Knowledge-Based Provision of Goods and Services
for People with Social Needs: Towards a Virtual Marketplace

Daniela Rosu, Dionne M. Aleman, J. Christopher Beck, Mark Chignell,
Mariano P. Consens, Mark S. Fox, Michael Grüninger, Chang Liu, Yi Ru, Scott Sanner
University of Toronto, Toronto, Canada
{drosu,aleman,jcb,chignell, consens,msf,gruninger,cliu, yiru,ssanner}@mie.utoronto.ca

Abstract

Traditionally, the needs of vulnerable populations have been addressed by a plethora of public and private agencies that rely on donations of money, goods and services which they distribute based on their perception of what is needed and where. This approach, however, lacks a comprehensive understanding of the demand side as well as the ability to coordinate between various suppliers of goods and services, identify latent supply and predict future demand. To help address these issues, we have developed a knowledge-based platform that harnesses advances in several AI fields for efficient and effective provisioning of goods and services.

1 Introduction

The ecosystem of social welfare agencies that tend to the needs of vulnerable populations is under increasing pressure and struggles to meet the growing demand for goods and services as poverty levels rise, driven by the continuous widening of the income gap (OECD 2016), and more and more migrants come to depend on social services in their re-settlement countries. There is, however, a great societal reservoir of good-will and resources that can help alleviate the strain on the established welfare system, given a cost-effective and accessible way to tap into it.

Consider the case of a family of Syrian refugees recently re-settled in Toronto who need a crib for a soon to be born baby. They are eligible to use furniture banks in the city, but this is of no help as furniture banks do not store cribs as a matter of policy. The family found on-line an offer for a free crib, but transportation is an issue: the donor cannot deliver the crib, commercial delivery is cost prohibitive and public transport is not an option. Meanwhile, a retired couple from a Toronto suburb is looking to downsize to a condominium in the city and offload some of their possessions, including a baby cot formerly used by their granddaughters. The couple would like to donate the cot, but the resettlement agencies they contacted lack the transportation capacity needed. A few other items the couple no longer needs would be of great help to other recent immigrants, but the couple is not aware of other avenues for donating them. Discovering and

tapping into this latent supply can be tremendously helpful, as furniture banks and other charities only accept a few types of household goods. In yet another part of the city, a contractor who works on a flexible schedule and owns a large pick-up truck would gladly help transport donated items while en route from one work site to another, but he does not know to whom to offer his services.

While various on-line platforms allow for listing requests and offers of free goods and services (e.g., Craigs' List, Kijiji), it is not possible, to our knowledge, to coordinate between demand and supply and secure additional services such as free transportation by a third party. There is no platform, to our knowledge, that would allow the refugees to place a request for a crib, the retired couple to offer the baby cot, match the two listings and arrange transportation with a volunteer driver, while also discovering and harnessing latent supply of goods and services.

In this paper we describe the Social Needs Marketplace (SNM) Portal, a knowledge-based platform that aims to provide an efficient way to (re)distribute goods and services. The SNM Portal enables the demand side to make known their needs and the supply side to make known what they can provide in a richly structured representation that further allows for the discovery of latent needs and potential supply.

The Portal employs methods from multiple disciplines such as artificial intelligence, operations research and human-computer interaction (e.g., ontologies, machine learning, constraint programming, human factors principles) to collect requests and offers, suggest matches between them, and address the logistical problems related to transporting goods from donors to the people in need and scheduling services within the time windows that people are available, using a Uber/TaskRabbit-like volunteer network.

The main research questions addressed in this work are: (1) representing the knowledge related to the users of the system, and the goods and services requested and offered (Section 3), matching the donated goods and services to the demand, and allocating them fairly and efficiently (Section 4), and scheduling the pick-up and delivery of the donated goods from donors to recipients (Section 5). We present additional research challenges in section 6 and discuss our conclusions in section 7.
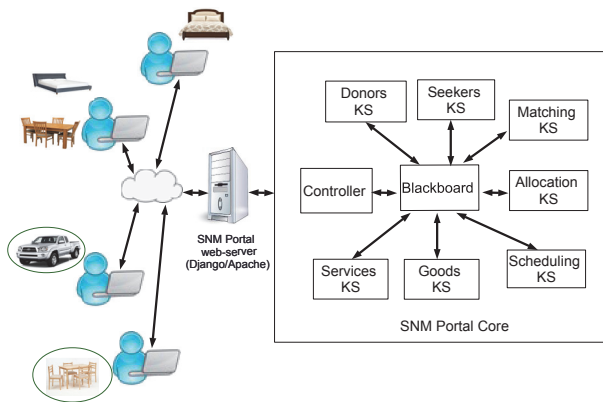
Figure 1: The SNM Portal - Architecture Overview

