Towards an Agent Coordination Framework for Smart Mobility Services

Andrea Sassi and Franco Zambonelli
Dipartimento di Scienze e Metodi dell'Ingegneria
Università di Modena e Reggio Emilia, Italy
name.surname@unimore.it

ABSTRACT

Smart and social mobility services will soon hit the streets of our cities. However, most of existing solutions so far are built through different operations that don't lie on the same processing flow, neither don't share with each others their input data streams. The understanding of how to design a general-purpose framework, supporting a variety of integrated services and promoting direct users involvement, is still missing. In this paper, we first show our conceptual vision of smart mobility services, focusing on the cooperation and interoperability of the actors involved. We then analyze the infrastructural requirements to enable such smart mobility services and present the characteristics of a general-purpose framework for the provisioning of smart mobility services, conceived as a distributed and open agent coordination infrastructure. To exemplify, we show how the framework can be applied in the context of an urban ride-sharing service.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Coherence and coordination, Multiagent systems

General Terms

Algorithms, Design

Keywords

Socio-technical System, Pervasive Computing, Smart Mobility Services, Agents Coordination, Ride-sharing

1. INTRODUCTION

The dramatic progress in embedded and mobile computing technologies, smart phones in primes, along with the pervasive diffusion of social networking tools, let us envision the emergence of a dense networked ICT infrastructure. In such infrastructure, coordinated human agents (i.e., the citizens) and software/hardware agents will interact with each other in such infrastructure so as to serve – at the same time – individual-level and urban-level goals, as if they were part of a single socio-technical system.

The overall behavior of such system will be driven by a variety of urban services which aim to improve the overall quality of life of individuals by providing them with tools to better interact with the urban environment, and also by shaping the activities of the urban environment itself, to suit their own needs.

One can consider a completely distributed software architecture deployed over individuals on their smart phones and over hardware sensors and actuators. However, a centralized entity able to continuously monitoring and redirecting the behavior of the agents will facilitate the dealing with city-scale problems.

The future pervasive urban services will be supported by bringing at work together the complementary sensing, computing, and actuating capabilities of the interconnected agents, and by closing them in a feedback loop (see Figure 1). After an initial learning phase in which raw data from sensors are collected, processed and classified, the agents will be skilled with context inference and anticipatory computing capabilities, as examples, and they will suggest tailored recommendations to the hardware actuators and to themselves. Closing these capabilities in a loop lets measure the goodness and the adoption rate of the suggested recommendations, by making clear their causal relation with the effects they generate. The process results in the generation of awareness, which can describe both individual and collective characters, related respectively to single agents and to a collection of those [3].

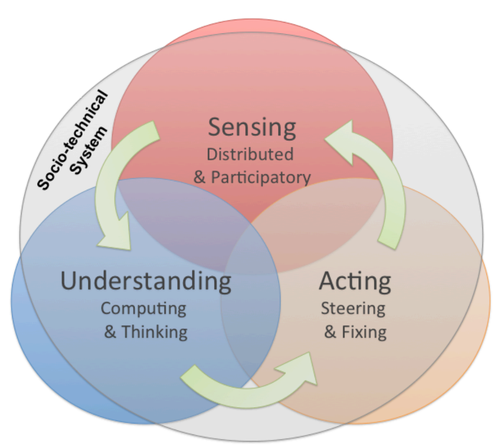

Figure 1: The sensing-understanding-acting feedback loop enabled by agents coordination.

This work focuses on smart mobility services enabled through the sensing-understanding-acting activities of the agents in the improvement of urban mobility. That is, to increase the effectiveness of individual mobility while at the same time improving the overall urban mobility (see Figure 2).

Human agents will play a fundamental role in the deployment of mobility services, since they can act both as consumers and as providers (e.g., via their private cars or simply by supplying information) of the services. Human social interactions can be pushed through a precise dynamic orchestration of the enabled data streams coming from both humans and ICT devices.

Our contributions are grounded on presenting how, in the scenario introduced above, the provisioning of integrated smart mobility services, can be effectively realized by a specifically suited coordination framework. Such coordination framework will be proposed as capable of supporting the iterative closed process of:

- Detecting mobility events related to the moving agents on the infrastructure, by harnessing the surrounding portion of the mobility data network shaped by the infrastructure itself, and also by processing the stream of incoming requests for mobility services;
- Identifying the possible solutions to satisfy expressed mobility needs based on the current state of things and of requests; anticipate future situations and future (or latent) mobility needs;
- Putting in act the necessary actions on actuator agents, or persuade human to act in certain ways, so as to end up realizing a coherent and sustainable set of services to satisfy the recognized needs.

