12]. Although paper-based data collection is time consuming and error prone, paper data source documents can be kept for further validation and investigation. It is however difficult during computerized data collection in systems that lack of provenance capacity. Even if clinical data is stored in centralized information systems, questions regarding where, how, by whom and in what circumstances it has been collected would still have to be answered in a validation and investigation event.

III. Related Work

A study on the state of the art in the relevant fields helps us to understand existing solutions and allows us to develop integrated solutions based on past research efforts without unnecessarily reinventing the wheel. In this synopsis, we review the recent efforts in developing pervasive computing applications in healthcare and clinical research domains.

Pervasive computing in healthcare focuses on improving the quality of healthcare services, automating patient condition monitoring and diagnosis, as well as optimizing data acquisition, storage and management. Conti et al. [13] suggest that the adoption of sensor networks in patient care has been the most successful area in pervasive healthcare research so far, while Bardram [3] argues that pervasive healthcare technologies have been designed and developed to support continuous well-being, treatment, and care of people rather than focusing on technologies for acute treatment and care. The later paper summarizes the commonly recognized challenges in pervasive healthcare and presents a set of research themes in related scientific disciplines, such as monitoring and body sensor networks, pervasive assistive technologies, pervasive computing for hospitals, preventive and persuasive technologies. The author advocates that information management in the healthcare environment is becoming increasingly decentralized. Pervasive computing can provide technologies to support patient self-care, allow clinicians to reach out to home-based patients, enable continuous monitoring for diagnosis, early detection, treatment of diseases, and optimize information acquisition, storage, integration, processing, and management. This paper presents two prototypical examples to illustrate the challenges in pervasive healthcare, discusses how pervasive computing technology can be designed to meet these challenges and recommends a "clinical proof-of-concept" approach for pervasive healthcare research.

A study conducted by Orwat et al. [2] also suggests that pervasive computing is increasingly influencing healthcare and medicine. The authors conclude that pervasive computing is loosely associated with many future technologies and its definition has not been clearly defined in the current literature. Their paper reviews recent pervasive computing systems in healthcare and highlights experiences in the development of pervasive computing systems. Their study has explored published research in pervasive computing and conducted a search with scientific databases and journals for the period of 2002 to 2006. The search results have revealed a great diversity of the pervasive computing systems in healthcare field. In total, 69 articles have been identified, in which the development of a pervasive computing system is reported. The identified articles describe 67 different systems, 84% of them are in their prototype or pilot stages, 9% have passed clinical trials, and 7% have been in regular operation. The majority of the systems (48%) have been developed for analytical and diagnostic support.

Pervasive computing is inherently context-aware. Black et al. [5] have envisioned a "smart hospital" environment in which context-aware applications are extensively adopted to assist in clinical care. The paper describes a scenario where

context middleware and pervasive computing technologies are used together for building an integrated enterprise environment in which wireless devices and wearable sensors are deployed around patients and their carers to collect contextual information and monitor their locations, conditions and behavior. An imaginative use case has been used in this paper to demonstrate how context-aware applications can help improve healthcare service delivery, nurse triage and information management. The paper reports a prototype of the enterprise smart space described in the use case and identifies requirements and challenges in developing context-aware pervasive healthcare systems. Similar work is described by Moran et al. [7], who report some empirical information about the movement pattern of hospital workers and how they interact with information while moving. The paper discusses how their study results can help in developing pervasive computing technologies for healthcare.

Bardram [6] describes a context-aware pervasive computing application for medical work in hospitals and elaborates on a context-awareness infrastructure in a hospital that supports various clinical applications. The paper uses some example applications to illustrate the conceived ideas. The use case scenarios demonstrate what context-awareness can mean in a healthcare environment, and explain how to design and develop a context-aware application to deal with the challenges. Finally, the paper concludes some key design principles for constructing a context-aware infrastructure in hospitals and developing context-aware clinical applications.

Floerkemeier et al. [4] propose the development of smart medication device pack. The device is equipped with sensors to detect the consumption of pills and can update information to the central information system via mobile network. The authors explain that the smart device pack can be used for monitoring the patient's compliance in clinical trials. The automatically collected information can assist the researchers to assess outcomes and determine whether adverse reactions or inadequate therapeutic responses are a result of the drug or a result of the patient not taking the medication. The compliance data can also be used for improving the management of clinical trials.

