Proceedings of the 2012 IEEE International Conference on
Cyber Technology in Automation, Control and Intelligent Systems
May 27-31, 2012, Bangkok, Thailand

THE HUMAN-ROBOT CLOUD: SITUATED COLLECTIVE INTELLIGENCE ON DEMAND

Nikolaos Mavridis
New York University AD
PO Box 129188
Abu Dhabi, UAE
nikolaos.mavridis@nyu.edu

Thirimachos Bourlai
West Virginia University
PO Box 6109
Morgantown, WV 26506-6109
thirimachos.bourlai@mail.wvu.edu

Dimitri Ognibene
Imperial College
Elec. And Electronic Eng. Dept
London SW7 2BT
d.ognibene@imperial.ac.uk

***Abstract*—The Human-Robot Cloud (HRC) is an innovative extension of Cloud Computing across two important directions: First, while traditional cloud computing enables transparent utilization of distributed computational as well as storage resources, the HRC enables, in addition to the above two, the utilization of (a) distributed sensing (sensor network technology) and (b) actuator networks (including robot networks). Thus, HRC extends the concept of cloud computing by connecting it to the "Physical World", through sensing and action. Second, while traditional cloud computing involves the usage of only electronic components, such as computers and storage devices, the HRC's capability is extended by the support of human physical and cognitive "components" as part of the cloud, which are neither expected to be experts nor to be engaged with the cloud full-time. Such components are primarily expected to interact with the system for only short periods of time (seconds), essentially providing crowd-servicing for the Cloud. Human components provide any or a mixture of the following: a) input arising from a number of sources through the usage of their sensory faculties (auditory, visual etc.), - thus, acting as "intelligent sensors" attached to the cloud; b) input that results from the usage of their cognitive faculties (pattern recognition, prediction, identification, planning etc.) – thus, acting as "intelligent systems" attached to the cloud; and c) actuation services to the Cloud (by moving around their bodies or other objects) – thus acting as "actuators" attached to the cloud. Thus, the proposed HRC is aiming to achieve the best of both worlds, i.e., either humans or machines, being able to carry out tasks which are very difficult or impossible for either humans or machines alone to carry out. Furthermore, the HRC enables the *construction* of situated agents exhibiting collective intelligence on demand, and the *transformation* of situated agency from a "capital investment" to a service, components of which can be provided by multiple providers, in a transparent fashion to the end user**

Keywords-robots; cloud computing; human computation; sensor networks; crowd-sourcing; crowd-servicing

I. INTRODUCTION

Computation, software, data access, and storage services, traditionally require end-user knowledge of the physical location, as well as ownership and configuration of the system that delivers the services. Users or organizations, requiring computational power or storage space, have to physically buy and usually own (or long-term lease) physical computers and storage devices. However, despite its simplicity, the main disadvantages of such a traditional concept are the following: a) a large initial capital investment; b) large amounts of potentially unused computational power and/or storage space during the lifetime of the system; and c) the concept's limitation to peak computational throughput and storage space due to hardware constraints. Hence, the concept of *cloud computing* has recently been proposed that overcomes the aforementioned limitations. Numerous cloud computing-based services are already operational [1-4], and considerable related research has and is taking place [5-8].

With cloud computing, the delivery of computing takes the form of a *service* rather than a *product*. In addition, shared resources (computational and storage power) as well as software and information, are provided to the end-users as a utility (like the electricity grid) over a network. Thus, the shared resources are provided to the end-user, without the need for him/her to know where the machines provided are physically located, or what the specifics of the configuration of such machines are. In essence, cloud computing aims to achieve *transparent sharing of distributed resources*, which can be utilized by multiple users.

Some of the main *advantages* of cloud over traditional computing are the following: First, for a large cloud, practically the peak computational throughput (that can be sustained for small periods of time) is high; and thus, bursts of activity can be accommodated. Second, depending on the Service-Level Agreement (SLA) between the user and the Ease of Use cloud, periods of zero or low activity can have little or no cost. Third, and quite importantly, there can usually be easy scalability to the cloud, as well as high robustness. As computational processes are not dispatched to faulty nodes, and replacement nodes are abundant, extra capacity can be added on demand, and reliability can be much higher.

