Chapter 32
ANTE: Agreement Negotiation in Normative and Trust-enabled Environments

Henrique Lopes Cardoso, Joana Urbano, Ana Paula Rocha, António J. M. Castro and Eugénio Oliveira

Abstract The ANTE framework encompasses results of research efforts on three main agreement technology concepts, namely *negotiation*, *normative environments* and *computational trust*. ANTE has been conceived as a general framework with a wide range of applications in mind. This chapter provides an overview of the main guidelines of this project, and explores two application domains for this framework: automated B2B electronic contracting, and disruption management in the context of an airline company operational control.

32.1 Introduction

Negotiation and task allocation have been in the multi-agent systems realm since its inception as a research field. More recently, social aspects of agenthood have received increasing attention, namely developments in the fields of normative and trust systems.

The ANTE[1] framework encompasses results of research efforts on three main agreement technology concepts, namely negotiation, normative environments and computational trust. ANTE is therefore the corollary of an ongoing long-term research project, which has been targeting the domain of B2B electronic contracting, although having been conceived as a more general framework with a wider range of applications in mind.

This chapter provides an overview of the main guidelines of this project, together with a brief description of its most important research contributions. Furthermore,

---

Henrique Lopes Cardoso, Joana Urbano, Ana Paula Rocha, António J. M. Castro and Eugénio Oliveira
LIACC, Dep. Eng. Informática, Faculdade de Engenharia, Universidade do Porto, Rua Dr. Roberto Frias, 4200-465 Porto, Portugal, e-mail: `{hlc,joana.urbano,arocha,antonio.castro,eco}@fe.up.pt`

[1] Agreement Negotiation in Normative and Trust-enabled Environments

two application domains for this framework are explored: automated B2B electronic contracting, and disruption management in the context of an airline company operational control.

Section 32.2 describes in broad lines the main concepts of the ANTE framework, and identifies the main research contributions in each of its main research areas: automatic negotiation, normative environments and computational trust. Section 32.3 describes the two application domains identified above and provides details regarding how ANTE has been exploited to fit those domains. Section 32.4 concludes the chapter.

32.2 The ANTE Framework

ANTE addresses the issue of multi-agent collective work in a comprehensive way, covering both negotiation as a mechanism for finding mutually acceptable agreements, and the enactment of such agreements. Furthermore, the framework also includes the evaluation of the enactment phase, with the aim of improving future negotiations.

Taking a broad perspective, an agreement can in this context be a solution obtained using a distributed cooperative problem solving approach. Therefore, a wide range of problems can be tackled. The agreement binds each negotiation participant to its contribution to the overall solution. It is therefore useful to represent the outcome of a successful negotiation process in a way that allows for checking if the contributions of each participant do in fact contribute to a successful execution of the agreement. A normative environment, within which agent interactions that are needed to enact the agreement will take place, takes care of this monitoring stage. Assessing the performance of each contribution is essential to enhance future negotiations. Computational trust may therefore be used to appropriately capture the trustworthiness of negotiation participants, both in terms of the quality of their proposals when building the solution (i.e. the practicability of the approach) and in terms of their ability to successfully enact their share.

In the following we provide some insight to the most important contributions of our developments in each of the aforementioned agreement technologies.

32.2.1 Negotiation

Negotiation is a form of decision-making where two or more parties jointly search a space of possible solutions with the goal of reaching a consensus. People use negotiation as a means of compromise in order to reach mutual agreements. In general, negotiation is defined as an interactive process whose goal is to achieve an agreement between interested parties. In competitive environments (as it is the case of e-business), self-interested agents have their own goals and are thus intrinsically competitive among each other; but even in this case it is also desirable for negotiating agents to have an incentive to cooperate in order to achieve efficient and mutually beneficial solutions. In cooperative environments, agents work together to find an optimal solution, e.g. by merging a set of multiple partial solutions. In ANTE, we have developed a negotiation protocol (Q-Negotiation) suitable for both competitive and cooperative environments that conducts to the selection of the best possible solutions [21]. Using this protocol, negotiation participants engage themselves in a sequential negotiation process composed of multiple rounds, by exchanging multi-attribute proposals and counter-proposals, trying to convince each other to modify the values for attributes they evaluate the most. The negotiation protocol selects the participants that, based on their capabilities and availability, will be able to make the best possible deal. However, since agents are autonomous entities, they are free to quit negotiation whenever they feel that no further concession is in their own interest.

It encompasses two important features:

- A multi-attribute evaluation to select the most favorable proposals at each round.
- A learning capability in order to enable agents to make the best possible deals even when faced with incomplete information and when operating in dynamic environments.