Mihailidis et al. [8] present a pervasive computing application for supporting clinical trials. The authors describe a system that can automatically remind nurses to perform scheduled tasks for a clinical trial. The system uses a range of smart devices to accurately acquire vital signs of a patient and transfer the data to a centralized information system. The authors argue that the system can help eliminate possible measurement and data entry errors that are commonly seen in clinical trials, improving clinical trial management and significantly reducing the amount of time spending on data entry compared to the conventional practice.

Finally, remote patient monitoring is becoming increasingly plausible in healthcare. The advancement of sensor and wireless mobile device technologies has enabled the pervasive healthcare approach for remote acquirement of information about a patient's clinical parameters and activity patterns, with some smart devices being able to automatically conduct remote medical analysis for risk prevention. For example, Chowdhury et al. [16] present a mobile phone-based system, called MediAlly, that has the ability to capture medical data streams from a remote subject. The system is context-aware, which in this case means that it can collect, store and process the medical data streams based on user-specified context rules. As an example, authors say that a doctor would find it useful to know if a data stream corresponding to "30 minutes of elevated heart rate", recorded a month ago, occurred while the subject was exercising at a health club or seated at his home. The current setup of MediAlly has only been tested in a laboratory environment and the question of the accuracy of context inference for specific application needs still remains open.

IV. A Clinical Trial Scenario

The review of the existing pervasive healthcare solutions suggests that most are lacking provenance capture functionality to specifically address the issues of data quality, integrity, and validation. The questions of how data has been collected, transformed and verified are crucial in the healthcare domain, especially when integrating multiple data sources (e.g. from several points of care) or taking same measurements in different settings (e.g. blood pressure of a standing or lying patient). For example, one clinical system can store patient's age, while another – his date of birth. Date format may also be different (e.g. including or not the time along with the date). Such variations could lead to fact distortion, which may be critical when selecting patients for participation in clinical trials; incorrect inclusion or exclusion of patients may seriously affect research findings.

Below, we propose a provenance-aware pervasive computing system that is designed to address the named issues. To assist us in demonstrating our solution (presented in the following section), we provide a running example that describes the data collection process in a clinical trial. This example illustrates how advanced pervasive computing technologies can be used for optimizing collection of clinical data from study subjects and improving the quality of the data.

In our scenario, a phase-I clinical trial is conducted for testing the safety of a new drug. The study is carried out in a smart hospital, in which a pervasive healthcare environment is implemented with methods similar to those we reviewed in the previous section. In the hospital, all patients wear Radio Frequency Identification (RFID) tags. The patient rooms are instrumented with portable vital sign monitors that can communicate to a central data server via wired or wireless networks. Site professionals are equipped with wireless PDAs that can download and upload data from and to centralized hospital information systems.

The study involves a number of patients. Each patient is assigned to a nurse. The study protocol specifies the dosage and data collection procedures. The procedure can be downloaded to a portable PDA which runs clinical trial application software to instruct the nurse how to perform the trial and collect data. The nurse needs to follow the procedures to distribute the drugs to her patients, and take blood samples and measure the patient's vital signs (blood pressure and body temperature). Finally, the nurse needs to complete case report forms (CRFs) and enter the data into the central data system.

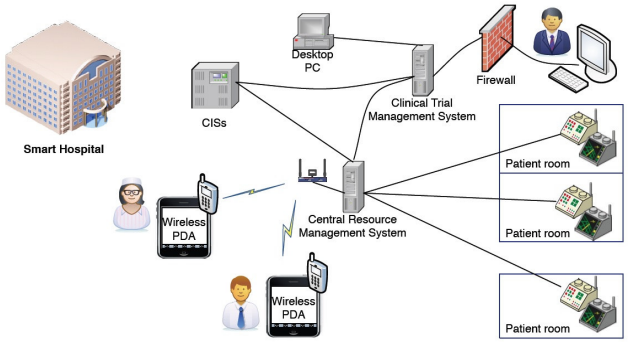

Fig. 1. A smart hospital with a pervasive healthcare environment.