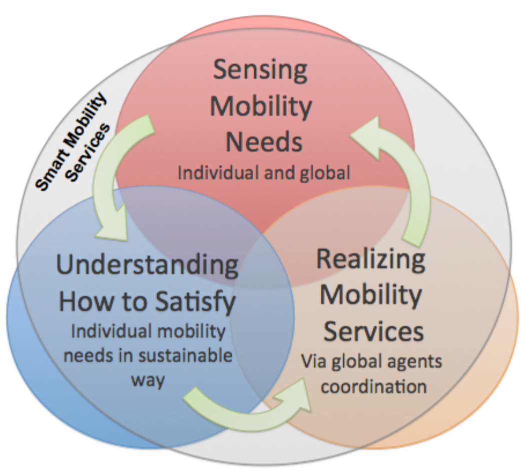

Figure 2: Smart mobility services enabled by agents coordination.

We believe, a general-purpose coordination framework that supports the shaping and the provisioning of smart mobility services will represent a powerful tool for urban designers and city administrators to make urban mobility services more efficient in terms of cost of the infrastructure (by harnessing the same sensors and actuators that self-reconfigure themselves upon specific requests, as well the same software architecture) and amount of data collected and processed (by sharing them among services with different purposes, instead of replicating similar operations for each service). Furthermore the provisioning of integrated solutions for different mobility needs can increase their individual adoption rate, towards the aim of reaching a critical mass of users, and thus can increase their effectiveness.

The contributions of this paper are to introduce our conceptual vision of what smart mobility services can be (Section 2), identify a set of infrastructural requirements for a general-purpose agent coordination framework (Section 3), sketch a conceptual model of the coordination framework for smart mobility services (Section 4), and introduce a use case in the area of ride-sharing (Section 5).

The paper also shortly discusses related works (Section 6) before concluding (Section 7).

2. SMART MOBILITY SERVICES

2.1 From ITS to smart mobility services

The recent dramatic progresses in ICT technologies, have led to the emergence of a very broad area of research in *Intelligent Transportation Systems* (ITS). ITS, in general, represent the most advanced way to establish a real-time transportation management, and consists in harnessing ICT technologies to better address users mobility needs and to support urban authorities decisions [1, 29].

ITS aim to improve urban transport performance, and can address in turns the problems and issues of pedestrians, cyclists, private vehicles, public transports, and roadside infrastructures. However, the application of ITS is often limited to the provisioning of on-demand web-services, with little or no interactions between users and contributions from user themselves. Furthermore, ITS do not offer a unified and integrated approach to support urban mobility in all its aspect, and often they own independent approaches for different mobility needs.

In general, the shift from ITS to *smart mobility services* must pursue the desired comfort for citizens and the satisfaction for urban authorities at the same level, by improving traffic efficiency and road capacity on the transportation network at an integrated, global, level. The services focus to impact on the development of increased social participation of citizens, where they are no longer simply requestors of mobility services, but can in turn play a role in the provisioning of services. Such an endeavor can feed cooperation and sharing practices with incentives and regulations.

Smart mobility services consist of all the mobility solutions enabled by pulling data from the available set of agents, generating higher information out of them, and enabling potential social interactions between a set of agents. The utility information is returned to them in such a way as to reinforce their interaction.

Citizens with mobility needs receive recommendations built on the matches with the services provided by other citizens, thanks to the supporting ICT infrastructure. Such recommendations can be strengthen if users have a similar profile, especially in terms of collaborative behavior. Data from social networks can detect social communities with same interests and mobility habits [5, 33]. The system will monitor the eventual adoption of the recommendation, and its effectiveness (was the service actually available?). Finally, it will update the profiles of the involved agents, to provide more useful recommendations in the future.

2.2 Example of smart mobility services

Let us now see some examples of such smart mobility services:

- **Parking Match.** A driver is approaching her destination and tries to find a vacant parking space. Some time earlier, another driver has left a parking lot in the same area. A parking match takes place and the driver is reached by a parking recommendation. Data involved in the matching process can come directly from the users involved, from the parking sensors installed on the infrastructure, or on users vehicles [22, 21].

- **Itinerary Match.** Consider the concurrent presence of the same users in a given set of locations at different times. When a spatio-temporal analysis on the data reveals that such co-location happens regularly (as seen in [8, 16]), it identifies a possible pool of commuters that make similar trips. The system should persuade them to switch to carpooling, making them aware of the benefits they have. Available carpooling services show how struggling is to reach a critical mass, hence social incentives are crucial (some carpooling issues are presented here [14]).

- **Taxi Match.** A taxi is hailed on the street by a person. While the driver is moving towards client's destination, he shares his route with other people that are looking for a ride (as described here [20]). If someone with a compatible trip ask for a ride, then taxi service becomes shared. Thus, its cost is lower for the clients and the revenue increase for the taxi driver. This service could seem similar to the previous one, but it mainly differs in terms of how the matches take place. The Itinerary Match mainly evaluates historical trips and habits, the latter considers real-time data.

