BSCA-P: Privacy Preserving Coalition Forming Among Rational Web Service Agents

Bastian Blankenburg, Matthias Klusch

In this paper, we propose a coalition formation protocol, called BSCA-P, that allows intelligent agents to negotiate game-theoretically stable coalitions for Web service trading with a maximum of individual monetary profit, while keeping certain kinds of financial data private. We show that no agent has to reveal its local service sales value and final payoff, and can achieve a certain degree of both agent and service anonymity while still successfully participating in rational coalitions.

1 Introduction

In today's increasingly networked and competitive world, the appropriate utilization of pay per use Web services are considered as one major key to the success of commercial service oriented business applications in domains such as e-logistics, tourism, and entertainment. In the near future, intelligent service agents are not only supposed to search for, interact with, and compose, but also negotiate access to, and execute such Web services on behalf of its user, or other agents. In fact, they may exhibit some form of economically rational cooperation by forming coalitions to share the created joint monetary value while at the same time maximizing their own individual payoff. According to classical microeconomics, means and concepts of cooperative game theory are inherently well suited to this purpose.
However, the public revelation of quantity and value of local service sales, and individual requests for particular services required to play cooperative games with complete knowledge could lead to an unsolicited competitive advantage in web service oriented business. The problem is, how can certain kinds of local financial data be kept private while still successfully participating in coalition negotiations to maximize individual profits? Research on privacy preserving coalition formation is in its infancies; first solutions to this problem have been presented in [1, 2][1].

The remainder of this paper is organized as follows. In section 2, we introduce basic notions of service agents, coalition theory, and negotiation used throughout this paper. Section 3 provides an analysis and examples of how certain types of financial information can be kept private during negotiation of coalitions, whereas in section 4, we prove that at the communication level, service requests can only be anonymized by means of an anonymous routing protocol. Finally, the overall coalition formation protocol BSCA-P is then presented with its computational and communication complexity in section 5. We conclude in section 6.

[1] This paper is an extended version of [2].

2 Coalitions of Service Agents

In this section, we introduce the basic notions of Web service agents and cooperative game theory that are required to understand the approach proposed in subsequent sections.

2.1 Service Agents

We consider a Web service to be any kind of task-oriented, XML-based business application software that is location transparent, i.e., network accessible from anywhere with one or multiple protocols of the IP suite, possibly enlarged with additional descriptive metadata to describe its semantics for service consumers, programmable via an API, and loosely coupled with other software applications to implement processes within, or across enterprises. It is supposed to be registered and located by means of web service registries, or intelligent middle agents [5]. Examples of ontology languages for describing Web services range from WSDL for the contemporary Web, to WSDL-S, OWL-S, and WSMO for the future semantic Web.

Unfortunately, in recent literature, the terms agent and Web service are often used interchangeably. An autonomous service agent is a special kind of intelligent information agent [4] that is supposed to pro-actively search for, interact with, and compose, but also negotiate access to, and execute atomic, or composed Web services on behalf of its user, or other agents. In contrast, Web services are considered passive in that they are not expected to be able to, for example, autonomously decide upon its invocation, or intelligently (re-)plan the composition of its own or other services either individually, or in joint cooperation with other services.

There are, in principle, three different ways of how an individual service consumer or provider agent can interact with an network accessible Web service, that is via (1) the service interface, or communication with another service agent that either (2) provides this and possibly multiple other services, or even (3) temporarily integrates parts of the service code into its own on demand, thereby changing the individual reactive agent behaviour accordingly. In this paper, we adopt the second perspective of interaction, and do not differentiate between the offering of atomic, or composite web services WS by service agents.

However, we do assume that each agent a is equipped with an individual model of monetary valuation $w_a(WS)$ of each local, or remote service WS it can deliver to its users. Besides, the local execution of its own services does produce a certain amount of costs $c_a(WS)$ per invocation. Any pair of service agents a, a' is interested to access or execute, respectively, a particular service WS provided by a' only if is possibly profitable to do so, i.e., $w_a(WS) > c_{a'}(WS)$. Since we further assume an individual service agent to act economically rational, it will try to negotiate a profitable joint agreement for cooperation with other service agents in a coalition to maximize its individual payoff. Such an agreement includes the commitment of each coalition member to deliver both relevant local services and those it is planning to compose jointly with other members, as well as the implementation of the negotiated payoff distribution among them. Such kind of rational cooperation between Web service agents can be described in terms of cooperative, or coalition games.

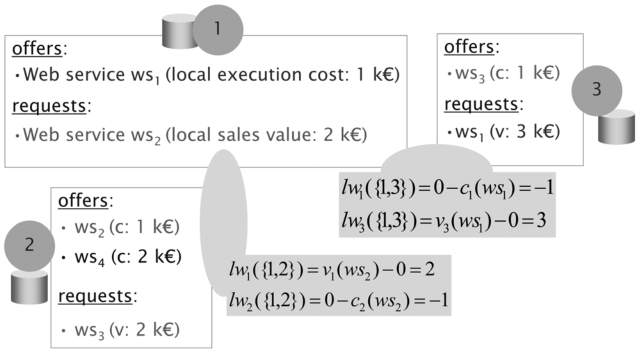

3-Agent Coalition Game

C	{i}	{1,2}	{1,3}	{2,3}	{1,2,3}
v(C)	0	1	2	1	4

Figure 1: Example coalition game for three web service agents.

