Sharing Policies for Multi-Partner Asset Management in Smart Environments

Christos Parizas, Diego Pizzocaro, and Alun Preece
School of Computer Science and Informatics
Cardiff University, UK
Email: {C.Parizas, D.Pizzocaro, A.D.Preece}@cs.cardiff.ac.uk

Petros Zerfos
IBM T.J. Watson Research Center
Yorktown Heights, NY, USA
Email: pzerfos@us.ibm.com

***Abstract*—Smart environments are ecosystems, which seamlessly embed IT assets into physical world's objects and hold promise for improving the services we receive from our social and economic ecosystems. The management of smart environment assets in multi-partner, dynamic collaboration scenarios where different sets of assets are owned and operated by different partners is a non-trivial problem, due to restrictive asset sharing policies applied by collaborating partners. In this work we formalize, evaluate and compare two asset sharing policies, investigating their impact on MSTA-P, a policy-regulated version of an existing asset-task assignment protocol. The first sharing policy is based on a traditional asset ownership model while the second is based on an *edge* model allowing asset sharing among collaborating partnes through cross-partner team formations. We find that while the traditional ownership model allows slightly better performance, the difference is only marginal, so a team-sharing model offers a viable alternative sharing approach.**

I. Introduction

Advances in information technology including the Internet, sensors and communication protocols, combined with the availability of cheaper computing power set the future for intelligent and interacting "smart environments". Smart environments are ecosystems composed of infrastructures that blend physical and IT assets, wherein sensors, network connectivity and data storage are embedded seamlessly in physical world's objects and hold promise for improving the efficiency of our social and economic ecosystems [1]. Smart environments come as a result of an effective integration of planning, construction and management methods prepared for both expected and unexpected events [2].

In unexpected scenarios such as humanitarian relief and emergency response operations, two or more organizations that own and operate disparate sets of smart environment assets merge forces and form coalitions to achieve common objectives. In such cases, facilitating the collaboration between multiple partners and establishing trust while considering security and privacy issues is a precondition for a successful operation. Typically, collaborating partners have their own inherent restrictions, which are stated as sets of policies on how to share their infrastructures with other organizations.

Recent examples of disaster situations, such as those unfolded during the Haiti earthquake[1] demonstrated the need for emergency responders affiliated with national and organizational groups to form cross-organizational teams and share assets in an ad-hoc manner, in order to provide humanitarian assistance. Within such stressed environments, assets such as sensors, and data analytics engines, collect, process and disseminate operational data and insights, which are then shared across organizations to enable quick decision making.

Sharing smart environment assets to support multiple concurrent and multi-organization missions is a non-trivial problem. Collaborating organizations have different backgrounds (e.g. area of expertise, cultural background) which reduce shared awareness and understanding of the mission, leading to different decisions about what assets can be shared, with whom, and when [3].

One of the most widespread asset sharing models is the static asset ownership approach that we refer to as *asset-centric* sharing, according to which assets belonging to a collaborating partner may or may not be shared with other partners [4], based on pre-defined policies. However, this approach assumes a centralized operations and control center that specifies a cumbersome access model. In this work we present a novel asset sharing scheme inspired from military operations, based on the *edge* model [5]. According to this, members from multiple organizations are grouped into cross-organizational teams and share ownership and management authority regardless of their organizational affiliation. The *team-centric* sharing model shifts decision making power regarding access and control of assets to the edge of organizations. allowing for a more dynamic assets sharing pattern to emerge.

In [6] we developed a distributed protocol called Multi-Sensor Task Allocation (MSTA) for addressing the problem of allocating heterogeneous bundles of sensing assets to a variety of different sensing tasks, with the goal of maximizing the usefulness of assets and satisfying the most critical task requests. While MSTA was initially designed for sensors, by considering sensors as sensing service providers and the sensor bundles as composite services, the algorithm is generally applicable to services-to-task allocation. In this work we formalize, evaluate and compare the two aforementioned asset sharing policies, investigating their impact on MSTA-P protocol. In particular, this paper makes the following contributions: (1) a formal representation of the two asset sharing models using predicate logic, which makes them amenable to analysis, (2) a

[1]Haiti Earthquake Response: Context Analysis - http://tinyurl.com/k8cffr7

novel asset sharing protocol, called MSTA-P, which integrates the evaluation of asset sharing policies, (3) evaluation and comparison of the sharing models using a discrete-time, multi-agent, simulation environment. We find that while the traditional ownership model allows slightly better performance, the difference is only marginal, so a team-sharing model offers a viable alternative sharing approach.