- **Multimodal Rides Match.** A person explicitly declares a destination from her starting location, asking for directions. A selection of a spatio-temporal portion of data streams occurs. Multimodal directions can be provided to reach that destination. Current traffic level and rides availability (from multiple means of transports) on the transport network is evaluated and several complex pattern matching mechanisms are put in place to shape the best multimodal way to reach the destination. Several approaches come from Operational Research [10, 4]. In [9], authors have considered ride-sharing as a complementary solution to usual means of transports in multimodal trip planning.

- **Chaperone Match.** Parents cannot bring their children to school every morning and they might find difficult to bring them back home when classes are over as well. When no other relatives or friends can look after a child, one can consider to share the path the child is going to follow, at a certain time, to look for someone that takes charge of assessing the presence of the child at intermediate checkpoints (e.g., a bus stop, a crossing, a public display, a store). Hardware sensors and reliable citizens located close the checkpoints can act as proximity probes and thus they can send actual feedback in real-time to the parents, and of course they send alerts when an unexpected event will occur.

The above examples in any case see a clear distinction between provider and requestor of a service, and consider that providers of a service are not influenced by the request. However, in a really integrated system, the mean to provide a service can be dynamically shaped upon the request, in a process of mutual influence. Indeed, those who provide a service is because they have a need to satisfy. It is thus possible to let the distinction between requestor and provider vanish, and dynamically adapt the shape of services depending on the need, also with some supra level objectives in mind behind the opportunistic self-interest of the involved parties.

3. INFRASTRUCTURAL REQUIREMENTS

Next generation smart mobility services should be pursued by settling some infrastructural requirements on its components. These requirements can determine the technical viability of smart mobility services deployment.

Interconnection. Based on the Internet of Things paradigm [2], the agents that populate the urban environment need to be connected and able to exchange messages each others. The distributed network of humans and ICT-devices will enable sensing, computing, and actuating capabilities only if information can flow seamlessly among a defined set of entities, despite network dynamics, and made ephemeral.

Heterogeneity. The inter-connected components of the ICT infrastructure are highly heterogeneous. This feature has not to be considered its weakness. We have to take advantage of their complementary role in knowledge mining. As example, one can consider a fixed entity on the roadside acting as a traffic sensor (e.g.: smart traffic light, CCTV camera). The data collected can be enriched with the one provided by mobile agents (e.g.: pedestrian, cars, buses), and hence its interpretation is made easier. Events detection and anticipation accuracy can improve as well.

Interoperability. Interoperable agents encourage combination of concurrent data streams from different locations, enabled in precise spatio-temporal patterns. Our coordination framework is based on the orchestration of such different data sources, dynamically selected due their complementary role, according to the incoming requests. Nevertheless, energy saving and classification accuracy should imply specific conditions that drive the concurrent activation of certain data sources and classifiers as well.

Individual tasks. Each agent has to share her knowledge among a collection of agents that provides complementary skills to her ones, in order to (i) "measure" the context of the surrounding environment, (ii) infer a certain situation, so become aware that is happening something relevant, (iii) and finally adapt the behavior of the actuators accordingly. Human actions and interactions are crucial during the whole process, and they can be tracked by explicit or implicit sensing of data through both personal devices like smart phones or smart vehicles, and through public interactive displays.

Collective intelligence. The brain of the system needs a software architecture designed by balancing a top-down and a bottom-up approach. The first usually results in very predictable and measurable systems that lack in reactivity in high dynamic contexts. The latter suits to cope with pervasive computing in decentralized systems, which their behavior is not always predictable, nor easy to be engineered. Collective intelligence can emerge from the reasoning and the collaborations among decentralized agents that aim to process individual and collective contents.

System safety. The system should own only a finite set of reachable states, which should be known in the design phase, and tested during the development. The aim here, is to avoid risks related to the eventual system's evolution towards uncontrolled situations. To enable this feature we need to own a deep understanding of system dynamics, and how to deal with them. In other worlds, citizens should feel safe to contribute in social collective intelligence initiatives, because they trust the system and its potentials, and find it useful in any circumstance.

Information propagation. The inferred information should pervade the nodes of the infrastructure till it can actually reach any potential agent that can be interested in it. Data should be packed in efficient structures, and routed via peer exchanges. A middleware architecture can be harnessed to reach these goals, and thus to support the purpose of the coordination framework, which is expected to become active supporter of agents interactions and facilitator of information propagation [7].

Data management. Big amount of spatio-temporally distributed and heterogeneous data will be concurrently evaluated by computing-enabled devices, at different stages. Thus, efficient storage, querying, and analysis practices are needed. Academic literature offers as many cues as many approaches it presents ([23], and [19] among the others), but a unified best practice is missing.

Users privacy. Discovering matches between needs and services implies computation on sensitive data coming from the set of agents. This task can contemplate the sharing of confidential information among them. Privacy concerns and sharing policies must be dealt on user agreements and should consider innovative practices to balance the value of the data shared with the value expressed by service enabled through the sharing of someone else [13]. One should be able to opt-out from collecting certain data once they could evaluate the purpose of that collection, the sharing rules, and the service(s) that could be enabled thanks to it.