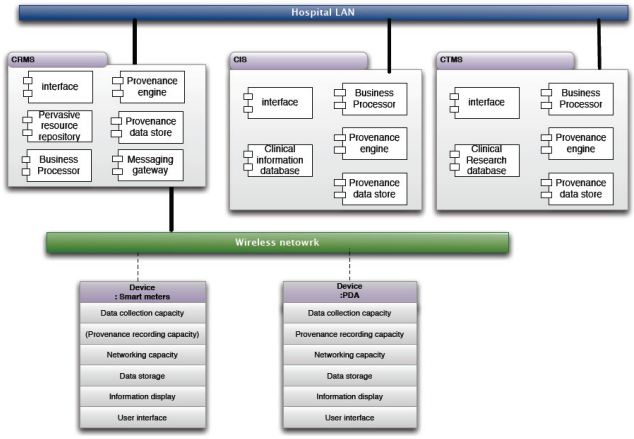

Fig. 2. The proposed system architecture

V. THE PROPOSED SYSTEM

In the following sections, we explain our generic, scalable, pervasive computing system architecture and a provenance-based approach to addressing the data collection, traceability and validation challenges in clinical applications. The system, as shown in Figure 1, can be seamlessly integrated with existing pervasive healthcare infrastructure to achieve automated clinical trial management and improve the quality and speed of clinical research data collection. In particular, the system employs provenance capacity in each of its subsystems, including the smart devices, to ensure the traceability of data collection. Figure 2 illustrates the basic functionality and capacities of the subsystems. It is worth noting that while it is plausible to employ the provenance recording capacity in all of the subsystems involved in the system, some smart devices may not have sufficient capacity to support such functionality (this issue is not addressed in this paper).

A. *System architecture*

As can be seen from Figures 1 and 2, our system consists of the following components.

Clinical Trial Management System (CTMS): The CTMS is a centralized information system that manages the clinical trial projects. It implements a web-based interface which allows authenticated users to access project information. The access control is role-based, so that users only have access to the information permitted by their roles. The CTMS is also a task scheduling system. A task schedule is implemented from the study protocol of a project. It includes a series of procedures which instruct a nurse on when to distribute doses to her patient and take blood samples and measure the blood pressure and cholesterol of the patient.

Central Resource Management System (CRMS): The CRMS is a central resource registry that manages the information about the pervasive devices. A pervasive device has to be registered before it can be used. Once registered, the device becomes a part of the pervasive healthcare infrastructure of the hospital. When the device is in use, it also needs to register the context information about its usage. For instance, when a nurse uses a wireless PDA, her login information will be sent to the CRMS, so the device is bound to the nurse until she logs out or someone else logs in on the device. The CRMS is also a communication gateway between the pervasive resources and other systems.

Clinical Information Systems (CISs): A CIS is a collection or integration of healthcare information systems of a hospital. A good example is the EHR system which is a centralized patient healthcare data repository. A hospital may have a number of heterogeneous CISs, therefore interoperability of various data sources has to be ensured (the issue we address elsewhere but not in this paper). Typically, data acquired from the CISs has to be anonymous when being integrated into clinical research data set.

Pervasive devices: This category includes the heterogeneous smart devices that form the pervasive healthcare infrastructure of the hospital. Examples include wireless PDAs which physicians and nurses carry with them and network-enabled blood pressure monitors being deployed in patients' rooms.

B. *Usage scenario*

Following our running example introduced in section IV, a specific usage scenario can be described as follows (Sarah is the name of the nurse).

1) Clinical trial schedule: When a scheduled task is due, the CTMS sends a reminder to Sarah's PDA via the CRMS. When Sarah becomes available, she logs into the CTMS to acknowledge the reminder and proceed to the patient room to perform the tasks.