2.2 Coalition Games

According to microeconomics, a *coalition game* $(\mathcal{A}, v)$ is constituted by a given set $\mathcal{A}$ of service agents, and the value v of every possible joint coalition $C \subseteq \mathcal{A}$ among them. Each coalition value $v(C)$

$$v(C) := \sum_{a \in C} lw_a(C) \tag{1}$$

is the maximum monetary gain that can be achieved by cooperation between the members of coalition C. This gain is defined by the sum of the so called *local worth* $lw_a(C)$ of each agent $a \in \mathcal{A}$ in C as its member.

Let $E_a(C)$ denote the set of services that are executed by $a \in \mathcal{A}$, and $R_a(C)$ the set of services of members of C which are accessed by a. The local worth of a in C

$$lw_a(C) := \sum_{WS \in R_a(C)} w_a(WS) - \sum_{WS \in E_a(C)} c_a(WS) \tag{2}$$

is its total monetary contribution to C (without sidepayments), that is the difference between the local income of the service agent by charging its users for relevant data produced by local, or remote services offered by another coalition member, and the cost of executing its local services as requested.

Example 1 *Consider a 3-agent coalition game as shown in figure 1. Service agent a_1, for example, offers its own web service ws_1 to any other known agent of the game, that are service agents a_2 and a_3. Each local execution of its service would cost a_1 an amount of 1ke, but produces no monetary income as it is of no relevance for its own users. Hence, its self value is zero.*

Agent a_3 is requesting access to service ws_1 from a_1, as it can charge its local users with an total amount of 3ke per use, but does not offer any service of interest for users of a_1 in turn. As a consequence, the local worth of a_1 in a joint coalition with a_3 is $lw_{a_1}(C_1) = -c_{a_1}(ws_1) = -1$ whereas that of a_3 is $lw_{a_3}(C_1) = -c_{a_3}(ws_1) = 3$. Summing up the local worths of all agents in every possible coalition yields the set of coalition values which is the cooperative game to solve by negotiation: What coalitions shall the agents form, and how then to distribute the coalition values to their members?

2.3 Stable coalitions

A solution (S, u) of a cooperative game (A, v) is a partition S of the set of agents, means a set of disjunct coalitions that have been formed together with a distribution u of their coalition values to each agent as member of respective coalitions. This payoff distribution is assumed to be efficient, that is the joint benefit is distributed completely without any loss, and individual rational, such that no agent gets less than it could obtain by staying alone.

As soon as the coalitions have been formed, the computed payoff distribution will be implemented by means of certain side-payments that are to be exchanged among the agents. In our case, each service agent a as a member of a certain coalition C may only claim for some sidepayment $sp_u(a, C)$ by other agents, if the difference

$$sp_u(a, C) := u(a) - lw_a(C); \tag{3}$$

$$sp_u(C^*, C) := \sum_{a \in C^*} sp_u(a, C), C^* \subseteq C \tag{4}$$

between its assigned payoff $u(a)$, that is the money it shall get, and its local worth $lw_a(C)$ in C, that is its local income based on charging its own users, is positive. Otherwise, it has to make a sidepayment of an amount of $|sp_u(a, C)|$ to other agents in C. If the payoffs u are distributed without loss, the same holds for its implementation by exchange of sidepayments between members of a coalition.

Corollary 1 *Let $C \in \mathcal{S}$, $(\mathcal{S}, u)$ be a solution of a game $(\mathcal{A}, v)$. If $C^* = C$, we write $sp_u(C)$. Then $sp_u(C) = 0$, if an only if u is efficient wrt. $\mathcal{S}$.*

A solution is called $stable$, in case no agent could have an incentive to leave its coalition due to its assigned payoff. There exist different stability concepts in game theory from which we adopted, for the work reported in this paper, an efficient variant of the Shapley value [7], the so called *bilateral Shapley value*.

Definition 1 *The union C of two disjoint coalitions $C_1, C_2 \subset \mathcal{A} \setminus \emptyset$ is called a* bilateral coalition, *with C_1 and C_2 called founders of C. A bilateral coalition C is called* recursively bilateral *iff it is the root node of a binary tree denoted by T_C for which (a) every non-leaf node is a bilateral coalition, and its* *founders and sub-coalitions are its children, and (b) every leaf is a single agent coalition. For the depth $d(C^*, T_C)$ of a node C^* in T_C with either $C^* = C$, or $C^* \subset C^{**}$, $C^{**} \in T_C$ it holds that*

$$d(C^*, T_C) = \begin{cases} d(C^*, T_C) = 0 & \text{if } C^* = C \\ d(C^*, T_C) = d(C^{**}, T_C) + 1 & \text{otherwise} \end{cases}$$

A coalition structure $\mathcal{S}$ for $(\mathcal{A}, v)$ is called (recursively) bilateral *if $\forall C \in \mathcal{S}$: C is (recursively) bilateral, or $C = a$, $a \in \mathcal{A}$. The* bilateral Shapley value *$\sigma_b(C, C_i, v)$,$C_i, i \in \{1, 2\}$ of the bilateral coalition C is defined as the Shapley value of C_i in the game $(\{C_1, C_2\}, v)$:*

$$\sigma_b(C_i, C, v) = \frac{1}{2}v(C_i) + \frac{1}{2}(v(C) - v(C_k)) \qquad (5)$$

with $k \in \{1, 2\}, k \neq i$.