Attaching utility values to different attributes helps to solve the problem of multi-attribute evaluation. Generally, an evaluation formula is a linear combination of the current attribute values weighted by their corresponding utility values. However, in some cases, it can be a very difficult task to attach absolute values to attributes' utilities. A more natural and realistic situation is to simply impose a preference order over attributes' values and/or attributes themselves. Q-Negotiation adopts a multi-attribute evaluation based on a qualitative, as opposed to quantitative, measure. A learning capability is included through a Reinforcement Learning algorithm. The choice of this kind of learning algorithm has two main reasons. First, reinforcement learning algorithms support continuous, on-line learning during the negotiation process itself by making decisions according to the environment reactions in the past. The history of a negotiation (past rounds) is a crucial piece of information to be considered when deciding what to do in the next round. Second, reinforcement learning includes not only exploitation but also exploration facilities. In dynamic environments or in the presence of incomplete information, exploration (i.e. trying out new different possibilities) becomes a powerful technique. Learning is done through a qualitative comment that an agent receives, concerning its last proposal, from negotiating partners. The negotiation process results in the selection of a set of agents that commit themselves to the issues discussed during negotiation. This agreement may be formalized into a contract and subject to monitoring by a normative environment, as discussed in the following section.

32.2.2 Normative Environment

The normative dimension of a multi-agent system may, in general, encompass two perspectives on the interactions that norms are supposed to govern. Norms regulating pre-established interactions apply to the agent population as a whole, e.g. by specifying appropriate interaction conventions for negotiation. On the other hand, run-time norms are those that come into force when agents negotiate or adopt them to govern subsequent interactions (e.g. negotiated contracts). Within ANTE we are mostly concerned with the latter case, i.e., with norms that are agreed upon through a negotiation process.

In the context of agreement technologies, the role of a *normative environment* [14] is twofold. Given the agreement on a possible solution as obtained from the negotiation phase, it is necessary to check if the partial contributions of individual agents make their way in enabling a successful overall resolution of the problem. In many cases, the execution of the solution is itself distributed, which requires agents to enact by themselves their part of the agreement. *Monitoring* this phase is therefore an important task. Furthermore, in non-cooperative or dynamic scenarios, it is possible that after successfully negotiating an agreement self-interested agents are no longer willing to fulfill their commitments. This puts in evidence the second role of a normative environment, that of *enforcing* norms by coercing agents to stand for their commitments.

The notion of norm has been used with different meanings. In ANTE, a norm is a rule prescribing some behavior that agents governed by that norm must meet in certain circumstances. Given that these norms will govern previously negotiated agreements, the normative environment should enable the run-time establishment of new normative relationships. The “normative shape” of the environment will therefore evolve and adapt to the actual normative relationships that are established. In order to make this feasible, we believe it is important to provide some infrastructure that facilitates the establishment of norm governed relationships: a supportive and extensible *normative framework* [15] for framing the agreements that might be achieved. The main aim of this infrastructure is to assist software agents in the task of negotiating and establishing agreements that need an explicit representation for monitoring and enforcement purposes.

32.2.3 Trust

In Section 32.2.2, we addressed the role of normative environments for agreement technologies. In fact, control, legal norms and monitoring are common governance mechanisms used to reduce opportunism in business transactions [18, 30, 22, 1]. However, the drafting and monitoring of detailed contracts is sometimes costly and ineffective, and trust is often seen as a complementary, or even supplementary, governance mechanism that helps reduce the risk associated with business transactions [31, 19, 6, 30, 8].

Although trust is typically associated with uncertainty (for instance, in business transactions), it is an ubiquitous social concept present in everyday life. In fact, trust has been studied in several research areas in distinct fields such as close relationships, political relationships between countries, social networks, and even cooperative relationships. Computer science scholars, especially in the area of multi-agent systems, have been proposing *computational models of trust* that can be used to assist the decision making of agents when negotiating agreements, particularly in the phase of resource allocation.

Current research on computational trust models addresses the estimation of the trustworthiness of the target entities (individuals, groups, institutions, or things) by aggregating past evidence on these entities. These models tend to focus on some particular problem of computational trust, such as the modeling of the dynamics of trust [9], the context in which the evidence was produced [27], the use of reputation as a trust antecedent [10], and the modeling of trust in a socio-cognitive perspective [2].

Our approach to computational trust has as it main desideratum the ability to compute adequate estimations of trustworthiness in several different environments, including those of high dynamicity, where evidence on the agent in evaluation is scarce or even inexistent. Our model is composed of two basic components. The first one is Sinalpha, which computes general values of trustworthiness by relying on different properties of the dynamics of trust. Sinalpha models the trustworthiness of an agent using a function of α that presents a sinusoidal shape (see Equation 32.2.3). By setting $\delta = +0.5$, the trustworthiness value is restricted to the range $[0,1]$.

$$trustworthiness = \delta \cdot (\sin\alpha + 1)$$

The trustworthiness score of the agent is minimum when $\alpha = 3\pi/2$ and maximum at $\alpha = 5\pi/2$. This score is updated using $\alpha_{i+1} = \alpha_i + \lambda \cdot \omega$, where λ reflects the outcome associated with the piece of evidence being aggregated (it assumes positive values for evidences with positive outcomes and negative values for evidences with negative outcomes, where $|\lambda^+| < |\lambda^-|$), and parameter ω is used to define the size of the ascending/descending step in the trustworthiness path. A detailed description of Sinalpha is given in [29].

