Published in OTM Conferences, pp. 90-108, 2015
Springer link: http://link.springer.com/chapter/10.1007%2F978-3-319-26148-5_6

Finding Collective Decisions: Change Negotiation in Collaborative Business Processes

Walid Fdhila[1], Conrad Indiono[1], Stefanie Rinderle-Ma[1], and Rudolf Vetschera[2]

[1] University of Vienna, Faculty of Computer Science
[2] University of Vienna, Faculty of Business, Economics and Statistics
email: {firstname.secondname}@univie.ac.at

Abstract. Change propagation has been identified as major concern for process collaborations during the last years. Although changes might become necessary for various reasons, they can often not be kept local, i.e., at one partner's side, but must be partly or entirely propagated to one or several other partners. Due to the autonomy of partners in a collaboration, change effects cannot be imposed on the partners, but must be agreed upon in a consensual way. In our model of this collective decision process, we assume that each partner that becomes involved in a negotiation has different alternatives on how a change may be realized, and evaluates these alternatives according to his or her individual costs and benefits (utilities). This paper presents models from group decision making that can be applied for handling change negotiations in process collaborations in an efficient and fair way. The theoretical models are evaluated based on a proof-of-concept prototype that integrates an existing implementation for change propagation in process collaborations with change alternatives, utility functions, and group decision models. Based on simulating a realistic setting, the validity of the approach is shown. Our prototype supports the selection of change alternatives for each partner during negotiation that depending on the group decision model used, provides solutions emphasizing efficiency and/or fairness.

1 Introduction

Collaborative business processes, in which a set of partners collaborate in order to achieve a common business goal G [2], have become an important way of coordinating economic activities. For example, in virtual factories different manufacturers collaborate in the production of goods like cars [25]. Technically, this is realized by a process choreography where the role of each partner is defined in the form of a public process. The latter describes the way a partner interacts with the other partners and the required data to be exchanged for the collaboration (inputs and outputs) [28]. Each role in the choreography; i.e., public process, is associated with a public goal G_i. In order to fulfill its role reflected by the public process, a partner develops its business logic based on a so called private process that is consistent with the public process. Due to confidentiality reasons, the details of the private process are hidden from the partners, only the public process as an abstraction of the private process is visible to the outside.

It should be noted that for fulfilling a role, a partner might develop alternative consistent private processes [27], each with an associated cost and possibly different private goal g_i[3]. Indeed, according to a partner business strategy, the private and public goals can diverge, but they should always align with the common goal G (of the collaboration). Since the private process represents the business logic of a given partner, it cannot be fully transparent to the other partners. Consequently, a partner can solely view the public process of the other partners, which serves as part of the SLA (Service Level Agreement) [17]; i.e. the contract. The SLA can also include non-functional requirements such as QoS or costs [17].

It is optimistic to say that once the collaboration is set, the business processes will not change. In fact, due to many factors; e.g., optimization, evolving business needs, changing laws (compliance), a process is often subject to change [6, 9]. Different alternatives of change formulations (process changes) can correspond to the same business change. For example, due to a change of the marketing strategy, different approaches can be adopted, and consequently different process configurations.

In a collaboration, a change rarely confines itself to a single partner, but might lead to knock-on effects on the associated partners; i.e., change propagation. As set out in [7], change propagation is realized by following these three steps:

- Private-to-public propagation (Pr2Pu): Changes are propagated from the private process to the corresponding public process of the same partner. This has consequences on the public goal of the partner G_i.
- Public-to-public propagation (Pu2Pu): Changes are propagated to the affected partners; i.e. the effects on their public processes. This has consequences on the common business goal G and the public goals G_i of the affected partners.
- Public-to-private propagation (Pu2Pr): Changes are propagated to the corresponding private processes of the affected partners. This has consequences on the goals g_i

Due to the autonomy and independence of partners in a collaboration, change effects cannot be imposed on the partners, but must be agreed upon in a consensual way. Hence, a propagation to a partner might result in a potentially costly and time-consuming negotiation process.

In a collaboration, two constellations are possible: (i) all partners know each other (which entails a negotiation that involves all partners) (ii) some partners are only visible to a subset of partners, for example, in supply chains. Specifically, constellation ii) entails P2P negotiations that involve only a subset of partner, or transitive negotiations. Consequently, with respect to the constellation, a full or P2P negotiation can be envisaged.

The arising research question is: How to find a collective decision on the concrete realization of a change propagation among the affected partners in a process collaboration in a fair and efficient way? A solution should specifically

[3] Compare to different goals for process variants as described in, e.g., [21].

support the affected partners to choose from their set of possible change alternatives associated with respective goals g_i the best one for them given a fair and efficient solution for the entire choreography and its associated common goal G. Addressing this research question requires a synthesis of research from the field of group decisions and negotiations on the one hand, and process change on the other hand, which poses new challenges for both fields. The present paper aims at this synthesis. Although the paper thus integrates known results from both fields (in particular, well known concepts in the area of group decisions), this synthesis is still important. Firstly, the concepts presented here were, to the best of our knowledge, previously not considered in the area of process management. Furthermore, concepts from group decisions usually deal with abstract decision alternatives. In the present paper, we show how models of cooperative processes, and changes made both to public and private processes, can be mapped to the level of abstraction required by group decision models.

This paper presents models from group decision making that can be applied for handling change negotiations in process choreographies in an efficient and fair way. The theoretical models are evaluated based on a proof-of-concept prototype that integrates existing implementation for change propagation in process choreographies with change alternatives, utility functions, and group decision models. Based on simulating a realistic setting, the validity of the approach is shown. Our prototype supports the selection of change alternatives for each partner during negotiation that depending on the group decision model used, provide solutions emphasizing efficiency and/or fairness.