Given a recursively bilateral coalition structure $\mathcal{S}$ for a game $(\mathcal{A}, v)$, a payoff distribution u is called recursively bilateral Shapley value stable *iff for each $C \in \mathcal{S}$, every non-leaf node C^* in T_C : $u(C_i^*) = \sigma_b(C_i^*, C^*, v_{C^*}), i \in 1, 2$ with $\forall C^{**} \subseteq \mathcal{A}$:*

$$v_{C^*}(C^{**}) = \begin{cases} \sigma_b(C_k^p, C^p, v_{C^p}) & \text{if } C^p \in T_C, \\ & C^* = C^{**} = C_k^p, k \in 1, 2 \\ v(C^{**}) & \text{otherwise} \end{cases} \qquad (6)$$

In other words, when merging two recursively bilateral coalitions into one its value will be distributed down the corresponding coalition tree to its members by means of recursively replacing the coalition value of the respective parent coalition with its payoff, that is the bilateral Shapley value.

Example 2 *Consider our example game, and the bilateral coalition $C_1 = \{a_1\} \cup \{a_3\}$. Since $v(\{a_1\}) = v(\{a_3\}) = 0$, it holds that $\sigma_b(\{a_1\}, \{a_1\} \cup \{a_3\}, v) = \sigma_b(\{a_1\}, \{a_1\} \cup \{a_3\}, v) = 0 + \frac{1}{2}(2 - 0) = 1$. Merging of C_1 with $C_2 = \{a_2\}$ ($C = C_1 \cup C_2$) yields $v(C) = 4$ and $v(C_2) = 0$, thus $\sigma_b(C_1, C, v) = 2 + \frac{1}{2}(4 - 2) = 3$ and $\sigma_b(C_2, C, v) = 0 + \frac{1}{2}(4 - 2) = 1$. Recursively replacing the coalition value $v(C_i)$ in (5) with the bilateral Shapley value of C_i then leads to the following payoff distribution (cf. figure 2): $u(a_1) = \sigma_b(\{a_1\}, \{a_1\} \cup \{a_3\}, v^*) = 0 + \frac{1}{2}(3 - 0) = 1.5$ and $u(a_3) = \sigma_b(\{a_3\}, \{a_1\} \cup \{a_3\}, v^*) = 0 + \frac{1}{2}(3 - 0) = 1.5$.*

2.4 Negotiation of stable coalitions

The BSCA protocol for negotiating such stable coalitions does restrict negotiation to pairs of voted leaders of coalitions of given maximum size, thereby reducing the communication complexity. Each coalition leader recursively distributes the potential joint coalition value to those agents that are members of its current coalition according to the bilateral Shapley values (cf. figure 2). Coalitions are formed bilaterally per round based on coalition proposals that are mutually accepted based on the expected maximum of individually rational payoffs for the agents involved. However, to determine these potential payoffs, the BSCA protocol requires each agent to reveal its local worth to every potential coalition partner per round.

From the knowledge about the local worth of an agent in some coalition, one could easily deduce, for example, its monetary self value, that is the local income of the agent from selling its own services exclusively to its own users. Further, from the distribution of service requests, and the known

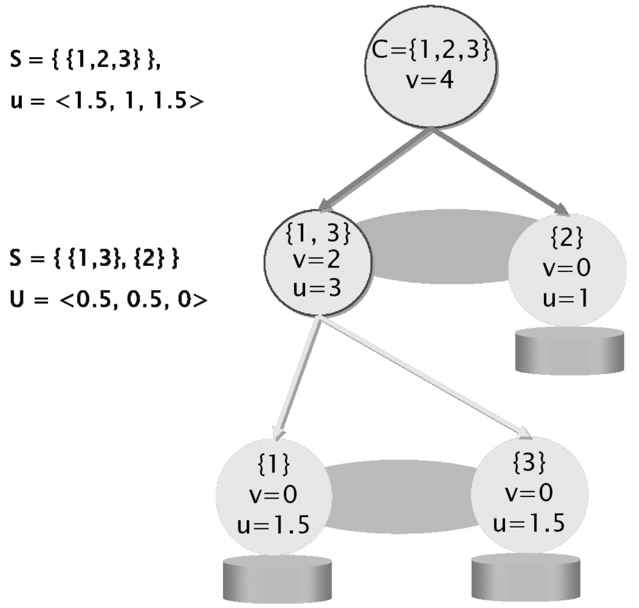

Figure 2: Binary tree of bilateral coalitions for the example game.

set of local worth values, any third party could easily deduce that some agent does apparently have a stronger interest in certain services offered by some agents than by others. These kinds of revelation could lead to an unsolicited competitive advantage of these parties in web service oriented business after, or in parallel to playing this particular coalition game.

In general, that is the problem of how to preserve data privacy in cooperative games playing: To what extent an individual service agent could keep its self values, and expected final payoffs private to other agents such that all agents are negotiating a solution of still the same game that is stable according to the bilateral Shapley value? More general, what is the trade off for any service agent between hiding certain kinds of private financial data from potential collaboration partners, and collectively rational profit making?