The second component of our computational model is Contextual Fitness, a situation-aware tuner that downgrades the trustworthiness scores computed by Sinalpha in cases where the agent in evaluation has proved to behave poorly in the situation in assessment. Its mode of operation is based on the dynamic extraction of *tendencies of failure* from the past evidence of the agent, using the *information gain* metric [20]. This approach differs from other situation-aware computational trust approaches by its flexibility and ability to reason in terms of context even when the evidence on the agent in evaluation is scarce [26, 27].

32.3 Application Domains

In this Section we will describe two application domains being addressed by the technologies that integrate the ANTE framework.

32.3.1 B2B E-contracting

The first scenario addressed by ANTE is that of B2B electronic contracting. With a strong automation perspective, the scenario envisages the use of software agents negotiating on behalf of their principals, which are buyers or suppliers in a B2B network. Negotiation is therefore used to select, among a group of potential suppliers, the best ones to fit a particular business opportunity. Contracts resulting from successful negotiations are validated, registered and digitally signed, before being handed to the normative environment for monitoring and enforcement purposes. Finally, the way agents enact their contracts provides important information for trust building. A repository of trust and reputation information may then complete the circle by providing relevant inputs for future negotiations. The integration of all these stages is depicted in Figure 32.1.

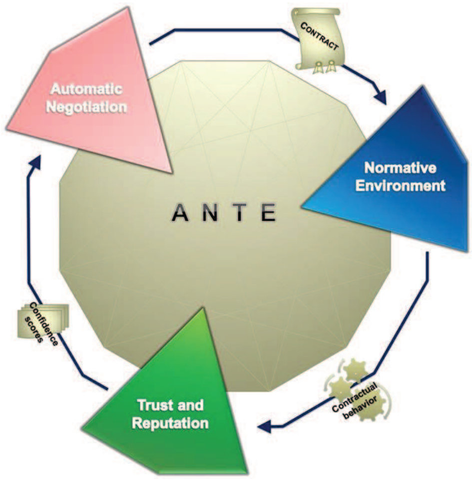

Fig. 32.1: ANTE.

Important synergies are obtained from the integration of the three main research domains identified in Figure 32.1. Negotiation is informed by trustworthiness assessments of negotiation participants. In ANTE, this may be put in practice in three different ways: using trust for preselecting the partners with whom to negotiate; evaluating negotiation proposals taking into account the trustworthiness of proposal issuers; or exploiting trust information when drafting a contract with a selected supplier, e.g. by proposing a sanction in case the supplier breaches the contract (thus trying to reduce the risk associated with doing business with a not fully trustworthy agent).

Connecting the monitoring facility of the normative environment with a computational trust engine means that we can use contractual evidences regarding the behavior of agents when enacting their contracts to build trust assessments. Our approach to modeling contractual obligations [17] allows for a rich set of possible contract enactment outcomes (fulfillments, delays, breaches, and so on), which in turn enables a trust engine to weight differently the possible sub-optimal states that might be obtained [23, 24].

As mentioned in Section 32.2.2, connecting negotiation with a normative environment that provides a monitoring service opens up the possibility of providing some normative infrastructure that facilitates contract establishment. For that, the normative environment should provide a supportive and extensible normative framework. Inspired by notions from contract law theory, namely the use of "default rules" [5], we have proposed a model for this normative structure based on a hierarchy of *contexts* [16], within which norms are created that may apply to subcontexts. The context hierarchy tries to mimic the fact that in business it is often the case that a B2B contractual agreement forms the business context for more specific contracts that may be created. Each contract establishes a new context for norm applicability. A *norm defeasibility* approach [15] is used to determine whether a norm should be inherited, for a specific situation, from an upper context. This feature allows the normative framework to be adapted (to better fit a particular contract case) and extended (allowing new contract types to be defined). The rationale behind this design is based on the assumption that "default rules" should be seen as facilitating rather than constraining contractual activity [11].

32.3.1.1 Prototype.

The ANTE framework has been realized as a JADE-based FIPA-compliant platform. In the case of the e-contracting application domain, as can be seen in Figure 32.2, there are three kinds of agents in the platform: those that provide contracting services (upper part of the figure), namely negotiator, computational trust, ontology mapping, notary and normative environment; external agents whose role is to make a connection to real-world contract enactment events (e.g. deliveries, payments); and users of the system (lower part of the figure), representing buyers and suppliers.

