JASPR Work Package 1: Requirements analysis

Nathan Griffiths, Simon Miles, Phil Taylor, and Lina Barakat

4 December 2015

1 Work Package Description (from Case for Support)

With BP [Black Pepper Software], we will specify logistics case studies that demonstrate the reputation assessment challenges, including coping with limited evidence, the requirement for detailed rationale, mitigating circumstances, and the need for personalised assessments. Since the effectiveness of reputation assessment is dependent on the data available, we will identify a set of data desiderata. Reputation indicators will be defined for reasoning about mitigating circumstances, provision of personalised assessments, and support for rationale generation. We will analyse what could be realistically captured by consumers and providers, both in general and with reference to the case studies. With BP, we will analyse the practically of augmenting existing systems with data capture mechanisms, explore scalability issues, and identify a minimum set of data capture requirements. Identity management is crucial to the effectiveness of any reputation mechanism, and although not the concern of this project, a suitable mechanism is required. We will determine which certification and security technologies can be used to support identity management in the architecture, for example to prevent whitewashing or the establishment of multiple identities. Deliverable: Case study definitions, analysis of data capture and identity management (report).

2 Overview of Logistics Case Study

Black Pepper Software (BP) has extensive experience of building software systems for a large multi-national logistics company. For reasons of commercial confidentiality the following description has been anonymized and we refer to the logistics company as Acme Logistics.

Acme Logistics is a commodity trading and logistics company that sources, stores, refines, and ships commodities globally. Activities are divided into four areas: oil and petroleum products, refined metals, non-ferrous concentrates, and bulk commodities. Black Pepper's experience with Acme Logistics is in the refined metals area, and that is the focus of the case studies considered in JASPR.

Conceptually, the business is divided into buying teams and selling teams, separating the purchase of materials from their sale. The logistics team responsible for arranging shipment of the commodities is another separate team. There are benchmarks for the expected cost of moving commodity X from A to B which is used when creating the trades. Key to the business is optimising the matching between the two sides: sourcing raw materials, refining them, selecting appropriate transport methods, routes, and storage. An example transactions might be the delivery of 100 tonnes of copper to a manufacturing plant, which first must be purchased as ore from a mine, refined through a smelter, formed, and delivered to the customer. Each of these stages could be outsourced or performed using Acme Logistics's own assets. Note that contracts typically express features such as quality (e.g. 99% purity) and have a tolerance (e.g. 100 tonnes $\pm 2\%$), with payment for what is actually received (potentially after remediation if, for example, the purity is found to be lower than specified).

Acme Logistics do not actually own the materials, but finance the purchase, e.g. borrow £10M for 100 tonnes of copper for 30 days at 2% rate (and so have to sell at more than 2% profit). The material itself is used as collateral to finance the purchase. Before the commodities are delivered to the customer, the finance is repaid. Shipping delays therefore have costs not only in terms of penalty payments to customers but also in terms of extra interest and the potential for additional warehouse time (and rent). Another factor for delay costs are delays in getting into port for undertaking the transshipment, or missing the pre-arranged slot for undertaking transshipment.

To avoid the profit on the sale of commodities being affected by the price moving in the markets, the price is hedged with paper trades, to compensate for price movements. This means that the profit is not impacted by market changes, but the cost of hedging is another factor that leads to delays incurring costs.

Acme Logistics own some ships, lease some long term, and buy passage for specific goods on a ship. They also have the potential to sell spare space on their ships. In owned and leased ships any delays have knock on effects for subsequent contracts.

Acme Logistics also own smelting plants inroducing the potential for using their own smelter if it is idle as this may be cheaper than another option, even if other option is usually cheaper. There is a need to understand the actual throughput compared to theoretical throughput and balance cost against loses incurred in smelting and quality of result. Acme Logistics may be tied into 3 year fixed contracts with third party smelting plants, and will need information at point of contract revision to make informed choice.