2 Overview of the SNM Portal

The overall architecture of the Portal is presented in figure 1. The Portal's users place requests and make offers for goods and services donations via a web based interface. All offers and requests are processed by the Portal's Core, whose architecture is based on the Blackboard paradigm (Hayes-Roth 1985; Nii 1986). As depicted in figure 1, a blackboard architecture is comprised of a shared data structure called a Blackboard, a Controller and a series of Knowledge Sources (KS). The blackboard acts as a central point of coordination of tasks (goals) to be performed by the knowledge sources, serving a number of purposes, including: (1) a means of managing SNM problem solving by representing all goals, sub-goals, goal assignments, goal status, and goal priorities, (2) a means to notify a KS that there is a goal to be satisfied, (3) a means for a KS to notify the status of their achieving goals.

The Controller is responsible for posting problems to be solved to the Blackboard, e.g., requests from users, and encapsulates control mechanisms such as prioritization of goals and conditions for terminating the system. The knowledge sources monitor the blackboard, respond to requests placed there by the web interface, and place responses back onto the blackboard. A Knowledge Source (KS) is a process that operates asynchronously. The current KSs are implemented in Python and C++. A KS can perform one or more tasks (e.g., matching requests for goods, allocation of goods, scheduling of pick-up and deliveries) and has its own data or knowledge base. A KS shares information with other KSs only through the blackboard. When a KS is activated, it receives all information on the Blackboard about the goal(s) it needs to solve. The KS then performs its computation using whatever information is available within its own knowledge base and posts results of its computation on the Blackboard as an update to its activating goal, and, possibly, additional new goals and elaborations of goals. The Core's architecture is rounded off by helper modules, which have functions that help maintain the blackboard and support the functioning of the knowledge sources.

User Interface. The design of the Portal's web-based interface reflects several assumptions made about the users of the portal: (1) *context of need*, e.g., urgent vs non-urgent, (2)*information-retrieval behavior*, e.g., preference for simple account features and interactions, (3) *literacy/proficiency in English*, e.g., heterogeneous, a large proportion of users have a limited command of English, (4)*digital literacy*, e.g., mostly limited literacy.

In order to accommodate users with reduced language skills, such as recent immigrants or refugees, we opted for simple vocabulary and phrasing throughout. The assumed level of digital literacy informs the amount of information to be displayed at any given point, and we implemented a conversational style interaction that asks for, or provides data in comparatively small chunks, in order to give users the ability to manage the information requested or given. The Portal uses a strict hierarchy model of content organization, i.e., pages are only made accessible from their parent page and users are guided through information and are not allowed to skip ahead.

Donors can register to provide service, donate goods, or both, and seekers can place requests for goods or services. After users navigates through the menu and select the type of service of interest, they are asked to provide information that will allow the portal to search for an appropriate match. Services identified as high demand, e.g., Goods, Food, Transportation, Housing, have their individual pictorial tabs on the interface, while services that are less frequently requested, e.g., Legal, Home Services, are grouped together and shown on a different page. For each individual type of request a user will need to fill in a series of related questions that will allow the portal to present them with a list of goods or services that best meet their criteria.

The system is designed to tailor the information requests at every step dynamically, based on the information received from the user at the previous steps. The data elicited from the users is stored in a rich semantic representation, which allows the Portal to use ontology-based reasoning.

Seekers are offered the opportunity to express their preferences for the goods or services that are matched to their requests by ranking them, information which is then made available to the allocation and scheduling knowledge bases.