However, despite all of the above advantages, traditional cloud computing has one important limitation: it is not effectively connected to the *physical world*, in the way that a situated robotic agent would be. This is the case, as traditional cloud computing does not include components that directly sense the physical world (cameras, microphones, temperature and motion sensors etc.), nor does it include components that can directly act upon the physical world (motors, robotic arms, mobile robots, UAVs, etc.).

We thus propose a *two-fold extension* of the traditional concept of cloud computing, that is expected to transform conventional cloud services to the concept of the *Human-Robot Cloud* (HRC). HRC extends the baseline cloud concept to include not only computational and storage resources, but also sensor as well as actuator networks – such as cameras and robots. An important aspect to these additional characteristics to the cloud concept is the capability for virtualization of such networks; the user requests a specific sensing service to be provided by the cloud, without the need for a detailed knowledge of its implementation, i.e., what type and which specific sensor will be used. Similarly, the user requests a motion / mobility or another physical action service, without caring about which robot, motor or relay will be activated. In practice, it is the responsibility of the cloud to recruit appropriate resources to carry out the physical action. Therefore, its *first extension* is the augmentation of the cloud with sensing and actuation resources, in order to effectively turn it from an abstract computational and storage entity to a physically situated agent, that can sense and act upon the real world.

But we also move beyond this extension, and propose a *second important extension*: Given the current state-of-the-art of Artificial Intelligence, there exist tasks for which machines surpass human performance (such as numerical calculations); and others for which humans are still much more competent than machines (such as object recognition and some types of pattern recognition – for example, CAPTCHA recognition [9]). Of course, there exist many other tasks for which the scale does not way so unilaterally to the machine or the human side. The question of whether the task should be assigned to a machine or a human, however, should depend on current availability, and constraints of cost, accuracy, response time etc. Thus, in order to be able to harvest the best of either the machine or the human world, we propose to extend the traditional cloud in order to include *not only electronic* computation, sensor and actuator networks *but also human* computation, sensing, and actuation (human computation such as in [10], human sensing such as in the DARPA Baloon Challenge [11], human actuation, i.e. physical mobility and manipulation of objects by humans).

Through the usage of these six types of components (electronic computation & storage, electronic sensors, electronic actuation networks, as well as human computation, sensing and actuation), the proposed HRC cloud effectively becomes a *distributed situated agent* of mixed (human and machine) nature. The aforementioned components of human or machine nature are capable of entering or exiting the agent on-the-fly. Also, the humans providing services to the cloud are not expected to be dedicated to it. They can be *micro-engaged* to the cloud for bursts of time (e.g., seconds), while they are engaged in other business- or entertainment-related activities.

Through the inclusion of humans, as well as sensor and robot networks, many new capabilities arise – moving the cloud far beyond the state-of-the-art of today's cloud computing. Also, many new complexities as well as *open research questions* are expected to be brought forward. In the next section, a sampler of some of the most important questions is presented, followed by a background section, a discussion, and a conclusion of this position paper.

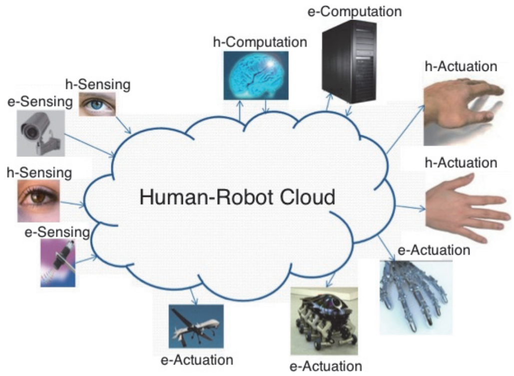

Figure 1. The 6 types of components of the Human-Robot Cloud: (e-) Electronic Sensing/Computation/Actuation, as well as (h-) Human Sensing/Computation/Actuation

II. AREAS OF OPEN RESEARCH QUESTIONS

As promising a concept as the Human-Robot cloud seems to be, its wider realization will be catalyzed by obtaining answers to a number of interesting groups of related research questions. Many of such questions are related to the specifics of including human elements within the cloud, but also to vital technical considerations, as well as issues of security and, quite importantly, ethics. In more detail:

A. *Human Interface and Recruitment:*

How can the cloud interface to humans (providing sensing, computation, or actuation services for the cloud), and how can humans be recruited ?

