A Hundred Schools of Thought Automatically Contending

Peter McBurney and Simon Parsons
Department of Computer Science
University of Liverpool
Liverpool L69 7ZF U.K.
{p.j.mcburney,s.d.parsons}@csc.liv.ac.uk

August 8, 2001

We envision a rational cacophony of "numberless infinities" of autonomous software agents engaged in dialogue and argument with each other, debating automatically every topic which humans now discuss and many we do not. This note presents the major elements we believe are necessary for the achievement of this vision, and outlines the progress already made — by ourselves and others — towards its realization. We begin with a discussion as to why we believe argumentation is essential to modeling of autonomous agent interactions.

Much research to date in interaction between autonomous software agents has focused on negotiation, and has used economic auction and game-theoretic mechanisms as the enabling technology [24]. Perhaps this is understandable given the recent public attention paid to electronic commerce and the computational and theoretical challenges involved even in designing simple auction mechanisms. Yet, most every-day human interactions do not involve negotiation; indeed, even for those which do, the participants typically spend more time engaged in relevant dialogue before, during and subsequent to the transaction than in actually exchanging offers with each other. Participants to a negotiation even engage in relevant dialogue with others not party to the potential transaction, a fact which has led marketing theorists to study what they call *word-of-mouth* influences on consumer decision-making, e.g. [8].

It is interesting to ask why this is so: why do human parties to a negotiation exchange more than the simple acceptances and rejections permitted in most auction mechanisms? A key reason is that human participants rarely, if ever, satisfy one of the central tenets of the classical economic models of decision-making [26]: they do not usually begin a decision process with pre-determined preferences and utility valuations, but instead these are formed (either partially or completely) in the very process of undertaking the transaction [45]. In addition, even those participants with predetermined beliefs and preferences are rarely perceived by other participants to hold their views fixedly, thus leading to attempts by those others to persuade them to change their views. These acts of attempted persuasion occur because one party believes the other able and willing to change his or her mind, and that, with such a change, there is a greater chance of successful resolution of the issue being discussed.

We may assume that rational participants to an interaction would not change their minds without receipt of additional information, that is, that they would not change their beliefs, preferences or intentions on the basis merely of whim or malice. With this assumption, the transmission of relevant information between the participants becomes a key determinant of the successful resolution of an interaction, and hence, information-transmission needs to be enabled by the interaction protocol used by the participants. A weakness of economic auction mechanisms as interaction protocols is that they typically permit only limited information to be transmitted between the parties.[1] Negotiation protocols based on argumentation, however, enable the transmission of considerably more information between the participants, and in a manner which is rational and coherent.

These ideas have greater importance if we extend the remit of interaction beyond negotiation, viewing the latter (following [49]) as dialogue concerning the division of some scarce resource. Agents (or humans) engaged in joint deliberation to decide a course of action, for example, may each have different knowledge, none of which is complete but all of which is germane to the problem; achievement of a common joint intention may only be possible by coherent sharing of this distributed knowledge, as is the case for the room-furnishing agents of [39]. Because theories of argumentation apply to a wider class of discourses than negotiation dialogues, they are capable of supporting these other types of interaction; auction mechanisms, in contrast, appear limited to negotiations.

To achieve our vision of a cyber-world of multitudes of automatically arguing agents, a number of distinct components will have to be created and assembled. The first of these will be languages for argument and dialogue, which specify locutions and define syntactical rules for their use. The development of standardised general agent communications languages has been the focus of considerable work, e.g. [13, 29]. We have instead approached the task by defining application-specific protocols for particular types of interactions, protocols which may be seen as examples of electronic *institutions* [38]. We have done this using formal dialogue games, taken from philosophy, in which two or more players "move" by uttering locutions according to some defined rules. Dating at least from the time of Aristotle, dialogue games have been used by contemporary philosophers to study fallacious modes of reasoning [20, 28] and as a game-theoretic semantics for intuitionistic, classical [27] and quantum [37] logic. Within Artificial Intelligence they have been applied to: modeling complex reasoning, e.g. in legal domains [6]; human-computer interaction [5]; designing software development processes [12]; and for the design of computational protocols for autonomous agent interactions. Examples of dialogue-game agent protocols include systems for: team-formation [9]; information-seeking dialogues, where one agent seeks the answer to a question from others believed to know it [23]; persuasion dialogues, where one agent seeks to persuade another to adopt some statement [10, 34]; negotiation dialogues, where agents discuss how to divide some scarce resource [2, 3, 36, 46]; deliberation dialogues, where agents collaborate to decide a course of action in some situation [22]; and chance discovery dialogues, such as identification of risks or opportunities, where agents together seek the answer to some question subject to value constraints on the possible answers [35].[2]

A second requirement for our vision will be mechanisms "in the heads" of participating agents for interpreting locutions uttered in a dialogue and choosing appropriate responses, both single locutions and strategically. For negotiation dialogues, we have

[1] Participants to negotiations may exercise considerable creativity in overcoming the limited capabilities for information transmission of an auction mechanism, as was seen in the allegedly-collusive behaviour of bidders in the US Federal Communications Commission PCS Spectrum Auctions of 1994-5 [7].