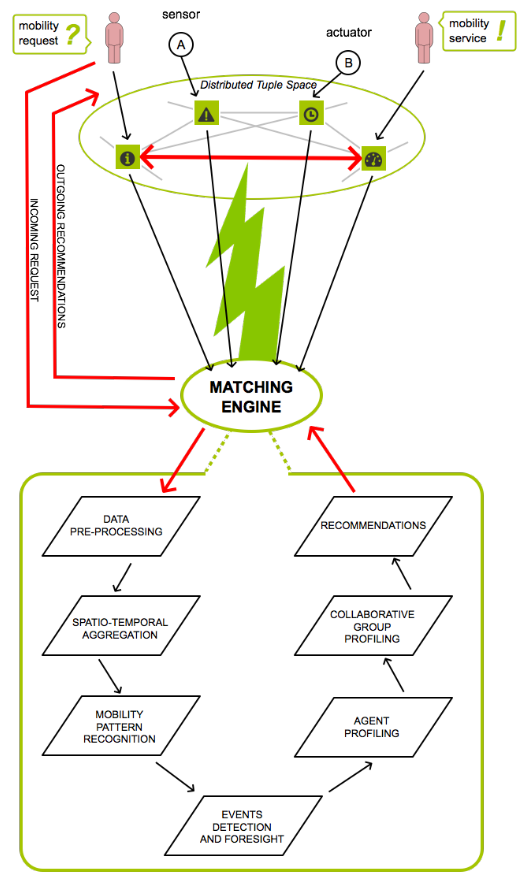

Figure 3: Conceptualization of the coordination framework that matches mobility services with mobility requests. Agents on the mobility network include humans and ICT sensors-actuators such as smart traffic lights (which count the approaching vehicles) and smart signals (which change the displayed information). Data streams are dynamically selected and processed through the matching engine.

4. CONCEPTUAL MODEL

As shown in Figure 3, the framework grounds on a matching engine that processes several data streams from a dense distributed tuple space, which is made of information concerning mobility status, requests, and services, generated by the agents on the mobility network. The rationale of the matching engine is triggered by incoming mobility requests, which in turns drive continuous processing steps. After several computing iterations on the available relevant information, the matching engine discovers and builds services on the mobility network, which finally result in mobility recommendations for the requesting user.

4.1 Distributed Tuple Space

Agents on the mobility network can implicitly or explicitly generate contextual information related to their mobility status, requests, and services. We believe a middleware infrastructure based on a set of networked tuple spaces [28] could represent a viable and suitable solution to store and share knowledge among all the agents interested in some particular generated contents, as well to properly feed the matching engine with the necessary information.

In particular, in the current demonstrative implementation of our infrastructure, we have built our coordination framework by exploiting the SAPERE tuple-based infrastructure [31, 32]. SAPERE has the following characteristics that make them suitable to implement our proposed coordination infrastructure;

- It integrates an advanced and semantic pattern matching mechanism which can act as the basic building block to realize advance matches between mobility requests and offers;

- SAPERE defines a context-aware and spatial tuple space model, where one can adopt context-aware and spatial rules to dynamically select, evaluate, and propagate information, which is particularly suited to the area of mobility;

- SAPERE can associate specific middleware agents to react to events occurring in the network of tuple spaces, which can be used (and has been used, indeed) to realize advanced and multifold matching mechanisms, as described in the following subsection.

4.2 Matching Engine

The matching engine concept can be described through the definition of a set of sub-activities, each of which has been implemented as a SAPERE middleware agents. A description of the iterative phases that compose the matching engine follows.

Data pre-processing and spatio-temporal aggregation. At a first place, incoming data is filtered, cleaned and aggregated. The process of course needs a considerable amount of data, collected over time, until this activity results in meaningful content for the engine. At further iterations, each incoming raw data will be filtered, cleaned and aggregated again, according to spatio-temporal constraints of the incoming request.

Data modeling supports upper-level meaning abstractions, by generating complex data structures useful to understand a special mobility pattern of the considered agents.

Mobility pattern recognition and events detection and foresight. Machine learning techniques enable regular patterns identification and anomalies detection on the aggregated input data. Really well trained classifiers can perform effective anticipatory computing [24] that can be crucial in dynamic environments.

Not only agents on the move own mobility patterns (inferred, as example, by mining their mobility routes from GPS data). Roadside sensors can shape the mobility status on the mobility network as well, and so they let creation of tuples that characterize the mobility context of a geofenced area in a specific time interval. So, it is clear that it will be possible to detect and anticipate the occurrence of significant mobility events and have a real-time distributed representation of them.

