Evaluating the Effect of Utility-based Decision Making in Collective Adaptive Systems

Vasilios Andrikopoulos[1], Marina Bitsaki[2], Santiago Gómez Sáez[1], Michael Hahn[1], Dimka Karastoyanova[1], Giorgos Koutras[2] and Alina Psycharaki[2]
[1]Institute of Architecture of Application Systems, University of Stuttgart, Stuttgart, Germany
[2]Transformation Services Laboratory, University of Crete, Heraklion, Greece

Keywords: Utility, Decision making, Collective Adaptive Systems, Choreography.

Abstract: Utility, defined as the perceived satisfaction with a service, provides the ideal means for decision making on the level of individual entities and collectives participating in a large-scale dynamic system. Previous works have already introduced the concept into the area of collective adaptive systems, and have discussed what is the necessary infrastructure to support the realization of the involved theoretical concepts into actual decision making. In this work we focus on two aspects. First, we provide a concrete utility model for a case study that is part of a larger research project. Second, we incorporate this model into our implementation of the proposed architecture. More importantly, we design and execute an experiment that aims to empirically evaluate the use of utility for decision making by comparing it against simpler decision making mechanisms.

1 INTRODUCTION

The concept of *Collective Adaptive Systems (CAS)* as defined by (Kernbach et al., 2011) and extended in (Andrikopoulos et al., 2013) encompasses systems that consist of heterogeneous entities, both physical and virtual, that collaborate towards satisfying their individual objectives, and the global objective(s) of the collective. Such systems underline the emergence, rather than the prescription of behavior through the interaction between entities that may be distributed geographically and organizationally. The dynamics of relationships between entities in CAS evolves over time as their constituting entities adapt to external or internal to the system environmental events. Large scale, complex, and dynamic systems like those required for the realization of smart cities can therefore be reflected and studied efficiently and effectively as collective adaptive systems.

Previous work proposed the use of the concept of *utility* from economics theory as the means to evaluate the satisfaction of the participating entities in a CAS, calculated as the perceived achievement of their individual objectives (Andrikopoulos et al., 2014b). Utility for each entity is expressed as a function over the preferences, context, state, and interactions of the entity with other entities in the system. The concept of utility is particularly interesting in the context of CAS since it provides the means for informed decision making on both local and global level. This is achieved by optimizing (that is, maximizing) the perceived satisfaction of the entities from a certain state of the system. Utility can be maximized for each participant entity individually, or as a group, introducing the notion of collective utility. This work focuses on the former case. A straightforward algorithm is proposed in (Andrikopoulos et al., 2014b) for this purpose, together with the necessary steps for implementing them in practice.

In this paper, we expand on the discussion in (Andrikopoulos et al., 2014b) by defining concrete utility functions for entities participating in the case study CAS used by the ALLOW Ensembles project[1]. We then show how the Cloud-enabled architecture proposed in the same work can be realized in order to support utility-based decision making in the CAS of the case study. The main focus of this work is on demonstrating the effect of this decision making mechanism in practice by emulating the behavior of the entities in the CAS using the architecture we developed in an experimental setting. This evaluation allows us to draw conclusions with respect to the efficacy of our

[1]http://www.allow-ensembles.eu/

Andrikopoulos, V., Bitsaki, M., Sáez, S., Hahn, M., Karastoyanova, D., Koutras, G. and Psycharaki, A.
Evaluating the Effect of Utility-based Decision Making in Collective Adaptive Systems.
In *Proceedings of the 6th International Conference on Cloud Computing and Services Science (CLOSER 2016) - Volume 2*, pages 39-47
ISBN: 978-989-758-182-3

proposal.

The contributions of this work can therefore be summarized by the following:

- A detailed utility model for participant-level decision making in an example CAS (Section 3).
- An experimental evaluation of the proposed utility model (Section 4) using the architecture proposed in (Andrikopoulos et al., 2014b).

In addition, the paper provides the necessary background (Section 2) which this work builds upon, discusses the State of the Art in utility theory and related works (Section 5), and presents our conclusions and plans for future work in Section 6.

2 BACKGROUND

As discussed in previous works (Andrikopoulos et al., 2013; Andrikopoulos et al., 2014b; Andrikopoulos et al., 2014a), the ALLOW Ensembles project uses the concepts of *entities, cells* and *ensembles* to model CAS. *Entities* represent both software and human actors by aggregating different functionalities performed or associated with this actor, expressed as reusable *cells*. Cells therefore encapsulate functionalities that the entity offers or requires from the system. The interaction of entities follows the interaction between their cells through their predefined functionalities and results into *ensembles*. Ensembles are therefore representing the emerging interactions between different entities in order to fulfill their individual goals. Despite the fact that each entity has its own objectives, the cooperation with each other should result into an increase of their perceived satisfaction in participating in an ensemble, expressed as an increase in each entity's *utility*. Deciding therefore whether to participate into an ensemble, or to which of one of the possible alternative ensembles that are available in the system, becomes a utility maximization problem for each entity. While in the context of the project we also discuss cooperative games (Wiese, 2010) where utility is aggregated into a group level, for the purpose of this work we scope the discussion to non-cooperative games, where each entity attempts to optimize for itself its utility.

