DCF-S: A Dynamic Coalition Formation Scheme for Rational Agents

Andreas Gerber
German Research Center for Artificial Intelligence GmbH,
Stuhlsatzenhausweg 3, 66123 Saarbrücken, Germany
agerber@dfki.de

Matthias Klusch
German Research Center for Artificial Intelligence GmbH,
Stuhlsatzenhausweg 3, 66123 Saarbrücken, Germany
klusch@dfki.de

ABSTRACT
We introduce the notion, issues, and challenges of dynamic coalition formation (DCF) among agents in open, heterogeneous and world widely distributed environments. Results achieved in the traditional field of static coalition formation are briefly discussed, as well as the desired results in the new research field of DCF. We propose a simulation-based DCF scheme, called DCF-S, to be applied to any multiagent system acting in environments in which agents face imperfect information on tasks and society. Using this scheme each agent attempts to form task-oriented coalitions with other agents it knows about. Coalition leading agents are continuously striving to improve their coalition by simulation of potential alternatives in case agents leave or enter the scene, or tasks are changed by the users. In this paper we outline and discuss the simulation-based scheme for dynamic coalition formation.

General Terms
Design, Economics.

Keywords
Agents, Game Theory, Co-operation.

1. INTRODUCTION
Self-interested, autonomous software agents on the Internet may negotiate rationally to gain and share benefits in stable (temporary) coalitions. This is to save costs by coordinating activities with other agents. For this purpose, each agent determines the utility of its actions and productions in a given environment by an individual utility function. The value of a coalition among agents is computed by a commonly known characteristic function which determines the guaranteed utility the coalition is able to obtain in any case. In a characteristic function game the agents may use imposed individual strategies to achieve a desired type of economically rational behaviour such as altruistic, bounded rational, or group rational.

Conference '00, Month 1-2, 2000, City, State.

In any case, the distribution of the coalition's profit to its members is de-coupled from its obtainment but is supposed to ensure individual rational payoffs to provide a minimum of incentive to the agents to collaborate. Rational agents are usually required to form beneficial coalitions in open, distributed and heterogeneous environments at any and in reasonable time. That includes scenarios in which dynamically occurring events may interfere with the running coalition.

Due to its nature dynamic coalition formation (DCF) methods promise to be particularly well suited for applications of ubiquitous and mobile computing. In m-commerce settings, for example, personalized information agents each representing a potential business partner may dynamically form temporary profit-oriented coalitions on demand at any time. That may increase customers' benefits of purchasing and negotiating sets of items at multiple electronic market places world wide in reasonable time. This appealing vision for the common Internet user appears to be not far away from being realised with respect to recent advances in appliances of wireless computing and communication. First research towards appropriate economic models for agent-based applications in this domain includes, for example, [20] [23] [16].

The remaining sections of this paper are structured as follows. Section 2 introduces to the traditional field of static coalition formation; readers who are familiar with the domain may skip this section. In section 3 and 4 we briefly describe the problem of dynamic coalition formation, and then discuss how selected work from different research disciplines may be coping with parts of the DCF problem. We propose a simulation-based DCF scheme in section 5, and conclude the paper with a brief outlook on future work in section 6.

2. Traditional Methods to Coalition Forming
According to [4] models of coalition formation may be classified into two main approaches: utility-based and complementary-based models dividing the societies of actors into ones following either the principle of 'bellum omnium contra omnes' as it is largely favoured, for example, by game theory [12], or ones which rely on the collaborative use of complementary individual skills to enhance the power of each agent to accomplish its goals, respectively.

Up to now, most classic methods and protocols for a formation of stable coalitions among rational agents follow the utility-based approach. They rely on derived concepts from co-operative game theory, economics, and operations research. Utilitarian coalition formation covers two main activities:

(1) the generation of coalition structures, that is partitioning or covering the set of agents into coalitions, so as to maximize the monetary value depending on the benefit of accomplishing tasks regarding used resources and time spent;

(2) the distribution of gained benefit among the participants of each of the coalitions.

These activities may be interleaved and are not independent. A comprehensive discussion and classification of relevant work on coalition formation is given, for example, in [11] [21]. In the following we provide the reader with a very brief introduction to basic concepts and notions of co-operative game theory.

2.1 Basic Concepts and Notions

A *co-operative game* (A, v) is defined by a set A of agents wherein each subset of A is called a coalition, and a real-valued characteristic function v assigning each coalition C in A its maximum gain, that is the expected total income of the coalition (the so-called coalition value). It is commonly assumed that the value of any coalition C is in money and the value $v(C)$ does not depend on the actions of agents outside the coalition. Furthermore, we may assume that any coalition C forms by a binding agreement (coalition contract) on the distribution of its coalition value $v(C)$ among its members, in particular no side-payments are allowed from C to any agents outside C within the game. Finally, the characteristic function v is assumed to be known to all agents in A.