3 Non-Disclosure of Financial Data

The basic idea to solve this problem is that each agent should not disclose its total local worth in a potential joint coalition to any other agents but the amount resulting from collaboration only. This so called *additional local worth* is the difference between its local worth in the merger C and its current coalition. In fact, any coalition (leader) C_1 can locally compute its bilateral Shapley value $u_C(C_1) = v(C_1) + \frac{1}{2}av(C_1, C_2)$ in a joint coalition C with some other coalition C_2 simply by means of its self value, and an equal distribution of the *additional joint coalition value* $av(C_1, C_2)$. The latter value is computed by summing up the additional local worths of the agents in each of the bilateral coalition founders. As a consequence, coalition C_1 could compute its expected payoff without knowing anything about the total local worth of its potential coalition partner C_2.

In more detail, (5) can be rewritten as

$$\sigma_b(C_i, C, v) = v(C_i) + \frac{1}{2} \cdot (v(C) - v(C_1) - v(C_2)) \qquad (7)$$

with $i \in \{1, 2\}$. Thus, the *additional coalition value*

$$av(C_1, C_2) := v(C_1 \cup C_2) - v(C_1) - v(C_2) \tag{8}$$

produced by forming coalition $C_1 \cup C_2$ is evenly distributed among C_1 and C_2. For recursively bilateral Shapley value stable payoff distributions, this means that each child node in the coalition tree gets half of the additional payoff of its parent node. The share of the total payoff that a node gets is thus directly dependent on its depth in the tree, which is shown by the following lemma.

Lemma 1 *Let $(\mathcal{S}_1, u_1)$ and $(\mathcal{S}_2, u_2)$ configurations for a game $(\mathcal{A}, v)$, with u_1 and u_2 being recursively bilateral Shapley value stable, and $\exists C_1, C_2 \in \mathcal{S}_1 : C = C_1 \cup C_2 \in \mathcal{S}_2$. Then*

$$\forall C^* \in T_C : u_2(C^*) = u_1(C^*) + \frac{av(C_1, C_2)}{2^{d(C^*, T_C)}}$$

Proof: Induction over $d(C^*, T_C)$. *The case $d(C^*, T_C) = 0$ is obvious because of the efficiency of σ_b and definition of av. For $d(C^*, T_C) = 1$, we have $C^* = C_i$, $i \in \{1, 2\}$ and $u_2(C_i) = \sigma_b(C_i, C, v) = v(C_i) + \frac{1}{2}av(C)$. Again because of the efficiency of σ_b, $v(C_i) = u_1(C_i)$, and thus $v(C_i) + \frac{1}{2}av(C) = u_1(C_i) + \frac{av(C)}{2^{d(C^*, T_C)}}$. In case $d(C^*, T_C) = k > 1$ and lemma 1 holds for all C^{**} with $d(C^{**}, T_C) < k$, we have $C^* = C_i^p$, $i \in \{1, 2\}$, $C^p \in T_C$, $d(C_i^p, T_C) = d(C^p, T_C) + 1$ and $u_2(C_i^p) = \sigma_b(C_i^p, C^p, v_{C_i})$ with $v_{C_i^p}(C^p) = u_2(C^p) = u_1(C^p) + \frac{av(C)}{2^{d(C^p, T_C)}}$. Applying 6 and 7, we get*

$$\begin{aligned} u_2(C_i^p) &= v(C_i^p) + \frac{1}{2}(u_2(C^p) - v(C_i^p) - v(C_k^p)) \\ &= v(C_i^p) + \frac{1}{2}(u_1(C^p) + \frac{av(C)}{2^{d(C^p, T_C)}} - v(C_i^p) - v(C_k^p)) \\ &= v(C_i^p) + \frac{1}{2}(u_1(C^p) - v(C_i^p) - v(C_k^p)) \\ &\quad + \frac{av(C)}{2^{d(C^p, T_C)+1}} \\ &= u_1(C_i^p) + \frac{av(C)}{2^{d(C_i^p, T_C)}} \end{aligned}$$

For the merge of C_1 and C_2 to form $C = C_1 \cup C_2$, we further define the *additional local worth* of agent $a \in C_i$, $i \in \{1, 2\}$:

$$alw_a(C_i, C) := lw_a(C) - lw_a(C_i), \tag{9}$$

and the summarized additional local worth for a subcoalition $C^* \in T_{C_i}$

$$alw(C^*, C_i, C) := \sum_{a \in C^*} alw_a(C_i, C) \tag{10}$$

Also, note that

$$\begin{aligned} av(C_1, C_2) &= \sum_{a \in C} lw_a(C) - \sum_{a \in C_1} lw_a(C_1) - \sum_{a \in C_2} lw_a(C_2) \\ &= alw(C_1, C_1, C) + alw(C_2, C_2, C) \end{aligned} \tag{11}$$

The following theorem shows that in order to compute its sidepayment when merging coalitions C_1 and C_2, each subcoalition $C^* \in T_{C_i}$ only needs to consider its sidepayment for the case without the merge and the additional local worths of C_1, C_2 and C^*:

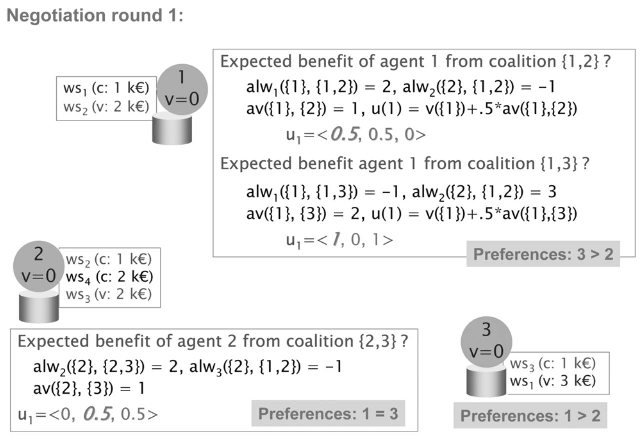

Figure 3: Privacy preserving negotiation of coalitions (round 1).