Multiple possibilities exist: cell phones, tablets, PC's – in the form of text messages, audio messages, pictures, videos, and so on. Also, the interface might be through a video game that a human is playing; one of the activities or the riddles that the user has to solve in the game, might actually be what the cloud needs at that moment.

B. *Human Incentivisation:*

Why should the human give some of his sensory, processing, or actuation cycles to the cloud?

There is a wide range of options, such as: 1) extrinsic financial incentive: for monetary return – where the human is explicitly being paid for his services in traditional currency, 2) intrinsic altruism: as a volunteer while the cloud is serving a worthy purpose, 3) extrinsic non-financial incentive: for example, by bartering so that the cloud can give services in return back to the service-providing human, in exchange for the services that he has provided to the cloud, or even 4) transparent parasitic: for example through the human providing services as part of the game-play of an online game, without explicit knowledge that the tasks he/she is carrying out in the game are useful for the cloud.

C. *Service-Level Agreements:*

What agreements should be made between the cloud and its users, and what among the cloud, humans and machine service providers?

Many open questions exist here and many simple answers as well as existing schemes can be used to bootstrap real-world HRC's. For example, how should pricing be dynamically scaled, on the basis of current supply and demand, as well as of contracted and estimated Quality of Service?

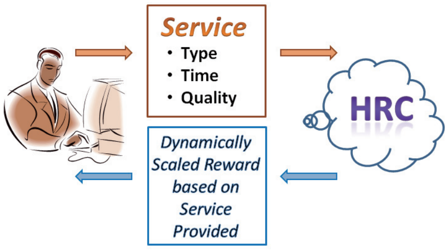

Figure 3. Hypothetical SLA for the Human-Robot Cloud, between a human service provider and the cloud. A complementary SLA structure also exists, for the case of a human service consumer (second possible relation to the cloud) and the cloud which provides the required service.

D. Human Quality Control:

How can one estimate the quality of the services supplied by the humans? Are the humans performing sensing (search for a red ball), pattern recognition (perform object recognition or OCR, or speech recognition in a noisy recording etc.), or actuation (move a ball) services, whose quality adheres to the specification of the service that was originally assigned to them?

Multiple possibilities exist, starting from the simple:

1) One approach is related to *redundancy* (a task is provided to more than one humans) and trusting majorities (i.e. an approach related to inter-annotator agreement measures in computational linguistics etc. [12]). The underlying idea here could be that, truth is approximated by the beliefs of the majority. Of course, such schemes require significant refinement, in order to be resistant to attacks (e.g, such as the sensor network versions of [13], for example against the node replication attack).

2) Another approach is based on testing the quality of the human services provided, by *comparing them with known ground truth*; for example a) assigning an image pre-labeled with very high confidence as a labeling task to a human, and comparing his answer to the known ground truth, or b) assigning an object sensing task to a human (find the red ball), for an object that was already placed by the cloud on purpose in a specific location in order to test the human's performance. Many combinations of the above as well as other approaches exist [14].

E. Typology of Services:

What are some basic building blocks for services that can ideally be generic and application independent? In other words, what are the basic sensing, computation, and actuation primitives that the system should support, and how can these become implementation-independent?