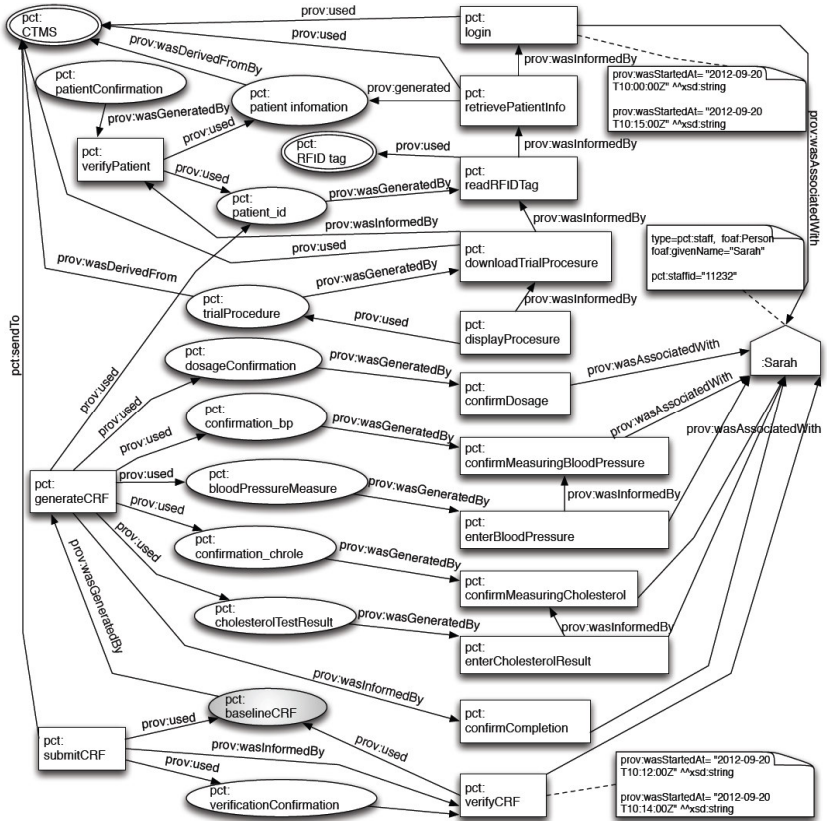

Fig. 3. The provenance of the baseline CRF in the PDA.

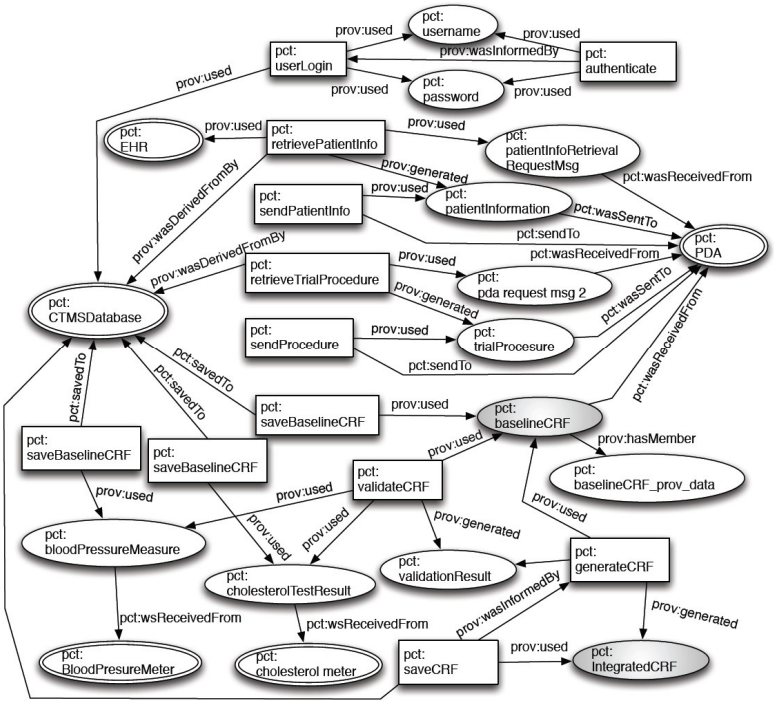

Fig. 4. The provenance of the final CRF by the CTMS.

2) Data collection: When Sarah logs in to the CTMS on her PDA, the procedure is displayed on the device. She needs to follow the procedure, tick each instruction upon completion, and enter data if required. The procedures include:

- First, Sarah needs to read the patient's RFID tag with her PDA to identify the patient. The software sends the information to the CTMS which in turn invokes the EHR to verify the patient.
- Once the patient is verified, Sarah needs to prepare the instrumentation, i.e. the blood pressure meter and the blood cholesterol meter in our case. These smart devices can also read the RFID tag the patient wears. So the measurements will be sent to the CISs and the CTMS with the patient's unique identifier. Sarah needs to verify the patient's ID against the one displayed on her PDA.
- When the patient is verified, the nurse needs to instruct the patient to take the doses as described. When Sarah ticks to indicate the completion of the task, the software logs the time.
- The procedure instructs the nurse to take the clinical measurements of the patient after ten minutes. Sarah needs to verify the clinical measures displayed on the meters and enter them on her PDA. The meters also automatically send the measurements to the EHR once the readings are confirmed by Sarah.
- When all tasks in the procedure are ticked, the PDA will automatically integrate the data, and display a CRF on the screen [7]. Sarah needs to check the form, and if everything is correct, she clicks the "Submit" button, upon which the PDA will send the form to the CTMS.
- Finally, Sarah logs out from the CTMS, and moves to her next patient.