Theorem 1 *Let $(\mathcal{S}_1, u_1)$ and $(\mathcal{S}_2, u_2)$ configurations for a game $(\mathcal{A}, v)$, with u_1 and u_2 being recursively bilateral Shapley value stable, and $\exists C_1, C_2 \in \mathcal{S}_1 : C = C_1 \cup C_2 \in \mathcal{S}_2$. Then $\forall C^* \in T_{C_i}$, $i \in \{1, 2\}$:*

$$\begin{aligned} sp_{u_2}(C^*, C) =& sp_{u_1}(C^*, C_i) - alw(C^*, C_i, C) \\ &+ \frac{alw(C_1, C_1, C) + alw(C_2, C_2, C)}{2^{d(C^*, T_C)}} \end{aligned}$$

Proof: *Remember that for any u, $sp_u(C^*, C) = \sum_{a \in C^*} u(a) - lw_a(C) = u(C^*) - \sum_{a \in C^*} lw_a(C)$ (see 4). Because of lemma 1, 9 , 10 and 11, we can rewrite*

$$\begin{aligned} sp_{u_2}(C^*, C) &= u_1(C^*) - \sum_{a \in C^*} lw_a(C) + \frac{av(C_1, C_2)}{2^{d(C^*, T_C)}} \\ &= u_1(C^*) - \sum_{a \in C^*} (lw_a(C_i) + alw_a(C_i, C)) \\ &\quad + \frac{av(C_1, C_2)}{2^{d(C^*, T_C)}} \\ &= sp_{u_1}(C^*, C_i) - alw(C^*, C_i, C) + \frac{av(C_1, C_2)}{2^{d(C^*, T_C)}} \\ &= sp_{u_1}(C^*, C_i) - alw(C^*, C_i, C) \\ &\quad + \frac{alw(C_1, C_1, C) + alw(C_2, C_2, C)}{2^{d(C^*, T_C)}} \end{aligned}$$

Please note that in case of $C^* = C_i$, it holds that $sp_{u_1}(C^*, C_i) = 0$, because of $C_i \in \mathcal{S}_1$ and corollary 1. Hence, in order to obtain recursively bilateral Shapley value stable payoff distributions by repeatedly merging coalitions, all subcoalitions have to inform each other only about their additional local worths. Absolute local worths as well as coalition values do not have to be revealed at all. This is in contrast to the traditional way of negotiating stable coalitions with complete prior knowledge about local worth and coalition values that constitute the game to be solved. We acknowledge that this does hold in particular for the bilateral Shapley value but not necessarily for other game-theoretic stability concepts.

Example 3 *Consider, again, our example coalition game (cf. fig. 1). During the first negotiation round, it turns out that agents a_1 and a_3 would prefer each other as a coalition partner, since both of them could obtain a higher individually rational payoff in a joint coalition than each could get in a separate*

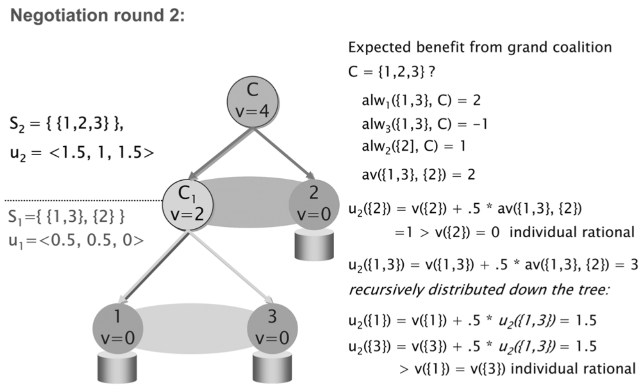

Figure 4: Privacy preserving negotiation of coalitions (round 2).

coalition with agent a_2 (cf. figure 3). Agent a_2 is even indifferent in respect to the coalition it would prefer.

More concrete, $\{a_1\}$ and $\{a_3\}$ form a coalition C_1, with $alw_{a_1}(\{a_1\}, C_1) = -1 - 0 = -1$ and $alw_{a_3}(\{a_3\}, C_1) = 3 - 0 = 3$. According to theorem 1 we get

$$sp_u(\{a_1\}) = 0 + \frac{(-1)+3}{2^1} - (-1) = 2$$

and

$$sp_u(\{a_2\}) = 0 + \frac{(-1)+3}{2^1} - 3 = -2.$$

Thus, the net amount received by a_1 and a_3 are

$$\begin{aligned} u(a_1) &= lw_{a_1}(C_1) + sp_u(\{a_1\}) = -1 + 2 = 1 \\ &= \sigma_b(\{a_1\}, \{a_1\} \cup \{a_3\}, v) \end{aligned}$$

and

$$\begin{aligned} u(a_3) &= lw_{a_3}(C_1) + sp_u(\{a_3\}) = 3 - 2 = 1 \\ &= \sigma_b(\{a_2\}, \{a_1\} \cup \{a_2\}, v). \end{aligned}$$

In the second round, agent a_2 negotiates with the leader of the newly formed coalition C_1 for joining as it is individually rational to do so: Its expected payoff in a potential grand coalition amounts to 1ke, that is it may obtain more by means of cooperation than it would by staying alone. On the other hand, forming of coalition C_1 is consent with this proposal for the same reason: Its bilateral Shapley value of 3ke, recursively distributed down the coalition tree to agents a_1 and a_3, yields a rational expected payoff for both members.