Figure 32.3 shows a buyer interface for specifying its needs (left side) and for configuring a multi-attribute negotiation (right side) to take place using the negotiator service. Options include how trust is to be used in each of the negotiation steps, as described earlier in this section. Also, the buyer may indicate the type of contract that is to be created should negotiation succeed; norms governing specific contract types are already available in the normative environment, thus making it easier to establish a contract.

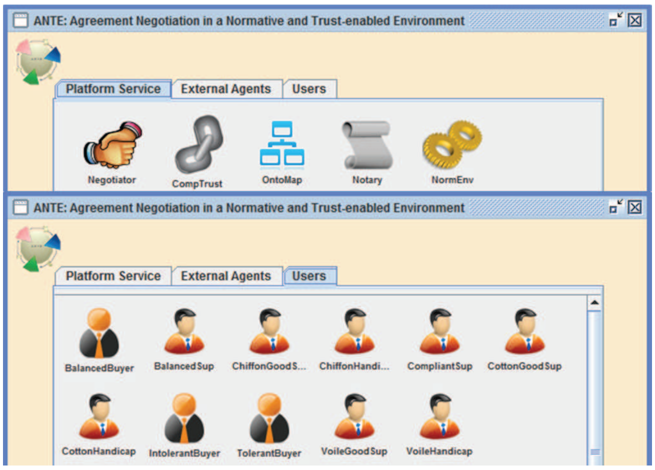

Fig. 32.2: ANTE main screen.

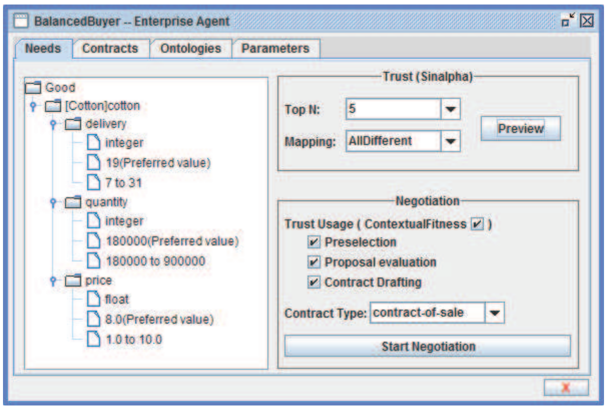

Fig. 32.3: Buyer's needs and negotiation configuration.

Figure 32.4 shows, on the buyer interface, the contracts it has already established (upper part) and a set of events related to their enactment (lower part). These events are automatically reported by the normative environment in the contract monitoring phase.

Turning to the supplier interface, in Figure 32.5 we can inspect the negotiations that took place, together with the messages exchanged using the Q-Negotiation protocol described in Section 32.2.1.

The negotiator interface (see Figure 32.6) shows the evolution of the proposals exchanged during a negotiation protocol in terms of their utility for the buyer that started the negotiation process.

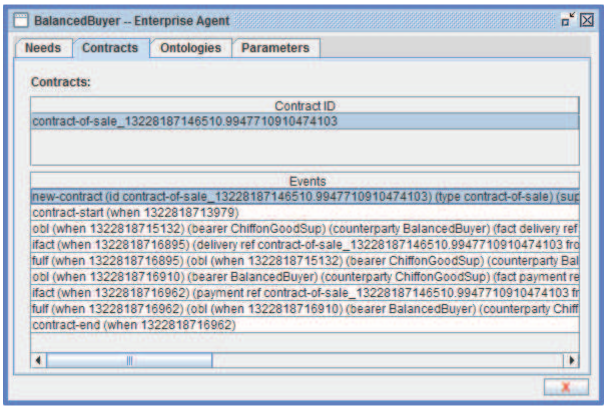

Fig. 32.4: Buyer’s established contracts and their enactment.

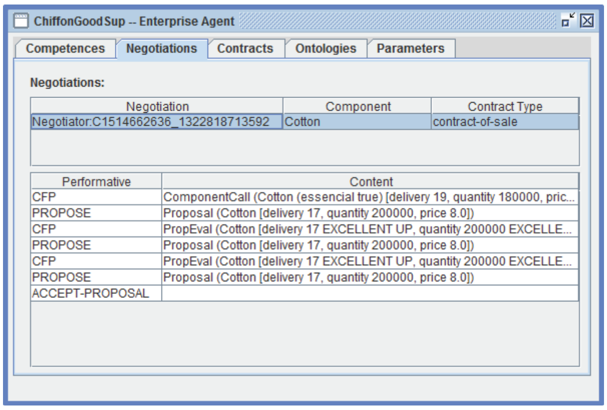

Fig. 32.5: Supplier’s negotiations.