The case study used in the ALLOW Ensembles project for demonstration and evaluation purposes, and therefore the one adopted by this work, is that of an *Urban Mobility System (UMS)* that acts as a smart city planner in a multi-modal transportation scenario. The city used as the model for this scenario is that of Trento, Italy. From the transportation means available in the scenario, of particular interest for the purposes of this paper is the use of utility for comparing between using one own's car versus the use of public transportation (that is, bus), which may also include some walking to reach the closest bus stop, and of course some waiting time until a bus arrives. The main entities under consideration in this case consist of the users of the UMS that want to travel from randomized points to other randomized points on the map of Trento in a given time interval. Depending on whether users decide to use the public transportation (and walk/wait as necessary) or to take their own car for this trip, the entities participate in either a bus route ensemble, or form their own ensemble, respectively.

In the following sections we discuss a) how the utility of the UMS users can be modeled, b) what is the necessary infrastructure for supporting the discussed scenario, and c) what kind of conclusions can be drawn about the effect of using utility in the decision making of the entities in this scenario.

3 UTILITY MODEL

In this section we extend the discussion in (Andrikopoulos et al., 2014b; Andrikopoulos et al., 2014a) to propose a utility model that describes the procedure for the calculation of the utility of an entity for participating in a specific ensemble (given as an alternative among various ones that fit the entity's request). The model comprises the following components:

1. Input/information required to calculate utility for an alternative (e.g. a route or a combination of routes):

   a. Variables that affect utility (e.g. travel time, travel cost, walking duration, etc.). They are specified by the designer of the utility model in accordance with the characteristics of the entity's request. These variables are random variables; their value is realized after the end of the execution of the ensemble. Their estimation is used in calculating values of utility beforehand. Let $\vec{v} = (v_1, \ldots, v_m)$ be the vector of variables that we are interested in. In our case study we consider the following variables: *travel time* (t), *travel cost* (c), *walking duration* (d), *and number of changes* (n).

   b. Entity's preferences including: upper and lower bounds of the values of variables (constraints) and weights for the variables that indicate their effect on the utility function (e.g. specify priorities with respect to variables: I prefer a cheaper trip than a faster trip). Preferences are specified by the entity and are part of the request. They are constant numbers for the whole life-cycle of the specific ensemble given by the user.

Let $\vec{a} = (a_1, \ldots, a_n)$ be the vector of preferences and $\vec{w} = (w_1, \ldots, w_m)$ be the vector of weights for the m variables defined above. In our case study we consider the following preferences: *maximum travel time* (t_{max})*, maximum travel cost* (c_{max})*, maximum walking duration* (d_{max}) *and maximum number of changes* (n_{max}).

2. Interdependencies between any two of the above variables in order to select the appropriate forms of utility functions. They are specified by the designer of the model. We consider two different ways for recognizing that two variables v_1, v_2 are interdependent:

   a. There is a known function f such that $v_1 = f(v_2)$. In this case we substitute v_1 for v_2 in the utility function and eliminate one of the two variables.

   b. For each value v_{1k} of variable v_1 there is an associated value v_{2k} of variable v_2. In this case, we calculate the utility function for pairs $(v_{1k}, v_{2k}), k \in K$ and apply a common weight for both variables.

   Let $\vec{v}_{ij} = (v_i, v_j)$ be the vector of the interdependent variables v_i and v_j.

   In our case study, the variables we use (see previous step) are independent of one another, thus, we have an additive model for the utility function as we will see in the next step.

3. Utility functions that are specified by the designer and reflect the entity's preferences. Let $u(\vec{v}, \vec{a}, \vec{w})$ be the utility function of the entity as a function of the variables, preferences and weights as defined above.

   a. In the case where no interdependencies exist then:

$$u(\vec{v}, \vec{a}, \vec{w}) = \\ = w_1 f_1(v_1, \vec{a}) + \ldots + w_m f_m(v_m, \vec{a}), \qquad (1)$$

   where $f_i, i = 1, \ldots, m$ is the component that contributes to the total utility related to variable v_i.

   b. In the presence of interdependencies, we have:

$$u(\vec{v}, \vec{a}, \vec{w}) = \\ = \sum_{k \in I_1} w_k f_k(v_k, \vec{a}) + \sum_{i,j \in I_2} w_{ij} f_{ij}(v_{ij}, \vec{a}) \qquad (2)$$

   where I_1 is the index for all independent variables and I_2 is the index for all pairs of interdependent variables. For example, given that $\vec{v} = (v_1, \ldots, v_7)$ and dependencies exist on pairs $v_{34} = (v_3, v_4)$ and $v_{67} = (v_6, v_7)$, then the utility is given by: $u(\vec{v}, \vec{a}, \vec{w}) = w_1 f_1(v_1, \vec{a}) + w_2 f_2(v_2, \vec{a}) + w_{34} f_{34}(v_{34}, \vec{a}) + w_5 f_5(v_5, \vec{a}) + w_{67} f_{67}(v_{67}, \vec{a})$.