Agent and collaborative group profiling. Each entity on the network is characterized by its own capabilities, which let it play specific activities with proper tasks, which are, in turns, driven by the nature of the entity itself. These conjectures bring the necessity to model agents behavior to the foreground (the Belief-Desire-Intention (BDI) model [26] is one of the approaches suggested by the literature in this field). The distributed tuple space should be populated with profiling contents related both to individual agents and to groups of them.

Interactions among group of agents is actually a crucial aspect to model. A survey with some proposals is presented in [6]. Humans interactions offers a good starting point in collaborative behavior understanding. The discovering of interaction reasons, modes, and effectiveness is pursued, in order to bring collaboration aspects to the shared knowledge. In order to motivate users in deeper collaborations, behavioral changes can be stimulated through tailored incentives and mechanisms taken from persuasion theory [11].

Recommendations and feedback impacts. Once the engine is able to infer the up-to-date context of the agents, the process goes on to the evaluation of which mobility events can be useful in addressing mobility requests. The set of identified alternatives is then sent to the requesting user, as recommendations, in the form of available services.

A similar mining can be performed when a reconfiguration of the ICT components on the mobility network is needed. Consider, as example, the increase of the sampling rate for a traffic sensor, according to the increase of variation in the traffic level measured. In that case, the granularity of the data collected should be increased. Hardware sensors and actuators have to be solicited with the optimal self-reconfiguration rules.

We believe the closing loop lets a profitable feature to come out from the coordination framework. It determines the continuous learning of the system, which becomes aware of how effective has been the mobility recommendations exchanged among the agents, and which benefits are generated thanks to them.

5. CASE STUDY EXAMPLE

To evaluate the effectiveness of the proposed coordination framework in the provision of smart mobility service, we have developed a set of algorithms that aims to reproduce the main conceptual activities involved in the matching engine described above.

We have focused our efforts on an Itinerary Match service, as described in Section 2. Our aim is to evaluate potential matches between mobility requests and offers. Commuters with a similar typical daily route should be detected and recommended to join ride-sharing opportunities. Of course, the framework should support the provisioning of integrated services, but our work is still on an initial stage and our testings have been delimited in shaping a single service.

Even if our framework is expected to collect real-time data, coming from the distributed tuple space, we have undertaken an offline experiment, by simulating the matching engine activities on a large dataset previously collected.

Raw data involved in our study covers one week of detections in the city of Turin, and it consists in Call Description Records (CDRs) collected by a mobile network operator, through the cellphone network. However, one can assume that data can be collected opportunistically from a set of drivers, through an application installed on their smartphones, and propagated on the nodes of the infrastructure.

Basically, each time a user performs data exchange on the Internet, starts a call, or sends a text message, a spatio-temporal record is created. Each occurrence contains the user's identifier (who makes it happen), the location of the antenna related to the network cell (where it has happened), and the timestamp (when it has happened).

5.1 Towards agents classification

According to the conceptual model of the coordination framework, the first step involves an initial **pre-processing of raw data**. In our case study, we have filtered data in a way that tries to exclude non-commuters. In particular, we define commuters as all the users that generate at least one event in both a pair of enough distant geographic zones (let us call them A and B), during working days. Furthermore, we have narrowed our definition of commuters by considering two particular regions to perform that filtering. We want to study urban mobility, so we have considered an area that covers the inner part of Turin as the zone A (about 100Km2 wide), and a geo-fence of a broader zone (about 3000Km2 wide), which surrounds the city center (suburban area), as the zone B. We have not made any consideration on the mean of transportation used by the users, because the input data is too fragmented and sparse. Best practice to succeed in this activity consists in excluding all the commuters that are used to move along railways, cycling paths, metro stations, or bus stops.

Next phase has involved the **spatio-temporal aggregation** of the selected CDRs into mobility traces.

We were interested in modeling data into upper-level meaning abstractions, useful to better understand mobility patterns of the considered users. We define a mobility trace as the conjunction of a pair of temporally adjacent events, which represents the origin-to-destination path covered by a certain user in a defined temporal interval. This process has resulted in the detection of sequential mobility traces (the destination of the first matches with the origin of the second) that can cover wide areas and time intervals.

Each user is characterized by a set of mobility traces that can be reduced in length by doing some further spatio-temporal aggregation. The aim here is to compress the amount of data linked to each user, by merging the mobility traces through their sequential relationships (in both spatial and temporal domains). Thus, they shape brand new, more extended, mobility traces. The amount of merging occurrences has also been stored in the resulting mobility trace. Spatial proximity has been computed on the pair of geographical points that characterize the origin or the destination on the pair of the involved traces. This is easy to compute with a point-to-point distance formula (e.g.: haversine, euclidean). Temporal closeness has been computed on the time interval associated to the same pair of mobility traces. This task is more tricky and it concerns the evaluation of several temporal relations. In our case we have followed the ones presented by Van Beek and Manchak [27].