[2] The categorization of dialogues here is based on that of [49].

proposed mechanisms based on models adopted from the marketing theory of consumer decision-making [36]. Similar generative mechanisms have been proposed in recent agent negotiation architectures which do not use argumentation, e.g. [4, 11]. An important question is the extent to which such semantic mechanisms can ever be verified, since a sufficiently-clever agent will always be able to simulate insincerely any semantic requirement [51]. Recent work has proposed the use of a social semantics [19, 48], effectively public statements of private mental states such as beliefs and intentions, to ensure all participants assign the same meaning to syntactical statements. Of course, duplicitous agents will still be able to make false declarations of such social commitments.

A third requirement will be for mechanisms to decide to initiate or to enter dialogues of particular types on particular subjects at particular times. For this task, research in computational linguistics linking dialogue locutions with the beliefs and intentions of participants is likely to be relevant. For example, the model of human dialogues of [18] identifies a purpose for the discourse and purposes for subsidiary segments of the discourse. As the dialogue proceeds, these purposes are jointly recognized by the participants (or not) and then achieved by means of discussion (or not), thus leading to closure or to subsequent dialogue segments and dialogues.

Fourthly, we will require formal rules of encounter and procedure, the argumentation equivalent of auction mechanisms. Examples of such systems of rules include: Robert's rules of order for public meetings [44]; Alexy's rules for discourses over ethical and moral questions [1]; and Hitchcock's principles for rational mutual inquiry [21]. Because, as we argued above, the value of argumentation in agent interactions lies essentially in the ability and willingness of participants to change their beliefs, desires and intentions, then argumentation-based systems must permit participants to undergo self-transformation [14]; i.e. agents must be permitted to change their mental states. But to do this rationally and coherently requires rules for making assertions, for questioning and contesting assertions, and for justifying and retracting prior assertions. Thus, a system of dialogue between rational agents requires an explicit theory of argumentation and rules of encounter. Some work has been undertaken to develop computational versions of such rules, e.g. [15, 40], but the only dialogue-game agent communications protocol known to us which has been shown formally to comply with explicit rules of encounter is our own model of persuasion in scientific domains [34].

A fifth requirement for our vision are coherent models of complex dialogues, such as dialogues undertaken in sequence and in parallel, and embeddings of one dialogue inside another. An essential step towards this was the partial typology of six basic types of human dialogues developed by Walton and Krabbe [49], which categorizes dialogues according to the dialogue purpose, the initial information known to each participant and their respective intentions. Building on this typology, Reed [41] proposed a computational formalism to represent complex combinations of the basic types. A second computational model for complex dialogues is our dynamic modal formalism [30]. Before such formalisms can be applied, a better understanding of their properties, such as their computational complexity, will be required.

A final requirement is a formal theory of dialogue protocols within which to compare one protocol with another. Currently, potential users of on-line auction systems are faced with choosing between auctions conducted under different rules and procedures. The same problem will confront intending users of dialogue protocols. When are two dialogue protocols the same? How similar are two different protocols? When is one protocol better than another? In what sense, and for what purpose? These questions will be increasingly important as agent system proliferate and more dialogue protocols are developed.

In order to answer such questions, we have recently proposed a geometric semantics for agent dialogue protocols [32], in which dialogues are interpreted as paths in a real, multi-dimensional space. The rules of the protocol may then be seen as delimiting the sub-space of possible paths, for example by the establishment of forbidden regions when specific combinations of locutions are not permitted. Using this semantics we have been able to show that two similar generic protocols are in fact not equivalent, in the sense that they may give rise to dialogues with qualitatively different commitments being uttered by the participants. Because this interpretation treats dialogues as directed execution traces, there is an apparent similarity to geometric models of concurrency in distributed computer systems [17], in which domain the methods of algebraic topology have been successfully applied. We are currently exploring the relevance of these ideas for a mathematical theory of dialogue protocols. Because auction mechanisms may be viewed as protocols for very simple dialogues, a formal theory of dialogue protocols should also encompass these mechanisms, and thus answer similar questions about the properties and relative performance of auctions.