The remainder of the paper is structured as follows: Section 2 introduces change negotiation scenarios and describes a motivating example. Section 3 formalizes the problem of change negotiations in process choreographies and Section 4 introduces three evaluation approaches of change alternatives. Section 5 evaluates and analyzes the approaches through a prototype proof of concept. Finally, Section 6 describes the related work and Section 7 concludes the paper.

2 Change Negotiation Scenarios and Motivating Example

This section explains different negotiation scenarios and introduces a motivating example to illustrate the change negotiation problem.

2.1 Change Negotiation Scenarios

As aforementioned, change alternatives are derived from a set of possible process changes across partners, which could result from different scenarios. In particular, we can distinguish three different scenarios:

1. One partner (say partner A) has to make a mandatory change to its (private) processes, for example because of a change in legal requirements, which renders the existing process impossible (e.g., because it now violates some new legal requirements). For this scenario, we assume that there is only one possible change to A's private process.

2. One partner initiates a change, either because of a change in the environment, or because the partners wants to change the process to improve it or to exploit a new business opportunity. In contrast to scenario 1 above, we assume for this scenario that there are several possible private changes of A, which are then propagated through the network. All possible change alternatives considered result from this propagation.
3. In addition to scenario 2, we now also assume that the other partner B responds to the propagated change proposals of A by proposing alternative changes to the public processes (which result from a propagation of changes to the private processes of B, that B develops in response to the changes proposed by A).

If each change in the private process of one partner results in exactly one possible change of the public process of that parter, and each change in the public process of one partner requires exactly one change in the public process of the other partner, then scenario 1 leads to a unique outcome at the group level and does not require any further negotiation between partners. The change in partner B's public process could still be implemented via several different changes in partner B's private processes. However, since these changes do not affect partner A, they do not need to be decided at the group level. As explained above, partner B can select the change to its private processes that on the one hand implements the required change to partner B's public processes, and on the other hand optimizes partner B's private goal.

In contrast, scenario 2 provides an opportunity for a real group decision. The different changes proposed by partner A will most likely not be exactly equivalent even to that partner. Although it is unlikely that one partner will propose a change which is extremely bad from that partner's perspective, one can expect partners to include change alternatives which are slightly worse to them than other proposals, in order to broaden the set of alternatives. Since the situation is not a zero sum game, it is possible to have some alternatives, which are better for all partners, and thus to have efficient alternatives and actually create value for both partners. However, since it can be difficult for one partner to estimate the full consequences of proposed changes for the other partner, it is quite likely that the range of outcomes present in the set of alternatives under discussion will be different for partners. Alternatives will likely be quite similar from the perspective of partner A initially triggering the change, but could vary considerably from the perspective of partner B (if some of these alternatives have very negative effects on partner B, of which partner A is initially not aware).

In particular in such a setting, the third scenario will likely take place, and partner B will also make some counter-proposals for changes to the public processes. These changes must then be propagated from B's public processes to A's public processes, and A must find changes in its private processes to implement them. Since B might also be unaware of the difficulties some of these changes will cause for A, it is to be expected that this set of alternatives will span a wider range of outcomes also for A compared to scenario 2 above.

If such counter-proposals are made, the entire set of process changes considered is the union of the set of changes proposed by A and those proposed by B. One can also imagine a situation in which the intersection of the proposals of both parties is considered. However, since this is a dynamic process, that would in fact mean that B is granted a veto on proposals made by A, he can eliminate them by not nominating them himself. In such a setting, it would not make sense for B to propose any change that has not already been proposed by A, since it would also not be contained in the intersection. Thus, using the intersection would lead to a quite narrow set of alternatives, which are very similar both from A's and from B's perspective. We thus consider an approach using the union of both sets of proposals to be more suitable to find creative solutions.

2.2 Motivating Example

Consider the collaboration scenario between a *bank* and an *insurance company* as shown in Fig. 1 to provide new services; e.g., retirement funds and financial consultancy. The collaboration combines expertise and customer base of both partners in order to increase the benefits in terms of reputation and marketshare; i.e., the **global goal** of the collaboration.

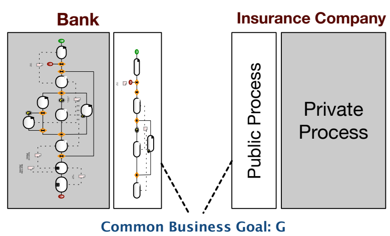

	Private Goal	Public Goal	Common Goal
Bank	Increase customer base by 10% p.a.	-Increase customer base by 3% p.a. -Increase product value by 4% p.a.	Increase reputation and market share
Insurance Company	Increase product value by 9% p.a.	-Increase customer base by 7% p.a. -Decrease financial costs by 5% p.a.	

Business Goals Table

Fig. 1. Process Collaboration: Example from Financial Domain

Each partner contributes to this global goal in a different way (through its public process), and aligns with a public goal. For example, the *bank* contributes through its financial expertise with the **public goal** of increasing its customer base by 3% p.a. and its product value by providing additional insurance offers. Similarly, the *insurance company* aims at increasing its customer base by 7% p.a. and reducing its financial costs through its partnership with the bank. Additionally, each of both partners holds a **private goal**, which is aligned with its private business process. For instance, one of the private goals for the *bank* constitutes an increase of the total customer base by 10% p.a., of which 3% is covered by this collaboration.

Now, we assume that the *bank* wants to apply a new marketing strategy through a categorization of its customers. This leads to a major change in the private process of the *bank*, by for example, creating a different procedure for each customer category; e.g., *gold category* customers receive additional insurance coverage and high saving interest rates. It should be noted that the implementation of this marketing strategy can be achieved in many ways (i.e., different alternatives). Each of these alternatives can directly affect the public model and consequently the collaboration with the *insurance company*. Similarly, each of them corresponds to a change alternative for the *insurance company*'s public process with a different impact on its public and private goals.