Both seekers and donors of goods and services are permitted to specify strict windows of availability, which are taken into account when scheduling pick-ups, deliveries and appointments for various services. At the end of each day, the Portal's Allocation KS produces an allocation of goods and sends emails to the seekers asking them to confirm that they accept the goods allocated to them, and to donors to confirm that the goods are still available. The Scheduling KS then produces a daily schedule that allocates the volunteer drivers available on that day for transporting the confirmed goods from donors to seekers. Cancellations at any stage of the process by donors, recipients of volunteer drivers result in recalculations of the daily schedule to take into account the canceled segments.

Usage Scenario. Said, his wife Sara and their two young children, aged 6 months and 24 months, have recently relo-

cated to Toronto after spending a year in a Turkish refugee camp. Said used to be a financial analyst in Syria, and Sara used to be a Biology teacher. Said's proficiency in English is intermediate, while Sara's is very basic.

The five Ontario families who sponsored them rented a small unfurnished apartment on their behalf, but could only partially furnish it. Within a few days of arriving in Toronto, Said logs into the SNM Portal and places a request for a dinner table for four, a sofa-bed and a crib. Upon receiving Said's request, the Portal inquires if the family would also need a stroller, a baby changing table and car seats. Said answers yes and indicates that the family needs help with transporting any donated goods and would be available to receive the requested items in the evenings, between 6pm and 9pm, on week days and at any time during the day, 8am to 9pm, on weekends, for the next month.

The Portal's Matching KS produces a list of available items, from local donors and furniture banks, that fit Said's requests and he is asked to rank them.

After the completion of the request, the Portal presents Said with a short questionnaire concerning proficiency in English. Upon learning that one family member's knowledge of English is basic, the Portal asks if the services of a language tutor would be wanted. Said answers yes to this question and also indicates that a female tutor with knowledge of Arabic will be preferred. The Matching KS produces a list of available female English tutors who also speak Arabic and Said is asked to make a selection.

Claire and Martin, a couple from Scarborough have registered with the Portal an offer to donate several items they no longer need, including gently used clothes suitable for children aged from 1 to 6 years, toys, a combo stroller, two convertible car seats and a crib that can be disassembled. They indicate that they are available for donation pick up on weekends, between 9:30 and 11am.

Robert is a contractor who lives in Markham and has registered with the Portal his offer to use his Ford F-250 pick-up truck to transport goods to/from Markham (and other neighborhoods in Toronto) between 9am to 2pm on weekends, for three months. Sonya is a professional accountant. She was born in Canada to Lebanese parents, speaks Arabic, and lives in Parkdale. She is volunteering with various organizations and is also registered with the Portal as a provider of English tutoring for children and adults, financial advice for newcomers (including income tax filing), and transport of (small) goods on Sunday afternoons, for six months.

The Portal matches Said's request for a crib with Clair and Martin's offer and schedules Robert to pick up and deliver the crib on a Saturday morning. It also puts Said's wife in contact with Sonya for English language tutoring.

3 Representing Demand and Supply: Ontologies for the SNM

The match of supply and need of physical objects requires a detailed description of used goods from multiple aspects, for the use of classification as well as transportation. Such a representation faces several daunting challenges. First, the terminology that people use for these descriptions is rife with

ambiguity. Second, it is often unclear whether or not the terminology used is adequate to provide complete descriptions of a given range of goods and services. A third challenge is that many of the terms that people use to describe goods and services are ad hoc and arbitrary.

We base our representation of knowledge of users, goods, services, logistics, and preferences on ontologies. Ontologies are computer-interpretable specifications of the meanings of the terminology used by software applications. They play a key role in advanced applications in artificial intelligence, such as search, decision support systems, and the semantic integration of software systems. Several ontologies for products and services have been developed. Foremost among these is the GoodRelations Ontology (Hepp 2006) and the Schema.org, (Krutil, Kudelka, and Snasel 2012) effort. One drawback of these ontologies is that they are primarily taxonomies, and consequently do not fully address the challenges of ambiguity, incompleteness, and arbitrariness. The first objective focuses on the design and evaluation of ontologies for goods and services that are correct and complete with respect to the semantic requirements that arise from the motivating scenarios for the Supply/Demand Module. This involves specifying taxonomies of classes of goods and services, logical definitions for properties of goods and services, the relationships among goods and services, and the identification of other related concepts, such as offers.