The solution of a co-operative game with side payments is a *coalition configuration (S, u)* which consists of a partition S of A, the so-called coalition structure, and a n-dimensional, real-valued payoff distribution vector which components are computed by a real-valued payoff or utility function u. The payoff distribution assigns each agent in A its utility $u(a)$ out of the value $v(C)$ of the coalition C it is member of in a given coalition structure S. It is commonly assumed that every coalition may form, including singletons or the grand coalition A. However, the number or size of coalitions to be formed using a coalition formation method is often restricted to ensure, for example, polynomial complexity of the formation process.

Individually rational payoff distributions are assigning each agent at least the gain it may get without collaborating within any coalition. For *group rational* coalitions it holds that the group of all agents is assumed to maximize its joint payoff. In coalition configurations with so-called *Pareto-optimal* payoff distributions no agent is better off in any other valid payoff distribution for the given game and coalition structure. A coalition configuration (S, u) is called *stable* if no agent has an incentive to leave its coalition in S due to its assigned payoff $u(a)$. Different characteristics and criterions of stability define different solution space for co-operative games (section 2.2).

Rational agents which are involved in a co-operative game (A,v) are supposed to negotiate a stable payment configuration (S,u) as a solution of the game by the use of an appropriate *coalition algorithm* (CA). The CA is completely decentralized, and should provide a stable coalition configuration for any co-operative game in the considered environment at any time. The latter property is, in general, hard to achieve in DCF environments.

A *coalition formation environment* (CE) for a given set of agents A is the set of assumptions and constraints which are valid for any kind of coalition forming activity between agents in A including propositions on:

1. The task-related functionality of each individual agent in A, including its set of tasks and goals to accomplish, set of appropriate actions, and methods to compute the individual utilities of task-related productions;
2. Valid methods for computing the values of coalitions, for example, by the sum of production utilities of all agents in a coalition;
3. Valid methods for determining coalition configurations, including methods for searching coalition structures, negotiation and payoff distribution schemes;
4. Commitments, obligations of and agreements between agents in A concerning the type of collaboration and interaction.

In a given coalition formation environment the agents particularly agree on (a) what kind of stable coalitions shall be negotiated, and (b) what particular coalition algorithm CA shall be used for the negotiation.

A coalition formation environment is called *super-additive* or *sub-additive* depending on the type of all co-operative games it allows, and *general* if it allows for at least one sub-additive game. Both general and sub-additive environments are often called non-super-additive environments. In sub-additive games at least one pair of potential coalitions is not better off by merging into one. This could be caused by, for example, communication and co-ordination overhead costs, etc.

A *coalition formation model* CFM = (CE, CA) is defined by both, the considered coalition formation environment CE and a given coalition algorithm CA for this environment. Interesting models are those where coalition formation is concerned with general and sub-additive environments.

The meaning of *stability of coalitions* depends on the considered discipline and application domain. Many if not most of the coalition formation algorithms today rely on chosen game-theoretic concepts for stable pay-off division within coalitions according to, for example, the Shapley-value, the Core, the Bargaining Set, or the Kernel [8]. For example, the *Kernel K* of a co-operative game (A,v) with respect to a given coalition structure is the set of *K-stable* configurations (S,u) in which all coalitions in S are in *equilibrium*. A coalition C is in such an equilibrium if each pair of agents in C is in equilibrium, i.e., any pair of agents in C is balanced, that is, none of both agents can outweigh the other in (S,u) by having the option to get a better payoff in coalition(s) not in S excluding the opponent agent. In other words, agents may argument each other like “Since I could obtain more without you in alternative coalitions than you without me, I deserve more, but without going to harm you.” For this purpose each agent has to compare its *surplus* with those of other agents; the calculation of the surpluses bases on that of the *excesses* of every alternative coalition possible. Thus, the kernel of a game is exponentially hard to compute unless, for example, the size of the coalition is limited by a constant. One attractive feature of this stability concept is that it is locally Pareto-optimal in the set K.

2.2 Limits of Traditional Approaches to Form Stable Coalitions

All traditional approaches to coalition formation remain static in the sense that they do not allow for any type of (non-deterministic) interference with the running coalition formation process. In addition, many results known today only hold for super-additive coalition formation environments. Regarding results achieved in the domain of fuzzy and stochastic co-operative games, further basic research on appropriate concepts and criterions of "vague" stability and corresponding coalition algorithms remains to be performed. Besides, most work on coalition formation rely on coalition formation environments in which all agents are assumed to be homogeneous, means they have the same task-related internal functionality; which in an open and changing world we are considering is not a feasible assumption. In the following section we discuss issues and problems of dynamic coalition formation environments and point out what kinds of results would be desirable in the new research field of DCF.