   In our case study, the additive model in which no interdependencies exist is appropriate and therefore we use Equation 1 for the definition of utility functions.

Based on the above, we propose the following types of utility functions for the variables of interest, namely, travel time, travel cost, walking duration, and number of changes:

1. $u_1(t, t_{max}) = e^{-k_1 t / t_{max}}, t \geq 0$
   The component of utility that corresponds to time is a decreasing function of time. Parameter $k_1 (k_1 > 0)$ is a constant number that reflects the rate of the decrease of utility with respect to t_{max}.

2. $u_2(c, c_{max}) = e^{-k_2 c / c_{max}}, c \geq 0$

   Utility is a decreasing function of cost. Parameter $k_2 (k_2 > 0)$ is a constant number that reflects the rate of the decrease of utility with respect to c_{max}.

3. $u_3(d, d_{max}) = e^{-k_3 (d/d_{max})^2}, d > 0$

   This is a decreasing function of walking duration. The utility decreases slower in low values of d than in high values.

4. $u_4(n) = \begin{cases} \frac{n_{max} - n + 1}{n_{max}} & n = 1, \ldots, n_{max} \\ 0 & n > n_{max} \end{cases}$

   This is a decreasing function in the number of changes.

The utility of each user entity (as discussed in Section 2) can therefore be calculated as the weighted sum of $u_1, \ldots, u_4$:

$$u(\vec{v}, \vec{a}, \vec{w}) = \\ = w_t e^{-k_1 t / t_{max}} + w_c e^{-k_2 c / c_{max}} + \\ + w_d e^{-k_3 (d/d_{max})^2} + w_n \frac{n_{max} - n + 1}{n_{max}} \qquad (3)$$

where $w_t + w_c + w_d + w_n = 1$ for normalization purposes.

4 EVALUATION

The purpose of evaluation in this work is twofold: on one hand, we want to evaluate the efficacy of the proposed utility model, as discussed in Section 3, in decision making in the context of the CAS outline in Section 2; on the other hand, we also aim to evaluate the appropriateness of the architecture discussed in the previous section as the underlying infrastructure supporting the same CAS. For this purpose we use an experimental evaluation that *emulates* the actual operation of the CAS in a smaller, and more manageable sample of the population under the assumptions discussed in the following.

4.1 Experiments Description

We designed two experiments for evaluation purposes. In both experiments we assess the impact of utility in decision making by comparing the decisions taken for a fixed set of users when attempting a trip from a random point in Trento's map to another also randomly generated destination in Trento. Users have profiles with different preferences allowing them to prioritize the available transportation modes (public transportation or car) and routes in different ways. The experiments are performed for a set of one thousand users distributed in three different profile types: *workers, students*, and *pensioners*. All trip requests take place within a fixed interval during the morning of a working day. *Experiment A* measures which transportation mode was selected by each user when *a)* the utility model of Section 3 (and in particular, Equation 3) was used to decide the best option for each user, *b)* only the duration of the trip is taken into consideration (shortest is better), and *c)* only the cost of the trip is used (cheaper is better). *Experiment B* measures the effect of the bus fare price to the choice between public and private transportation by users when utility is used for decision making. For this purpose we reduce the bus fare in fixed decrements and we measure the selected transportation mode by each user as before. The setup for both experiments is described in the following.

4.2 Experimental Setup

In terms of infrastructure for our experimental evaluation, we extend the system architecture proposed in (Andrikopoulos et al., 2014b) and add the necessary components for generating and driving the resulting system with a representative load. The resulting system is summarized by Fig. 1. More specifically, the system discussed in (Andrikopoulos et al., 2014b) distinguishes between a *Modeling Environment* and a *Runtime Environment*. The Modeling Environment, implemented as an Eclipse Graphical Editor[2] allows the definition of cells and ensembles (as discussed in Section 2) as a set of *service orchestrations* and *choreographies*, respectively. WS-BPEL is used for the former, and the BPEL4Chor language for the latter, as discussed in (Andrikopoulos et al., 2014b). For the purposes of our evaluation, we used the Modeling Environment to design the ensemble and the respective cells allowing a passenger to inquire the UMS for traveling options between two points in a city. UMS replies with a set of route alternatives that include both *public*, i.e. buses and possibly walking, and *private*, i.e. car driving transportation modes. Decision making based on different policies, i.e. utility, trip duration, or cost is also modeled as a cell in the system, allowing the automation of the experiment.