More concrete, their additional local worths in the grand coalition C are

$$\begin{aligned} &alw_{a_1}(\{a_1\}, C) = 1 - (-1) = 2, \\ &alw_{a_3}(\{a_3\}, C) = 2 - 3 = -1 \\ &alw(C_1, C_1, C) = alw_{a_1}(\{a_1\}, C) + alw_{a_3}(\{a_2\}, C) = 1 \\ &alw(C_2, C_2, C) = 1 - 0 = 1 \end{aligned}$$

The additional coalition value is thus

$$av(C_1, C_2) = alw(C_1, C_1, C) + alw(C_2, C_2, C) = 2$$

Applying theorem 1 again, we get the new payoff distribution u^* with

$$sp_{u^*}(C_1) = 0 + \frac{1+1}{2^1} - 1 = 0$$

$$(= sp_{u^*}(C_2))$$

The net payoffs of C_1 and C_2 are equal to their bilateral Shapley values:

$$\begin{aligned} u^*(C_1) &= lw_{a_1}(C) + lw_{a_3}(C) + sp_{u^*}(C_1) \\ &= 1 + 2 + 0 = 3 = \sigma_b(C_1, C, v) \\ u^*(C_2) &= lw_{a_2}(C) + sp_{u^*}(C_2) \\ &= 1 + 0 = 1 = \sigma_b(C_2, C, v) \end{aligned}$$

For sidepayments within C_1, we again apply theorem 1:

$$\begin{aligned} sp_{u^*}(\{a_1\}, C) &= sp_u(\{a_1\}, C_1) + \frac{1+1}{2^2} - 2 \\ &= 2 + 0.5 - 2 = 0.5 \\ sp_{u^*}(\{a_3\}, C) &= sp_u(\{a_3\}, C_1) + \frac{1+1}{2^2} + 1 \\ &= -2 + 0.5 + 1 = -0.5 \end{aligned}$$

Consequently, the net payoffs of a_1 and a_3 are equal to their recursively bilateral Shapley value stable payoffs:

$$\begin{aligned} u^*(a_1) &= lw_{a_1}(C) + sp_{u^*}(a_1) \\ &= 1 + 0.5 = 1.5 = \sigma_b(\{a_1\}, C, v) \\ u^*(a_3) &= lw_{a_3}(C) + sp_{u^*}(a_2) \\ &= 2 + (-0.5) = 1.5 = \sigma_b(\{a_3\}, C, v) \end{aligned}$$

4 Anonymity of Service Requests

Another issue of privacy concerns the non-disclosure of private non-financial information of an individual service agent such as the number and kind of services it does request from some other agent. Even if service agents were enforced to negotiate stable coalitions based on the exchange of their additional local worths only, the question is whether they still would be able to deduce such kind of knowledge about service oriented interests of potential competitors from the set of additional local worths?

Unfortunately, it turns out that this indeed is possible. For example, consider a bilateral coalition $C = C_1 \bigcup C_2$ with $alw(C_1, C) > 0$. From this information, one can deduce that $lw(C_1, C) > lw(C_1, C_1) = v(C_1)$, which implies that agents in C_1 produce more value, and/or less costs than in C. That, in turn, means that at least one agent in C_1 did request services that are offered by agents in C_2. This kind of reasoning chain can be recursively applied to every sub-coalition of C_1 in the coalition tree. In particular, the first coalition partners of an agent, that are its direct siblings in the coalition formation tree, know that it did request some services from them. There is no way to hide this fact other than by committing each of them to keep it private, and trust them to do so.

Though the existence of service requests of any individual agent in coalition C_1 can be detected by other agents in coalition C_2, it turns out that they can be anonymized, thus providing the agents with a weaker notion of privacy at least. To measure degrees of anonymity, different notions have been proposed in the literature, such as total, or group anonymity, under possibilistic or probabilistic interpretations [6, 3]. In fact, if some agent in C_1 does request some service offered by another agent in C_2, the rest of the agents in C_1 could readily observe that, but do not know what kind

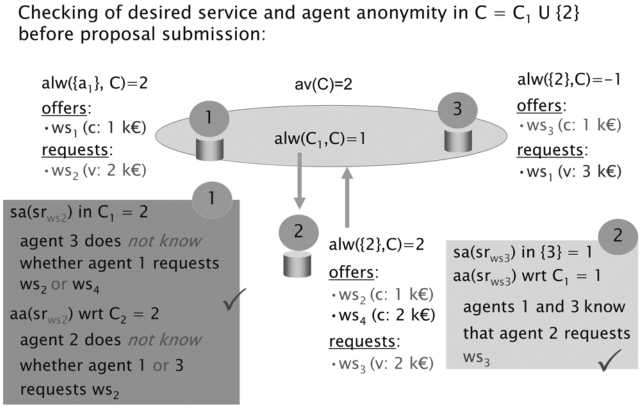

Figure 5: Individual service request anonymities.