With respect to the multitude of choices; i.e., change alternatives, a negotiation process is required to decide on the changes to be implemented. The negotiation involves the *bank* and the *insurance*, which collectively try to find an alternative that is aligned with the common goal and the respective private and public goals. An evaluation of the alternative costs helps selecting a fair and efficient solution.

3 Problem Formulation

This section formalizes the problem of change negotiations in process choreographies.

3.1 Process changes

We model possible process changes as discrete alternatives. Each such alternative represents a coherent change made to a public or private process model, and alternatives are exclusive in the sense that exactly one of the alternatives can be implemented. This implies that different variants of changing the process (even if they are different only on a small part of the total change) are considered as two different alternatives, and that maintaining the current process is also one of the alternatives (although it might have rather negative consequences, e.g., if due to a change in regulations the current process is no longer allowed and would lead to high fines).

We denote the j-th alternative for a change in partner i' *private* process by δ_i^j and the j-th alternative for a change in that partner's *public* process by σ_i^j. In the present paper, we only consider two actors $i \in \{A, B\}$.

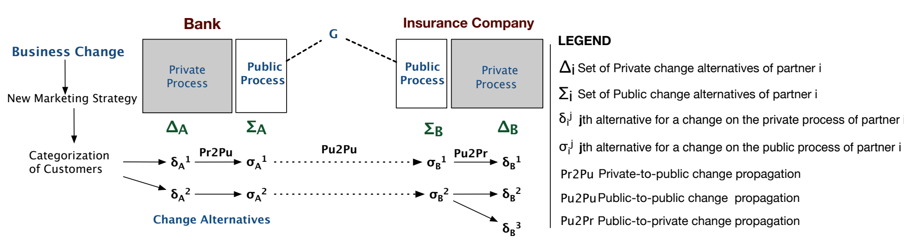

Fig. 2. Example from Financial Domain: Propagation and Negotiation Scenario

As described in the example from financial domain (cf. Figure 2), changes in the public and private processes of the same partner are not independent of each other. A change to the public process can possibly be implemented by different changes to the private process. We denote the set of possible changes to the private process, which can be used to implement change σ_i^j to the public process by $\Delta_j(\sigma_i^j)$.

3.2 Goals and preferences

We distinguish three levels of goals in the model:

1. *Common goals* G refer to the goals of the collaboration. They depend only on the public process, and are shared by all partners.
2. *Public goals* G_i of partner i are goals which are relevant only to partner i, but of which all partners are aware, and which are also influenced only by the public processes.
3. *Private goals* g_i of partner i are also only relevant to partner i, and furthermore they are only known to that partner. Their fulfillment depends on the private and public processes of partner i, and could also be influenced by the public processes of the other partners.

For simplicity, we consider only one common goal, and one public and private goal per partner in this paper. Extension to multiple goals at each level is straightforward, but would increase the complexity of notation without providing much additional insight for the purpose of the present paper.

The dependence of these goals on changes to the public and private processes can thus be expressed as

$$G = G(\sigma_1^j, \sigma_2^j, \ldots, \sigma_N^j) \tag{1}$$

$$G_i = G_i(\sigma_i^j, \sigma_{-i}^j) \tag{2}$$

$$g_i = g_i(\delta_i^j, \sigma_i^j, \sigma_{-i}^j) \tag{3}$$

where $-i$ refers to all partners except partner i.

Only the private goal of a partner is influenced by changes to the private process of that partner. Thus we can always assume that a partner selects the change to its own private process, that best fulfills the private goal of the same partner.

$$\begin{aligned} \delta_i^j = & \arg\max_{\delta_i^j} \; g_j(\delta_i^j, \sigma_i^j, \sigma_{-i}^j) \\ & \text{s.t.} \\ & \delta_i^j \in \Delta_j(\sigma_i^j) \end{aligned} \tag{4}$$

Although this implies a sequence between changes to the public and the private process in which the public process is changed first, this does not imply that the entire change process cannot be triggered by a change to the private process of one partner, only that final adjustments are to be made to the private process after fixing the public processes (cf. Section 2.1 above for details).

Consequently, the impact of changes on all goals (common goal, public and private goals of each partner) can be seen as being dependent on the changes to the public processes of all partners. The common decision problem of the partners thus consists in agreeing on a new set of public process models. Each design alternative, which can be considered by the partners, therefore consists of a vector of changes to each of the partner's public process model, which can be written as $(\sigma_i^j, \sigma_2^j, \ldots, \sigma_N^j)$, or more specifically in the case of two partners considered here as (σ_A^j, σ_B^j). For simplicity, we write this entire vector of changes as $\boldsymbol{\sigma} = (\sigma_i^j, \sigma_2^j, \ldots, \sigma_N^j)$

When evaluating possible process changes, each partner takes into account the common goal as well as its public and private goals. We represent this simultaneous overall evaluation of all goals via a multi-attribute utility function [14]

$$u_i = w_i^G v_i^G(G) + w_i^p v_i^p(G_i) + w_i^r v_i^r(g_i) \tag{5}$$

where w_i^G, w_i^p, and w_i^r refer to the weights which partner i assigns to the common goal, its own public goal and its own private goal, respectively, and $w_i^G + w_i^p + w_i^r = 1$. The functions v_i^G, v_i^p and v_i^r are the partial utility functions of that partner which transform these goals onto a common utility scale. For simplicity, we assume that these functions are linear and are scaled so that the worst possible outcome in each goal is assigned a partial utility value of zero and the best outcome in each goal is assigned a partial utility value of one. Thus, for the common goal, the partial utility function (which in that case is identical for all partners) is

$$v^G = \frac{G - \underline{G}}{\overline{G} - \underline{G}} \tag{6}$$

where $\overline{G}$ and $\underline{G}$ are the best and worst possible values of the common goal, respectively. Partial value functions for the other goals can be defined analogously.