The uses the methodology for the design and verification of first-order logic ontologies originated in (Gruninger and Fox 1995) (Uschold and Gruninger 1996), and (Katsumi and Gruninger 2010), which outlines the lifecycle leading to the development of a correct ontology. Automated reasoning plays a critical role in the specification of requirements, design, and verification of the ontology. Competency questions are queries that impose demands on the expressiveness of the underlying ontology. Intuitively, the ontology must be able to represent these questions and characterize the answers using the terminology. The relationship between competency questions and the requirements is that the associated query must be provable from the axioms of the ontology alone. Since a sentence is provable if and only if it is satisfied by all interpretations, competency questions implicitly specify the intended semantics of the ontology.

Evaluation is at the heart of the ontology design methodology, allowing a rigorous way to measure success and progress. The specification of competency questions is a formal approach to capturing the requirements for the goods and services ontologies. The design of an ontology is successful if the axiomatization of the ontology is sufficient to entail the competency questions, and within this project, the demonstration of this capability will be shown through the use of automated theorem provers. In this way, the evaluation of the ontologies arises from the support for search capabilities through automated reasoning with descriptions of goods and services. Existing theorem provers can be used to solve queries related to the matching of goods/services offerings and requests within the Supply-Demand Module, thus supporting intelligent meaning-based search and retrieval.

To name a few motivating situations, the system needs to

know the level of mobility in the case of matching appropriate chairs for disabled individuals; in order to obtain eligibility for transportation, the matching algorithm needs to be told if a furniture is able to be assembled and reassembled or not; both the demand side need to have the complete knowledge of damage and condition of a good to make the decision of acceptance, for example a dining furniture set with a missing chair and a table with a chipped surface.

We have identified five ontologies to be developed in order to realize the full functionality of the SNM project: parthood, material, product condition, shape and mobility.

From the parthood ontology perspective, when describing a mereological relationship, the use of the words part and whole hides the intended semantics of the relations. Existing studies in mereology usually either present part of as a general top level mereology relation and all other mereological relations are sub-relations under it, or having a part of relation to summarize all parthood relations. Such as Winston et al. (Winston, Chaffin, and Herrmann 1987) presented a taxonomy of part-whole relations with including different part-whole relations as specializations of a single, general part-of relation satisfying the basic axioms of mereology. In more recent researches, upper ontology SUMO introduces the relation part as a spatial relation, Keet (Keet 2016) introduces a taxonomy with summarization of Odells types of part-whole relations (Odell 1998), Bittner and Donnelly (Bittner and Donnelly 2005) also have a P relation which stands for part of. We present an alternative approach, in which all relations that are associated with mereology are formalized in different modules of the ontology.

Firstly, we introduce tests to validate whether a relation is mereological and evaluate relations from previous work. Secondly, we propose a new relationship among the mereological relations that is different from an abstract conceptual hierarchy relationship. Thirdly, using a bottom-up technique, we identify four mereological relations based on use cases developed for the Social Need Marketplace. These relations are capable of representing the physical structure of used household goods, including the description of damage and relationship between multiple goods. We provides a standard ontology of mereology that is more simple, comprehensive and semantically rich.

The parthood module ontology is capable of describing the different containment and connection relations, as well as incorporate the representation of damaged products which has not been explored before in the field, for example describing paint scratch and tearing of silk material. In the situation of accommodating possible industrial use, the ontology could be applied through a domain extension steered into incorporating concepts required for enabling the automatic eligibility identification for special assembly, delivery and transportation handling activities according to the parthood characteristics of the product, for example a product can be processed in the disassembling activity only if it is a whole object or a member in a collection but not a dense thing and has externally connected components.

There are many different relationships that can be close to parthood, portion, component part, component, constituent, piece, region, share and percentage can be perceived as part.

However, in our approach these are treated as multiple distinct mereological relations, each of which is axiomatized in its own module of the ontology rather than being sub-relations of part. Our claim is that there is no super-relation of the mereology relations, but each mereology relation is comprehensive by its own.