TABLE I. SOME BASIC TYPES OF E-SERVICES AND H-SERVICES THAT COULD BE PART OF A TYPOLOGY.

e-Services (Provided by Machines)	
Sensing	High-level active sensing primitives (location-, object-, human- or feature-directed), with or without addressing sensor genre and timing window
Cognition	Pattern Recognition, AI Planning and Scheduling, Inference, Database Operations, etc.
Actuation	High-level motion & mobility primitives (specifying target configuration at multiple levels of abstraction, with or without trajectory details / constraints, in a body- or object-directed manner, with or without addressing actuator genre
h-Services (Provided by Humans)	
Sensing	Reports of what is being seen, heard, etc. (location-, object-, human- or feature-directed, etc.)
Cognition	Pattern recognition, Human-Assisted Planning & Scheduling, Translation Services, Personal and Communal memories, Commonsense knowledge etc.
Actuation	High-level motion & mobility primitives, similar to e-service specs, at human-achievable accuracy granularities

F. Discovery and Composition of Services:

There are a number of important questions implicated here: 1) How can the cloud deal with connections and disconnections of humans and devices to the cloud? 2) How can the cloud keep track of the resources that are available at a particular moment in time? 3) How can it dynamically discover a service that fulfills the requirements of the task at hand? And 4) How can the cloud compose micro-services to create larger building blocks (macro-services)?

G. Task Interpretation, Planning, Goal Selection:

The cloud could either be 1) executing a task based on a list of potentially concurrent micro-services (serial execution with pre-defined parallel threads), or 2) it could be given a task described at a more macro level (without an explicit specification of the micro-services required to achieve it), in which case it needs to assemble micro-services in order to achieve it, in a compiled or interpretive manner. Alternatively, 3) it might be given a goal specified in terms of desired world- and/or internal-state. In such as case it needs to perform planning and re-planning, in order to reach the achieved goal-state. Planning will essentially consist of taking a sequence of actions (physical or computational), which would correspond at the lowest level to execution of the relevant micro-services from the cloud. Furthermore, notice that action selection (which e-service or h-service should be selected?) and planning in the cloud could be given to a mixture of machine (planning algorithms) and human (human planners) components. Finally, one could envision the cloud to be operating with multiple goals assigned to it; and these goals might be coming either from cloud users, or might even be coming from its human components (service providers), which are playing the role of cloud users accessing the cloud in return to their services. In both cases, goal arbitration and/or selection is required, in order to prioritize, parallelize or serialize execution, and resolve conflicts between the top-level goals that the cloud, viewed as a distributed monolithic situated agent, is trying to achieve.

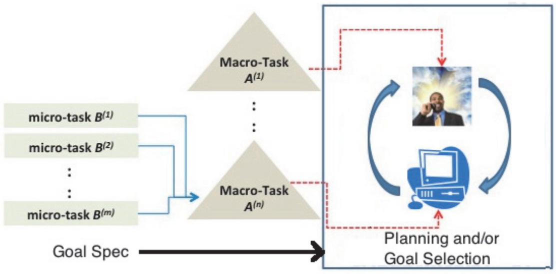

Figure 4. Figure 5: Multiple avenues for specifying what the cloud should perform: through a combination of micro- and macro-services, as well as single and multiple Goal specs. The actual goal selection as well as planning might be performed by either AI programs and/or humans providing appropriate services to the cloud. Furthermore, the goal specs and/or tasks to be executed could either be coming from humans or machines requesting services, or from the human service providers, as a rewarding return to their contributions to the operation of the cloud.

H. Security:

How can the cloud as well as the information flowing or stored in it be secured against potential intruders or malignant use? What types of attacks can be envisioned, and what are the respective countermeasures? How can one guarantee the cyber- as well as physical security of the components of the cloud, but also of the humans that might reside nearby where the actuation components of the cloud are (robotic arms etc.)?

I. Ethics, Privacy, and Legal Responsibility:

What are the important ethical issues that arise from the Human-Robot Cloud? How can one protect privacy? Also, how can one at the same time assign legal responsibility for cases of intentional or unintentional, foreseeable or unforeseeable misconduct of users, humans providing services, machine components, or the system as a whole?

Thus, a wide range of important research questions exist, towards the effective wider-scale realization of the Human-Robot Cloud. However, many important advantages of HRC are foreseen, which we will explicate in the discussion that follows, after providing some background.