The execution of the cells takes place in the *Execution Engine* component of the Runtime Environment, implemented based on the Apache ODE[3] Engine, an open source implementation of BPEL. The *Utility Module* in Fig. 1 implements the utility model discussed in Section 3 as a set of Web Services that are interacting with the Execution Engine through an *Enterprise Service Bus (ESB)*. More specifically, the Utility Model services accept a list of route alternatives and a unique entity identifier, representing each user of the UMS, and return the route alternative list ordered by their calculated utility for the specific entity. The entity identifiers, together with their profiles consisting of their preferences with respect to maximum traveling time, cost, etc. are stored in the *Entity Management System* component of the system, also implemented as a set of Web Services on top of a database for persistence purposes. The *Adaptation Manager* and *Monitoring* components in Fig. 1 are out of the scope of this evaluation, and therefore they will be omitted from the rest of the discussion; here they are presented for completeness.

The main difference between the system discussed in (Andrikopoulos et al., 2014b), and the one presented in Fig. 1 is the addition of a third aspect, that of *Load Generation & Driver*. This contains the components of *Entity Generator*, responsible for generating entities and their profiles for experimental use, the *Route Generator*, which produces the route alternatives for the possible trips, and the *Load Driver* that generates a load for the system emulating the behavior of the entities generated by Entity Generator using the routes produced by the Route Generator. In the following we discuss how we implemented these latter components, as well as the steps necessary for preparing and executing our experiment. W.r.t. the deployment of the infrastructure depicted in Figure 1, the components are distributed in on-premise, private cloud facilities of both institutions. The Modeling and Load Generation & Driver components, and Execution engine are distributed in separate machines in the University of

[2]Eclipse Graphical Editing Framework: https://eclipse.org/ gef/

[3]Apache ODE: http://ode.apache.org/

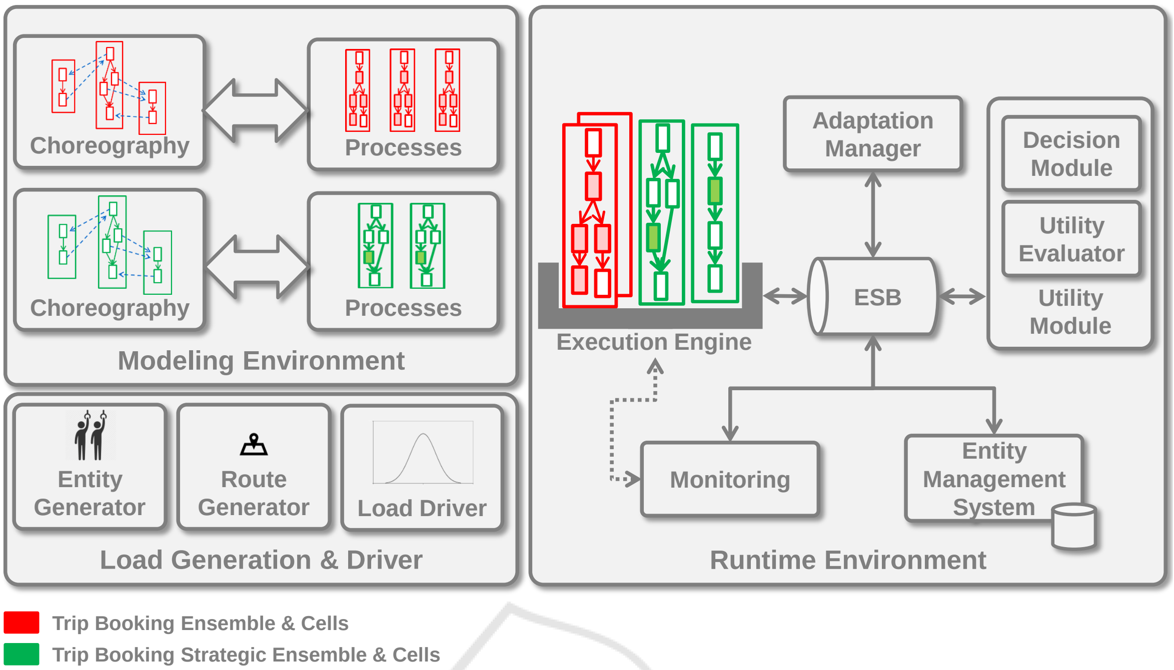

Figure 1: Architecture of the Experimental Setup (adapted from (Andrikopoulos et al., 2014b)).

Stuttgart[4]. The Utility Module, as well as the Entity Management System, reside in an on-premise infrastructure in the University of Crete[5].