Since outcomes in all goals depend on changes to the public process models, the total utility for each partner will also depend on these changes:

$$u_i(\boldsymbol{\sigma}) = w_i^G v_i^G(G(\boldsymbol{\sigma})) + w_i^p v_i^p(G_i(\boldsymbol{\sigma})) + w_i^r v_i^r(g_i(\boldsymbol{\sigma})) \tag{7}$$

Note that in (7), the first two terms refer to goals which are known to all partners, while the third term refers to partner i's private goal, which is not known to the other partner. Information provided by one partner about the utility it will achieve when a certain change to the public processes is implemented therefore is based on some information which the other partner cannot verify. This could create incentives to misrepresent the effect of changes on one's private goals (for example, to avoid a certain change, one could indicate that this change would be even more damaging to one's private goals than it really is). In this paper, we do not take into account this form of strategic behavior by the partners and assume that for the sake of maintaining a long term partnership, partners will refrain from such behavior.

Figure 3 illustrates the relationship between changes to private processes, changes to public processes, and utilities. The axes represent utilities of both partners. Consider the change to public processes indicated by (σ_A^1, σ_B^1). For

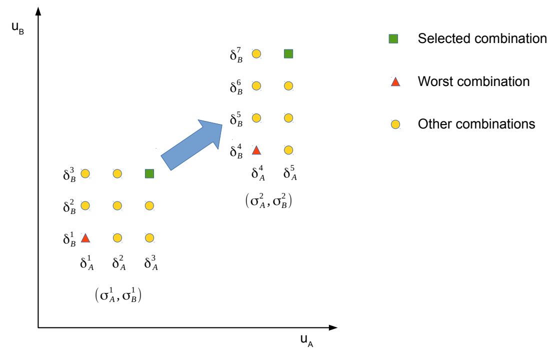

Fig. 3. Public and Private Process Changes and Utilities

this change in public processes, there are three possibilities to implement it in the private processes of A indicated as δ_A^1, δ_A^2 and δ_A^3. Since changes in A's private process only affect the utility of partner A, consequences of these changes are all located on a horizontal line. Similarly, we assume that there are also three possible changes to the private processes of B labeled δ_B^1, δ_B^2, and δ_B^3. In utility space, the consequences of these changes are located on a vertical line, so all possible combinations of changes to private processes are located in the grid shown in the lower left part of Figure 3. Assuming that each partner implements the change to its private processes which optimizes its own utility, the change to public processes indicated by (σ_A^1, σ_B^1) will lead to the upper right corner point of that grid, where the changes δ_A^3 and δ_B^3 are implemented. Similarly, another change to public processes (σ_A^2, σ_B^2) leads to another grid. Obviously, the number of changes to private processes of the partners which can implement a given change in public processes does not have to be the same neither across partners, nor across different changes to the public processes, so we show here only two private changes for partner A, but four for B. Still, assuming optimizing behavior of all partners, this set of changes to the public processes will lead to an evaluation corresponding the the upper right corner of the respective grid.

With respect to the example introduced in Section 2, Figure 4 summarizes the different change dependency graphs; i.e., propagation alternatives, as well as the corresponding cost graphs. A cost graph, uses the cost functions defined previously and a change propagation alternative in order to generate a cost alternative. The evaluation of an alternative cost depends on the adopted negotiation method; e.g., additive, Nash-bargaining, as will be defined in the next section.

4 Collective decision

The partners have to jointly agree on a change $\boldsymbol{\sigma} = (\sigma_A, \sigma_B)$ to the public processes. This is a decision situation in which the decision of each partner (the public process change of each partner) affects the outcome for all partners. Furthermore, these decisions are not independent of each other, since changes

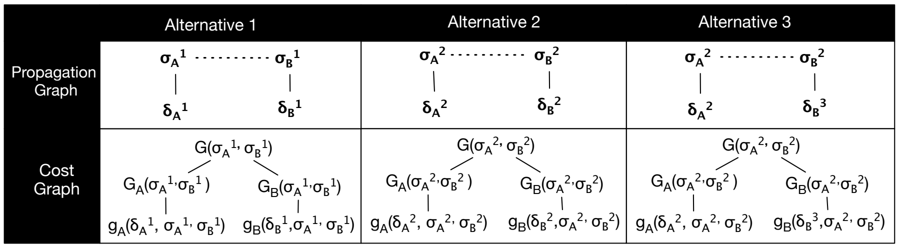

Fig. 4. Example from Financial Domain: Change Alternatives Table

to the public processes of all partners must be compatible with each other. For example, if partner A changes its public process so that it no longer generates an output which is required from the entire process, then partner B must change its process so that it now provides that output. Such decision problems are often referred to as group decision problems, and are also studied in the field of cooperative game theory. In the following paragraphs, we will present several solution concepts from these fields for this type of problems.

In general, solutions to collective decision problems can be evaluated using the criteria of efficiency and fairness [29]. The approaches we survey below are based on systems of axioms, that describe the precise criteria of fairness and efficiency used, as well as the necessary trade-off between the two concepts.

Efficiency in the most general sense can be defined via the concept of Pareto-optimality. A solution is Pareto-optimal if no other solution exists, which would make at least one partner better off, and no partner worse off, than the solution under consideration [1]. Pareto-optimality can thus be defined if individual preferences are represented by rankings of alternatives. However, if preferences are only specified as rankings, Arrow's well known impossibility theorem [1] holds. This theorem indicates that no mapping from individual rankings to a group ranking exists, that in addition to Pareto optimality fulfills the conditions of universal domain (i.e., a group ranking is obtained for any profile of individual rankings), independence of irrelevant alternatives (i.e., the ranking of two alternatives in the group ranking does not depend on the availability of some other alternative), completeness and transitivity of the group ranking as well as the non-dictatorship condition (no group member uniquely determines the group's ranking).