For each user, the **mobility pattern recognition** phase has contemplated the inference of the most visited mobility path described by the mobility traces. In particular, this process has first resulted in the application of a K-means clustering algorithm on the spatial dimension of the mobility traces. The evaluation of the clustered points has been done on the amount of occurrences related to them. Only the two most populated clusters have been evaluated (origin and destination candidates). The typical daily route of each user has been discovered. Figure 4 shows an extract of the daily routes in a 1-hour time lapse.

We formally define a daily route as the most frequent outward plus the most frequent inward mobility traces generated by the same user from/to an origin to/from a destination. Actually, our daily routes mining has returned a significant result only for the 10% of the considered users, since most of the results have revealed the same amount of occurrences on multiple candidates in the same cluster (too much ambiguity on the data). We think this sudden loss of significance can be tackled by evaluating more temporally distributed data (one week of CDRs collection does not provide enough significance to our study).

The user's typical daily route can be useful **to detect and anticipate mobility events**. As examples:

- it describes the **daily journey** the user is used to perform, and so it represents a daily event itself;
- it reveals the **expected presence** of an agent on the underlying road network, during a certain time span;
- it can be harnessed to anticipate any **expected traffic congestion** on the underlying road network at a certain time;
- it lets to locate a moving probe that can be queried just in case in the future to detect **mobility status and alerts**.

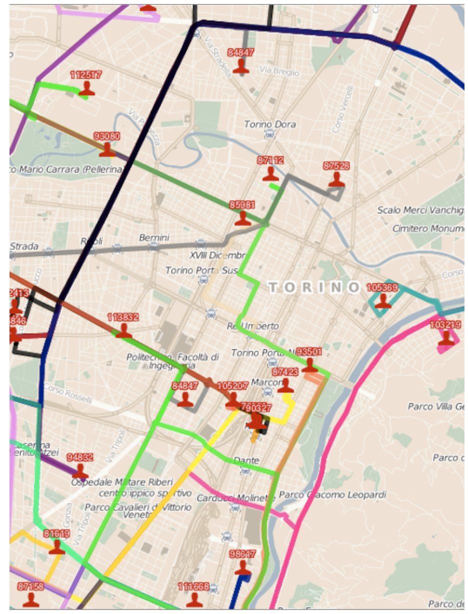

Figure 4: Partial representation of the users daily routes in the city center of Turin at a given time interval.

The contributions provided by each agent, wether it is a requesting agent or that it is a potential service provider, have been used to classify them, by creating their **agent profile** (in the distributed tuple space) with classifying labels and higher information contents.

5.2 Towards agents recommendation

Each agent profile is related to a commuter, and it initially contains only information about its daily route. In a real case scenario it can additionally include agent's personal details (such as demographics and interests) that can be collected through online social networks, its mobility preferences (such as its usual mean of transportation, and its willingness to do ride-sharing), the most likely home and work locations, its belonging to a same group of agents (due their commute similarities), and the most relevant historical events detected. Furthermore, one can think to add a ranking information to the agent profile that quantifies how much it has been involved in crowdsourcing and collaborative initiatives. This data can narrowly reflect social interactions among agents and their resulting benefits.Thus, it can outline the rise of **collaborative group of agents**.

As our next step, mining pool of users with similar daily routes has been done through an exhaustive search on all the users, by assuming that they were currently moving alone in a private car with 5 seats capacity. Each user should express at the same time its availability in offering ride-sharing services, and its necessity to find more efficient mobility solutions (in terms of vehicle occupancy rate). Our study has been limited to consider uniformed users that are characterized by the same mobility desiderata. However, on mining pool of users, one should consider maximum detour distance and time admitted as individual factors of the driver and each one of the passengers. A further improvement of the algorithm should contemplate the evaluation of the existing collaborative groups, in order to prefer users to rely on.

For any driver, we have selected the pair of mobility traces that composes its daily route, and we have compared each one of those with the whole set of concurrent mobility traces (within a confidence time interval) generated by other users. All the compared users with a mobility trace detected along the one generated by the selected driver represent potential passengers in ride-sharing pools, which can outline new potential collaborative groups of agents. The resulting pools contain information about the most suitable sequence of timed-stops to pick-up and drop-off passengers.

Once the pools are detected, ride-sharing **recommendations** can be sent out. Their should push social interactions between the involved users, and let new collaborative communities to be shaped. The system can track how they affect users mobility behavior and update both their individual and collaborative profiles.

6. RELATED WORK

Finding new approaches to enable mobility services has recently received a lot of attention. However, most of the studies are far from reaching effective and integrated solutions (from the collection of the requests to the provision of the services).

As discussed in Section 2, most of current ITS approaches do not offer a unified and integrated approach to support urban mobility in all its aspect, and often they own independent approaches for different mobility needs [1]. Also, in our proposed framework, and unlike most of ITS proposals, citizens are active agents of the overall infrastructure, by collaborating implicitly and explicitly towards the provisioning of smart mobility services.