Once the CTMS receives the baseline CRF from Sarah's PDA, it automatically verifies and integrates the data received from the smart meters to generate the final CRF. The clinical data set and the report can then be inspected by the clinical investigator for compliance checks and further validation. In certain systems, the clinical researcher from the trial sponsor side may be allowed access to the CTMS under appropriate security and privacy arrangements.

C. Provenance awareness

Traditional validation of clinical data involved someone manually comparing the paper document with the computerized data. Validation is one inherent issue in computerized data collection. This issue is especially prominent in a pervasive healthcare environment where data collection and integration are computerized. In the previous section, we discussed an architectural solution and its usage scenario to demonstrate how to eliminate errors in data collection in an appropriately designed clinical application. However, manual validation processes with paper-based source documents are still considered the gold standard by many. We propose a provenance-based approach to addressing the data validation challenges in clinical research projects in a pervasive healthcare environment. We argue that the provenance data of the clinical data can sufficiently replace the paper-based data validation approach in many circumstances. In addition, it will remarkable improve the traceability and auditability of the data collection process.

We use the emerging W3C PROV [17] standard to illustrate the form of provenance documentation captured in the data collection processes during a clinical trial. We provide examples of provenance graphs of the scenario described above to illustrate the provenance data available to answer the following validation questions:

1) How was the data collected? The question can be further decomposed into several more specific questions such as who collected what data, what procedure was followed during the data collection, and so on. We expect to use the provenance of the final CRF to ensure the quality of the data.

2) How was the data integrated? During the data collection process, datasets are transformed and integrated into a structured dataset. These integration processes are largely automated by software. We want to use the provenance information of the transformation and integration processes to inspect the correctness of the system functions.

3) How was the data verified? As we discussed in the previous section, data is repeatedly verified by the nurse in the patient room and the CTMS. We want to use the provenance information of these validation processes to affirm the validation.

Provenance documentation: The W3C PROV [17] specification proposes a collection of interoperability solutions, such as the PROV data model, the PROV ontology, and provenance accessing and querying mechanisms. During the data collection, processes involved need to be recorded in a pro-actively defined template to facilitate provenance data integration.

The systems involved in the data collection process, such as the PDA, CTMS, and the CISs, employ provenance recording engines to capture provenance information regarding the data the systems handle and the processes they enact. During the execution of the processes, provenance information such as the input, output and execution time are recorded to generate provenance data. The provenance data from the systems can be integrated to generate a provenance graph, as shown in Figures 3 and 4, which can then be visualized for validating the final CRF and the entire data collection process.

In the graphs (Figures 3 and 4), the visual notation employed by PROV is followed. An oval denotes an *entity*, which is a digital or physical thing, in this case used for data items or devices. A rectangle denotes an *activity*, corresponding to the execution of an action or process in an individual instance. A pentagon (envelope-shape) is an *agent* denoting something or someone who has responsibility for an activity taking place, commonly referring to a person or organization. Some attributes of activities are given, specifically timestamps of when they started and ended. As indicated by arrows, entities and activities are causally related, so the graph states where an activity used an entity in its execution, where an entity was generated by the execution of an activity, where one entity was derived from another, or one

activity received data or was triggered by ('was informed by') another.

The PROV specifications [17] are application-agnostic. The PROV Ontology (PROV-O) specification defines a collection of vocabularies in OWL2 Web Ontology Language to encode the PROV Data Model (PROV-DM). Domain specific profiles can be created to extend the PROV-O to facilitate interoperable provenance modeling in a particular domain or an application. In our system, we define a Pervasive Clinical Application (PCA) PROV profile to ensure the interoperability of the provenance capacities in systems and devices forming the pervasive healthcare infrastructure. We use the namespace prefix *pct* to denote the terms defined in the profile. Those vocabularies are used in combination with the PROV-O vocabularies, with the namespace prefix *prov*, to annotate the relations between entities, activities and agents. The PCA profile specific annotation vocabularies are defined in Table I.