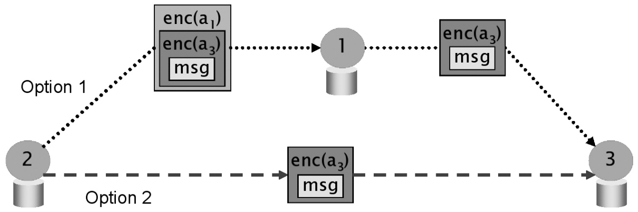

Figure 6: Options of encrypted service request message "onion" routing from agent a_2 to agent a_3.

of service it is. Likewise, the service provider agent in C_2, with $|C_2| \geq 2$ knows that its service has been requested by an agent in C_1 but not which one (cf. figure 5).

We can quantify these kinds of possibilistic anonymity for each service WS requested by an individual agent $a \in C_1$ in terms of

- *service anonymity* $sa(WS, C_1) = |\bigcup_{a \in C_2} OS_a|$ within C_1 in terms of the number of services offered by members of C_2, such that, in the extreme, no agent knows which of its coalition partners does access what specific service, and
- *agent anonymity* $aa(WS, C_2) = |C_1|$ with respect to C_2 in terms of the size of its actual coalition C_1, since from the perspective of agents in C_2, any agent in C_1 might be the originator of the service request.

Assuming that each agent specifies its desired (default) minimum degrees of service and agent anonymity for each web service WS it is interested in, any request and coalition proposal to potential cooperation partners will be submitted, i.e., $WS \in R_a(C)$, if and only if these requirements are met.

To maintain the above mentioned types of anonymity also at the communication level, we adopt the simple onion routing protocol [8] to anonymize the exchange of service request messages between the service agents. In essence, each service request message gets routed between sender and receiver via randomly selected intermediate agents each of which encrypting the message with its individual public key (cf. figure 6). This way, for communication paths consisting of at least three agents, no intermediate agent is able to determine both the origin and the receiver of a service request message nor to decrypt its content to some extent as guaranteed by the underlying encryption protocol.

5 Coalition Formation Protocol BSCA-P

In this section, we finally propose the coalition formation protocol BSCA-P that makes use of all concepts and means that have been introduced in the previous sections. We assume that service offers along with service execution costs are known in prior.

Algorithm 1 *For a game* $(\mathcal{A}, v)$, $\mathcal{S}_0 := \{\{a\}|a \in \mathcal{A}\}$, $r := 0$ *and* $\forall C \in \mathcal{S}_0 :\ sp_0(C) := 0$. *In every coalition* $C \in \mathcal{S}_r$, *every agent* $a \in C$ *performs:*

1. *Let* $C \in \mathcal{S}_r,\ a \in C$ *and* $\mathcal{S}^* := \mathcal{S} \setminus C$.
2. Communication:
   (a) *For all* $C^* \in \mathcal{S}^*$ *do:*
      i. *Determine set* $R_a(C^*)$ *of requests, subject to the sets* OS_{a^*} *of offers for each* $a^* \in C^*$, *costs and minimum anonymity degrees.*
      ii. *For each service request which is in* $R_a(C) \cap R_a(C^*)$ *keep the one with minimum costs.*
      iii. *Set* $alws_a(C^*) := alw_a(C, C^*)$.
      iv. *For each bilateral coalition* $C^a,\ C^a \in T_C,\ a \in C^a,\ a = Rep(C_a)$, *wait for a message from* $Rep(C^a_i),\ i \in 1, 2, a \notin C^a_i$ *containing* $alws_{Rep(C)}(C^*)$ *and set* $alws_a(C^*) := alws_a(C^*) + alws_{Rep(C)}(C^*)$.
      v. *If* $a = Rep(C)$ *then send* $alws_a(C^*)$ *to* $Rep(C^*)$; *else send* $alws_a(C^*)$ *to* $Rep(C^+)$ *with* $C^+ \in T_C,\ a = Rep(C^+_i),\ i \in 1, 2, a \neq Rep(C^+)$.
   (b) *If* $a = Rep(C)$ *then receive* $alws_{Rep(C^*)}(C)$ *and set* $alws(C^*) := alws_{Rep(C^*)}(C) + alws_a(C^*)$ *for all* $C^* \in \mathcal{S}^*$; *else go to step 3i.*
3. Coalition Proposals:
   (a) *Set* $Candidates := \mathcal{S}^*$, $New := \emptyset$ *and* $Obs := \emptyset$
   (b) *Determine a coalition* $C^+ \in Candidates$ *with* $\forall C^* \in Candidates :\ alws_a(C^+) \geq alws_a(C^*)$.
   (c) *Send a proposal to* $Rep(C^+)$ *to form coalition* $C \cup C^+$.
   (d) *Receive all coalition proposals from other agents.*
   (e) *If no proposal from* $Rep(C^+)$ *is received and* $Candidates \neq \emptyset$, *set* $Candidates := Candidates \backslash \{C^+\}$ *and go to step 3b.*
   (f) *If a proposal from* $Rep(C^+)$ *is received, then form the coalition* $C \cup C^+$:
      i. *If* $o(Rep(C)) < o(Rep(C^+))$ *then set* $Rep(C \cup C^+) := Rep(C)$; *else set* $Rep(C \cup C^+) := Rep(C^+)$.
      ii. *Inform all other* $Rep(C^*), C^* \in \mathcal{S}^* \setminus C^+$ *and all* $a^* \in C,\ a^* \neq a$ *about the new coalition and* $Rep(C \cup C^+)$
      iii. $New := \{C \cup C^+\}$, $Obs := \{C, C^+\}$
   (g) *Receive all messages about new coalitions. For each new coalition* $C_1 \cup C_2$ *and* $Rep_{C_1 \cup C_2}$, *set* $Candidates := Candidates \setminus \{C_1, C_2\}$, $New := New \cup \{C_1 \cup C_2\}$ *and* $Obs := Obs \cup \{C_1, C_2\}$.
   (h) *Send the sets* New *and* Obs *to all other coalition members* $a^* \in C,\ a^* \neq a$
   (i) *If* $a \neq Rep(C)$ *then receive the sets* New *and* Obs *from* $Rep(C)$.
   (j) *Set* $r := r + 1$, $\mathcal{S}_r := (\mathcal{S}_{r-1} \setminus Obs) \cup New$.
   (k) *For each (sub-)coalition* $C^* \in T_C$ *with* $Rep(C^*) = a$, *determine* $sp_r(C^*)$ *according to theorem 1 (using* $sp_{r-1}(C^*)$ *instead of* $sp_u(C^*)$*).*
   (l) *If* $C_r = C_{r-1}$ *then stop; else go to step 2*