In the preceding section, we have defined the evaluation of the partners for each possible package of changes to the public processes in terms of cardinal utilities to avoid this problem. In this setting, efficiency can be defined as maximizing the total output to the group, i,e. the sum of utilities. If, as defined above, utilities are all scaled between zero and one, then

$$Eff = \frac{1}{N} \sum_{i=1}^{N} u_i(\boldsymbol{\sigma}) \tag{8}$$

is an indicator of efficiency, which is scaled so that one indicates perfect efficiency.

Fairness refers to the balance of payoffs between the partners. In the case of two partners, fairness can be measured by contract imbalance, which is the difference in utilities between the partners:

$$F = 1 - |u_A(\boldsymbol{\sigma}) - u_B(\boldsymbol{\sigma})| \tag{9}$$

It can be shown that cardinal utilities can be aggregated in a way that is compatible with very similar requirements as Arrow's axioms [5, 13]. A group utility function which fulfills these axioms is the additive function

$$\max U(\boldsymbol{\sigma}) = \sum_i w_i u_i(\boldsymbol{\sigma_i}) \tag{10}$$

where w_i is a weight assigned to member i's preferences in the group. The additive structure of the group utility function implies that utilities of different members are perfectly substitutable. Additive group utilities, which are sometimes also referred to as utilitarian solutions [15], therefore exclusively focus on efficiency.

In contrast, solution concepts from cooperative game theory also take fairness into account. The best known solution concept of cooperative game theory is the *Nash bargaining solution* [18], which selects the alternative that maximizes

$$\max N(\boldsymbol{\sigma}) = \prod_i (u_i(\boldsymbol{\sigma}) - d_i) \tag{11}$$

where $u_i(\boldsymbol{\sigma})$ is again the utility which partner i assigns to the proposed set of process changes, and d_i is partner i's utility for the *disagreement point*, i.e. the solution that would obtain if the partners did not agree on some alternative. For the problem at hand, the disagreement point can be interpreted as a situation in which no change takes place and thus the current processes continue to be used.

The Nash bargaining solution is based on the following set of axioms [18]:

1. Independence of linear transformation of utilities:
   The solution remains the same, when any individual utility function is changed by a linear transformation.
2. Pareto condition:
   The solution is not dominated by another feasible solution.
3. Symmetry:
   If both partners are symmetric (i.e. for each possible alternative, there exists another alternative in which payoffs to partners are exchanged), they receive the same payoff in the solution.
4. Independence of irrelevant alternatives:
   If some alternatives other than the solution are eliminated from the set of alternatives, the solution stays the same.

The axiom of symmetry introduces some notion of fairness to the Nash bargaining solution. Still, it can be shown that among comparable solutions, the Nash bargaining solution puts a comparatively high weight on efficiency [22]. Another important characteristic of the Nash solution for our problem is independence of linear transformation of utilities. Although one might argue that consequences of process changes for firms can usually be evaluated in terms of money, even monetary consequences could have a different relevance for the partners involved. The same absolute amount might be a significantly higher financial burden for a comparatively small firm than for a large corporation. In the additive model (10), such differences would have to be taken into account by assigning different weights w_i to partners.

In literature, independence of irrelevant alternatives is often seen as a problematic axiom [11, 16]. One can argue that the bargaining power of a partner depends on which alternatives are available. Consider a situation in which the solution gives to each partner a payoff which is approximately in the middle of the range of possible payoffs for that partner. Then suddenly, all alternatives which would be better than that solution for one partner become unavailable. Because of independence of irrelevant alternatives, the Nash bargaining solution would still give both partners the same payoff as before, which is now the best outcome for one partner, but only an average outcome for the other one. This could be considered unfair. Since the axiom does not cover the case that new alternatives are added, it is even possible to construct examples in which alternatives are added which are better for both parties, but one partner receives less than in the solution of the original problem.

The Raiffa-Kalai-Smorodinski (RKS) solution [11, 16] replaces the axiom of independence of irrelevant alternatives by the axiom of monotonicity. This axion states that if a problem is extended by alternatives which dominate the solution of the original problem, then the new solution must also dominate the previous solution (i.e. no partner must be worse off, and at least one partner must be strictly better off in the new solution). The RKS solution maximizes the payoff to the partner whose utility is smallest. It therefore follows the principle of egalitarism by considering the partner who is worst off [15].

Considering just the minimum of utilities would introduce a dependence on the scaling of utilities. This problem can easily be avoided by scaling the utilities of all partners so that for each partner, the best feasible outcome is assigned a utility of one. The RKS solution can thus be represented as [16]

$$\max K(\boldsymbol{\sigma}) = \min_i \frac{u_i(\boldsymbol{\sigma})}{\max\limits_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} u_i(\boldsymbol{\sigma})} \tag{12}$$

A similar transformation can also be applied to the additive model (10). Note, however, that it leads to a violation of the axiom of independence of irrelevant alternatives, since the transformed utilities of each partner now depend on the availability of the best solution for each partner.

All three solutions can be interpreted as special cases of a general α-fair social welfare function [4]

$$A(\boldsymbol{\sigma}) = \sum_i \frac{u_i(\boldsymbol{\sigma})^{1-\alpha}}{1-\alpha} \tag{13}$$

For $\alpha = 0$, (13) corresponds to (10), for $\alpha \to 1$, it approaches the Nash bargaining solution (11) and for $\alpha = \infty$, it corresponds to (12).