Given a formal theory of dialogue-game agent interaction protocols, we should be able to use this to articulate a set of evaluation criteria for protocols, providing guidance to both designers and potential users. We have already presented a tentative list of evaluation criteria [31, 42], drawing on our protocol design experience, and informed by criteria proposed for the design and assessment of automated auction mechanisms [24, 47]. However, the economic problems for which auction mechanisms have been designed can usually assume the existence of a common unit of inter-personal utility comparison, namely money. Consequently, the criteria proposed for assessment of auction mechanisms have emphasized the assessment of outcomes rather than the assessment of procedures used to reach those outcomes. In contrast, many decisions, especially those in public policy domains, have possible outcomes which are not comparable, as comparison requires inter-personal utility assessments. For example, to assess whether banning cigarette smoking would be a desirable policy would require comparison of the disutilities to one set of individuals, current and future smokers, with the utilities to others, non-smokers; to add to the comparison difficulty, some of these individuals are as yet unborn. That many decisions result in incommensurable outcomes leads us to believe that we also require formal theories of *decision-process rationality* to complement the *decision-outcome rationality* defined by maximum-expected-utility in classical decision theory [26]. Recent research in argumentation, e.g. [25, 50], and in political science, e.g. [14, 43], may lead to a theory of rational decision-processes, and may guide the design of dialogue protocols for public policy domains, as in [16, 33].

As is clear from this list of requirements, the realization of our vision will draw on contributions from argumentation, computer science, decision theory, economics, linguistics, marketing theory, pure mathematics, philosophy, political science, rhetorical theory and sociology. This complex mix of large parts of human thought is one reason why research towards the vision we have outlined is so exciting, although we are still a long way from its achievement.

References

[1] R. Alexy. A theory of practical discourse. In S. Benhabib and F. Dallmayr, editors, *The Communicative Ethics Controversy*, pages 151–190. MIT Press, Cambridge, MA, USA, 1990. (Published in German 1978).

[2] L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In E. Durfee, editor, *Proc. 4th Intern. Conf. on Multi-Agent Systems (ICMAS-2000)*, pages 31–38, Boston, MA, USA, 2000. IEEE.

[3] L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proc. 14th European Conf. on AI (ECAI 2000)*, pages 338–342, Berlin, Germany, 2000. IOS Press.

[4] M. Barbuceanu and W-K. Lo. A multi-attribute utility theoretic negotation architecture for electronic commerce. In C. Sierra, M. Gini, and J. S. Rosenschein, editors, *Proc. 4th Intern. Conf. on Autonomous Agents (AA-2000)*, pages 239–246, New York, NY, USA, 2000. ACM Press.

[5] T. J. M. Bench-Capon, P. E. Dunne, and P. H. Leng. Interacting with knowledge-based systems through dialogue games. In *Proc. 11th Intern. Conf. on Expert Systems and Applications*, pages 123–140, Avignon, France, 1991.

[6] T. J. M. Bench-Capon, T. Geldard, and P. H. Leng. A method for the computational modelling of dialectical argument with dialogue games. *Artificial Intelligence and Law*, 8:233–254, 2000.

[7] P. Cramton and J. A. Schwarz. Collusive bidding: Lessons from the FCC Spectrum Auctions. *Journal of Regulatory Economics*, 17:229–252, 2000.

[8] J. A. Czepiel. Word of mouth processes in the diffusion of a major technological innovation. *Journal of Marketing Research*, 11:172–180, May 1974.

[9] F. Dignum, B. Dunin-Kęplicz, and R. Verbrugge. Agent theory for team formation by dialogue. In C. Castelfranchi and Y. Lespérance, editors, *Proc. 7th Intern. Workshop on Agent Theories, Architectures, and Languages (ATAL-2000)*, pages 141–156, Boston, 2000.

[10] F. Dignum, B. Dunin-Kęplicz, and R. Verbrugge. Creating collective intention through dialogue. *Logic Journal of the IGPL*, 9(2):305–319, 2001.

[11] P. Faratin. *Automated Service Negotiation Between Autonomous Computational Agents*. PhD thesis, Queen Mary and Westfield College, University of London, UK, 2000.

[12] A. Finkelstein and H. Fuks. Multi-party specification. In *Proc. 5th Intern. Workshop on Software Specification and Design*, Pittsburgh, PA, USA, 1989. ACM Sigsoft Engineering Notes.