3. The DCF Problem

The research domain of *dynamic coalition formation* (DCF) can be defined by the set of co-operation methods, schemes, and key enabling technologies to cope with the problem of dynamically building beneficial coalitions among agents in open, distributed, and heterogeneous environments. This DCF problem has to be solved in any collaboration environment and scenario in which

(1) agents may enter or leave coalition formation processes at any time,

(2) the set of tasks to be accomplished by and the (computational) resources of individual used may change dynamically, and

(3) the information, network, and user environment of each of the agents and the system as a whole may change dynamically as well.

Co-operation scenarios inducing uncertain, time-limited, context-based utilities and coalition values may exacerbate the DCF problem. In dynamic coalition formation environments the following classes of events may non-deterministically occur.

- *Tasks*: The set of tasks, goals to accomplish and corresponding plans to pursue may change for each individual agent at any time. Such changes concern, for example, the volume of tasks, utilities, and costs of task execution as well as the frequency of such changes. General task allocation problems are known as at least NP-hard problems. Real-time issues and requirements to perform planning under time-dependent uncertainty [22] may even exacerbate these kinds of problems.
- *Agents*: Agents may leave or enter the agent society at any time; some agents may even temporarily hide their existence to parts of the society for different reasons.
- *Information*: Information and data may change, or become corrupt either intentionally by the sender or via its transmission over the network. User preferences may continuously change as well as network connections to agents in the considered agent society.

One hard challenge for agents to negotiate stable coalitions in such dynamic settings is how to flexibly react on different kind of changes preferably in real-time without having to restart the complete negotiation process. This in particular requires the agents to individually handle uncertain environment knowledge via appropriate adaptation mechanisms. Basic research is needed to clarify in which kinds of dynamic settings and to what extent available algorithms for the static formation of stable coalitions may be adopted such as the meaning of the chosen stability criterion may be retained. In particular, efficient methods have to be invented allowing the agents to deliberately restart their coalition negotiations at any time depending on the kind of possible changes in the considered environment.

4. Relevant Work for DCF

The development of DCF schemes may in particular benefit from adopting appropriate methods for quantitative or qualitative decision making in face of imperfect information. Reasonable solutions for fuzzy [30] and stochastic co-operative games [29] may be adopted for the development of co-operation schemes which enable the agents to deal with different types of imperfect information such as the vagueness of expected coalition values, payoffs, or membership of agents in coalitions. In both cases such uncertainties may be induced in dynamic coalition formation environments. Other relevant work for developing co-operation schemes for dynamic environments include, for example, utility-based schemes for dynamically re-organising organisational structures [1], and exception tolerant reasoning and multi-criteria decision making under uncertainty [2].

Social reasoning mechanisms are considered as essential building blocks suitable to situations where agents may dynamically enter or leave the society, without any global control. Therefore advances in social reasoning have a clear impact on the development of DCF schemes. Social reasoning mechanisms are often based on the notion of social dependence [3], or aim at reputation and trust management.

In order to acquire and use dependence knowledge on the considered agent society, each agent has to explicitly represent some properties of the other agents, which may change dynamically; exploit this representation thereby optimising its behaviour according to the evolution of the society; and to monitor and revise its representation to avoid inconsistencies to an acceptable degree, without any pre-established global control.

Reputation management aims at avoiding interaction with undesirable participants and may complement other security technologies for authentication and authorisation. Mechanisms for building, propagating, measuring and maintaining reputation and trust [24] are useful to apply, for example, to settings for coalition formation among self-interested agents in e-commerce applications where trusted third parties are required but not available. Negotiation schemes for uncertain games with trusted third party are proposed, for example, in [18].

Rational agents may face many potentially beneficial choices related to the timing of events, which may occur during the individual decision process, and the negotiation with other potential coalition partners. Regarding the use of social reasoning mechanisms for the formation of temporary coalitions in continuously changing environments temporal dependence networks and adequate temporal social reasoning mechanisms may be applied to DCF schemes. Work on real-time issues in the context of agent-based online auctions on a single auction server suggesting a design for maximal asynchrony and robustness to network delay includes, for example, [22]. In general, any problem of performing time-constrained reasoning for coalition formation may be viewed naturally as a constraint satisfaction problem [6]. Therefore one might adopt approaches for solving dynamic CSPs [13] to cope with this part of the DCF problem.

5. DCF-S: A Simulation-Based DCF Scheme

The simulation-based dynamic coalition formation scheme DCF-S is designed to enable agents for effectively reacting on changes of their set of goals and agent society. This DCF-S scheme may be instantiated by, for example, the utilization of different computational methods and negotiation protocols with respect to the considered application and chosen criterion for coalition stability. Each instantiation yields a particular DCF-S based coalition algorithm. In the following we briefly outline and discuss the DCF-S scheme.