3 Information Recording

Records are generated throughout the process, in particular every time materials are moved (e.g. into warehouse, onto boat, into dock) their details are recorded. Specific information recorded includes time, where it arrived from, claimed actual weight on arrival (provisional weight), quality, and time of arrival at customer. Note that the actual weight may not be checked and recorded for several days after arrival.

Currently the information recording process is a mixture of online and paper based, and there is no integrated repository of information. However, this information and recording process could be used to create provenance records. Once the information is recorded in a structured way, it becomes possible to check for patterns, e.g. whether weight always goes up (e.g. ore mixed with lower grade) or down (e.g. poor handling or theft) in a particular warehouse.

4 Overview of entities/processes

4.1 Goods

Goods have several characteristics:

- Grade (quality). Goods have different qualities depending on supplier, e.g. mines in South Africa and Australia may produce ore with different purity.
- Shape (rod, sheet, cathode etc.).
- Brand (Producer X or Producer Y etc.), which is related to reputation and provenance.
- Weight. Some materials may be shipped in bundles, with bundle weights recorded, and each bundle having an identifier for tracking. Weight can vary with environment (e.g. wet ore weighs more, but metal content is unchanged).

On receipt of goods Acme Logistics typically must start transportation to onward destination immediately. At the same time an assay is sent to UK for testing. At a later point may find that quality is below what was specified and so have to refer back to contract and enter into remediation regarding payment and must also take remedial action for downstream customer supply. For example, if a delivery of goods at 99% purity was expected, but only received 92% purity this can have a knock on effect for a customer requiring 99%, and may introduce the need to divert goods to a customer who has lower requirements and source an alternative supply of 99% purity goods.

There are a number of risks to consider, including:

- Accidental damage, e.g. sheet metal falling and being misshapen, with the result that it must be resold, potentially as scrap.
- Environmental degradation (e.g. rusting if left in rain).
- Contamination/loss, e.g. some forms of material are moved by digger (e.g. ore), and some can be lost off bucket, or may be mixed (accidentally or deliberately) with another material; should consider how often it needs to be assayed.

4.2 Customer

Customers vary according to whether they only buy specific goods, or whether they buy a range of items, e.g. copper refinery only purchases copper ore, while a manufacturer may purchase copper, tin and steel. Where a customer purchases several materials they may have separate contracts for each. Goods may only be useful in combination, i.e. delays of one may impact use of another. Customers may interact via a distribution centre (dependent on commodity), e.g. distributor may buy large quantity of steel and sell small amounts, or may aggregate lots of small purchases and sell in large amounts. Of prime importance to a distribution centre is selection of the right suppliers.

4.3 ACME Logistics

ACME Logistics is responsible for the overall process of managing the physical movement and storage of material throughout the chain (purchase of materials from suppliers, refinement, delivery to customers).Typically use third party warehouses for intermediate storage, paying rent for each day of storage. ACME Logistics also own warehouses, and when goods are in one of these detailed information may be available, e.g. pallet x contains y kg of good and is stored in location z.

In making decisions ACME Logistics must consider a range of risks, including:

- Route options have associated risk (rebel held territories, piracy risk etc.)
- For road transport there is RTA risk, e.g. if shipping overland in 20 trucks might lose one and so only be able to deliver 19/20 of load.
- Risk of damage or loss (shrinkage) when goods in third party warehouse.
- General in-transit risks:
  - Vehicle (lorry/boat/train) crash or mechanical failure.
  - Adverse weather, causing delay, damage or loss.
  - Environmental events, e.g. ash cloud from volcano, causing delay, damage or loss
  - Civil unrest causing delay, damage or loss
  - Strikes causing delay (may have predictable risk associated with a geographical region and time of year).

5 Circumstance patterns

Where there is little data from which to assess reputation, individual pieces of evidence can carry great weight and, where negative, may cause a provider rarely to be selected, and never given the opportunity to build their reputation. While the honesty of a reviewer can be tested from past behaviour and dishonest negative reviews ignored, it is possible for a review to be accurately negative, because a service was provided poorly, but for this not to be an accurate predictor of future behaviour. These are cases of mitigating circumstances, where the context of a past service provision rather than the ability of the agent meant that it was poorly provided, but that context was temporary.