III. BACKGROUND

There exist a number of domains of existing research that are highly relevant to implementing the HRC. These include: Cloud computing; Service Discovery and Composition; Sensor Networks; Robot Networks; Crowd-Sourcing; Crowd-Servicing, Human Computation; Incentivisation; Security; Ethics; Quality Control. Furthermore, one of the many promising application areas of the HRC is in enhancing real-world biometrics and surveillance applications, and thus we will also provide relevant background towards this direction.

Regarding *Cloud Computing*, the main existing platforms include Amazon web services [1], Google App Engine [2], Microsoft's Azure [3], as well as open-source solutions such as Eucalyptus [4]. An overview of the current state of research related to the cloud can be found at [5], while some primers regarding practical issues include [6-7], and terminology [8]. *Web Service Discovery and Composition* is a subject with more than a decade of coverage, a sampling of which can be achieved through [15-18], with a basic overview included in the textbook [19].

On *Sensor Networks*, there is already a wide literature, including introductory textbooks such as [20-22], and also the classic but somewhat outdated review [23], as well as topics highly relevant to the HRC such as simple query processing systems for sensor networks [24]. Also of particular importance to the Human-Robot Cloud are early attempts to integrate Sensor and *Robot Networks* for real-time situated tasks, such as in emergency response scenarios [25]. The Robot Networks literature traditionally concentrates on decentralized multi-robot motion planning and communications-related aspects [26-28], and is complemented with research termed as Distributed Robotics, with real-world examples including the Scout [29] and COTS robots [30], among others.

Crowd-sourcing, crowd-servicing, and Human Computation have a much more recent history, but are already booming in applications and importance. A prime example of real-world platforms is the Amazon Mechanical Turk (MTurk), one of the suites of Amazon Web Services, that enables computer programmers (Requesters) to co-ordinate the use of human intelligence (Workers) to perform tasks that computers are unable to do yet, which has been used already for a wide variety of tasks, including speech transcription, translation as well as evaluating translation quality [31], annotation [32], user studies [33], and much more. The pioneering work of Luis Von Ahn on Human Computation also offers a good illustration of real-world crowdsourcing, for tasks such as labelling towards automated object recognition [34], as well as on exploring the potential of game-like activities for harvesting human cycles towards collectively solving large-scale problems [35]. Of particular relevance to the Human-Robot Cloud are a number of special issues, such as appropriate task routing to people with relevant abilities, which can be implemented for the case of relying on local-knowledge of the agents involved through the approach of [36], ingenious approaches towards achieving simple but effective quality control such as the iterative dual-pathway structure used for speech transcription in [37], and consideration regarding high response speed real-time crowd-powered systems [38].

Finally, cyber- and physical- *security* as well as *ethics* are important existing background areas, which must be reconsidered and extended in order to account for the new challenges created by the Human-Robot Cloud. Of particular importance is existing work in security for sensor networks, as well as for security for robotic systems in close proximity to humans. Furthermore, information ethics as well as robot ethics [39] [40], can provide a strong starting basis for building upon for the Human-Robot Cloud.

IV. DISCUSSION

Given sufficient progress in the above research questions, a number of main advantages of the proposed HRC are foreseen. These include the following:

A. Revolutionary Change in Economic Aspects

Separation of infrastructure owners from users becomes possible for the case not only of computation and storage but also for sensor networks (for example, camera installations) as well as robots (for example, mobile robots). Also, depending on pricing and delay considerations, the ability the offsource human computation to considerable physical distances from the locus of situatedness is possible; as well as the ability to recruit human sensing and actuation whenever their electronic counterparts are not available or not suited for the task, become possible. Bartering and Dynamic pricing schemes for effective situated agency can be implemented on the Human-Robot Cloud, too.

B. Best of Human and Machine Worlds

If Mechanical Turk-inspired approaches and crowdservicing aim to provide the best of both worlds when it comes to computation, the HRC extends the abilities of the resulting system to the best of both worlds, not only for computation, but also for sensing and actuation.