5 Implementation and Analysis

This work has been integrated into an existing change propagation framework [6] that is able to simulate various choreography settings for stress testing change propagations. This simulation framework served as the basis for ensuring the correctness of the change propagation algorithms [6], as well as the data source for devising a change mining algorithm for change prediction [8]. In the context of this work, we have extended the change propagation framework by designing and implementing a new package for simulating change negotiation scenarios using the mathematical models described in the previous section. Concretely, we have implemented the *additive*, *Nash bargaining* and the *RKS* solutions.

As a sample scenario we have implemented the example shown in figure 1 consisting of the two partners: *bank* and *insurance*. Although this example consists of only two partners, we have kept the implementation as generic as possible, being able to generate change propagation scenarios consisting of an unlimited number of partners. The restriction to two partners is kept in this paper to be able to graphically illustrate the results of the resulting negotiation outcome.

Using the change propagation framework, several change scenarios have been generated. The effects of each change on the different partners is calculated through the same framework. Each change and its effects is considered as an alternative and is associated with a cost. The basis of choice during the change negotiation is represented by a single change propagation alternative, as can be seen in the *Change Alternatives Cost Table* in Fig. 4. An alternative is defined as a unique pair of private changes, with their associated public changes. Since goals and costs, especially private ones, require intimate details of partners during a negotiation example, we have opted to estimate both private and public costs in our simulation by randomizing their values in the range [0..1]. Each partner's utility is then derived based on these private as well public random costs. As we want to compare the group ranking functions themselves, even random values should allow correct rankings, and highlight the most fair and effective change negotiation alternatives.

Bank (U)	Insurance (U)	Scoring Function	Score Value
0.3729	0.4785	Additive	0.4257
0.3729	0.4785	Nash-Bargaining	0.1784
0.3729	0.4785	RKS	0.7793
0.3729	0.3861	Additive	0.3795
0.3729	0.3861	Nash-Bargaining	0.1439
0.3729	0.3861	RKS	0.9658

Table 1. Negotiation Log sample

The following table 1 illustrates parts of the extracted log file after a complete change negotiation simulation. Each row represents the ranking of a single alternative, where the alternative is represented by the partner utilities: *Bank Utility* and *Insurance Utility*. The associated score value shows the outcome of applying the respective ranking function. In this particular instance, the highest score is attained with the alternative where the *Bank*'s utility is measured at 0.3729 and the associated *Insurance*'s utility at 0.3861 with the final RKS score of 0.9658. This alternative is higher ranked compared to the next best alternative (with *Insurance*'s utility at 0.4785) due to the fairness criteria, which prefers alternatives where both partner utilities do not stray away too far from each other.

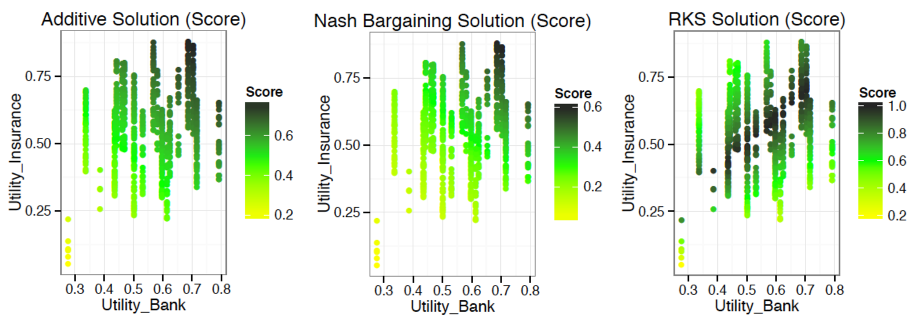

Fig. 5. Decision options for multiple propagation scenario

For evaluating the group ranking functions, we have taken the complete negotiation log as described in table 1 and visualized in Fig. 5. The charts are split into the three group ranking functions, each highlighting the final ranking score of all alternatives. A change negotiation scenario is represented as one vertical line with the corresponding alternatives. The x-axis depicts the *Bank*'s utility while the y-axis depicts the *Insurance* utility values. The final ranking score is represented in the hue of the color of a unique point representing one alternative. The darker the color, the better ranked are the alternatives. As can be seen in the charts, the *additive* solution seems to prefer *efficient* solutions, while disregarding fairness. This can be identified by the grouping of the black points as alternatives on the upper half of the chart for the *additive* ranking function chart. This is due to the behaviour of the *additive* function, which ranks alternatives with increasing partner utilities higher, without considering the difference between the utilities; i.e., fairness (cf. Equation 10). Similarly, the *Nash Bargaining* solution also prefers *efficient* solutions, and disregards those that are considered unfair. This behavior can be identified by the dark colors grouped in the upper-right hand corner of the respective chart. This can be explained by the multiplication of partner utilities (cf. Equation 11). Finally, the *RKS* solution seems to prefer more balanced alternatives to the other solutions, as the black colored points form a diagonal line from the bottom left to the upper-right corner of the chart. Indeed the *RKS* function chooses alternatives with less spread between partner utilities (cf. Equation 12).

Figure 6 shows the behavior of the best alternatives of each negotiation scenario generated by each negotiation function. In particular, it classifies these alternatives with respect to their efficiency and fairness (cf. Equations 8, 9). Both the *additive* and *Nash bargaining* functions pick up similar best alternatives per scenario for both *efficiency* and *fairness*.