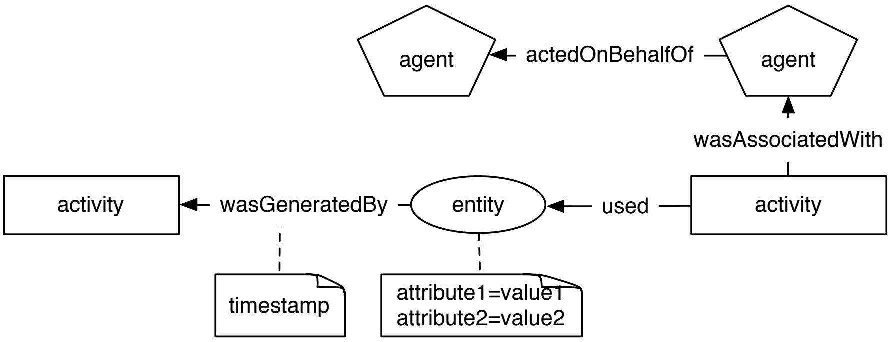

Figure 1: PROV graph illustrating the key elements

For example, with Acme Logistics, sea transportation of goods by a shipping company may be delayed by an unexpected severe storm. The fact that such a delay is out of the shipping company's control should be considered when assessing this company's reputation. Similarly, a low reputation of a shipping company (combining road, rail and sea transportation) due to the past failures of its rail transportation sub-provider to meet consumers' constraints on delivery time might give an unfair view if the company has changed the sub-provider to avoid such failures re-occurring.

We thus argue for the circumstances of past interactions to be recorded and taken into account more explicitly. This raises the question of what form these records should take, and who should record them. Also, in order to share interaction records between different parties, as many reputation methods do, they need to be recorded in a commonly interpretable format. In 2013, the W3C standards body published the PROV standard for modelling, serialising and accessing *provenance information*, the history of processes [10]. A PROV document describes in a queryable form the causes and effects within a particular past process of a system (such as agents interacting, the execution of a program, or enactment of a physical world process), as a directed graph with annotations. A visualisation of such a graph, showing PROV's key elements, is shown in Figure 1. In summary, an *activity* is something that has taken place, making *use* of or *generating entities*, which could be data, physical or other things. *Agents* are parties that were responsible for (*associated with*) activities taking place, and one agent may have been *acting on behalf of* another in this responsibility. Activities, entities and agents (graph nodes) may be annotated with key-value *attributes* describing features that the elements had. *Timestamps* can also be added to show when entities were used or generated by activities.

In this section, we specify three mitigating circumstances patterns that could be detected in provenance data. These examples are not intended to be exhaustive, but illustrative of the approach.

5.1 Unreliable sub-provider

In the first mitigating circumstance, poor service by a particular provider on a past occasion was due to their reliance on a poor sub-provider for some aspect of the service. If the provider has changed sub-provider, the past interaction should not be considered relevant to their current reputation (such a situation may indicate poor judgement and so may still have a degree of relevance, especially if the provider repeatedly selects poor sub-providers).

In other words, Provider A's reputation should account for the fact that previous poor service was due to Provider A relying on Provider B, who they no longer use. The provenance should show all of:

1. Provider B was used where there was poor service provision,
2. Provider B's activities were the likely cause of the poor provision, and
3. Provider A no longer uses Provider B (not necessarily shown through provenance.)

A provenance pattern showing reliance on a sub-provider in a particular instance can be defined as follows. For reference, activities are labelled with An (where n is a number) and entities are labelled with En. Figure 2 illustrates this pattern, along with some of the specific cases below.

