

Dynamic Coalitions: A position paper

The Dynamic Coalitions Coalition*

14 Feb 2005

1 Context of the work

This work has been motivated by work within the GOLD[2] project, which seeks to build an architecture to facilitate the creation and maintenance of Virtual Organisations within the Chemical Engineering sector, and DSTL.

Within the School of Computing Science at Newcastle, a group has been working to develop a common understanding of Dynamic Coalitions. This document seeks to set out our common position, insofar as one has been achieved¹.

We begin by describing some possible contexts in which Dynamic Coalitions may arise. In Section 3 we discuss our progress in formal modelling, and in Section 4 we discuss the tools we intend to build.

2 Dynamic Coalition situations

Dynamic Coalitions emerge when the individual interests of a number of parties are considered to be best served by co-operation with each other. Within a political context, the OED defines *coalition* as

An alliance for combined action of distinct parties, persons, or states, without permanent incorporation into one body.

We can step from politics into many other contexts, by allowing these “parties, persons or states” to also

*Budi Arief, Jeremy Bryans, John Fitzgerald, Carl Gamble, Michael Harrison, Nigel Jefferson, Cliff Jones, Igor Mosolevsky, Thea Peacock, Peter Ryan.

¹Although this work is not being performed within a formal DIRC Targeted Activity, we believe that it may of interest to many within the DIRC project.

include companies, software or hardware agents, web-services and military units, among others.

The word “dynamic” points to the nature of the alliance mentioned above. A dynamic coalition may be formed spontaneously, members may join and leave without warning, and the nature of the relationships between participants may vary dramatically across the coalition as well as throughout its lifetime.

Dynamic coalitions are distinct from a collection of interacting parties by having a common goal or intention. Each party may have several goals, and some of these may be in conflict with those of other parties in the coalition, but the common goal is sufficiently strong for them to choose to form a coalition together.

Dynamic coalitions occur in many settings, for example:

- In response to an emergency, where dealing with the immediate emergency becomes the overriding interest of each party. This will be a dynamic grouping, because as the focus of the problem changes the necessary capabilities will be better supplied by different parties. For example, in the event of an earthquake, immediate evacuation of the area may be crucial and require the rapidly deployable capabilities of the army; later the health of the refugees may become the focus, requiring the intervention of the Red Cross.
- In response to commercial market forces, where each company believes that profitability will increase if they cooperate. In this context, they are often referred to as Virtual Organisations or Virtual Enterprises. For example, in the Chemical Engineering sector, the patent-holder of a

drug will have only a finite amount of time to manufacture and sell the drug before any other company may also exploit the same drug. If a company does not have all the necessary manufacturing skills “in-house”, it may reduce its time-to-market by forming a suitable coalition to compensate for the missing techniques. It will then have more time to exploit its monopoly.

- In a programmed-agent context, where the agents have been designed to use input from other agents or sensors around them. These are called multi-agent systems. As the communications topology changes, agents will have to continually reconfigure their “dynamic coalition” of communications with the other agents in order to maintain some desired functionality.
- In a military context, where achieving an objective may require the cooperation of a large number of military units. Here the dynamic nature of the cooperation could be due to new objectives becoming necessary as the conflict evolves.

3 Formal modelling

We are interested in the flow of information around these models of coalitions. For example, we are interested in identifying states of formal models in which information has reached the “wrong” actor, or where information has not reached the “right” actor. The purpose of the formal modelling will be to identify these states.

We therefore intend to model a theory of knowledge distribution in a distributed environment in which dynamic coalitions form, change and disperse.

To do this we will develop a formal model, within which we can specify a number of overlapping dynamic coalitions. We will begin with as simple a model as possible and continue by building an increasingly refined set of formal models.

Our formal model will at least be able to describe the following properties:

- Parties may know (or believe) different things.

- Information will be “time-valued” — items of information may become more or less valuable to parties over time.
- Different parties may make different inferences from the same pieces of information.

We have begun work on a simple formal model which will contain these elements.

We have said above that the purpose of our formal modelling will be to identify potential interesting or problematic information flow properties within coalitions. This will lead us to identify the behaviours of the individual parties that led to these properties.

After this, we may be able to make some deductions about the motives of the parties for these behaviours. For example, trust is an important aspect of dynamic coalitions. Differing trust relationships may lead to differing behaviours which may lead to differing information flow outcomes. It may be possible at a later stage of development to ascribe trust models or policies to the actors. These would essentially act as a set of constraints on the possible behaviours of the coalition.

We discuss below a set of open questions which we have discussed for the formal model, and our current thinking on each. We do not necessarily expect any one model to be able to describe all of these situations.

- Should we represent knowledge as atoms, or should relations between these atoms also be a component of knowledge? Atomic representation of knowledge may be sufficient to throw up interesting problems to do with information transfer, but it seems cumbersome to allow agents to reason about and manipulate these atoms in a meaningful way. It therefore seems sensible to continue to consider both.
- How do the knowledge and/or beliefs of individual actors map to the real world? Actors themselves may have some internal model of the real world. This will include facts which they consider to be true.

We could include in our model an “oracle” component to allow the user of the model to state

whether or not these facts actually are true. This has proved a thorny subject, and at this stage it is still not clear if it is a good idea or not.

- We should be able to allow actors associate meta-data with different “bits” of knowledge. This could include degree of certainty, secrecy, allowed recipients and provenance of data, among other things.

4 Analysis Tools

We have introduced the idea of building formal models that will allow us to explore the properties of dynamic coalitions, particularly in relation to the flow of knowledge/belief among agents. We also aim to develop tool support that will allow a range of different analyses to be performed on the formal models.

One relatively simple tool is an animator that supports the exploration of an executable version of a formal model showing individual agents and the coalitions of which they are members. Figure 1 shows a possible sketch of what this might look like. The model contains agents (circles) linked with possible communication channels. The scenario on the left describes some constraints on how these agents may evolve and communicate. We can query the model at stages of this evolution, to learn the answers to questions such as “when does agent A know this fact?”; “what does agent B know at this stage?”; “from whom did agent C get this piece of info?” etc.

A further form of automated analysis is the use of model checking to generate scripts that lead to undesirable states. The model of agents and coalitions needs to be represented as a state-transition system and encoded in a suitable model checking tool. The property characterising states of interest is presented to model checker, which generates a trace leading to such a state. This would need to be converted back to a script which could be executed through the tool interface. We are exploring the use of Spin [1] and CTL as a basis for such a tool. The challenging question is what properties characterise the states we wish to search for.

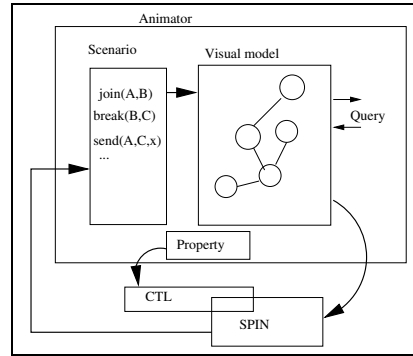


Figure 1: A possible view of the animator

Acknowledgements: We would like to thank Fred Schneider for his insightful comments on early versions of this work. Also to EPSRC projects DIRC, GOLD, and DSTL.

References

- [1] Edmund Clarke, Orna Grumberg, and Doron Peled. *Model Checking*. MIT press, 1999.
- [2] The GOLD project. www.gigamesh.ac.uk.