C. Increased Utilization, Efficiency, Reliability

The dead time of ownership-based sensor and actuator networks can now be "rented out" at competitive prices, thus enabling much higher utilization. Through reuse of sensing and actuation resources, as well as through re-use of intermediate results and synergistic/opportunistic planning, increased efficiency becomes possible. Also, through the redundancy and reconfigurability of the components of the Human-Robot Cloud, increased reliability becomes achievable.

D. Potential Performance Increments by Scaling

There are many cases for which, as the number of users of the cloud increases, the intermediate results of sensing, computation, or actuation, might become applicable to more than one tasks which are running on the cloud. Thus, the component utilization time vs. tasks completed ration is foreseen to potentially improve, as well as other aspects of performance.

E. User and Human-Service-Provider Satisfaction

By belonging to a wider team, and also by seeing their goals achieved in return to their contribution to the cloud, in the case of a "bartering", for which time dedicated to the cloud by a human service provider is returned to him in terms of credit for specific types of services that the human-robot cloud can in turn perform according to his requests. Thus, the cloud can also effectively become an agent which exhibits aspects of partial auto-telicity, as its goals are partially dictated by the goals of some of its contributing components, i.e. the humans contributing services to the cloud.

V. CONCLUSION

In this position paper, we have introduced the concept of the "Human-Robot Cloud" (HRC). The HRC is a two-fold extension of cloud computing, which includes not only computation and storage, but also sensing and actuation elements, of both human as well as machine nature. Thus, it transforms "situated agency" (interconnected sensory, cognitive, and actuation elements) from a resource to a utility; and provides numerous noteworthy advantages of traditional solutions.

The most important characteristic of the proposed HRC is its dual nature, i.e., the combination of both human and electronic elements. Thus, the advantage over the conventional cloud concept is that sensing, cognition, and actuation is provided both by machines as well as humans, chosen on the basis of availability, cost, as well as depending on the nature of the task carried out. Our proposed concept is expected to achieve the best of both worlds, i.e., either humans or machines, being able to carry out tasks which are very difficult or impossible for either humans or machines alone to carry out.

Two of the main benefits of the proposed HRC concept are: first, the utilization of a dynamic mixture of machine and human resources in order to carry out specific tasks such as search and identification, tracking, co-ordinated operations involving physical mobility and situated communication etc.); and second, the fact that human components operators that are connected to the Cloud (Cloud service providers) do not need to be experts, (i.e. they need not (a) have to be physically close to an operations center and (b) dedicate a lot of personal time and (c) have expert knowledge and abilities.

Thus, an important consequence of the proposed HRC concept is the fact that it revolutionizes the economics of sensor and robot networks, as well as the economics of crowd-serviced human computation. This is achieved by enabling dynamic pricing and time-sliced utilization of sensor / robot network resources instead of settling for the traditional most expensive and inflexible options of long-term leasing or purchase of dedicated lines, computers, sensors, actuators and human resources, which furthermore often need to be physically relocated. Thus, the concept of situated distributed intelligence that we propose does not require an inflexible physical system that needs to be purchased and installed at a huge capital cost, but effectively becomes a flexible rentable "service on demand", with much higher "instantaneous peak capacity" as well as efficiency and robustness.

In order to reach a wider implementation of the HRC concept, there exist a number of areas of important *open research questions*, which need to be addressed in the near future. Such questions include aspects not only of human interface and recruitment, incentivisation and quality control, but also of being able to derive typologies of services at multiple levels of abstraction, as well as perform service discovery and composition, simple- and composite-task execution, as well as planning and goal selection. Quite importantly, numerous open research questions regarding security as well as ethics also exist.

Answers to these open questions require further research in the future, but are also based on existing background areas, which were briefly reviewed. However, the prospective benefits of a wide implementation of the Human-Robot cloud are multiple, and a sampling of them was provided in the discussion section.

Last but not least, we foresee that through appropriate design considerations, the human-robot cloud might also serve as a vehicle for partial auto-telicity, by potentially propagating the goals of the humans which contribute to it, and for a deeper sense of communal participation, collaboration, and harmony.

REFERENCES

[1] F. P. Miller, A. F. Vandome, and J. Mcbrewster, “Amazon Web Services”, Alpha Press, ISBN-13: 978-6131788369, 2010.