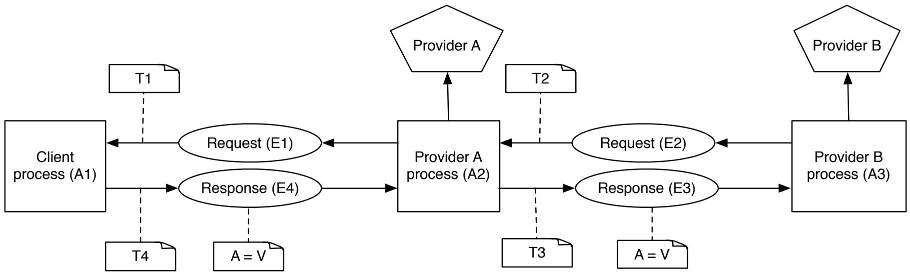

Figure 2: Provenance graph pattern for unreliable sub-provider circumstance

Step 1 A client process, A1, sends a request, E1, for a service to a service process, A2, for which Provider A is responsible. In the PROV graph, this means that E1 wasGeneratedBy A1, A2 used E1, and A2 wasAssociatedWith Provider A.

Step 2 A2 sends a request, E2, to a service process, A3, for which Provider B is responsible. In the PROV graph, this means that E2 wasGeneratedBy A2, A3 used E2, and A3 wasAssociatedWith Provider B.

Step 3 A3 completes the action requested, and sends a result, E3, back to A2. In the PROV graph, this means that E3 wasGeneratedBy A3, and A2 used E3.

Step 4 A2 completes the service provision, sending the result, E4, back to A1, so that the client has received the service requested. In the PROV graph, this means that E4 wasGeneratedBy A2, and A1 used E4.

We can then distinguish different cases in which Provider B would be the likely cause of poor quality service provision. Each case corresponds to an extension of the above provenance pattern.

Case 1. An aspect of the result of provision is poor, and that aspect is apparent in the contribution provided by Provider B. The extensions to the original pattern are as follows.

- The service provision result, E4, has an attribute A=V, which is a reason for the result being poor.
- The intermediate result from Provider B, E3, has this same attribute A=V.

Case 2. The poor provision may not be due to eventual outcome but due to the time taken to provide the service, and this can be shown to be due to the slowness of Provider B. The extensions to the original pattern are as follows.

- The sending of the service request (i.e. the relation E1 wasGeneratedBy A1), is timestamped with T1.
- The receipt of the service result (i.e. the relation A1 used E4), is timestamped with T4.
- The sending of the delegated request (i.e. the relation E2 wasGeneratedBy A2), is timestamped with T2.
- The receipt of the delegated service result (i.e. the relation A2 used E3), is timestamped with T3.
- $T4 - T1 > X$, where X is the reasonable upper limit for the service to be provided, and $T3 - T2 > Y$, where Y is some significant portion of X.

The final criterion required for the above patterns to affect Provider A's reputation assessment is to show that Provider A no longer uses Provider B. This could be through *(i)* the fact that recent provenance of Provider A's provision show no use of Provider B, or *(ii)* Provider A's advert for their service specifying which sub-provider they currently use.

We also note that a variation of this pattern is also useful, namely to identify situations in which successful service provision was due to a *good* sub-provider who is no longer used. In this variation the same pattern is used but with poor provision replaced by good provision.

5.2 Freak event

In the second circumstance, the operation of Provider A in providing a service was affected by a one-off substantial event, e.g. ash from an erupting volcano, flooding blocking roads, etc. The freak event can

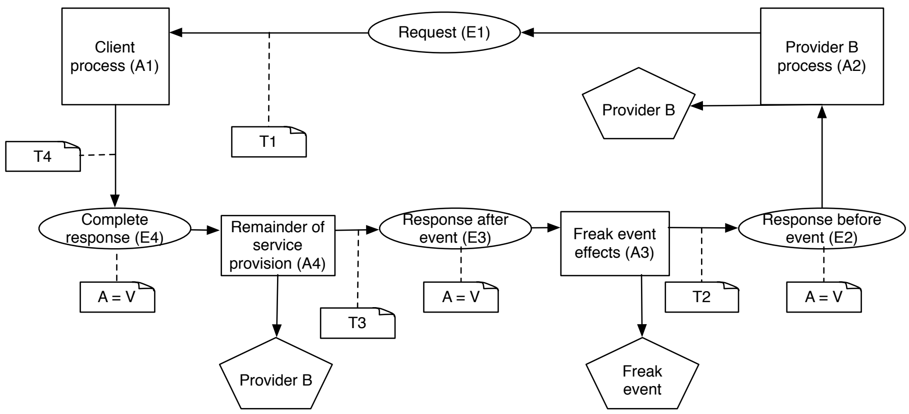

Figure 3: Provenance graph pattern for freak event circumstance

helpfully be considered to be an agent in the provenance, as it is an autonomously acting entity. The provenance should show all of:

1. the effects of a known freak event were part of the process of Provider A providing the service, and
2. the part of the process affected by the freak event was the likely cause of the poor service.

The pattern should show that the effects of the freak event were part of the service provision process, illustrated in Figure 3.

Step 1 A client process, A1, sends a request, E1, for a service to a service process, A2, for which Provider B is responsible. In the provenance graph, this means that E1 wasGeneratedBy A1, A2 used E1, and A2 wasAssociatedWith Provider B.

Step 2 A2 begins providing the service by producing entity E2. E2 wasGeneratedBy A2.

Step 3 The relevant effects, A3, of the freak event affect the service provision, so we distinguish what is provided before those effects, E2, and that provided after the effects, E3. A3 used E2, E3 wasGeneratedBy A3, A3 wasAssociatedWith the freak event.

Step 4 The remainder of the service provision process, A4, completes from the state after the freak event has affected the process, E3, and produces the final provision result, E4. A4 used E3, E4 wasGeneratedBy A4.

Step 5 Finally, the service provision is completed and returned to the client. A1 used E4.

Similar to the first circumstance above, we can then distinguish the cases in which the freak event is the likely cause of eventual poor service. The attributes can indicate that the product before the event (E2) was high quality, while the entity after (E3) was not, e.g. water damage affecting a parcel. The period between the request and response could be primarily due to the length of the effects of the freak event (A3).

5.3 Poor organisation culture

In the third case, Provider A may be an individual within Organisation B. In such cases, the culture of the organisation affects the individual and the effectiveness of the individual affects that of the organisation. If Provider A leaves the organisation, this past relationship should be taken into account: Provider A may operate differently in a different organisational culture. The provenance should show:

1. Provider A provided poor service while working for Organisation B, and
2. Provider A is no longer working for Organisation B.

First, a provenance pattern showing provision of a service within an organisation in a particular instance could be as follows (illustrated in Figure 4).

Step 1 A client process, A1, sends a request, E1, for a service to a service process, A2, for which Provider A is responsible. In the provenance graph, this means that E1 wasGeneratedBy A1, A2 used E1, and A2 wasAssociatedWith Provider A.

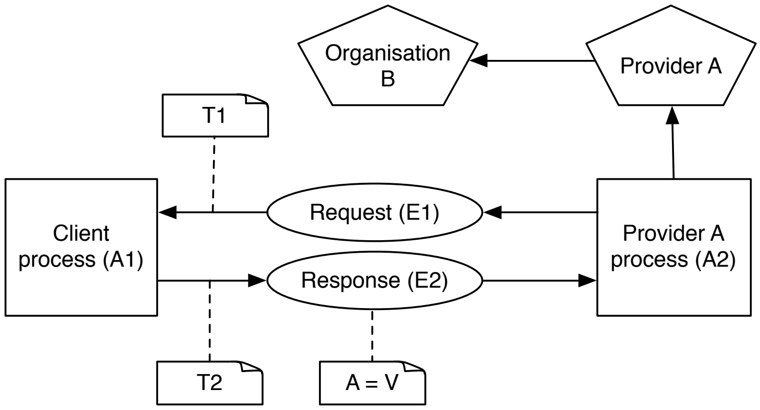

Figure 4: Provenance graph pattern for poor organisation culture circumstance

Step 2 Provider A is acting on behalf of Organisation B in performing A2. In the provenance graph, this means that Provider A actedOnBehalfOf Organisation B in its responsibility for A2 (the last part not depicted in the figure to retain clarity).