The interface for the computational trust service (shown in Figure 32.7) allows us to inspect how trustworthiness assessments are being computed, including the contractual evidences that are used as input for each agent. It also allows us to choose the mapping method that associates different weights to each of the possible contract enactment outcomes.

The scenario that is illustrated throughout this sequence of screenshots is from the textile industry domain. We have run several experiments with the aim of trying to figure out the best ways of integrating negotiation, norms and trust [28, 24, 25].

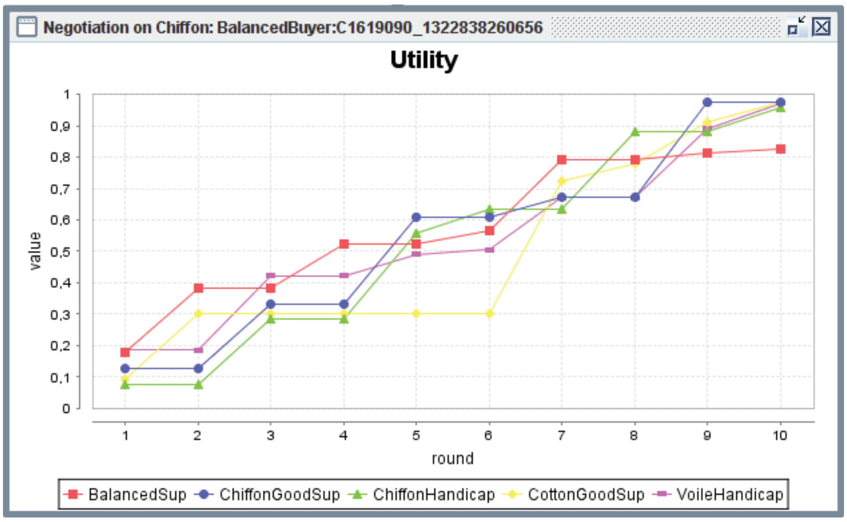

Fig. 32.6: Negotiator: proposal evolution in a particular negotiation.

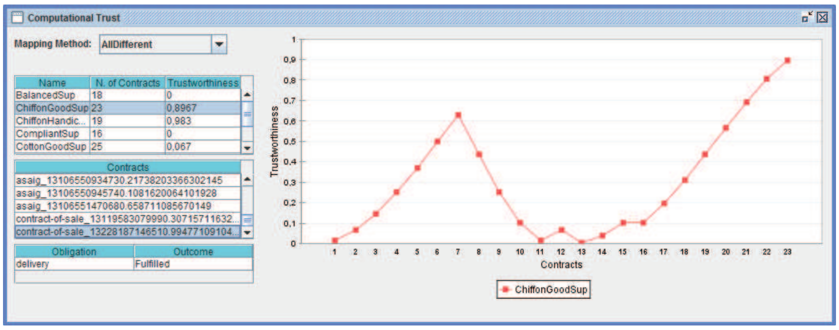

Fig. 32.7: Computational trust: computing trustworthiness assessments from contractual evidences.

32.3.2 MASDIMA

The second scenario addressed by ANTE is related to disruption management in Airline Operations Control. MASDIMA [2] is an agent-based application that represents the Airline Operations Control Centre (AOCC) of an airline company. The AOCC is the organization responsible for monitoring and solving operational problems that might occur during the execution of the airline operational plan. It includes teams of human experts specialized in solving problems related to aircrafts and flights, crewmembers and passengers, in a process called Disruption Manage-

[2] Multi-Agent System for DIsruption MAnagement

ment. In this section we will briefly introduce the Airline Operations Control Problem (AOCP) and we will present our solution to this problem, i.e., the agent based application MASDIMA [3]. Although we present a high level view of the system architecture we will give more emphasis on how we used the ANTE Framework (described in Section 32.2) to implement this application.

Airline companies developed a set of operations control mechanisms to monitor the flights and check the execution of the schedule. During this monitoring phase, several problems may appear related to aircrafts, crewmembers and passengers [4]. According to Kohl et al. [13], disruption management is the process of solving these problems. To be able to manage disruptions, airline companies have an entity called Airline Operations Control Centre (AOCC). This entity is composed of specialized human teams that work under the control of an operations supervisor. Aircraft, Crew and Passenger teams are amongst the most important ones. Although each team has a specific goal (for example, the crew team is responsible for having the right crew in each flight), they all contribute to the more general objective of minimizing the effects of disruption in the airline operational plan. During the execution of an operational plan, several events or problems might appear, e.g., aircraft malfunctions, enroute and/or departure or destination weather conditions, crewmembers not reporting for duty, passengers not reporting at gate, and so on. These problems, if not solved, might cause flight departure and/or arrival delays. AOCCs have a process to monitor the events and solve the problems, so that flight delays are minimized with the minimum impact on passenger and, preferably, with the minimum operational cost. Typically, the main costs to consider are: (1) Crew Costs, (2) Flight Costs and (3) Passenger Costs. There is also a less easily quantifiable cost that is also included: the cost of delaying or cancelling a flight from the passenger point of view. Most airlines use some kind of rule-of-thumb when they are evaluating the impact of the decisions on passengers. Others just assign a monetary cost to each minute of delay and evaluate the solutions taking into consideration this value. When faced with a disruption, the AOCC needs to find the best solution that minimizes the delay and costs of the flights, crewmembers and passengers affected returning, as soon as possible, to the previous operational plan.