4.3 Experiment Preparation & Execution

Entity Profiles

The Entity Generator component of Fig. 1 was created using the capabilities offered by the R language[6]. More specifically, for each randomly generated user identified by name/surname and a generated UUID, a profile was created by randomizing their preferences within acceptable value ranges for the variables under consideration (trip duration, trip cost, number of changes, walking duration, and their respective weights). Table 1 summarizes the ranges that we defined for this purpose, chosen to reflect the different characteristics of the defined user groups. As shown in the table, users were assigned to the student, worker, or pensioner group in a 0.3/0.5/0.2 ratio, respectively, and values for each of their preference components was generated based on a uniform distribution model within the ranges defined in Table 1. The created user profiles were then imported into the Entity Management System of Fig. 1 for persistence and retrieval purposes.

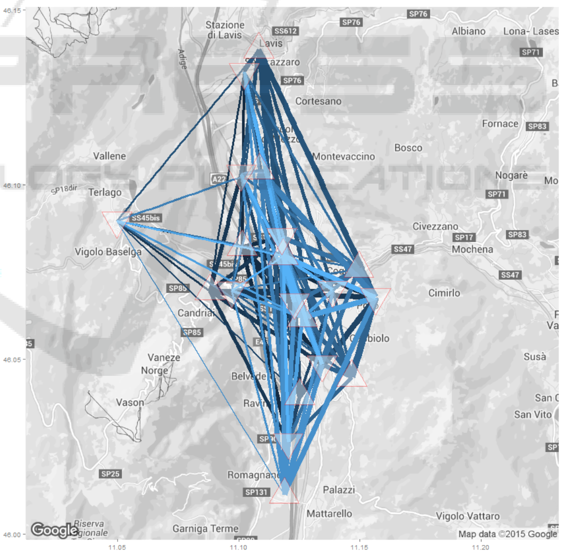

Figure 2: Trip distribution organized by population cluster center.

Trips

The next step for setting up the experiments was to distribute randomly the generated users on the map

[4]IAAS - University of Stuttgart: http://www.iaas.uni-stuttgart.de/

[5]TSL - University of Crete: http://www.tsl.gr/

[6]The R project for Statistical Computing: https://www.r-project.org/

Table 1: Value ranges for user profile generation.

	User Profile Type		
Preference	Student	Worker	Pensioner
Upper bound of duration t_{max} (minutes per km)	$[3,5]$	$[2,4]$	$[3,4]$
Upper bound of travel cost c_{fixed} (euros)	$[1,3]$	$[1,4]$	$[2,3]$
Maximum number of changes $n_{changes}$ (integer)	$\{2,3\}$	$\{2,3\}$	$\{1,2\}$
Maximum walking duration w_{max} (minutes)	$[10,20]$	$[10,15]$	$[5,10]$
Weight of travel time w_t	$[0.3,0.5]$	$[0.5,0.7]$	$[0.2,0.4]$
Weight of cost w_c	$[0.5,0.9]$	$[0.4,0.8]$	$[0.3,0.5]$
Weight of walking duration w_d	$[0.1,0.3]$	$[0.2,0.5]$	$[0.4,0.6]$
Weight of number of changes w_n	$[0.2,0.4]$	$[0.3,0.6]$	$[0.2,0.3]$
Ratio in population	30%	50%	20%

of Trento, roughly following the population distribution in the actual city, and by choosing an origin and a destination point for each. In order to reduce the complexity of the experimental setup, remove possible performance perturbation due to uncontrolled network latencies, and avoid unnecessary for our purposes development efforts we reduced the possible trips into a smaller possible set and pre-calculated the route alternatives for each trip. For this purpose we started by clustering (more specifically: k-means clustering) the origin and destination points for each user into two sets of 10 clusters each. The resulting origin and destination cluster centers are shown in Fig. 2 as upward- and downward-facing triangles, respectively. The width of the straight lines between cluster centers in Fig. 2 denotes the relative density of this trip in numbers of entities per trip (cluster center pair).

We then created the Route Generator component of Fig. 1. For this purpose, we used the integration offered by the ggmap package of R with the Google Maps API[7] to retrieve, and then re-format appropriately and consequently persist route alternatives for the 10×10 possible trips between cluster centers. For most of the trips, 4 route alternatives combining bus and walking, and 3 route alternatives for using a car were produced in this way. A total of 611 possible routes were produced in this way. The estimated time and distance provided by the Google Maps service was also attached to each route alternative. This allowed us to calculate an estimate for the cost of each trip, as follows:

1. For routes by car, we used the gas consumption of a reference popular car model (Fiat Panda 2014[8]) and an indicative price for gas for the Trento area[9] to calculate the cost as
$$c_{driving} = \frac{\text{distance(km)} \times \text{gas price per liter}}{\text{kilometers per liter}}.$$
2. For routes combining buses and walking, we amortized the cost of public transportation tickets in Trento (0.70€ for 70' of traveling, as of September 2015) into a linear cost model and used only the time spent on bus segments to calculate the cost of the route alternative as $c_{public} = \sum_{i}^{\#segments} \text{duration (sec)} \times \frac{0.7}{70 \times 60}$.