II. Related Work

Well known asset sharing approaches in multi-organizational environments are represented by [7]–[9]. In particular, [7] proposes a model where new collaborating members can only have access to a specific resource if they are first invited by authorized partners, [8] proposes a role-based framework that combines users' characteristics with parameters such as time and user IP address, allowing or denying access to resources accordingly, while [9], based on an automated trust negotiation approach focusing on a type of "contract" in which collaborating partners agree to share their resources over a given time period. Although all the above models cope with resource sharing in a secure and confidential manner, they are likely to fail or encounter difficulties in being applied to highly dynamic environments. In order for all of them to comply with situational changes, an extra overhead is needed due to the spatial or chain of command distance between the decision making center, and therefore the policy making centre, and the place where the changes take place. Differently, the edge, team-based sharing model presents much lower overhead regardless the situational changes frequency. The teams on the edge of the network – being event-driven entities – are formed, reformed or disassembled as a response to environmental changes therefore, sharing policies based on this model are always up-to-date to the unfolding operations.

III. Asset Sharing Policy Models

The asset sharing models that we propose and experiment with are binary; that is they either give or not access to assets services. We acknowledge the existence of finer-grained asset sharing models, which using techniques such as obfuscation, can grant access to subsets of services capabilities. The investigation of finer-grained sharing approaches are outside the scope of this work. However, we assume that ours and fine-grained models complement each other and that one can provide multi-level asset sharing management by combining the two. Consider for example the case where there are three sharing grades of a service (e.g. gold, silver, bronze). Using the sharing grade parameter as input in our models we can support multi-level sharing patterns (i.e. canAccess(U, A, silver) see Algorithm 1).

The first sharing policy (*asset-centric sharing*) is based on the traditional ownership approach. It considers a model making resources either available for any partners to use or alternatively reserving them for the exclusive use of the owning partner. We experiment with different sharing levels by allowing collaborating organizations to share different proportions of the assets they own.

In typical multi-partner operations usually there are a number of smaller, more focused formations, which are dynamically created in response to an on the field event and execute missions for only a short time [10]. In the second sharing model (*Team-centric sharing*), we assume collaborating partners share none of their owning resources. Instead, following the *edge* model we introduce a mechanism of cross-partner formations (small, focused formations), which we call *teams* and allow users participating in the same team to share assets freely; therefore team members have access to all assets owned by any organization represented in the team. In this case we experiment with a variety of sharing levels by applying different degrees of homogeneous (comprise members from a single partner) & heterogeneous (comprise members from two or more partners) teams.

Below we present a formal representation of the sharing policies using the following predicates. Suppose U, U' are users affiliated with different partners, A is an asset, P is a coalition partner, and T is a team.

canAccess(U, A) == true **if** U can access asset A
hasPartner(U, P) == true **if** U belongs to partner P
ownsAsset(P, A) == true **if** partner P owns asset A
hasTeam(U, T) == true **if** U is member of team T

Asset-centric sharing policy:

canAccess(U, A) **if**
hasPartner(U, P) $\wedge$ ownsAsset(P, A)

Team-centric sharing policy:

canAccess(U, A) **if**
hasTeam(U, T) $\wedge$ hasTeam(U', T) $\wedge$ hasPartner(U', P)
$\wedge$ ownsAsset(P, A)

IV. MSTA-P Protocol

The initial MSTA distributed protocol (the reader is referred to [6] for more in depth description of protocol's algorithms) runs on two main entities, (1) the user devices (e.g. smart phone or tablet) and (2) the smart environment assets and it consists of two main stages:

The initial negotiation stage: the user devices respond to user generated tasks requests, compute the best set of assets which may satisfy the request, and distribute the generated bids to this optimal set of assets.

The bundle formation stage: the assets decide upon which bundle to join in order to serve a particular task, giving priority to the most important tasks to which they can provide the highest average utility

Each task in the protocol is characterized by an expiration time (i.e. a deadline within which the task must be supported by an assets bundle or alternatively must be dropped) and a duration time (i.e. time during which the task remains active on the field). Provided that available resources are scarce, we assume that a subset of created tasks will not be supported, which implies the need for dropped tasks. A task is considered dropped if there are no available resources to satisfy the task

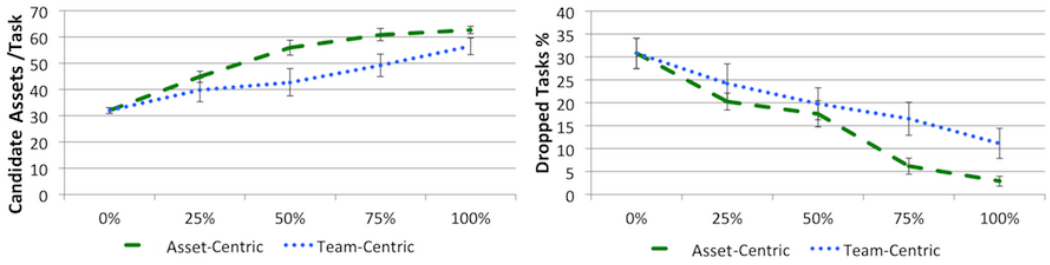

Fig. 1. Asset-centric vs. team-centric sharing models: effect on candidate assets per task and dropped tasks %.