In the shape and mobility ontologies, we need ontologies that go beyond textual description or taxonomy of terms for shapes to support reasoning. A shape ontology to represent the complicated shapes with logical relations with the least number of primitive or predefined concepts is in need. We also need a formal mobility ontology to represent the mobility of products or physical objects for the use of intelligently identifying the functionality and installation requirements of a product. Examples include deciding the space needed for the sway of a swing, the eligibility to support disable personals with the rotation and lifting of an office chair or wheelchairs, and the handling requirement causing from internal part movement of adornments.

Other modules supplementing the product ontology include the study of material, the indication of product condition and representing an "ideal product" to determine the completeness of an object.

4 Matching and Allocating Goods and Services

A request placed by a seeker of goods or services is processed by the Matching KS, which identifies using the rich semantic representation of offers and demands, a list of matches. The identification of potential matches involves a combination of database look-ups and ontology-based reasoning, provided by Z3 (de Moura and Bjorner 2008) an open source SMTsolver.

Requests placed by clients can have several matches, each with various associated preference rankings, and similarly, offers registered with the Portal can be matched to several requests. In order to establish a feasible distribution of goods and services, the Allocation knowledge source employs several existing techniques.

The goods allocation problem in our framework has been described in the literature as the house allocation, or the assignment problem (Abdulkadiroglu and Sonmez 1998), for which several algorithms exist. The default mode of the Allocation KS uses the Serial Dictatorship Method (Svensson 1999), which requires that the preferences expressed by users be strict. If a user declares that he or she is indifferent between two or more of the goods presented as potential matches, the allocation of goods to clients is computed using a variant of the Top Trading Cycles algorithm (Shapley and Scarf 1974).

Currently, the system does not take into account information related to past requests made by clients, i.e., each allocation is performed in isolation. We are currently exploring extending the framework to deal with repeated allocation problems and include information about past requests, e.g., if a client has made several requests that have not been satisfied, or the goods he or she was allocated have not been among his or her top choices. In essence, we have a re-

peated fair division combined with a routing problem, whose solving requires the introduction of a social welfare function that captures social fairness without compromising computational efficiency.

5 Pick-up and Delivery Scheduling

Given a set of allocations, the Transportation Scheduler KS seeks to assign and route vehicles and drivers from an existing volunteer pool to deliver the goods during the times at which the driver and the source and destination clients are available. As we assume we have volunteer drivers, it is likely that not all the goods can be delivered in one day and that each driver is only willing to perform one or two deliveries per day while incurring a small additional distance on a trip that they were already planning. For example, a volunteer driver may be willing to perform a pickup-and-delivery on his/her way to work.

Problem Definition Let $G = (\mathcal{V}, \mathcal{A})$ be a complete directed graph with vertex set $\mathcal{V} = \mathcal{D} \cup \mathcal{C}$ where $\mathcal{D}$ represents the start and end locations of each driver and $\mathcal{C}$ represents the source and destination client vertices. Each arc $(i, j) \in \mathcal{A}$ has a non-negative routing time $T_{i,j}$ satisfying the triangular inequality.

Let $\mathcal{R} = \{1, ..., N\}$ represent the set of pickup-and-delivery requests. Each request i is paired with a positive weight, W_i reflecting its importance. A request $i \in \mathcal{R}$ has associated pickup and delivery vertices $i^+, i^- \in \mathcal{C}$ and a positive load size, Q_i. In addition, each request i is associated with two time windows, $[E_i^+, L_i^+]$ and $[E_i^-, L_i^-]$ and two service times, S_i^+ and S_i^-, representing the time windows and durations of the pickup and delivery activities, respectively.

Let $\mathcal{K} = \{1, ..., M\}$ represent the set of vehicles. Each vehicle $k \in \mathcal{K}$ is associated with a start and end vertices $k^+, k^- \in \mathcal{D}$ which may represent the same geographical location, a capacity P_k, and an availability time window $[E_k, L_k]$ representing the time within which the driver must travel from start to end, including any additional time for deliveries.