These costs were persisted together with the rest of the trips information also in the Entity Management System of Fig. 1. For the purposes of Experiment B, we recalculated the cost of the route alternatives involving public transportation by reducing the original bus fare by 10% in each decrement until 50% of the original fare is reached. The cost for the respective trips was updated in the Entity Management System before each iteration of the experiment as appropriate.

System Load

For the system load used by the Load Driver (Fig. 1) that drives our experiments we decided that it needs to be as close as to the real-word as possible. For this purpose, we modeled the load of the system using a Poisson distribution with arrival rate $\lambda = 5$, which allows us to spread the 1K possible trips (one trip per person) into 15 time intervals using the rpois command of R for generating sample values of this distribution. Figure 3 illustrates the number of requests per time interval. We created intervals of 10 seconds in between each group of concurrent requests (meaning that e.g.

[7]Google Maps API: https://developers.google.com/maps/?hl=en

[8]Based on the statistics available in http://www.statista.com/statistics/417585/italy-leading-car-brand-sales/

[9]That is, 1.54€ per liter on average in September 2015, according to http://www.numbeo.com/cost-of-living/

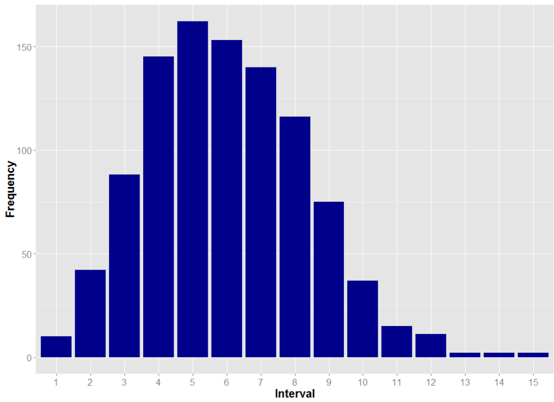

Figure 3: Trip requests by time interval ($\lambda = 5$).

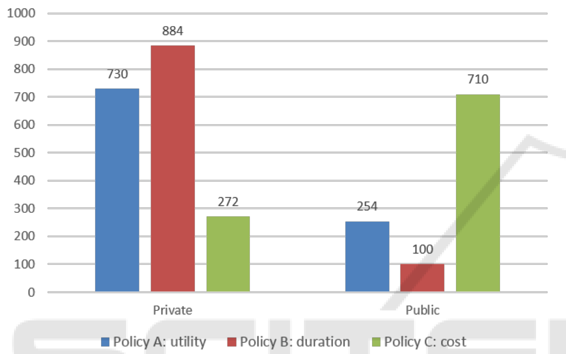

Figure 4: Experiment A — Distribution of transportation modes per policy.

10 requests were sent simultaneously at the beginning of the experiment, 42 after 10 secs, etc.), and that of 15' in the emulation starting from 06:00 on the 1st of October 2015 (i.e. 10 users sent their trip request at 06:00, 42 at 06:15, etc.).

Apache JMeter[10] was used as the load driver for the system. We created a JMeter *Test Plan* comprising the collection and processing functionalities for the load provided by the Load Driver, and the aggregation of the results obtained from the execution of each ensemble. System performance metrics were also measured and plan to be reported in future reports.

4.4 Findings

Figure 4 summarizes the results of Experiment A in terms of the decisions between route alternatives based on private or public transportation for each entity in the load, and for the different policies applied. As it can be seen from the figure, when entities make decisions solely based on the cost of the trip, the majority (72.3% of the population) opts for using public transportation. While this result is sensitive to the cost model attached to the routes entailing bus usage, it confirms the intuition that public transportation is indeed cheaper on average. On the other hand, selecting transportation mode based on trip duration has the opposite result, with almost 90% of the population opting for driving their own cars to their destination. Finally, as shown in Fig. 4, using utility essentially balances the effect of cost and duration to decision making with 25.8% of the population choosing for public transportation. However, cost still seems to play the most important role in the decision making. This can be largely attributed to the generated profile values for the weight coefficients in Equation 3 in comparison to the available route alternatives. Evaluating the effect of different profiles on the decision making is currently ongoing work.