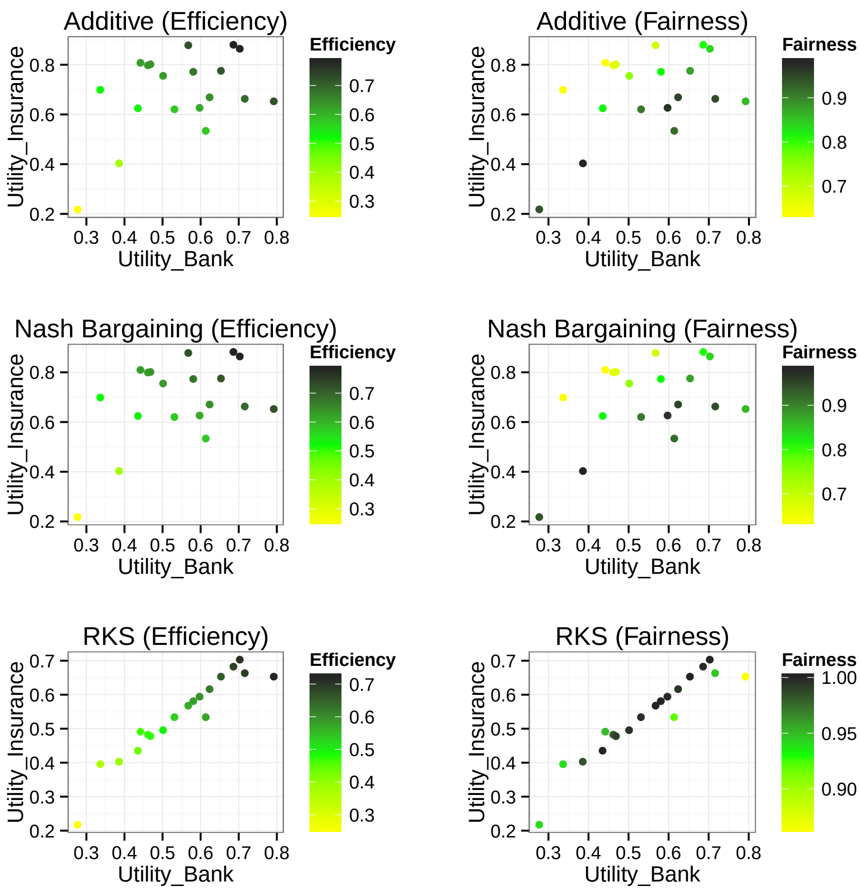

Fig. 6. Efficiency vs Fairness

In contrast, the *RKS* function selects the most fair solutions as best alternatives, which are not necessarily efficient. This is shown through the color difference between the *efficiency* and the *fairness* graphs. For example, for $utility\_bank = 0.5$ and $utility\_insurance = 0.5$ the fairness in the *RKS* is superior to 0.95, while the efficiency is less than 0.5. The same alternative is not picked up as best solution by both the *additive* and *Nash bargaining* functions.

6 Related Work

The research presented in this paper integrates three areas, i.e., process choreographies, change and change propagation, and collective decision making through negotiations. In each of the areas, a multitude of research questions and approaches exist. Hence, the following discussion will concentrate on the interfaces between the areas.

Change propagation in process choreographies has been addressed by different approaches as shown in a recent survey [7]. In [6, 8], the process of changing a process choreography is outlined and mentions negotiations as building block of this process. However, the approaches [6, 8] focus on structural correctness of changes and the prediction of change impacts, but have not addressed the negotiation of changes so far. [27, 9] deals with correctness and consistency during change propagation as well, but does not focus on negotiation aspects.

In the web service area in general, negotiation is part of the service discovery phase [23]. Negotiation has also been identified as part of building process choreographies, specifically for contracting [20] and the interface design [26, 19]. Moreover, negotiation protocols and strategies have been offered in the context of service level agreements between different partners, e.g., [10].

Although the solution concepts we presented in this paper, and the trade-off between fairness and efficiency which they represent, are well known concepts in the areas of collective decision making, the present problem also provides some innovative aspects for that field. Models of collective decision making usually assume that group members have a direct evaluation of the alternatives under discussion, or in some cases describe the underlying preference structure of group members by referring to risk or to separate attributes of the decision alternatives. e.g.,[12, 24]. The situation we are considering here is different because of the hierarchical relationship between public and private process models. This hierarchical structure allows partners to adapt their private processes, within the boundaries of the public process, to optimally fulfill their own private goals. Although we have shown that the model can ultimately be mapped back to one in which the collective decision determines the outcome for each partner, this is still an important extension to standard models of collective decision making.

In [3], State Charts are proposed to capture different e-negotiation models as business processes. Moreover, it is described how these models can be mapped onto BPEL models for web service orchestrations. The difference to the work at hand is that neither choreographies nor changes have been considered. Moreover, [3] focuses on e-negotiation protocols such as Dutch auction.

Overall, to the best of our knowledge, an approach that addresses the interfaces between all three areas, i.e., process choreographies, change, and negotiation, has not been proposed so far.

7 Conclusion

In a collaborative environment, changes to business processes are not confined to one partner, but have to be propagated through the processes of several partners. Since such changes affect the performance of processes of all involved partners, a collective decision problem arises.

In the present paper, we have raised the question whether models and methods from various theories of collective decision making can help to find a joint solution to this problem. The methods we studied in this paper clearly show the central dilemma of collective decision making, that is to find a balance between efficiency and fairness. Depending on how this trade-off is resolved, different solution concepts can be recommended.

One important assumption made in this paper is that all partners provide correct and truthful information about the possible impact of process changes on their public goals and utilities. While this assumption seems reasonable in a cooperative environment characterized by a high level of mutual trust, there is still the possibility that not all partners behave in that way, and some might at least slightly distort information in order to reach a solution that better fits with their own interest. Analyzing the effects of such strategic behavior thus leads to additional questions, that need to be addressed in future research.