A *route* for vehicle k is a sequence of vertices, $[k^+, ..., k^-]$. A request is *served* when it is part of a route. For served requests, the set of routes must satisfy the following constraints:

1. The pickup and delivery vertices of any request must be on the same route;

2. The pickup vertex must precede the delivery vertex;

3. A vertex is only visited by at most one vehicle;

4. The load of a vehicle k cannot exceed its maximum capacity P_k at any point;

5. A route must start and end within the vehicle availability window;

6. No subtours are allowed in any route;

7. The vertex is served within its specified time window.

The objective is to maximize the weighted sum of served requests.

Our problem requires the selection of time-window constrained pickup-and-delivery requests to serve, the assignment of requests to vehicles, and the routing of each vehicle. Assignment and routing are common aspects of the standard dial-a-ride problem (DARP) (Cordeau and Laporte 2007) and pickup-and-delivery problems (Parragh, Doerner, and Hartl 2008a; 2008b). However, selectivity is unusual for these problems as the objective is typically to serve all requests with minimum total travel cost. Selectivity can be seen in other problems such as the traveling purchaser problem (Booth, Nejat, and Beck 2016b) and the team orienteering problem (Vansteenwegen, Souffriau, and Van Oudheusden 2011).

Though these three components have been studied extensively, their combination has received little attention in the literature. Baklagis et al. (Baklagis, Dikas, and Minis 2016) proposed a branch-and-price framework to tackle this problem and Qiu et al. (Qiu, Feuerriegel, and Neumann 2016) proposed a graph search and a maximum set packing formulation that is specially tailored for homogeneous fleets. However, these papers do not consider multiple depots and heterogeneous fleets. Furthermore, the volunteer drivers tend to create problems where each vehicle serves a small number of requests (e.g., one or two) in a day, leading to solutions that look substantially different than vehicle routing problems with full-time drivers.

Solution Approaches. The time window and routing aspects of these problems make them challenging to solve to optimality. While many of the approaches seek heuristic solutions, the state of the art for exact approaches has typically been based on sophisticated mathematical programming such as the branch-and-price framework of Baklagis et al. (Baklagis, Dikas, and Minis 2016). Even with such approaches, the problem size for which optimal solutions can be found is typically in the range of 100 to 200 requests (Parragh, Doerner, and Hartl 2008a; 2008b).

Building on our previous work (Beck 2010; Tran, Araujo, and Beck 2016; Booth, Nejat, and Beck 2016b; 2016a) on constraint programming and decomposition methods for scheduling and routing problems, we are developing pure constraint programming approaches as well as a hybrid of mixed integer programming and constraint programming based on logic-based Benders decomposition for this problem. The current KS implements our initial constraint programming model which we expect to evolve as both our research on problem solving and our experience with the pickup-and-delivery problems within SNM develops.

6 Additional Research Challenges

The additional research challenges we are exploring include, but are not limited to, extending the support provided by the Portal for the semantic matching of supply and demand, identifying latent supply, ensuring information privacy and effective fraud detection, as well as developing methodologies for solving the repeated allocation problems that arise in the context of the Social Needs Marketplace.

Identifying Latent Supply. In any population, valuable and reusable goods such as furniture are often needlessly sent

to the landfill or held unused in storage areas for extensive periods of time. These are goods that could be allocated to vulnerable populations if only one knew they existed or could canvas areas (with flyers, mailings, door-to-door visits) likely to have a latent supply of needed goods.

This need to identify latent supply raises a number of important research questions that are largely unaddressed in the literature. These questions fall into to two distinct subcategories:

1. How can we estimate latent supply if we don't observe it? While we cannot observe all latent supply, we can observe a sample from online classified websites such as craigslist.org or kijiji.ca. Under the simplifying assumption that data not in this sample is missing at random, we can train geospatial-temporal models to generalize to future latent supply predictions as well as areas for which observation data was sparse. As a further extension, we could attempt to obtain true ground truth for latent supply through active surveys and then attempt to correct for the sampling bias observed on the classified websites.