Figure 5 presents the results of Experiment B on the effect of reducing the bus fare cost by 30% and 50% to the decision making of entities in the system, when utility is used as policy. As expected, reducing the bus fare results into more entities choosing for public transportation; however this increase into public transportation is linear in nature (Fig. 5(a)). Furthermore, the rate of increase in terms of utility (Fig. 5(b)) for entities choosing public transportation actually decreases as the cost is further reduced, since the other parameters in the model (total duration, walking duration, number of changes) start affecting the entity's utility mode more. As a result, when considering the trade-off between passenger utility gains and transportation company revenue loss it can be concluded that the 30% reduction to bus fare would be preferable for all parties involved.

5 RELATED WORK

Utility-based decision making is a well-known economic study field typically used in decision making mechanisms, e.g. in multi-attribute utility theory (Keeney and Raiffa, 1976). Focusing on the scope of this work, we leverage the usage of utility as a the fundamental mechanism for measuring the satisfaction and the decisions taken by travelers in a multi-modal transportation system, e.g. based on their commuting time, city pollution, etc.

In the scope of multi-modal transportation, several studies have been conducted for modeling commuting time and analyzing congestion management strategies, including travelers departure time choice, route choice or mode choice (Lam and Small, 2001; B. Johansson and Olsson, 2003; Li and Huang, 2005). In (Lam and Small, 2001), a method to measure how travel time and its reliability are valued by travelers is proposed, e.g. taking into account the value of time for a trav-

[10] Apache JMeter 2.10: http://jmeter.apache.org/

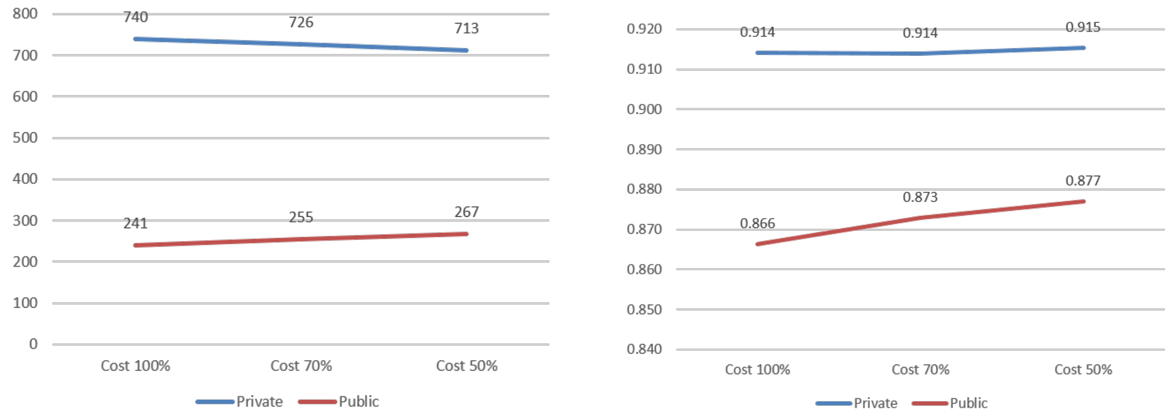

(a) Number of private/public transportation mode selections as a function of bus fare cost

(b) Average utility as a function of bus fare cost

Figure 5: Experiment B — Effect of bus fare cost reductions to decision making.

eler as the rate of change in utility. In (B. Johansson and Olsson, 2003), the labor market commuter behavior is analyzed in terms of the number of traveled short, medium and long time distances. In (Li and Huang, 2005), the reliability of morning commuting in congested and uncertain transport networks is investigated. Focusing on non-cooperative games, i.e. how commuters independently choose their optimal routes, (Sun and Gao, 2007), (Anas and Berliant, 2010)), and (Sun and Gao, 2007) examine how commuters choose their optimal routes and trip modes using non-cooperative games. Our work defers w.r.t. current state of the art in the following aspects: (i) we provide services that are shared among various users in the system taking into consideration the interdependencies among them, (ii) the decision making mechanism in our system is a distributed process performed by the users themselves using private information (value of utility) and sharing common parts of information (route characteristics such as travel time and cost), and (iii) our approach deals with the evaluation of services offered to users that are not predefined, but rather customized according to user preferences.

The definition of such services and decision making mechanisms as interdependent distributed processes, however, requires further analysis on existing modeling, provisioning, execution, and adaptation techniques in large scale CAS. For instance, typical interactions between multiple participants in a choreography can be modeled following the interaction or the interconnection modeling approaches (Barker et al., 2009). In the former, communication between participants is modeled using atomic *interaction activities*. The WS-CDL (Kavantzas et al., 2005) and the Savara[11] project are approaches that support the explicit specification of the *interaction activities*. On the other hand, the interconnection modeling approach consists on interconnecting the communication activities of each participant of the choreography. This approach is supported in the CHOReOS Integrated Development and Runtime Environment[12], in the Open Knowledge European project[13], and in BPEL4Chor (Decker et al., 2007). As BPEL4Chor enables the choreography specification atop of WS-BPEL and decouples the choreography behavior specification from the technical communication details, it is considered as an appropriate extensible point in CAS modeling and specification. The CoDyCo framework has been developed in the scope of the ALLOW Ensembles project supporting the modeling, provisioning, and execution of large scale interactions among multiple entities (Gómez Sáez et al., 2015). Such a framework has been completely reused and extended towards integrating performance and utility aspects for the experiments depicted in this paper.