5.1 Assumptions on the Environment

We consider a coalition formation environment in which agents may continuously receive a set of goals. Furthermore, any agent may freely enter or leave the society at any time. Due to this dynamic nature of the environment the agents are situated in their mutual knowledge and possible benefits to share in coalitions may be vague and incomplete. That requires each agent to utilize appropriate mechanisms to cope with these uncertainties and gradually adapt its decision making with respect to non-deterministically occurring changes in the environment. Furthermore, we assume an additional set of special agents, called *world-utility agents* (WUA) to the environment. Any WUA may receive, compile and maintain information about each of its registered agents. This information includes, for example, statements on the problem solving capability of an individual agent and the evaluation of its quality of service by other agents. Such evaluation may concern, for example, the reliability and trustworthiness of an agent with respect to its co-operation with other agents. It is assumed that these evaluation records are safe against possible manipulation and securely distributed to and updated by the networked world-utility agents. However, it is noteworthy to emphasise that each agent in the considered agent society is free to request its nearest world-utility agent. Though this set of WUA is not essential it may be quite useful as an additional source of information for the agent society.

We define a goal-oriented co-operative game $(A, v)|G$ as a co-operative game (section 2.1) with respect to a given goal G. Such a game is determined by a given set A of agents and a real-valued function v assigning each coalition C in A its total expected outcome with respect to the accomplishment of the goal G. In particular, the computation of the individual utility of the set of productions of coalition members in C is restricted to the set of productions related to G. We further assume that any coalition may be represented to the outside world by an appropriate *coalition leading agent* (CLA). Therefore, we consider each coalition as one entity or agent. Since one agent may also be considered as a single-agent coalition the term agent and coalition may be used interchangeably. Initially, the set of all possible, non-empty coalitions is the set of single-agent coalitions, and each agent is a CLA for a coalition which it has initiated to form for one of its goals. Any CLA is supposed to reliably and trustworthy act on behalf of the members of its coalition. This includes, for example, the responsibility to negotiate and control the distribution of resources and payoffs among the coalition members according to the coalition contract.

5.2 The DCF-S Scheme

In the DCF-S scheme each coalition leading agent concurrently simulates, adaptively selects, and negotiates coalitions each of which is able to accomplish one of its goals with an acceptable ratio between estimated risk of failure and individual profit to gain in the coalition. In the following we first summarize the main steps of the scheme, then present the scheme in more detail in pseudo-code.

5.2.1 Outline of the DCF-S Scheme

The main steps of the DCF-S scheme are executed by each coalition leading agent (CLA).

(1) *Preparation*: The CLA determines the set of goals to be accomplished in co-operation with other agents. It periodically updates its knowledge on the environment. The local knowledge base includes information on (partially) known problem-solving capabilities of other agents as well as individual evaluations of past collaborations with these agents with respect to their reliability and trustworthiness in co-operation. To obtain this information it may in particular request its nearest world-utility agent. Since this environment knowledge may be incomplete or vague the agent is assumed to utilize appropriate learning mechanisms for approximating the needed information.

(2) *Simulation*: The CLA simulates the formation of coalitions each of which may be able to accomplish a given goal with an acceptable ratio between the estimated individual profit and risk of forming the coalition. This simulation consists of the following steps.

- The agent performs a capability-based matching to determine the set of possible candidates for jointly accomplishing a goal in some coalition.
- For each goal and corresponding set of candidates the agent then randomly simulates coalitions of limited size until it finds a coalition which appears to be significantly better than the currently valid coalition with respect to its estimated individual profit to gain and risk of forming the coalition. In particular, each of these coalitions is simulated by randomly adding and removing candidates. The individual decision on whether a candidate will be added to or removed from the currently valid coalition is made by the agent with respect to its estimated risk of co-operating with this candidate as well as the marginal contribution of the candidate to the coalition. Both estimation values may be derived from information available in the updated local knowledge base of the agent. The individual risk for the agent to form the simulated coalition is related to the sum of all individual penalties the agent will have to pay to those of its partners in the currently

valid coalition which are not supposed to remain in the new restructured coalition.

- If no events did occur so far this may endanger the formation of at least one of the simulated coalitions the CLA proceeds with the following negotiation step to realise these coalitions. Otherwise it continues the simulation deliberately taking the detected changes in the environment into account trying to avoid a complete restart. This is achieved by keeping the already committed agents in the initial hypothetical coalition for the new simulation.

(3) *Negotiation*: The CLA negotiates all coalitions it has determined in the previous simulation step. Each goal-oriented coalition is bilaterally negotiated by the CLA with each potential member of the coalition in sequence; the complete set of negotiation sequences may be performed concurrently. The result of a successful negotiation is a binding agreement between agents on the constraints and attributes of their co-operation in the new coalition. As mentioned above, bilaterally negotiated coalition configurations may be stable according to, for example, the bilateral Shapley-value transformed to general environments.