2. The second research question is how to actually perform geospatial-temporal modeling of goods availability for future prediction based on historical data. There are a number of practical issues, with perhaps the most important one being that existing geospatial predictive models (Kriging or Gaussian Processes, nearest neighbor approaches) tend to spatially smooth too much when predicting sparsely observed urban phenomena. What is needed for latent goods modeling is a model that is spatially aware of demographic and socio-economic regions of the city and to generalize from these underlying traits. A secondary but equally important issue is modeling temporal trends in availability (e.g., more heaters may become available at the end of spring) and taking advantage of this for future goods allocation planning.

In summary, being able to accurately predict latent supply allows us to do long-term planning for goods allocation as well as targeted advertising and canvassing to make these predicted goods available to the populations that need them.

Privacy and Fraud Detection. The protection of individuals' privacy is a significant challenge faced in many application areas (healthcare applications been a prime example), and the Social Needs Market Portal is not an exception.

An important aspect of privacy protection is information accountability, that is ensuring that the policies and configured preferences that govern the flow of personal information are respected by the parties that collect, use, and share users' data.

Ensuring compliance with a variety of privacy policies is an increasingly complex task involving multiple participants including watchdogs and enforcement agencies that prosecute privacy violators.

While privacy auditing is complementary to other privacy enforcement mechanisms (e.g., data encryption and access control), the presence of comprehensive auditing is also a significant deterrent against violations. Privacy auditing in the SNM Portal presents multiple challenges, and foremost among them is the need to support auditing by the multiple participants involved. The information captured by auditing logs also provides an important source for analysis services that aim at detecting fraud in the marketplace. These fraud detection services can be executed online (e.g., providing warnings or blocking suspicious transactions) and offline (e.g., alerting relevant authorities of potential abuses and identifying the marketplace participants involved).

We are developing a privacy and security auditing framework for the marketplace that can provide the basis for ensuring privacy compliance and the information needed to implement online and offline fraud detection analytic capabilities.

The current infrastructure to support privacy auditing is ad-hoc, and not reinforced by widely adopted technology standards, organizations usually creating local logs to address their privacy audit needs. As a consequence, the logs generated by individual organizations are heterogenous and privacy auditing across organizations requires translation of the logs on an application by application basis. Recent work (Samavi and Consens 2012) has proposed a linked data based privacy framework designed to facilitate the auditing task of privacy auditors when multiple participants are dealing with individuals' personal information. The research we will carry out will build on these recent proposals and develop new solutions to deal with the specific challenges of the SNM environment. In addition to framework extensions, we will develop prototypes that demonstrate the feasibility and value of publishing privacy log data. We will also carry experimental validations of the suitability of the data captured in logs for completing fraud detection analytics.

7 Conclusion

We present the Social Needs Market Portal, a knowledge-based platform for the efficient and effective provisioning of goods and services to people in need. The Portal enables the consumer side to make known their needs and the supply side to list their available goods and services. It also enables the discovery of latent supply and assists with solving the logistical problems related to the provisioning of goods and services by supplying the means to schedule and route the capacity provided by a network of individual and agency-based volunteers. The Portal also provides NGOs and government-sponsored organizations with the capability to gain a better picture of the needs people have, post what they are able to provide, when and where, as well as reach a larger segment of the vulnerable populations and achieve better outcomes by being able to seamlessly combine and coordinate their efforts.

A major benefit of this project is the impact on the quality of life of the vulnerable. The flexible and nimble provisioning of basic needs reduces the stress of many segments of our population. Reducing stress provides a better quality of life, and a better quality of life, in turn, results in better work place performance and productivity and reduced social costs.

References

Abdulkadiroglu, A., and Sonmez, T. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66(3):689–701.

Baklagis, D.; Dikas, G.; and Minis, I. 2016. The team orienteering pick-up and delivery problem with time windows and its applications in fleet sizing. *RAIRO-Oper. Res.* 50(3):503–517.

Beck, J. C. 2010. Checking-up on branch-and-check. In *Proceedings of the Sixteenth International Conference on the Principles and Practice of Constraint Programming (CP2010)*, 84–98.

Bittner, T., and Donnelly, M. 2005. Computational ontologies of parthood, componenthood, and containment. In *Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI*.