As in the B2B scenario presented in Section 32.3.1, important synergies are obtained from the integration of the three main research domains identified in Figure 32.1. It is important to point out that in the disruption management scenario we have a closed and cooperative environment that contrasts with the open and competitive environment of the B2B scenario. Figure 32.8 shows the MASDIMA architecture. The agents *Tracking* (keep log of negotiation messages), *Learning* (increase robustness of future plans) and *Event Information* (system that registers the event information on the environment), although implemented, are not relevant for the scope of this chapter. The agent *Data Visualization* is responsible to update the user interface (Figure 32.9) with information so that the users can interact with the system. The *Monitor* agent is responsible for the runtime execution of the problem and is the counterpart of the normative environment as modeled in the scenario presented in Section 32.3.1.

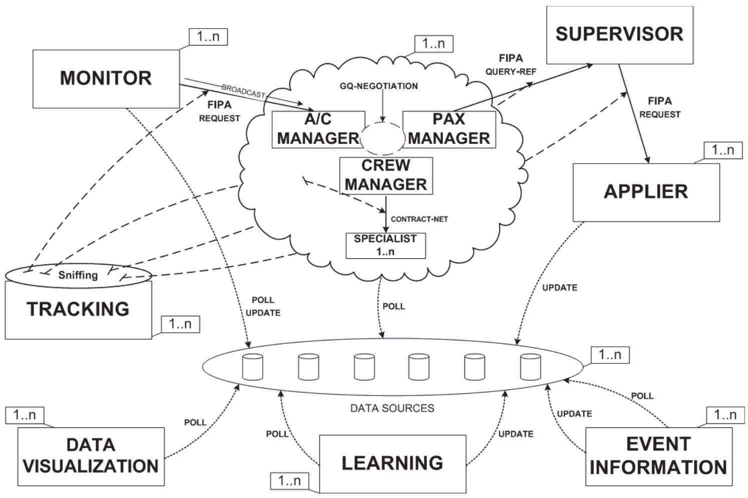

Fig. 32.8: MASDIMA Architecture.

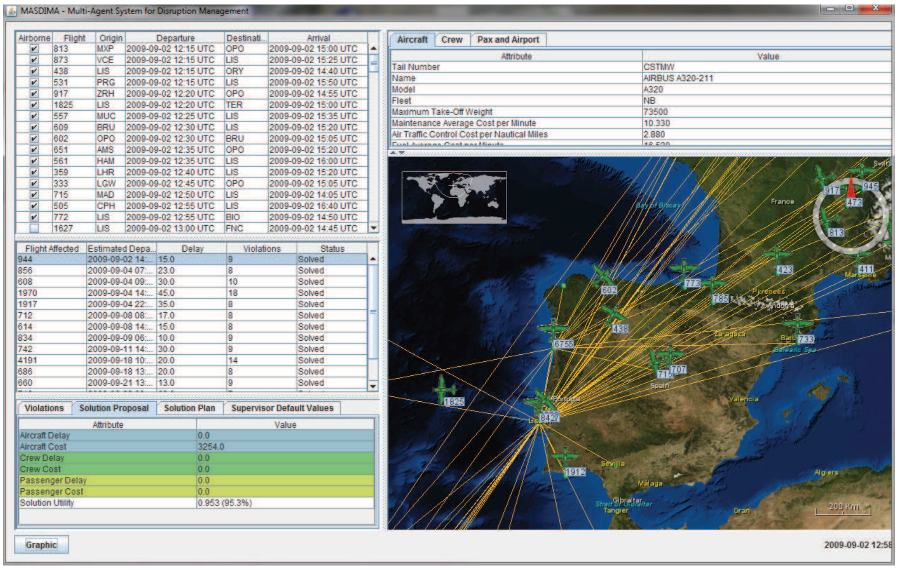

Fig. 32.9: MASDIMA User Interface.

The main negotiation takes place between the *Supervisor* and the *A/C, Crew and Pax* manager agents and the negotiation protocol used has the characteristics identified in Section 32.2.1. The Supervisor acts as the organizer agent and the managers as respondents. Since we are in a cooperative environment, each manager does not possess the full expertize to be able to propose a solution to the supervisor. As such, each manager needs to start an inter-manager negotiation to be able to complete their proposal and participate in the main negotiation. Although we are in a cooperative environment each manager wants to maximize its own utility and act according to its preferences. In this scenario the number of participants is defined according to the problem we want to tackle. As a minimum we need to have at least one manager for each part (or dimension) of the problem. Nevertheless, we can have more than one agent with the same expertize in the same dimension of the problem.