[2] E. Ciurana, “Developing with Google App Engine”, First Press, ISBN-13: 978-1430218319, 2009.

[3] D. Chappell, “Introducing the Azure Services Platform”, available online at: http://download.microsoft.com/download/e/4/3/e43bb484-3b52-4fa8-a9f9ec60a32954bc/Azure_Services_Platform.pdf, 2008.

[4] D. Nurmi et al., “The Eycalyptus Open Source Cloud Computing System”, CCGrid: The 9th IEEE Intern. Symposium on Cluster Computing and the Grid, 2009.

[5] Q. Zhang, L. Cheng, and R. Boutaba, “Cloud computing: state-of-the-art and research challenges”, Journal of Internet Services and Applications, 1(1), 2010 , pp. 7-18.

[6] M. Armburst, et al., “Above the clouds: A Berkeley view of cloud computing”, Tech. Report EECS-2009-28, EECS Dept, UCB, 2009.

[7] R. Buyya, C. Yeo, and S. Venugopal, “Market-Oriented Cloud Computing: Vision, Hype, and Reality for Delivering IT Services as Computing Utilities”, in Proceedings of High Performance Computing and Communications, HPCC, 2008, pp. 5-13.

[8] P. Mell, and T. Grance, “The NIST definition of cloud computing”, NIST, Special Publication 800-145, 2011.

[9] H. D. Truong, C. F. Turner, and C. C. Zou, “iCAPTCHA: The Next Generation of CAPTCHA Designed to Defend Against 3rd Party Human Attacks”, in Proc. IEEE International Conference on Communications, pp. 4758-4763, 2011.

[10] Luis von Ahn, B. Maurer, C. McMillen, D. Abraham, and M. Blum, “reCAPTCHA: Human-Based Character Recognition via Web Security Measures”, Science, 12:321(5895), 2008, pp. 1465-1468.

[11] Tang, J.C., et al., “Reflecting on the DARPA Red Baloon Challenge”, Communications of the ACM, vol 54 no. 4 , 2011, pp. 78-85.

[12] M. Banerjee, M. Capozzoli, L. McSweeney, and D. Sinha, "Beyond Kappa: A Review of Interrater Agreement Measures". Canadian Journal of Statistics, 27 (1), 1999, pp. 3–23.

[13] H. Yu, M. Kaminsky, P.B. Gibbons, and A. D. Flaxman, “SybilGuard: Defending Against Sybil Attacks via Social Networks”, IEEE/ACM Transactions on Networking, 16(3), 2008, pp. 576-589.

[14] R. Kern, H. Thies, C. Bauer, and G. Satzger, “Quality Assurance for Human-Based Electronic Services: A Decision Matrix for Choosing the Right Approach”, Current Trends in Web Engineering, Lecture Notes in Computer Science, 6385, 2010, pp. 421-424.

[15] K. Sycara, M. Paolucci, A. Ankolekar, and N. Srinivasan, “Automated discovery, interaction and composition of semantic web services”, Journal of Web Semantics, 2003, 1(1):27–46.

[16] X. Su, and J. Rao, “A Survey of Automated Web Service Composition Methods”, in Proceedings of First International Workshop on Semantic Web Services and Web Process Composition, 2004.

[17] A. Ragone, et al., “Fully Automated Web Services Discovery and Composition through Concept Covering and Concept Abduction”, International Journal of Web Services Research (JWSR), 4 (3), 2007.

[18] T. Weise, et al., “Different approaches to semantic web service composition”, In A. Mellouk et al. in Proceedings of The Third IEEE International Conference on Internet and Web Applications and Services, 2008, pp.90–96.

[19] M. P. Papazoglou, “Web Services: Principles and Technology”, Prentice Hall, 2007, ISBN: 978-0321155559.

[20] W. Dargie, and C. Poellabauer, “Fundamentals of Wireless Sensor Networks: Theory and Practice”, Wiley, ISBN 978-0470997659, 2010.