In case one bilateral negotiation fails or an event changing the value or structure of the considered coalition is detected the negotiation process for that coalition is immediately halted. The CLA then evaluates the negotiation process for this coalition so far in the next step and restarts the simulation of potential coalitions for the particular goal. Please note that, for the restart it keeps those agents in its coalitions with which it has already successfully negotiated and takes the current situation of the environment into account. This way the CLA may avoid a complete restart, thereby avoiding possible penalty payments and a corresponding decrease of its own reliability in co-operation for the rest of the agent society.

Evaluation: The CLA may evaluate its recent negotiations and report these evaluations to the nearest world-utility agent for distribution. Concurrently, it controls the distribution of payoffs and resources to members of the newly formed coalitions according to the successfully negotiated contracts.

5.2.2 DCF-S Scheme: Prerequisites

In the following we introduce the terms, sets, values, and functions used in the description of the DCF-S scheme.

The *local knowledge base* of an agent *a* consists of the following components.

- The set *GS(a)* of goals the agent has to accomplish. Interleaved goals are assumed to be aggregated by *a* into one goal.
- The list *CL|G* (*BestCL|G*) contains the (best) *candidates* with which agent *a* may coalesce to accomplish a goal *G* in *GS(a)*.
- The list *ACL|G* of *agent information records* contains information of agent *a* on the capabilities of other agents *a'* with respect to *G*. Each record stores a finite dimensional vector of real-valued attributes of an agent *a'* with respect to its estimated value of contribution to the accomplishment of goal *G* in *GS(a)*. These attributes are as follows.
- *Goal-related* attributes of agent *a'* concern, for example, the estimated amount of its available resources, costs, quality, and efficiency with respect to goal *G*.
- Other attributes of *a'* concern its *reliability and trustworthiness in co-operation*.
  - The real value *crv(a', ACL|G, C)* in [0,1] denotes the risk of agent *a* to cooperate with agent *a'* in coalition *C* for goal *G* in *GS(a)* with respect to the information on *a'* in the list *ACL|G*.
  - The real value *crl(a', C)* denotes the worst acceptable risk of agent *a* to cooperate with agent *a'* in coalition *C*.
  - The real value *rrl(a' C)* in [0,1] denotes the worst acceptable risk of agent *a* to remove agent *a'* from valid coalition *C* the agent *a* is leading. This risk value might be computed with respect to the implied payment of trust penalty *tp(a', a)* and the penalty payment limit *ppl(a)*.
  - The real value *tp(a', a)* is the trust penalty to be paid by *a* to *a'* in case *a* breaks a coalition agreement with *a'*.
  - The real value *ppl(a)* denotes the upper limit of penalty payments by agent *a*.

In addition, we define the following values, sets, and functions.

- Integer *MaxSim* denotes the *maximum number of simulation steps* per simulation round.
- *CC|G* set of *candidates for the formation of a coalition with respect to goal G*. These candidates are taken from the current set *CS* of all valid coalitions and determined by the function Match.
- Match(*CS*, *G*, *ACL|G*) determines the set of agents in the set *CS* of all agents (individual agents and valid coalitions which are actually known to agent *a*) each of which is capable of contributing to the accomplishment of goal *G*. The *capability-based matching* [9] determines to what degree the agents' capability descriptions in *ACL|G* match the description of the goal *G*. In addition, other agent attributes concerning trust and reliability in co-operation are matched. The returned set may be limited in size by including only few top-ranked matching candidates.
- Request(*ACL|G*, WUA), Update_Agt_Information (*ACL|G*, RecentAC) concern the request of the nearest WUA for information to (periodically) update the list *ACL|G*. The function Update_Agt_Information utilizes the results of an appropriate *learning mechanism* to approximate the needed information; this mechanism is assumed to be part of the functionality of the agent *a*. We suggest to use an appropriate reinforcement learning method [19].
- SelectAgt_MinRisk_MaxValue(*CC*, *C*, *ACL|G*) returns an agent *a'* from set *CC* of agents with estimated minimum risk of co-operation in and maximum value of contribution to the coalition

$HC \cup \{a'\}$ with respect to goal G regarding the attributes of a' stored in the agent information list $ACL|G$. The payoff for a' in C is individual rational for a'.

- SelectAgt_MaxRisk_MaxValue($CC|G$, $HC|G$, $ACL|G$) returns an agent a' from set $CC|G$ with estimated maximum risk of co-operation in $HC|G$, maximum value of the coalition $C\backslash\{a'\}$ regarding the information on a' in $ACL|G$.
- Events($BestCL|G$) returns the set of events occurred which have an impact on the formation of the coalition which consists of all agents in the list $BestCL|G$. These events concern the modification, deletion, or adding of a goal in *GS(a)*, and the leaving and entering of agents. Any processed event is automatically removed from the set.
- Value($CL|G$) determines the value v(C) of the coalition C which consists of all agents in the list $CL|G$.
- BilateralNegotiation(a, a', Value($BestCL|G$), $C \cup \{a'\}$) returns true if the *bilateral negotiation of agent a with agent a'* on a joint coalition $C \cup \{a'\}$ with respect to its value is successful, otherwise it returns false.