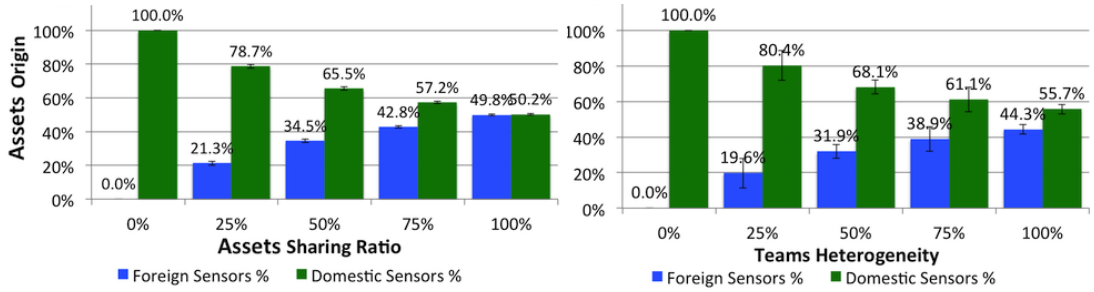

Fig. 2. Asset-centric vs. Team-centric sharing models: effect on assets origin.

utility demand in the initial negotiation stage, or if no resource can provide support to the task on time during the bundle formation stage. We refer to the set of these tasks as *dropped tasks*.

Algorithm 1 MSTA-P

```
for all A within SR do
  if canAccess(U, A) then
    addCandidateAsset(A)
  end if
end for
for all candidateAsset[A] do
  if canServe(A, T) then
    addBundle(B_AT)
    calculateUtility(B_AT)
  end if
end for
distributeBundle(B_AT)
```

Algorithm 1 performs the MSTA-P *initial negotiation stage* steps. The policies evaluation is the first process executed by the protocol after the task's creation. When the users create a new task, their devices query the assets within a sensing range SR of the area of interest in order to create assets bundle able to serve it. The sharing policies are considered at this step taking into account if the tasks' creators can access specific assets based on the sharing policies set by coalition partners (i.e. if *canAccess(U, A)* == true). As a result of the policies' evaluation is the creation of a list of assets that could be accessed by the task's creator. We call this list *candidate assets per task*. Therefore, the sharing policies affect the assets bundles creation by limiting the number of assets a user can access based on the applied sharing policies. In next section we use the *candidate assets per task* and the *dropped tasks %* variables as metrics for the sharing policies' evaluation.

V. Policies Evaluation & Simulation Results

For the implementation of MSTA-P protocol and the evaluation and comparison of the formalized asset sharing models we use REPAST Symphony[2], an open source agent-based, discrete time simulation environment. Implementing a multi-partner operation scenario we simulate a smart environment network, composed of 250 heterogeneous static "smart" assets and 50 mobile users which are randomly deployed on a 2D grid of 500m x 500m. Users and assets are equally distributed to each partner, while mobile users are free to move with no constraints following a *random waypoint mobility model*. In the experiments where the team-centric sharing model is tested, teams of users are formed at the first timestep and they do not change until the end. Each team has a team leader and consists of minimum two members while there is no members upper bound. At the simulations first timestep 25 out of 50 users are randomly selected (using java.util.Random, 48-bit seed) as candidate team leaders. Each one of them consecutively queries the nearby users within a *Team Range* (TR) radius of 100m in order to create a team. If the queried users are free do not belong to any team, the team leader adds them to their team. Each user belongs to maximum 1 team at a time, while there are users who do not belong to any team. In the experiments we use the most effective version of the MSTA-P the *Cost-driven Preemption* The task creation (i.e. users demand for asset services) rate is stable at arrival rate equal to 1 task per timestep and we repeat each simulation 10 times for 10000 timesteps averaging the measurements.

In order to have a complete picture in evaluating and comparing the two asset sharing models, we experiment in

[2]repast.sourceforge.net - last checked December 10th, 2013

asset-centric model by linearly increasing its sharing level starting from minimum 0% sharing ratio (none of the assets are shared with non-owning partners) and increasing it by 25% for each experiment until maximum 100% sharing ratio is reached. In team-centric model we start with the minimum sharing level 0% heterogeneous teams (all the teams in the field are homogeneous) and we increase it linearly by 25% for each experiment until reaching the maximum sharing level 100% heterogeneous teams. In essence, by increasing the degree of teams heterogeneity, indirectly we increase the overall shared assets but unlike the first model we do so through teams.