6 CONCLUSIONS AND FUTURE WORK

The previous sections use the case study of an urban mobility system discussed in an EU project to demonstrate the use of the concept of utility in decision making in large-scale systems like this of the case study. Utility as the perceived satisfaction with a service is used here to quantify the degree of satisfaction of the

[11]http://www.jboss.org/savara

[12]CHOReOS: Large Scale Choreographies for the Future Internet: http://www.choreos.eu/

[13]Open Knowledge: http://www.openk.org/

entities with the system, i.e. the users asking the system to provide them with the optimal option for transportation between points in the city. In this case, utility is expressed as a function over a set of preferences and context attributes related to e.g. the duration and cost of each route alternative that the system can produce. As it is shown by the experimental evaluation based on an emulation of the actual urban mobility system using synthetic data but the actual infrastructure for it, utility-based decision making produces results that depend heavily on the preferences of the entities in the system.

Future work aims to further evaluate the observed effect of user preferences to utility calculation in the same experimental scenario. Furthermore, we intend to introduce and evaluate similarly a cooperative game formulation to the case study, taking into account, for example, the delays to the trip depending on the number of passengers joining a bus route.

ACKNOWLEDGEMENTS

This work is partially funded by the FP7 EU-FET project 600792 ALLOW Ensembles and the German 610872 DFG Project SitOPT.

REFERENCES

Anas, A. and Berliant, M. (2010). The commuting game. *Unknown Journal*.

Andrikopoulos, V., Bitsaki, M., Bucchiarone, A., Gómez Sáez, S., Karastoyanova, D., Leymann, F., Nikolaou, C., and Pistore, M. (2014a). A Game Theoretic Approach for Managing Multi-Modal Urban Mobility Systems. In *Proceedings of the 5th International Conference on Applied Human Factors and Ergonomics (AHFE 2014)*, Kraków, Poland. CRC Press/-Taylor & Francis.

Andrikopoulos, V., Bitsaki, M., Gómez Sáez, S., Karastoyanova, D., Nikolaou, C., and Psycharaki, A. (2014b). Utility-based Decision Making in Collective Adaptive Systems. In *Proceedings of the Fourth International Conference on Cloud Computing and Service Science (CLOSER'14)*, pages 308–314. SciTePress.

Andrikopoulos, V., Bucchiarone, A., Gómez Sáez, S., Karastoyanova, D., and Mezzina, C. A. (2013). Towards Modeling and Execution of Collective Adaptive Systems. In *Proceedings of WESOA'13*, pages 1–12. Springer.

B. Johansson, J. K. and Olsson, M. (2003). Journal of geographical systems. *Commuters? non-linear response to time distances*, 5:315–329.

Barker, A., Walton, C. D., and Robertson, D. (2009). Choreographing web services. *IEEE Transactions on Services Computing*, 2(2):152–166.

Decker, G., Kopp, O., Leymann, F., and Weske, M. (2007). Bpel4chor: Extending bpel for modeling choreographies. In *Proceedings of ICWS'07*.

Gómez Sáez, S., Andrikopoulos, V., Hahn, M., Karastoyanova, D., and Weiß, A. (2015). Enabling Reusable and Adaptive Modeling, Provisioning & Execution of BPEL Processes. In *Proceedings of the 8th International Conference on Service-Oriented Computing and Applications (SOCA'15)*, Rome, Italy. IEEE Computer Society.

Kavantzas, N., Burdett, D., Ritzinger, G., Fletcher, T., Lafon, Y., and Barreto, C. (2005). Web services choreography description language version 1.0.

Keeney, R. and Raiffa, H. (1976). Decisions with multiple objectives: preferences and value tradeoffs. *Cambridge University Press*.

Kernbach, S., Schmickl, T., and Timmis, J. (2011). Collective adaptive systems: Challenges beyond evolvability. *ACM Computing Research Repository (CoRR)*.

Lam, T. C. and Small, K. A. (2001). The value of time and reliability: measurement from a value pricing experiment. *Transportation Research Part E: Logistics and Transportation Review*, 37:231–251.

Li, Z. and Huang, H. (2005). Fixed-point model and schedule reliability of morning commuting in stochastic and time-dependent transport networks. In *WINE 2005*, pages 777–785.

Sun, L. and Gao, Z. (2007). An equilibrium model for urban transit assignment based on game theory. *European Journal of Operational Research*, 181:305–314.

Wiese, H. (2010). Applied cooperative game theory. http://www.uni-leipzig.de/∼micro.