[21] A. Willig, and K. Holger, “Protocols and Architectures for Wireless Sensor Networks”, Wiley, ISBN 978-0470519233, 2007.

[22] I. F. Akyildiz, and M. C. Vuran, “Wireless Sensor Networks”, Wiley, ISBN 978-0470036013, 2010.

[23] I. Akyildiz, S. Weilian, Y. Sankarasubramaniam, and E. Cayirci, “A Survey on Sensor Networks,” IEEE Comm. Mag., 40(8), 2002, pp. 102–114.

[24] S. R. Madden , M. J. Franklin , J. M. Hellerstein, and W. Hong, "TinyDB: An Acquisitional Query Processing System for Sensor Networks", ACM Trans. Database Systems, 30(1), 2005, pp. 122-173.

[25] V. Kumar, D. Rus, and S. Singh, “Robot and sensor networks for first responders”, IEEE Pervasive Computing, 3(4), pp. 24-33.

[26] P. Basu, and J. Redi, “Movement control algorithms for realization of fault-tolerant ad hoc robot networks”, IEEE Networks 2–10, July/August, 2004.

[27] S. Das, H. Liu, A. Kamath, A. Nayak, and I. Stojmenović, “Localized Movement Control for Fault Tolerance of Mobile Robot Networks”, Wireless Sensor and Actor Networks, vol. 248, 2007, pp. 1-12.

[28] C. A. Luciano, “Simultaneous Coverage and Tracking (SCAT) of Moving Targets with Robot Networks”, Springer Tracts in Advanced Robotics, vol. 57, 2009, pp. 85-99.

[29] L. J. Pierce, P. E. Rybski, S. A. Stoeter, and N. Papanikolopoulos, “Dispersion behaviors for a team of multiple miniature robots”, in Proccedings of the IEEE International Conference on Robotics and Automation, 2003, pp. 1158-1163.

[30] B. Kratochvil, et al., “Heterogeneous implementation of an adaptive robotic sensing team”, In Proc. Of the IEEE International Conference on Robotics and Automation, 2003, pp. 4264-4269.

[31] C. Callison-Burch, “Fast, cheap, and creative: evaluating translation quality using Amazon's Mechanical Turk”, in Proceedings of the Conference on Empirical Methods in Natural Language Processing, 1(1), 2009.

[32] A. Sorokyn, and D. Forsyth, “Utility Data Annotation using Mechanical Turk”, in Proccedings of First IEEE Workshop on Internet Vision at CVPR, 2008.

[33] A. Kittur, E. H. Chi B. Suhand, “Crowdsourcing user studies with Mechanical Turk”, in Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems, 2008.

[34] Luis von Ahn, and L. Dabbish, “Labelling images with a computer game”, in Procccedings of the SIGCHI conference on Human Factors Computing Systems, 2004.

[35] Luis von Ahn, “Games with a purpose”, IEEE Computer, 39(6), 2006, pp.92-94.

[36] H. Zhang, E. Horvitz, Y. Chen, and D. C. Parkes, “Task Routing for prediction tasks”, in Proceedings of EC'11 Workshop on Social Computing and User Generated Content, 2011.

[37] B. Liem, H. Zhang, and Y. Chen, “An iterative dual pathway structure for speech-to-text transcription”, in Proceedings of the third Human Computation Workshop HCOMP, 2011.

[38] M. S. Bernstein, J. Brandt, R. C. Milleret, and D. R. Karger, “Crowds in Two Seconds: Enabling Realtime Crowd-Powered Interfaces”, in Proceedings of the User Interface Software and Technology Symposium UIST, 2011.

[39] G. Veruggio, “The birth of Roboethics”. in Proccedings of the IEEE International Conference on Robotics and Automation Workshop on Robo-Ethics, 2005.

[40] P. Asaro, "Remote-Control Crimes: Roboethics and Legal Jurisdictions of Tele-Agency," Special Issue on Roboethics, Gianmarco Veruggio, Mike Van der Loos, and Jorge Solis (eds.), IEEE Robotics and Automation Magazine, 18(1), 2011, pp.68-71.