Figure 1 compares the two sharing models in terms of their effects on MSTA-P performance. In both models the starting point is the same because we start from minimum sharing levels. In asset-centric model when the sharing ratio is at its maximum (i.e. 100%) the number of candidate assets per task is twice as much as when the sharing ratio is at its minimum level (i.e. 0%). The difference in dropped tasks proportion is even larger where the total dropped tasks proportion is almost 8 times smaller. We also notice that by increasing linearly the sharing ratio, the number of candidate assets has quasi-logarithmic increase, while symmetrically the number of dropped tasks decreases quasi-logarithmically. Moreover, we observe that the difference of candidate assets per task and dropped task proportion moving from 75% to 100% assets sharing ratios is significant smaller compared to when we move to higher sharing ratios at lower sharing levels.

As for the team-centric model, in 100% teams heterogeneity case, the number of candidate assets per task almost doubles and the dropped tasks is three times smaller in comparison to when the degree of team heterogeneity is 0%. Moreover, in the team-centric sharing model we observe that by increasing linearly the team heterogeneity ratio we obtain a quasi-linear increase in the number of candidate assets per tasks and symmetrically a quasi-linear decrease in the number of dropped tasks. Overall, asset-centric seems to be more effective than team-centric model, especially in terms of dropped task %.

In the second set of experiments, represented by Figure 2 we measure the average proportion of "domestic" and "foreign" assets that are overall accessed by users. Domestic assets represent assets that have the same origin as the user to whom they provide services; foreign assets instead, represent those that are owned by different partner. We assume that the larger the proportion of the foreign assets is, the more efficient the sharing model is as well. The graph on the left presents the accessed assets' origin in asset-centric sharing model when we increase the assets sharing ratio and the one on the right presents the accessed assets' origin in the team-centric model when we increase the degree of team's heterogeneity. Both sharing models start with 100% domestic assets because we apply the minimum sharing level. Once again, the asset-centric seems to perform slightly better than the team-centric model in terms of foreign assets in use.

VI. Discussion & Conclusion

Although the asset-centric model performs better, there are additional reasons that might make team-centric model a more attractive option for multi-partner asset sharing. Through asset-centric approach users can only share their assets at partners level while the team-centric approach is a more agile sharing strategy giving them the ability for sharing at team level. Thus, by decreasing the overall number of users that can eventually access a resource, they also increase their privacy level. Furthermore, team-centric approach also increases the "quality" of users that can access a specific asset. In fact, as mentioned before teams are entities in which users are all focused towards a narrow specific objective; thus, by sharing resources with same team users it is usually beneficial provided that team members will be able to better carry on their work towards a common goal, and that they will most likely have complementary expertise.

We conclude that: (a) the asset-centric does not outperform the team-centric model's efficiency with a large margin, (b) we identify a cut-off threshold at 75% of assets sharing ratio, above which the MSTA-P protocol seems to perform without significant changes, (c) the team-centric model presents a more focused sharing approach and (d) the team-centric model presents a low-overhead sharing formation mechanism compared to asset-centric. Thus, the team-centric model seems to be an efficient asset sharing approach for highly dynamic multi-partner operations.

Acknowledgment: This research was sponsored by the US Army Research Laboratory and the UK Ministry of Defense and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the US Government, the UK Ministry of Defense or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] M. Weiser, R. Gold, and J. Brown, "Origins of ubiquitous computing research at parc in the late 1980s," *IBM Systems Journal*, vol. 38, no. 4, pp. 693–695, 1999.

[2] K. Su, J. Li, and H. Fu, "Smart city and the applications," 2011, pp. 1028–1031.

[3] "Abca: Coalition operations handbook," *ABCA Publication 332*, 2008.

[4] T. Frisanco *et al.*, "Infrastructure sharing and shared operations for mobile network operators from a deployment and operations view," in *NOMS. IEEE*, 2008, pp. 129–136.

[5] R. Hayes and D. Alberts, "Power to the edge: Command and control in the information age," *CCRP*, 2003.

[6] D. Pizzocaro *et al.*, "A distributed architecture for heterogeneous multi sensor-task allocation," ser. 7th IEEE, DCOSS'11.

[7] S. Ellison, C. Dohrmann, "Public-key support for group collaboration," *ACM TISSEC*, vol. 6, no. 4, pp. 547–565, 2003.

[8] C. Nita-Rotaru and N. A. Li, "Framework for role-based access control in group communication systems." in *In Proc. of International Workshop on Security in Parallel and Distributed Systems*, 2004.

[9] S. e. a. Firozabadi, "A framework for contractual resource sharing in coalitions," 2004, pp. 117–126.

[10] T. Pham *et al.*, "Intelligence, surveillance, and reconnaisance fusion for coalition operations," in *Proc 11th Fusion*, 2008.