Step 3 A2 completes the service provision sending the result, E2, back to A1, so that the client has received the service requested. In the provenance graph, this means that E2 wasGeneratedBy A2, and A1 used E2.

We can then distinguish the cases in which the culture of Organisation B may be a mitigating factor in Provider A's poor provision. Poor performance is identified as described above: either an attribute indicating low quality, a part that is of low quality, or too long a period between the request and response. A variation on the circumstance is to observe where agents were, but are no longer, employed by organisations with a *good* culture.

6 Managing Circumstance Related Failures and Deviations

While some failures and deviations from agreed QoS (quality of service) may be out of the provider's control, the ability to manage and deal with such potential violations is an important factor when making judgments of future interaction partners. In particular, the following additional criteria may need to be considered when assessing a provider's reputation.

- Are there *preventive measures* in place against failures and violations caused by mitigating circumstances, and to what degree of circumstance severity do such measures remain effective? Here, relevant points include:
  - Looking at how others cope in similar circumstances
  - How to judge coping with severity of circumstances (versus competence)
- Are there *recovery plans* in place in the case of failures and deviations from agreed QoS?
  - For example, how well a provider copes with a poor sub-provision: are there backups?
  - Does the provider offer the original service with an additional emergency support service?
- Are there *compensation policies* in place in the case of failures and deviations from agreed QoS, and to what degree of failure/deviation severity are such policies valid? Here, relevant points include:
  - Learning the actual penalty payment of a provider when the one specified by the provider is vague, e.g. the provider will pay a compensation, but this is only up to a particular limit, and is only applicable in specific cases.

- What is the relationship between the severity and frequency of circumstances? This enables risk assessment, and capturing various client preferences in this regard.

7 Inputs to reputation

Reputation systems are generally driven by ratings given to experiences of agents. The direct experience of the assessing agent is often bolstered by those of other advisor agents. Ratings of experiences can be of various aspects of an interaction, such as speed, quality, or cost. In many reputation systems, the recency of experiences affects how much they contribute to the overall reputation assessment. The intuition of this is that recent experiences more representative of the current environment and agents than older experiences. In FIRE, ratings are combined in a weighted sum, where the weights are determined by the recency of the rating [3]. In other models such as TRAVOS, ratings older than a threshold are discarded [9]. BLADE and HABIT aim to maximise the useful information extacted from ratings, both aiming to re-interpret ratings from witnesses, maximizing the information used in reputation computation [6, 8]. BLADE and HABIT are most easily described as Bayesian networks, the latter with a hierarchical structure.

To include circumstances into reputation assessment, the reputation models must accept circumstance patterns as inputs. The ratings can then be weighted in a similar way to recency, whereby ratings that were made in more similar circumstances to the current ones contribute more to the overall reputation assessment. The recency weighting function in FIRE can be changed to a circumstance weighting function, where ratings are weighted based on their circumstance similarity. Likewise, in models such as TRAVOS the ratings made in dissimilar circumstances can be discarded.

For reputation systems to be driven by such experiences, they should be represented with all details of the interaction. The experiences might have the form $(a, b, t, \mathbf{C}, \mathbf{R})$, which describes an interaction between agent a and agent b at time t. $\mathbf{C}$ represents the circumstances of the interaction, and $\mathbf{R}$ are the ratings given to the interaction for each aspect. The reputation system will then process these interaction tuples by extracting circumstance patterns, and combine their ratings to produce an overall reputation assessment.

8 Benefits of recording more information

It is likely to be beneficial to record all events of interactions, as the information required to produce the best reputation assessments in a domain is unknown in general. For instance, if a provider is often rated highly but is given very poor ratings in a small number of cases, this may mean their reputation is not high. If a ship is usually on time and gains high ratings, but may be late on occasion, this may negatively affect its reputation. It may be that the ship is only delayed in stormy weather, however, which if recorded can be used to discount these poor ratings in the reputation assessment.