Theorem 2 *Let $n = |\mathcal{A}|$ and $m := max_{a \in \mathcal{A}}\{|R_a|\}$. The computational complexity of the protocol BSCA-P is in $O(n^3 m^2)$. The communication complexity in terms of the number of exchanged messages per agent is in $O(n^2)$.*

Proof: cf. [2]

After stable coalition configurations have been negotiated among service agents following the BSCA-P protocol, it will be implemented by exchange of actual sidepayments. For this purpose, each leader of a (sub-)coalition C makes or receives payments sp to, or from other leaders of immediate parent and child coalitions in the binary coalition tree. this way, only leaders of 2-agent coalitions get informed about individual sidepayments, that are its own, and that of the other agent. As a consequence, only the very first coalition partner of an individual agent a, that is its direct neighbour leaf in the coalition tree, might ever know a's exact sidepayment, though its individual utility value still remains private. To ensure anonymous service requests and access, we require each agent to follow the simple onion routing protocol.

6 Conclusions

We proposed a protocol for privacy preserving and stable coalition formation among rational web service agents. In particular, the payoffs and utilities of these agents can almost or even completely be kept private, respectively, during bilateral negotiations of recursively bilateral Shapley value stable coalitions. Further, following the BSCA-P protocol, there is no need to reveal absolute coalition values to successfully participate in coalition negotiations at all. However, we showed that the existence of service requests might not be hidden in general but anonymized to a specified degree. In summary, the negotiation protocol BSCA-P allows service agents in the Internet to keep personal financial data private, while increasing their individual profits by means of rational cooperation with others in coalitions.

Acknowledgement: This work has been supported by the German ministry for education and research (BMBF) under project grant SCALLOPS-01-IW-D02.

References

[1] B. Blankenburg and M. Klusch. On safe kernel stable coalition forming among agents. In *Proc. 3rd Int. Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), New York, USA*, 2004.

[2] B. Blankenburg and M. Klusch. Bsca-p: Privacy-preserving coalition formation. In F. Kluegl et al., editor, *Proc. 3rd German Conference on Multi-Agent System Technologies (MATES), Koblenz, Germany*. Springer, LNAI, 3550, 2005.

[3] Joseph Halpern and Kevin O'Neill. Anonymity and information hiding in multiagent systems. *Journal of Computer Security*, Special Edition on CSFW 16:75–88, 2003.

[4] M. Klusch. Information agent technology for the internet: A survey. *Data and Knowledge Engineering*, 36(3), 1980.

[5] M. Klusch and K. Sycara. Brokering and matchmaking for coordination of agent societies: A survey. In A. Omicini et al., editor, *Coordination of Internet Agents*. Springer, 2001.

[6] A. Pfitzmann and M. Köhntopp. Anonymity, unobservability and pseudonymity: a proposal for terminology. In *International Workshop on Designing Privacy Enhancing Technologies*, pages 1–9, New York, 2001. Springer-Verlag.

[7] L. S. Shapley. A value for n-person games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games II*, volume 28 of *Annals of Mathematics Studies*, pages 307–317. Princeton University Press, Princeton, 1953.

[8] P F Syverson, D M Goldschlag, and M G Reed. Anonymous connections and onion routing. In *IEEE Symposium on Security and Privacy*, pages 44–54, Oakland, California, 4–7 1997.

Contact
Bastian Blankenburg, Dr. Matthias Klusch
DFKI GmbH
Stuhlsatzenhausweg 3, 66123 Saarbrücken
Tel.: +49 (0)681 302-64823
Email: {blankenb, klusch}@dfki.de

Bastian Blankenburg is a PhD student of the university of the Saarland and researcher of the multi-agent systems group of the German research center for Artificial Intelligence (DFKI). His research interests include cooperation, risk, trust and privacy in multi-agent systems. He received a diploma from the University of the Saarland, Germany.

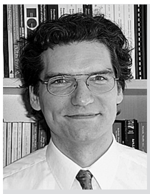

Matthias Klusch is a research fellow and senior researcher of the German research center for Artificial Intelligence (DFKI) where he is leading a research team on intelligent information agents and systems. His research interests include agent-based and service-oriented computing, intelligent information search and management, innovative applications of AI, rational cooperation and decision-making, and the semantic Web. He received a PhD in computer science from the University of Kiel, Germany.