Booth, K. E. C.; Nejat, G.; and Beck, J. C. 2016a. A constraint programming approach to multi-robot task allocation and scheduling in retirement homes. In *Proceedings of the Twenty-Second International Conference on Principles and Practice of Constraint Programming*, 539–555.

Booth, K. E. C.; Nejat, G.; and Beck, J. C. 2016b. Logic-based decomposition methods for the travelling purchaser problem. In *Proceedings of the Thirteenth International Conference on Integration of Artificial Intelligence and Operations Research Techniques in Constraint Programming*, 55–64.

Cordeau, J.-F., and Laporte, G. 2007. The Dial-a-Rde Problem: Models and Algorithms. *Ann Oper Res* 153:29–46.

de Moura, L., and Bjorner, N., eds. 2008. *Z3: An efficient SMT solver*.

Gruninger, M., and Fox, M. S. 1995. Methodology for the design and evaluation of ontologies. In *Proceedings of the Workshop on Basic Ontological Issues in Knowledge Sharing, IJCAI-95, Montreal, Canada*.

Hayes-Roth, B. 1985. A Blackboard Architecture for Control. *Artificial Intelligence* (26):251–321.

Hepp, M. 2006. Products and services ontologies: A methodology for deriving owl ontologies from industrial categorization standards. *International Journal on Semantic Web and Information Systems (IJSWIS)* 2(1):72–9–9.

Katsumi, M., and Gruninger, M. 2010. Theorem proving in the ontology lifecycle. In *Knowledge Engineering and Ontology Design*.

Keet, M. C. 2016. Introduction to part-whole relations: Mereology, conceptual modelling and mathematical aspects. Technical report.

Krutil, J.; Kudelka, M.; and Snasel, V. 2012. Web page classification based on schema.org collection. In *Proceedings of the IInternational Conference on Computational Aspects of Social Networks, CASon*.

Nii, H. P. 1986. Blackboard Application Systems and a Knowledge Engineering Perspective. *AI Magazine* 7(3):82–107.

Odell, J., ed. 1998. *Six different kinds of composition*, Advanced Object-Oriented Analysis and Design using UML. Cambridge University Press.

OECD. 2016. Society at a Glance 2016: OECD Social Indicators . Technical report, The Organisation for Economic Co-operation and Development (OECD).

Parragh, S.; Doerner, K.; and Hartl, R. 2008a. A Survey on Pickup and Delivery Problems Part I: Transportation between Customers and Depot. *Journal fr Betriebswirtschaft* 58:21 – 51.

Parragh, S.; Doerner, K.; and Hartl, R. 2008b. A Survey on Pickup and Delivery Problems Part II: Transportation between Pickiup and Delivery Locations. *Journal fr Betriebswirtschaft* 58:81 – 117.

Qiu, X.; Feuerriegel, S.; and Neumann, D. 2016. Making the most of fleets: A profit-maximizing multi-vehicle pickup and delivery selection problem. *European Journal of Operational Research*.

Samavi, R., and Consens, M. P. 2012. L2TAP+SCIP: an audit-based privacy framework leveraging linked data. In Pu, C.; Joshi, J.; and Nepal, S., eds., *8th International Conference on Collaborative Computing: Networking, Applications and Worksharing, CollaborateCom 2012, Pittsburgh, PA, USA, October 14-17, 2012*, 719–726. ICST / IEEE.

Shapley, L., and Scarf, H. 1974. On cores and indivisibility. *Journal of Mathematical Economics* 1(23):23–37.

Svensson, L. G. 1999. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare* 16:557–567.

Tran, T. T.; Araujo, A.; and Beck, J. C. 2016. Decomposition methods for the parallel machine scheduling problem with setups. *INFORMS Journal on Computing* 28(1):83–95.

Uschold, M., and Gruninger, M. 1996. Ontologies: Principles, methods, and applications. *Knowledge Engineering Review* 1:96–137.

Vansteenwegen, P.; Souffriau, W.; and Van Oudheusden, D. 2011. The Orienteering Problem: A Survey. *European Journal of Operational Research* 209:1–10.

Winston, M.; Chaffin, R.; and Herrmann, D. 1987. A taxonomy of part-whole relations. *Cognitive Science* 11:417–444.