Knowledge of reactions to circumstances may also be useful in reputation assessments. One ship in a storm may stop and wait for a storm to pass, whereas another may proceed through it. The ship that proceeded may be on time for the delivery, but the cargo may be damaged due to the rough seas. Depending on the context the delivery was requested in (whether or not time or damage was more of a concern), the ratings of the interactions will change. Ideally therefore, all aspects of a service provision and interaction, including preferences and circumstance, should be recorded to be used in reputation assessment.

Recording experiences in such high detail can be difficult to process, as not all information in records may be relevant to reputation assessments. As well as increasing the computational complexity of processing records, this may lead to flaws in reputation assessments. Further, some information may be redundant or add very little information to use in assessing reputation. Extracting the correct information is often a difficult task, and feature selection techniques may be appropriate.

9 Subsidiary use cases

In the JASPR project, given the collaboration with BP, our primary focus is on the logistics use case described above. However, to ensure that the project outputs are transferable to other domains and application areas it is useful to consider alternative use cases. To enable this, we have two subsidiary use cases: (i) on-demand transport in collaboration with Czech Technical University (CTU), and (ii) car leasing in collaboration with Circle Leasing.

9.1 On-demand transport

As investigated by the Agent-Based Computing for Intelligent Transportation Systems group at the Czech Technical University in Prague, the popularity of the market-based approach for on-demand transport services (e.g. adopted by Uber, and also other local transport services such as Lyftago in Prague) has grown rapidly over only a few years. Within this approach, a passenger requests a transport service from a provider, which in turn communicates with potential drivers (bidding for the service), and selects a suitable match accounting for passenger and driver preferences and requirements. Specifically, information and financial exchanges among the three parties are conducted via a multi-stage process: transport request (passenger-provider); driver selection and commission (driver-provider); passenger pricing (provider-passenger); passenger transport (driver-passenger); and driver payment (provider-driver). Passengers are typically interested in truthful providers and drivers, fair prices, on-time driver arrival, safety, car quality, etc. Similarly, when bidding for the service, criteria of interest to drivers include truthfulness of the provider and the passenger, payment (on time and in appropriate amount), passenger punctuality, passenger behaviour, etc. Assessing the reputation of the three parties with respect to such aspects is thus essential to decrease risk and increase satisfaction for both passengers and drivers. Accounting for and reasoning about mitigating circumstances is especially important to achieve accurate and fair assessments. For example, a passenger or a driver should not be penalised for lateness if it is due to events beyond their control, e.g. a late flight or an unexpected bad weather, respectively. Likewise, a low reputation of a provider due to the failure of a driver may not be appropriate if the provider is no longer dealing with this driver.

9.2 Car leasing

After an accident where a vehicle is damaged, an insurance company recommends a body-shop to a policy holder for repairing it, and a courtesy car may then be provided as short-term replacement. The processing of the courtesy car is performed by the leasing company, who provide year-long leases of cars to body-shops for this purpose. This sub-provision relationship is disclosed, but policy holders have little say in general over which body-shop or leasing company they interact with as body-shops are chosen by the insurance company and courtesy cars are selected by body-shops. A small number of body-shops charge damage of courtesy cars and fines to the policy holder through the car leasing company. These extra charges incur a handling fee which although agreed prior to the courtesy period can often aggravate customers into providing unwarranted negative feedback, which could be vindicated if the context was known. Reputation of the car leasing company can be provided by policy holders and body-shops, and similarly the reputation of body-shops from policy holders and insurance companies. Currently a net promoter score for the leasing company is used as a form of reputation, calculated via a yearly questionnaire asking body-shops if they would use or recommend the service and is reported as a ratio for use in promotional material when attracting new clients. There are other forms of reputation available, which may be computed from data recorded during all interactions between the car leasing company, body-shop, insurance company and policy holder. It is thought that timeliness, convenience, and cost have main influence in the car leasing reputation, although a more detailed questionnaire is required to further investigate this. A car leasing company with high reputation will enhance the reputation of a body-shop that uses its services, and the reputation of both will influence the choices of body-shops that insurance companies recommend.