In this scenario, trust is used when the supervisor is evaluating negotiation proposals from the managers. The trust information is built from the application of the winner solution on the environment through the *Applier* agent. In Figure 32.8 we can see that the manager agents are not responsible for applying the winning solution to the environment. That is a task for the Applier agent, which checks the successful execution of the solution. Connecting the monitoring facility with the trust engine enables the Supervisor agent to use evidence regarding the quality of solutions proposed in previous problems by the managers and applied by the Applier agent.

A final word regarding the *Specialist* agents that appear in the MASDIMA architecture in Figure 32.8. In this scenario, in order for the manager agents to present a proposal, they first need to find a candidate solution using the resources that exist on the operational plan. For example, a candidate solution for the aircraft dimension could be to swap the aircrafts between two flights. Likewise, a candidate solution to the crew dimension could be to use crewmember A instead of B and for the pax dimension a new itinerary from the departure airport to the destination. To find these candidate solutions, each manager might have a team of problem solving agents (the specialists) that, implementing algorithms like Simulated Annealing [12] or Dijkstra shortest-path [7] are able to find those solutions. It is from these candidate solutions that the managers take the attribute values necessary to present a proposal during the main negotiation.

MASDIMA is being tested at TAP Portugal (the major Portuguese airline) using real data and the results compared with the ones provided by the human operators in the Airline Operational Control Centre, using current tools and expertize. Results show that with MASDIMA it is possible to have less flight delays and lower operational costs.

At present we are integrating MASDIMA with the current end systems of the airline company and we are planning to enrich the negotiation protocol with arguments.

32.4 Conclusions

Real world applications of agreement technologies are better addressed by taking an integrative approach. The ANTE framework seeks to provide an environment where the interdependencies between different research domains – namely negotiation, norms and trust – can be experimented with. Although not addressed in this chapter, other areas of agreement technologies, such as semantics (ontologies) and argumentation, are also being addressed within the same research environment.

The quite disparate application domains described here demonstrate the effort that is being put into applying our research results in different areas. Having a strong initial focus on the formation of virtual enterprises, and later to general B2B electronic contracting, part of the framework (mostly negotiation) is being used to address the problem of disruption management in an airline operational control, together with all the issues that this problem raises.

Acknowledgements This research is supported by Fundação para a Ciência e a Tecnologia (FCT), under project PTDC/EIA-EIA/104420/2008, and by QREN under project RESPLAN. The authors would also like to thank Pedro Brandão for his work on the implementation of the ANTE prototype.

References

1. Bachmann, R.: Trust, power and control in trans-organizational relations. Organization Studies **22**(2), 341–369 (2001)
2. Castelfranchi, C., Falcone, R.: Trust Theory: A Socio-Cognitive and Computational Model. Wiley Series in Agent Technology. John Wiley & Sons Ltd., Chichester (2010)
3. Castro, A., Oliveira, E.: A new concept for disruption management in airline operations control. Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering **3**(3), 269–290 (2011). DOI 10.1243/09544100JAERO864
4. Clausen, J., Larsen, A., Larsen, J.: Disruption management in the airline industry - concepts, models and methods. Technical report, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby (2005). URL `http://www2.imm.dtu.dk/pubdb/p.php?3763`
5. Craswell, R.: Contract law: General theories. In: B. Bouckaert, G. De Geest (eds.) Encyclopedia of Law and Economics, vol. III, pp. 1–24. Edward Elgar, Cheltenham, UK (2000)
6. Das, T.K., Teng, B.: Between trust and control: Developing confidence in partner cooperation in alliances. Academy of Management Review **23**(3), 491–512 (1998)
7. Dijkstra, E.W.: A note on two problems in connexion with graphs. Numerische Mathematik **1**, 269–271 (1959)
8. Ireland, R.D., Webb, J.W.: A multi-theoretic perspective on trust and power in strategic supply chains. Journal of Operations Management **25**(2), 482 – 497 (2007)
9. Jonker, C.M., Treur, J.: Formal analysis of models for the dynamics of trust based on experiences. In: MultiAgent System Engineering, 9th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, pp. 221–231. Springer-Verlag (1999)
10. Josang, A., Ismail, R., Boyd, C.: A survey of trust and reputation systems for online service provision. Decision Support Systems **43**(2), 618 – 644 (2007)
11. Kaplow, L.: General characteristics of rules. In: B. Bouckaert, G. De Geest (eds.) Encyclopedia of Law and Economics, vol. Volume V: The Economics of Crime and Litigation, pp. 502–528. Edward Elgar, Cheltenham, UK (2000)