TABLE I. PCA *(PCT:)* ANNOTATION VOCABULARIES

SentTo	annotates a message or a document was sent to a remote system or device.
SavedTo	annotates a piece of data is saved into a local data repository or a database system.
wasReceivedFrom	annotates the source of a data or document.

During the data collection process described in the scenario, the provenance engines in both Sarah's PDA and the CTMS record the activities taken place on the system (device). Figure 3 presents the provenance graph produced by Sarah's PDA. The graph describes how the baseline CRF was produced. It shows the activities happened in the PDA, starting from Sarah's login, until the baseline CRF was generated and being sent to the CTMS. Figure 4 presents the provenance graph produced in the CTMS to reflect the activities happened in the system during the data collection process. These activities include the communication and data exchange processes between the CTMS and the smart devices, i.e. the PDA and the smart meters.

These graphs share a common node, the baseline CRF *(pct:baselineCRF)*, and therefore the two can be readily integrated into a single graph providing a richer history of the data collection than either individual device could provide. Furthermore, the provenance graphs produced by the other systems involved in the process can be queried to expand the integrated provenance graph.

In general, provenance graphs from the provenance data can be generated for different users and different purposes with different levels of information granularity.

Provenance-based validation: So far, we have discussed how to use provenance documentation to record the data collection processes and generate the provenance graphs. The provenance graph can trustfully reveal the collection of processes and resources involved in generating a result, i.e. the baseline CRF and the final CRF in our examples. The processes are not only recorded, but also annotated, so the provenance data and graphs can be used for examining the behaviors of the systems and devices, as well as the responsibilities of personnel involved in the data collection process. With the rapid advancement of ICT and pervasive computing technologies, it is foreseeable that the data collection and integration processes in clinical applications will become increasingly computerized. The scale and complexity of such applications will easily exceed what we discussed in the example scenario. However, the scalability of our provenance-based solution will ensure the efficacy and feasibility of our approach. In fact, as the complexity of the application increases, the benefits that our provenance-based validation solution can bring will make our approach even more prominent and appealing.

VI. CONCLUSION

In summary, the paper proposes provenance-aware system architecture for pervasive healthcare computing applications. Our generic and scalable solution allows a hospital to integrate its pervasive healthcare capacities and resources in order to improve the efficiency and reduce the cost of clinical trials. The pervasive computing-powered computerized data collection approach can significantly improve the quality and eliminate data entry errors in collection of clinical data. Our provenance-based approach demonstrates the feasibility of using provenance documentation and provenance graphs for validation of clinical data.

Our research is still in an early stage. We intend to apply a prototype of the system and the provenance tools proposed in this paper to real-world applications such as the EHR4CR project [11]. We also want to further investigate the challenges that exist in developing provenance-aware pervasive solutions for clinical applications, such as: (i) how to ensure the provenance information is accurate in a pervasive environment where many devices may fail; (ii) how to secure provenance information so that it is not corrupted by illegitimate users or innocent mistakes; (iii) how to efficiently store all of the provenance information in a way that scales and yet is accessible to all pervasive devices.

ACKNOWLEDGMENT

This research was partly supported by the IMI-funded collaborative project EHR4CR (Electronic Healthcare Records for Clinical Research), Grant agreement no.: 115189, and the National Institute for Health Research (NIHR) Biomedical Research Centre at Guy's and St Thomas' NHS Foundation Trust and King's College London. The views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or the Department of Health.

REFERENCES

[1] M. Weiser, "The Computer for the Twenty-First Century," Scientific American, vol. 265, no. 3, pp. 94–104, 1991.

[2] C. Orwat, A. Graefe, and T. Faulwasser, "Towards pervasive computing in health care - a literature review," BMC Medical Informatics and Decision Making, vol. 8, no. 1, p. 26, 2008. [Online]. Available: http://www.biomedcentral.com/1472-6947/8/26

[3] J. E. Bardram, "Pervasive healthcare as a scientific discipline," Methods of information in medicine, vol. 47, no. 3, pp. 178–185, 2008. [Online]. Available: http://view.ncbi.nlm.nih.gov/pubmed/18473081

[4] C. Floerkemeier and F. Siegemund, "Improving the Effectiveness of Medical Treatment with Pervasive Computing Technologies," Workshop on Ubiquitous Computing for Pervasive Healthcare Applications at Ubicomp 2003, Seattle, Washington, Oct. 2003.