10 Identity management

Identity management is crucial to the effectiveness of any reputation mechanism, and although not the concern of this project, a suitable mechanism is required. We require a mechanism that is able to determine identity, such that when agents interact or share information there is a high degree of confidence regarding their identities. Such a mechanism is crucial to prevent issues including whitewashing and the establishment of multiple identities to falsely build a positive reputation.

In this project we assume the existence of an underlying Public Key Infrastructure (PKI) that enables public keys to be associated with individual identities. A PKI includes one or more Certificate Authorities (CAs) that can digitally sign and publish public keys associated with agents. Several PKI models have been proposed, that define the PKI architecture and the relationships between entities and CAs [1, 4, 5]. For the purposes of this project we do not require commitment to a single approach, leaving the choice to system developers. Note that there is still a risk associated with the PKI approach to identity management, since different credential chains and certification paths have different risks of compromise [7]. Approaches such as Huang and Nicol's calculus of trust enable comparison of the trustworthiness of alternative certification paths [2], although consideration of such reasoning is beyond the scope of this project.

11 Personalisation

Many reputation assessment mechanisms, such as FIRE and TRAVOS, do not consider the subjectivity of agents' experiences, goals, and objectives. More recent methods, such as BLADE and HABIT, begin to consider such subjectivity by aiming to re-interpret the ratings received from witnesses in a way that is most relevant to the assessing agent. For example, if a witness has a history of giving a particular provider a strongly negative rating in situations where the assessing agent's experience is strongly positive this negative correlation can be learnt and exploited in reputation assessment (to suggest that when the witness gives a provider a poor rating, the assessing agent is likely to have the opposite assessment). The incorporation of such subjectivity can be thought of as a primitive form of personalisation. More advanced personalisation mechanisms, combined with the rich representation of circumstances enabled by provenance records, may enable reputation to better reflect the subjectivity that exists in interactions. For example, it is desirable for reputation assessment to take into account the factors that are currently of importance for the assessing agent (e.g. timeliness, quality, or cost). Current methods require the manual specification of the relative importance of such characteristics. It may be possible to incorporate methods from recommender systems and general personalisation techniques to enable this subjectivity to be managed automatically.

References

[1] Samia Bouzefrane, Khaled Garri, and Pascal Thoniel. A user-centric PKI based protocol to manage FC 2 digital identities. *International Journal of Computer Science Issues*, 8(1):74–80, 2011.

[2] J. Huang and D. Nicol. A calculus of trust and its application to PKI and identity management. In *Proceedings of the 8th Symposium on Identity and Trust on the Internet*, pages 23–37, 2009.

[3] Trung Dong Huynh, Nicholas R Jennings, and Nigel R Shadbolt. FIRE: An integrated trust and reputation model for open multi-agent systems. In *16th European Conference on Artificial Intelligence*, pages 18–22, 2004.

[4] John Linn. Trust models and management in public-key infrastructures. Technical report, RSA Laboratories, 2000.

[5] Radia Perlman. An overview of PKI trust models. *IEEE Network*, 13(6):38–43, 1999.

[6] Kevin Regan, Pascal Poupart, and Robin Cohen. Bayesian reputation modeling in e-marketplaces sensitive to subjectivity, deception and change. In *Proceedings of the National Conference on Artificial Intelligence*, pages 1206–1212, 2006.

[7] Michael K. Reiter and G. Stubblebine, Stuart. Authentication metric analysis and design. *ACM Transactions on Information and System Security*, 2(2):138–158, 1999.

[8] WT Luke Teacy, Michael Luck, Alex Rogers, and Nicholas R Jennings. An efficient and versatile approach to trust and reputation using hierarchical bayesian modelling. *Artificial Intelligence*, 193:149–185, 2012.

[9] WT Luke Teacy, Jigar Patel, Nicholas R Jennings, and Michael Luck. Coping with inaccurate reputation sources: Experimental analysis of a probabilistic trust model. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 997–1004, 2005.

[10] W3C. PROV model primer. http://www.w3.org/TR/prov-primer/, 2013.