  The negotiation should respect the individual rationality of payoffs derived from the value of the negotiated coalition and given, appropriate criterion of coalition stability such as the bilateral Shapley-value with equal or proportional share between coalition members [17] [5]. Regarding the negotiation of other attributes we suggest to use an appropriate multi-attribute negotiation method such as proposed in [7].
- Evaluate($ACL|G$) updates the agent information list $ACL|G$ according to the local *evaluation* of the recent negotiation processes of the agent and returns the updated list $ACL|G$. This evaluation gives input to the internal learning mechanism of the agent for adapting to changes in its environment.
- StopNegotiation($BestCL|G$) stops all running negotiation processes with all agents in $BestCL|G$ on a coalition for the goal G and updates the list $CL|G$ by keeping those agents with which the executing agent has already successfully negotiated.
- RedundancyCheck($CL|G$) returns a list of tuples (agt, op) in which (agt, add) and (agt, remove) are deleted.

Each coalition leading agent a executes the following steps. Consider agent a as leader of valid coalition $C|G$ for goal G in *GS(a)*.

```
for each G in GS(a) do concurrently, until external termination
{ halt:= false;
  while not halt do {
  (1) Preparation
      CL|G, BestCL|G :=null; op:=''; z = penalties:=0;
      if periodic(date, ACL|G) then
      {RecentAC := Request(ACL|G, WUA);
       ACL|G:=Update_Agt_Information(ACL|G,
       RecentAC);}
  (2) Simulation
      CC|G:= Match(CS, G, ACL|G);
      HC|G:= C|G;
      for z:= 1 to MaxSim do
      { op := Random({noop, add_agent,
      remove_agent});
          if op = add_agent then
        { agt:= SelectAgt_MinRisk_MaxValue(CC|G,
              HC|G, ACL|G);
          if crv(agt, ACL|G, HC|G) ≤ crl(agt, HC|G)
      then
            {CL|G:= CL|G + (agt, add);
             HC|G:=HC|G∪{agt};}
        }
        else
        if op = remove_agent then
        { agt:= SelectAgt_MaxRisk_MaxValue(HC|G,
              ACL|G);
          if crv(agt, ACL|G, HC|G) > rrl(agt, HC|G)
      then
          {CL|G := CL|G + (agt, remove);
           HC|G:=HC|G \{agt};
           penalties:=penalties + tp(agt, a); }
        }
      }
      if Value(CL|G) > Value(BestCL|G) then
         BestCL|G:= RedundancyCheck(CL|G);
      if Value(BestCL|G) >> v(C|G) && penalties <
      ppl(a) && Events(BestCL|G) = ∅ then halt:= true;
  }
  (3) Negotiation
      halt:= false;
      for each (a', op) in BestCL|G  do concurrently
      { try
          if op = add then{
            if BilateralNegotiation(a, a',
                 Value(BestCL|G), C|G∪{a'})
            then {GS(a)=GS(a)\{G}; C|G:= C|G∪{a'} }}
          else { C|G:= C|G \ {a'};
                 Payment( tp(a', a), a'); }
        catch(event) {
          GS(a) = GS(a) ∪{G}
          if Events(BestCL|G) <> ∅then {
          StopNegotiation(BestCL|G); Goto (4)};
        }
      }
  (4) Evaluation
      EvalRes:= Evaluate(ACL|G);
      [if desired then Send(EvalRes, WUA);]
}
```

In the following section we briefly discuss some issues of the outlined DCF-S scheme and then compare it with an alternative DCF scheme proposed in [17].

5.3 Discussion of the DCF-S Scheme

In the DCF-S scheme each agent concurrently simulates, adaptively selects, and negotiates coalitions each of which is able to accomplish one of its goals with an acceptable ratio between estimated risk of failure and individual profit to gain in the coalition. In other words, the agents strive to concurrently solve a set of single goal-oriented co-operative games $(A, v)|G$ by forming potentially overlapping coalitions with stable payoff distributions regarding their local environment knowledge and given stability criteria. The problem of assigning tasks and goals to coalitions of agents has been approached in [14] [15]. It is noteworthy that each of these goal-oriented co-operative games may change at any time subject to different kinds of non-deterministically occurring events such as agents leaving or entering the society, as well as deleting or modifying the considered goal. Each detected change may induce new co-operative games to be solved by the agents.