Of particular interest to our work is the role of a middleware, which supports interactions and information exchange among the agents on the socio-technical system, and its involved in the generation of distributed intelligence. As far as we know, the best examples in this field that deal with the underlying infrastructure are the work of Harnie et al. [15], which aims to specify urban-area applications with tuple spaces abstraction, and the work of Julien and Roman [18], which proposes a middleware to enable context-aware mobile applications. The former enables intelligence through moving buses that carry the tuples, the latter propagates intelligence through vehicle-to-vehicle short range communications.

The works of Yang et al. and of Qu et al. [30, 25] introduce the concept of Intelligent Transportation Spaces as the integration of various ITS modules, vehicles, and roadside infrastructure. They mainly analyze safe and effective communication technologies to enable pervasive intelligence without impacting too much on drivers workload. However, neither of the works mention social interactions in matching mobility needs and services.

To the best of our knowledge, existing works do not give their contributions on proposing new approaches that could enhance social interactions.

Most of the mobility services presented in literature (e.g., [17, 12]) merely offer tailored solutions, without worrying about the creation of a coordinated methodology that deals with the dynamic orchestration of heterogeneous data streams.

We believe that a unified framework that models sensing, computing, and actuating capabilities of a socio-technical system of mobility agents is currently missing.

7. CONCLUSIONS AND FUTURE WORK

Social interactions among humans and ICT devices could strengthen the awareness of what urban mobility needs are, and how they can be addressed with smart mobility services. Social collective intelligence can be enabled, and so its utility can hit citizens and convince them to collaborate and cooperate each others through innovative sharing practices regulated by suitable incentives.

In the future, we will reshape our case study based on ride-sharing recommendations over a longer collection period, in order to reduce data ambiguity. Then, we will experiment with a larger set of mobility services, and will attempt at integrating them towards the realization of composite multimodal mobility services through our coordination framework.

8. ACKNOWLEDGEMENT

Work partially supported by the Emilia Romagna SPINNER2013 project MUCCA (Cooperative Urban Mobility in Smart Cities) and by the EU Project ASCENS (Autonomic Service Component Ensembles).

9. REFERENCES

[1] S.-H. An, B.-H. Lee, and D.-R. Shin. A survey of intelligent transportation systems. In D. Al-Dabass, Suwarno, J. Yunus, I. Saad, D. Giriantari, and A. Abraham, editors, *CICSyN*, pages 332–337. IEEE, 2011.

[2] L. Atzori, A. Iera, and G. Morabito. The internet of things: A survey. *Computer Networks*, 54(15):2787–2805, 2010.

[3] N. Bicocchi, A. Cecaj, D. Fontana, M. Mamei, A. Sassi, and F. Zambonelli. Collective awareness for human-ict collaboration in smart cities. *2012 IEEE 21st International Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises*, 0:3–8, 2013.

[4] J. Booth, P. Sistla, O. Wolfson, and I. F. Cruz. A data model for trip planning in multimodal transportation systems. In *Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology*, EDBT '09, pages 994–1005, New York, NY, USA, 2009. ACM.

[5] C. Brown, V. Nicosia, S. Scellato, A. Noulas, and C. Mascolo. Social and place-focused communities in location-based online social networks. *CoRR*, abs/1303.6460, 2013.

[6] G. Cabri, L. Leonardi, L. Ferrari, and F. Zambonelli. Role-based software agent interaction models: a survey. *Knowledge Eng. Review*, 25(4):397–419, 2010.

[7] G. Castelli, M. Mamei, and F. Zambonelli. The changing role of pervasive middleware: From discovery and orchestration to recommendation and planning. In *PerWare Workshop at the 9th IEEE International Conference on Pervasive Computing and Communications, Seattle (WAS)*, pages 214–219, March 2011.

[8] C.-W. Cho, Y.-H. Wu, C. Yen, and C.-Y. Chang. Passenger search by spatial index for ridesharing. *2012 Conference on Technologies and Applications of Artificial Intelligence*, 0:88–93, 2011.

[9] B. J. Coltin and M. Veloso. Towards ridesharing with passenger transfers. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 1299–1300. International Foundation for Autonomous Agents and Multiagent Systems, 2013.

[10] B. de Jonge and R. H. Teunter. Optimizing itineraries in public transportation with walks between rides. *Transportation Research Part B: Methodological*, 55:212–226, 2013.

[11] B. Fogg. Persuasive computers: Perspectives and research directions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '98, pages 225–232, 1998.

[12] A.-J. Fougères, P. Canalda, T. Ecarot, A. Samaali, and L. Guglielmetti. A push service for carpooling. In *Green Computing and Communications (GreenCom), 2012 IEEE International Conference on*, pages 685–691. IEEE, 2012.

[13] Z. Gao, M. Li, S. Du, and H. Zhu. Fairness-aware and privacy-preserving friend matching protocol in mobile social networks. *IEEE Transactions on Emerging Topics in Computing*, 1(1):192–200, 2013.