12. Kirkpatrick, S., Jr, C.G., Vecchi, M.: Optimization by simulated annealing. Science **220**(4598), 671–680 (1983)
13. Kohl, N., Larsen, A., Larsen, J., Ross, A., Tiourline, S.: Airline disruption management: Perspectives, experiences and outlook. Technical Report CRTR-0407, Carmen Research (2004)
14. Lopes Cardoso, H.: Electronic institutions with normative environments for agent-based e-contracting. Ph.D. thesis, Universidade do Porto (2010)
15. Lopes Cardoso, H., Oliveira, E.: Norm defeasibility in an institutional normative framework. In: M. Ghallab, C. Spyropoulos, N. Fakotakis, N. Avouris (eds.) Proceedings of The 18th European Conference on Artificial Intelligence (ECAI 2008), pp. 468–472. IOS Press, Patras, Greece (2008)
16. Lopes Cardoso, H., Oliveira, E.: A context-based institutional normative environment. In: J. Hubner, E. Matson, O. Boissier, V. Dignum (eds.) Coordination, Organizations, Institutions, and Norms in Agent Systems IV, LNAI 5428, pp. 140–155. Springer (2009)
17. Lopes Cardoso, H., Oliveira, E.: Directed deadline obligations in agent-based business contracts. In: J. Padget, A. Artikis, W. Vasconcelos, K. Stathis, V. Torres da Silva, E. Matson, A. Polleres (eds.) Coordination, Organizations, Institutions, and Norms in Agent Systems V, LNAI 6069, pp. 225–240. Springer (2010)
18. Luhmann, N.: Trust and Power. John Wiley & Sons, New York (1979)
19. Mayer, R.C., Davis, J.H., Schoorman, F.D.: An integrative model of organizational trust. The Academy of Management Review **20**(3), 709–734 (1995)
20. Quinlan, J.R.: Induction of Decision Trees. Mach. Learn. **1**, 81–106 (1986)
21. Rocha, A., Lopes Cardoso, H., Oliveira, E.: Virtual Enterprise Integration: Technological and Organizational Perspectives, chap. Contributions to an Electronic Institution supporting Virtual Enterprises' life cycle, chapter XI, pp. 229–246. Idea Group Inc. (2005)
22. Sako, M.: Does trust improve business performance? In: C. Lane, R. Bachmann (eds.) Trust within and between Organizations: Conceptual Issues and Empirical Applications. Oxford University Press (1998)
23. Urbano, J., Lopes Cardoso, H., Oliveira, E.: Making electronic contracting operational and trustworthy. In: 12th Ibero-American Conference on Artificial Intelligence. Springer, Bahia Blanca, Argentina (2010)
24. Urbano, J., Lopes Cardoso, H., Oliveira, E., Rocha, A.: Normative and trust-based systems as enabler technologies for automated negotiation. In: F. Lopes, H. Coelho (eds.) Negotiation and Argumentation in MAS. Bentham Science Publishers Ltd. (2012)
25. Urbano, J., Lopes Cardoso, H., Rocha, A., Oliveira, E.: Trust and normative control in multi-agent systems: An empirical study. In: 10th International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS 2012), Special Sessio on Trust, Incentives and Norms in open Multi-Agent Systems (TINMAS). Springer, Salamanca, Spain (2012)
26. Urbano, J., Rocha, A., Oliveira, E.: Trustworthiness tendency incremental extraction using information gain. In: Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on, vol. 2, pp. 411–414 (2010). DOI 10.1109/WI-IAT.2010.151
27. Urbano, J., Rocha, A., Oliveira, E.: A situation-aware computational trust model for selecting partners. Transactions on computational collective intelligence V pp. 84–105 (2011)
28. Urbano, J., Rocha, A., Oliveira, E.: Trust-based selection of partners. In: C. Huemer, T. Setzer, W. Aalst, J. Mylopoulos, M. Rosemann, M.J. Shaw, C. Szyperski (eds.) E-Commerce and Web Technologies, *Lecture Notes in Business Information Processing*, vol. 85, pp. 221–232. Springer Berlin Heidelberg (2011)
29. Urbano, J., Rocha, A.P., Oliveira, E.: Computing confidence values: Does trust dynamics matter? In: L. Seabra Lopes, N. Lau, P. Mariano, L.M. Rocha (eds.) 14th Portuguese Conference on Artificial Intelligence (EPIA 2009), vol. LNAI, pp. 520–531. Springer, Aveiro, Portugal (2009)
30. Wathne, K.H., Heide, J.B.: Opportunism in interfirm relationships: Forms, outcomes, and solutions. The Journal of Marketing **64**(4), 36–51 (2000)
31. Williamson, O.E.: Transaction-cost economics: The governance of contractual relations. Journal of Law and Economics **22**, 233–261 (1979)