[5] J. P. Black, W. Segmuller, N. Cohen, B. Leiba, A. Misra, M. R. Ebling, and E. Stern, "Pervasive computing in healthcare: Smart spaces and enterprise information systems," in Proceedings of the MobiSys 2004 Workshop on Context Awareness, Boston, June 2004.

[6] J. E. Bardram, "Applications of context-aware computing in hospital work: examples and design principles," in Proceedings of the 2004 ACM symposium on Applied computing, ser. SAC '04. New York, NY, USA: ACM, 2004, pp. 1574–1579. [Online]. Available: http://doi.acm.org/10.1145/967900.968215

[7] E. B. Moran, M. Tentori, V. M. Gonzlez, J. Favela, and A. I. Martnez-Garcia, "Mobility in hospital work: towards a pervasive computing hospital environment." IJEH, vol. 3, no. 1, pp. 72–89, 2007. [Online]. Available: http://dblp.uni-trier.de/db/journals/ijeh/ijeh3.html#MoranTGFM07

[8] A. Mihailidis, J. Bardram, and D. Wan, Eds., Pervasive Computing in Healthcare. CRC Press, 2006.

[9] M. Bloomrosen and D. Detmer, "Advancing the framework: Use of health data report of a working conference of the American medical informatics association," Journal of the American Medical Informatics Association, vol. 15, no. 6, pp. 715–722, 2008. [Online]. Available: http://jamia.bmj.com/content/15/6/715.abstract

[10] [B. Delany, A. Taweel et al., "TRANSFoRm: Translational Medicine and Patient Safety in Europe," in Proceedings of the AMIA 2010 Annual Symposium, 2010.

[11] D. Kalra, A. Schmidt, H. Potts, D. Dupont, M. Sundgren, and G. Moor, "Case report from the ehr4cr project – a european survey on electronic health records systems for clinical research," iHealth Connections, vol. 1, no. 2, pp. 108– 113, 2011.

[12] R. G. Marks, "Validating electronic source data in clinical trials," Controlled Clinical Trials, vol. 25, no. 5, pp. 437–446, 2004. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0197245604000510

[13] M. Conti, S. K. Das, C. Bisdikian, M. Kumar, L. M. Ni, A. Passarella, G. Roussos, G. Tr¨oster, G. Tsudik, and F. Zambonelli, "Fast track article: Looking ahead in pervasive computing: Challenges and opportunities in the era of cyber-physical convergence," Pervasive Mob. Comput., vol. 8, no. 1, pp. 2–21, Feb. 2012. [Online]. Available: http://dx.doi.org/10.1016/j.pmcj.2011.10.001

[14] D. Estrin, D. Culler, K. Pister, and G. Sukhatme, "Connecting the physical world with pervasive networks," IEEE Pervasive Computing, vol. 1, no. 1, pp. 59–69, Jan. 2002. [Online]. Available: http://dx.doi.org/10.1109/MPRV.2002.993145

[15] M. Satyanarayanan, "Pervasive computing: Vision and challenges," IEEE Personal Communications, vol. 8, pp. 10–17, 2001.

[16] A. R. Chowdhury, B. Falchuk, and A. Misra, "Medially: A provenance-aware remote health monitoring middleware," pp. 125–134, 2010. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5466985

[17] Y. Gil, S. Miles et al., "PROV Model Primer - W3C Working Draft 24 July 2012," http://www.w3.org/TR/prov-primer/ , accessed September 25, 2012.

[18] J. Bohn, F. Grtner, and H. Vogt, "Dependability issues of pervasive computing in a healthcare environment," in Proceedings of the First International Conference on Security in Pervasive Computing, volume 2082 of Lecture Notes in Computer Science. Springer-Verlag, 2003, pp. 53–70.

[19] S. Chetan and R. Campbell, "Towards fault tolerant pervasive computing," in Pervasive 2004 Workshop on Sustainable Pervasive Computing, 2004, pp. 38–44.

[20] Taweel, A et al. (2011): Service and Model-driven Dynamic Integration of Health Data, Proceedings of the first international workshop on Managing interoperability and complexity in health systems, pp. 11-17, ACM, New York, USA.

[21] Chen Y, Bache R, Miles S, Cuggia M. and Taweel A, A SOA-based Platform for Automating Clinical Trial Feasibility Study, to appear e-Health conference, Prague, July 2013.

[22] Oxford English Dictionary, Oxford University Press.