[14] D. Graziotin. An analysis of issues against the adoption of dynamic carpooling. *arXiv preprint arXiv:1306.0361*, 2013.

[15] D. Harnie, T. D'Hondt, E. G. Boix, and W. De Meuter. Programming urban-area applications for mobility services. *ACM Transactions on Autonomous and Adaptive Systems*, 9(2), 2016.

[16] W. He, D. Li, T. Zhang, L. An, M. Guo, and G. Chen. Mining regular routes from gps data for ridesharing recommendations. In *Proceedings of the ACM SIGKDD International Workshop on Urban Computing*, UrbComp '12, pages 79–86, 2012.

[17] X. Hu, V. Leung, et al. Vssa: a service-oriented vehicular social-networking platform for transportation efficiency. In *Proceedings of the second ACM international symposium on Design and analysis of intelligent vehicular networks and applications*, pages 31–38. ACM, 2012.

[18] C. Julien and G.-C. Roman. Egospaces: Facilitating rapid development of context-aware mobile applications. *IEEE Transactions on Software Engineering*, 32(5):281–298, 2006.

[19] Y. Kwon, D. Nunley, J. Gardner, M. Balazinska, B. Howe, and S. Loebman. Scalable clustering algorithm for n-body simulations in a shared-nothing cluster. In M. Gertz and B. LudÃd'scher, editors, *Scientific and Statistical Database Management*, volume 6187 of *Lecture Notes in Computer Science*, pages 132–150. Springer, 2010.

[20] L. M. Martinez, G. Correia, and J. Viegas. An agent-based model to assess the impacts of introducing a shared-taxi system in lisbon (portugal). In *Proceedings of the 7th International Workshop on Agents in Traffic and Transportation*, 2012.

[21] S. Mathur, T. Jin, N. Kasturirangan, J. Chandrasekaran, W. Xue, M. Gruteser, and W. Trappe. Parknet: Drive-by sensing of road-side parking statistics. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services*, MobiSys '10, pages 123–136. ACM, 2010.

[22] S. Nawaz, C. Efstratiou, and C. Mascolo. Parksense: A smartphone based sensing system for on-street parking. In *Proceedings of the 19th Annual International Conference on Mobile Computing & Networking*, MobiCom '13, pages 75–86, 2013.

[23] A. Pavlo, E. Paulson, A. Rasin, D. J. Abadi, D. J. DeWitt, S. Madden, and M. Stonebraker. A comparison of approaches to large-scale data analysis. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data*, SIGMOD '09, pages 165–178, 2009.

[24] V. Pejovic and M. Musolesi. Anticipatory Mobile Computing: A Survey of the State of the Art and Research Challenges, June 2013.

[25] F. Qu, F.-Y. Wang, and L. Yang. Intelligent transportation spaces: vehicles, traffic, communications, and beyond. *IEEE Communications Magazine*, 48(11):136–142, 2010.

[26] A. S. Rao and M. P. Georgeff. Modeling rational agents within a BDI-architecture. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 473–484. Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, 1991.

[27] P. van Beek and D. W. Manchak. The Design and an Experimental Analysis of Algorithms for Temporal Reasoning. *Journal of Artificial Intelligence Research*, pages 1–18, 1996.

[28] M. Viroli, M. Casadei, S. Montagna, and F. Zambonelli. Spatial coordination of pervasive services through chemical-inspired tuple spaces. *TAAS*, 6(2):14, 2011.

[29] X. Yan, H. Zhang, and C. Wu. Research and development of intelligent transportation systems. In *Distributed Computing and Applications to Business, Engineering & Science (DCABES), 2012 11th International Symposium on*, pages 321–327. IEEE, 2012.

[30] L. Yang and F.-Y. Wang. Driving into intelligent spaces with pervasive communications. *IEEE Intelligent Systems*, 22(1):12–15, Jan. 2007.

[31] F. Zambonelli, G. Castelli, L. Ferrari, M. Mamei, A. Rosi, G. D. M. Serugendo, M. Risoldi, A.-E. Tchao, S. Dobson, G. Stevenson, J. Ye, E. Nardini, A. Omicini, S. Montagna, M. Viroli, A. Ferscha, S. Maschek, and B. Wally. Self-aware pervasive service ecosystems. *Procedia CS*, 7:197–199, 2011.

[32] F. Zambonelli, G. Castelli, M. Mamei, and A. Rosi. Programming self-organizing pervasive applications with sapere. In *Intelligent Distributed Computing VII - Proceedings of the 7th International Symposium on Intelligent Distributed Computing*, volume 511 of *Studies in Computational Intelligence*, pages 93–102. Springer, 2014.

[33] A. X. Zhang, A. Noulas, S. Scellato, and C. Mascolo. Hoodsquare: Modeling and recommending neighborhoods in location-based social networks. *CoRR*, abs/1308.3657, 2013.