Although the models we have discussed in this paper provide normatively appealing and axiomatically founded solutions to group decision problems, they do not describe the actual negotiation process by which such solutions can be reached. Future research thus needs to address this problem from a more dynamic perspective and include the actual bargaining process between partners that might be involved in finding a solution. Addressing these issues could eventually extend the first steps taken in this paper towards a directly applicable framework for collectively evaluating and choosing changes to process choreographies in an efficient and fair way.

Acknowledgments

The work presented in this paper has been conducted within the C^3Pro project I743 funded by the Austrian Science Fund (FWF).

References

1. Kenneth J. Arrow. *Social Choice and Individual Values.* Yale University Press, New Haven, 2 edition, 1963.
2. Alistair Barros, Marlon Dumas, and Phillipa Oaks. Standards for web service choreography and orchestration: Status and perspectives. In *Business Process Management Workshops*, pages 61–74. Springer, 2006.
3. Morad Benyoucef and Stefanie Rinderle. Modeling e-negotiation processes for a service oriented architecture. *Group Decision and Negotiation*, 15(5):449–467, 2006.
4. Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. On the efficiency-fairness trade-off. *Management Science*, 58(12):2234–2250, 2012.
5. James S. Dyer and Rakesh K Sarin. Group preference aggregation rules based on strength of preference. *Management Science*, 25:822–832, 1979.
6. Walid Fdhila, Conrad Indiono, Stefanie Rinderle-Ma, and Manfred Reichert. Dealing with change in process choreographies: Design and implementation of propagation algorithms. *Information Systems*, 49(0):1 – 24, 2015.
7. Walid Fdhila, David Knuplesch, Stefanie Rinderle-Ma, and M. Reichert. Change and compliance in collaborative processes. In *Services Computing*. IEEE, 2015.
8. Walid Fdhila, Stefanie Rinderle-Ma, and Conrad Indiono. Memetic algorithms for mining change logs in process choreographies. In *ICSOC*, volume 8831, pages 47–62. 2014.
9. Georg Grossmann, Shamila Mafazi, Wolfgang Mayer, Michael Schrefl, and Markus Stumptner. Change propagation and conflict resolution for the co-evolution of business processes. *International Journal of Cooperative Information Systems*, 24(01), 2015.

10. Ahmad Fadzil M Hani, Irving Vitra Paputungan, and Mohd Fadzil Hassan. Renegotiation in service level agreement management for a cloud-based system. *ACM Computing Surveys (CSUR)*, 47(3):51, 2015.
11. Ehud Kalai and Meir Smorodinsky. Other solutions to Nash's bargaining problem. *Econometrica*, 43(3):513–518, 1975.
12. Ralph L. Keeney. Foundations for group decision analysis. *Decision Analysis*, 10(2):103–120, 2013.
13. Ralph L. Keeney and Craig W. Kirkwood. Group decision making using cardinal social welfare functions. *Management Science*, 22:430–437, 1975.
14. Ralph L. Keeney and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. J. Wiley & Sons, New York, 1976.
15. James Konow. Which is the fairest one of all? A positive analysis of justice theories. *Journal of Economic Literature*, 41:1188–1239, 2003.
16. Zvi A. Livne. Axiomatic charaterizations of the Raiffa and the Kalai-Smorodinsky solutions to the bargaining problem. *Operations Research*, 37(6):972–980, 1989.
17. Heiko Ludwig, Alexander Keller, Asit Dan, Richard King, and Richard Franck. A service level agreement language for dynamic electronic services. *Electronic Commerce Research*, 3(1-2):43–59, 2003.
18. John F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
19. Michael Papazoglou. *Web services: principles and technology*. Pearson Education, 2008.
20. Shamimabi Paurobally, Valentina Tamma, and Michael Wooldridge. A framework for web service negotiation. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 2(4):14, 2007.
21. Karthikeyan Ponnalagu, Nanjangud C. Narendra, Aditya Ghose, Neeraj Chiktey, and Srikanth Tamilselvam. Goal oriented variability modeling in service-based business processes. In *Int'l Conf. on Service-Oriented Computing*, pages 499–506, 2013.
22. Shiran Rachmilevitch. The Nash solution is more utilitarian than egalitarian. *Theory and Decision*, pages 1–16, 2014.
23. Dumitru Roman, Uwe Keller, Holger Lausen, Jos de Bruijn, Rubén Lara, Michael Stollberg, Axel Polleres, Cristina Feier, Christoph Bussler, Dieter Fensel, et al. Web service modeling ontology. *Applied ontology*, 1(1):77–106, 2005.
24. Paula Sarabando, Luis C. Dias, and Rudolf Vetschera. Mediation with incomplete information: Approaches to suggest potential agreements. *Group Decision and Negotiation*, 22(3):561–597, 2013.
25. Stefan Schulte, Dieter Schuller, Ralf Steinmetz, and Sven Abels. Plug-and-play virtual factories. *IEEE Internet Computing*, (5):78–82, 2012.
26. Paolo Traverso, Marco Pistore, Marco Roveri, Annapaola Marconi, Raman Kazhamiakin, Pierluigi Lucchese, Paolo Busetta, and Piergiorgio Bertoli. Supporting the negotiation between global and local business requirements in service oriented development. In *Proceedings of the 2d International Conference on Service Oriented Computing, New York, USA*, 2004.
27. Wil MP van der Aalst. Inheritance of interorganizational workflows: How to agree to disagree without loosing control? *Information Technology and Management*, 4(4):345–389, 2003.
28. Mathias Weske. *Business process management: concepts, languages, architectures*. Springer Science & Business Media, 2012.
29. Moshe Zukerman, Musa Mammadov, Liansheng Tan, Iradj Ouveysi, and Lachlan L.H. Andrew. To be fair or efficient or a bit of both. *Computers and Operations Research*, 35:3787–3806, 2008.