According to the DCF-S scheme each coalition leading agent reacts on these changes during negotiation via a partial rather than complete restart of necessary simulation and negotiation of alternative size-bounded coalitions. The agent tries to keep those agents in the affected coalitions with which it has already successfully reached a coalition agreement. However, the DCF-S scheme does not guarantee in general that an optimal solution of these games will be found. Rather the agents are supposed to continuously approximate best solutions given their current knowledge on the dynamic environment by simulation and adaptation. One problem is that changes in the environments may occur that rapidly that it is not possible to realize any of the continuously simulated coalitions. However, that is a general problem for any type of DCF algorithm.

The DCF-S scheme assumes the existence of a set of networked world-utility agents, which each agent in the society is free to contact for obtaining needed information on the environment. In addition, the update of local knowledge by an agent is assumed to utilise results from a continuous adaptation process to approximate the needed information on its environment. That may improve the quality of its decision making independent from the world-utility agents, and reduce the overall complexity in computation and communication. As mentioned above, the complexity of any DCF-S based coalition algorithm as an instantiation of the DCF-S scheme largely depends on the complexity of the implemented methods.

5.4 Comparison With Other DCF Scheme

To date there exists only one alternative DCF scheme which has been suggested in [17]. It is to some extent similar to our DCF-S scheme. This scheme (ST-DCF) consists of the following steps.

1. An agent that detects a problem to be solved by some proper coalition first hastily forms an initial coalition by selecting neighbouring agents that it considers having high potential utilities for this task. This selection is based on its individual profile of its neighbourhood including information on the capability of each neighbour and its respective past inter-agent relationships. Agents may use a particular learning (case-based reasoning) mechanism for learning to form coalitions better by adjusting the weighted potential utility of neighbours, driven by the failure and success rates of past coalition formation.
2. The agent finalizes the coalition via concurrently performed bilateral negotiations with only selected top-ranked neighbours of high potential utility, during which constraints and commitments are exchanged in a special (argumentation) setting. The selection enables to conserve both computational resources and communication usage. When a coalition is formed or no longer can be formed due to some negotiation failures, all remaining negotiations are terminated.
3. The initiating agent acknowledges the status of a coalition, a responsible act that seals the validity of a planned coalition.

In principle, both ST-DCF and DCF-S schemes are opportunistic and of high risk since the formation of an optimal coalition cannot be guaranteed. However, both schemes significantly differ in the way coalitions are formed. Unlike the DCF-S scheme, the ST-DCF scheme does not allow agents to further improve a particular coalition for a given task if the corresponding information is not available at the time of its creation. This implies that sub-optimal solutions may not be improved to optimal ones even when the required information to compute such solutions becomes available to the agents. Moreover, it remains unclear how agents using the ST-DCF scheme may discover its neighbourhood; the DCF-S scheme allows agents to contact world-utility agents to obtain such information.

Both schemes allow for overlapping coalitions such that one agent may participate in multiple single goal-oriented coalitions. But the ST-DCF scheme prohibits agents to leave their coalition voluntarily; coalitions may only be disrupted just in case of the occurrence of external events like network failures. In addition, unlike the DCF-S scheme no penalties have to be paid by agents for leaving a valid coalition. The ST-DCF scheme pays less attention to avoiding possible deception and fraud via distribution of information on individual agent’s trust and reliability in co-operation with other agents.

6. Conclusions

We introduced the notion, selected issues, and challenges of dynamic coalition formation (DCF) among rational software agents. Using the DCF-S scheme each agent may concurrently simulate, adaptively select, and negotiate coalitions each of which is able to accomplish one of its goals with an acceptable ratio between estimated risk of failure and individual profit to gain in the coalition. Application-specific instantiations of this scheme may lead to a variety of DCF-S coalition algorithms. For this purpose, many relevant approaches and theoretical work stemming from different disciplines are available to date including work on temporal social reasoning, machine learning, and fuzzy and stochastic co-operative games. DCF algorithms promise to be particularly well suited for applications of ubiquitous and mobile computing, including mobile commerce in wireless network environments. However, further basic research is needed to investigate the potential of the new research field of dynamic coalition formation, which still remains in its very infancies to date.

7. REFERENCES

[1] Barber, K.S. and Martin, C.E. (2001) “Dynamic re-organisation of decision-making groups” Proceedings 5th International Conference on Autonomous Agents, Montreal, Canada, ACM Press, May/June, 2001.

[2] Benferhat,S., Dubois, D., Prade, H., and Sabbadin, R (2001) “Exception–tolerant reasoning and decision under uncertainty using qualitative possibility theory” Journal Künstliche Intelligenz, Vol. 3, arendtap Verlag, Beremen, 2001.

[3] Castelfranchi, C., and Micelli, M., and Cesta, A. (1992) “Dependence Relations Among Autonomous Agents” W Eric and Y Demazeau (eds.), Decentralized A. I. 3, Elsevier Science Publishers, pp 215-227, 1992.

[4] Conte, R., and Sichman, J.S., “DEPNET: How to benefit from social dependence” Journal of Mathematical Sociology, 20(2-3), 161-177, 1995.

[5] Contreras, J., Wu, F., Klusch, M. and Shehory, O., “Coalition formation in a power transmission planning environment”, Proceedings of International Conference on Practical Applications of autonomous Agents and Multi-Agent Systems (PAAM-07), London, UK, April 1997.

[6] Eaton, P.S., Freuder, E.C., and Wallace, R.J. (1998) “Constraints and agents – Confronting ignorance” AI Magazine, Vol. 19(2), 1998.

[7] Jonker, C.M., and Treur J., “An agent architecture for multi-attribute negotiation”. Proceedings of the seventeenth International Joint Conference on Artificial Intelligence (IJCAI 2001), (ISBN 1-55860-777-3, Volume 2):1195-2001, 2001.

[8] Kahan, J.P., and Rapoport, A., “Theories of coalition formation” Lawrence Erlbaum Associates, New Jersey. 1984

[9] Sycara, Widoff, Klusch, Lu, “LARKS Dynamic matchmaking among heterogeneous agents in cyberspace”, J Autonomous Agents and Multiagent Systems, 5(3), 2001

[10] Klusch, M., “Intelligent information agents: A survey”, J Data and Knowledge Engineering, Elsevier, 36, 2001

[11] Kraus, S., “Negotiation and cooperation in multi-agent environments”, Journal of Artificial Intelligence, Vol. 94(1-2), pp. 79-98. 1997

[12] Luce, R. D. and Raiffa, H., “Games and Decisions: Introduction and Critical Survey”, John Wiley & Sons Ltd., 1957

[13] Schiex, T. and Verfaillie,G. (1993) “No-good recording for static and dynamic constraint-satisfaction problems” Proceedings of International Conference on Tools with AI, pp. 48-55, Washington DC, IEEE CS Press, 1993.

[14] Shehory, O. and Kraus, S., “A Kernel-oriented model for autonomous-agent coalition-formation in general environments: implementation and results” Proceedings of AAAI Conference, pp. 134-140, Portland, Oregon, August 1996.

[15] Shehory, O. and Kraus, S., “Methods for task allocation via agent coalition formation”, Journal of Artificial Intelligence, Vol. 101 (1-2), pp. 165-200, May 1998.

[16] Shehory, O., “Optimality and Risk in Purchase at Multiple Auctions“ Proceedings of 5th International Workshop on Cooperative Information Agents (CIA 2001), M Klusch, F Zambonelli (eds.), Lecture Notes in Artificial Intelligence, Vol. 2182, Springer, 2001.

[17] Soh, L.K & Tsatsoulis, C., “Real-Time Satisficing Multiagent Coalition Formation” Proceedings of First International Joint Conference on Autonomous Agents and Multi-Agent Systems, Bologna, Italy 2002

[18] Soo, V.W. (2000) “Agent negotiation in trusted third party mediated uncertain games” Proceedings of 4th International Conference on Autonomous Agents, Barcelona, June, ACM Press, 2000.

[19] Sutton, R.S., and Barto, A.G., “Reinforcement Learning: An Introduction”, MIT Press, Cambridge, MA, Bradford Book, 1998.

[20] Tsvetovat, M. and Sycara, K., “ustomer coalitions in the electronic marketplace” Proceedings of 4th International Conference on Autonomous Agents, Barcelona, June, ACM Press, 2000.

[21] Vauvert, G., and El Fallah-Segrouhni, A., Coalition formation among strong autonomous and weak rational agents”, Proceedings of the 10th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW), Annecy, France, May 2001.

[22] Wellman,M.P., Ford, M., and Larson, K., “Path planning under time-dependent uncertainty” Proceedings of 11th Conference on Uncertainty in AI (UAI-95), Montreal, Canada, August, 1995

[23] Wellman, M.P. and Wurman, P.R. (1998) “Real time issues for Internet auctions” Proceedings 1st Workshop on Dependable and Real-Time E-Commerce Systems (DARE-98), Denver, CO, USA, June 1998.

[24] Yamamoto, J. and Sycara, K., “A stable and efficient buyer coalition formation scheme for e-marketplaces” Proceedings 5th International Conference on Autonomous Agents, Montreal, Canada, ACM Press, May/June, 2001.

[25] Yu, B., Singh, M.P. (2000) “A social mechanism of reputation management in electronic communities” Proceedings of International Workshop on Cooperative Information Agents IV, M Klusch & L Kerschberg (eds.), Lecture Notes in Artificial Intelligence, Vol. 1860, Springer, July 2000

[26] Suijs, J., Borm, P., and De Waegenaere, A., and Tijs, S. (1999) “Cooperative games with stochastic payoffs” European Journal of Operational Research, Vol. 113, pp. 193-205, 1999.

[27] Mares, M. (2001) “Fuzzy cooperative games – Cooperation with vague expectations” Physica Verlag, 2001.