[13] FIPA. Specification part 2 – agent communication language. Technical report, Foundation for Intelligent Physical Agents, 23 October 1997.

[14] J. Forester. *The Deliberative Practitioner: Encouraging Participatory Planning Processes*. MIT Press, Cambridge, MA, USA, 1999.

[15] T. F. Gordon. The Pleadings Game: An exercise in computational dialectics. *Artificial Intelligence and Law*, 2:239–292, 1994.

[16] T. F. Gordon and N. Karacapilidis. The Zeno argumentation framework. In *Proc. 6th Intern. Conf. on AI and Law*, pages 10–18, New York, NY, USA, 1997. ACM Press.

[17] E. Goubault. Geometry and concurrency: a user's guide. *Mathematical Structures in Computer Science*, 10(4):411–425, 2000.

[18] B. J. Grosz and C. L. Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204, 1986.

[19] F. Guerin and J. Pitt. Denotational semantics for agent communication languages. In J. P. Müller, E. Andre, S. Sen, and C. Frasson, editors, *Proc. 5th Intern. Conf. on Autonomous Agents*, pages 497–504, New York, NY, USA, 2001. ACM Press.

[20] C. L. Hamblin. *Fallacies*. Methuen, London, UK, 1970.

[21] D. Hitchcock. Some principles of rational mutual inquiry. In F. van Eemeren et al., editor, *Proc. 2nd Intern. Conf. on Argumentation (ISSA-91)*, pages 236–243, Amsterdam, 1991. SICSAT.

[22] D. Hitchcock, P. McBurney, and S. Parsons. A framework for deliberation dialogues. In H. V. Hansen, C. W. Tindale, J. A. Blair, and R. H. Johnson, editors, *Proc. 4th Biennial Conf. Ontario Soc. Study of Argumentation (OSSA-2001)*, Windsor, Ontario, Canada, 2001.

[23] J. Hulstijn. *Dialogue Models for Inquiry and Transaction*. PhD thesis, Universiteit Twente, Enschede, The Netherlands, 2000.

[24] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. Wooldridge, and C. Sierra. Automated negotiation: prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.

[25] R. Johnson. *Manifest Rationality: A Pragmatic Theory of Argument*. Lawrence Erlbaum Associates, Mahwah, NJ, USA, 2000.

[26] D. V. Lindley. *Making Decisions*. John Wiley and Sons, London, UK, second edition, 1985.

[27] P. Lorenzen and K. Lorenz. *Dialogische Logik*. Wissenschaftliche Buchgesellschaft, Darmstadt, Germany, 1978.

[28] J. D. MacKenzie. Question-begging in non-cumulative systems. *Journal of Philosophical Logic*, 8:117–133, 1979.

[29] J. Mayfield, Y. Labrou, and T. Finin. Evaluating KQML as an agent communication language. In M. J. Wooldridge, J. P. Müller, and M. Tambe, editors, *Intelligent Agents II*, LNAI 1039, pages 347–360, Berlin, Germany, 1996. Springer.

[30] P. McBurney and S. Parsons. Agent ludens: games for agent dialogues. In S. Parsons and P. Gmytrasiewicz, editors, *Game-Theoretic and Decision-Theoretic Agents: Proc. 2001 AAAI Spring Symposium*, pages 70–77, Menlo Park, 2001. AAAI Press.

[31] P. McBurney and S. Parsons. Desiderata for dialogue-game protocols for agent interactions. In A. Hunter and S. Parsons, editors, *Proc. Workshop on Adventures in Argumentation, ECSQARU-2001*, Toulouse, France, 2001. ECSQARU.

[32] P. McBurney and S. Parsons. A geometric semantics for dialogue-game protocols for autonomous agent interactions. In M. Raussen, editor, *Third Workshop on Geometric and Topological Methods in Concurrency (GETCO-01)*, Aalborg, Denmark, 2001. Basic Research in Computer Science (BRICS).

[33] P. McBurney and S. Parsons. Intelligent systems to support deliberative democracy in environmental regulation. *Information and Communications Technology Law*, 10(1):33–43, 2001.

[34] P. McBurney and S. Parsons. Representing epistemic uncertainty by means of dialectical argumentation. *Annals of Mathematics and Artificial Intelligence*, (in press), 2001.

[35] P. McBurney and S. Parsons. Chance discovery using dialectical argumentation. In Y. Ohsawa, editor, *Selected Proc. Intern. Workshops of 15th Annual Conf. AI Society of Japan (JSAI-2001)*, Advanced Information Processing, Berlin, Germany, 2001 (In press). Springer.

[36] P. McBurney, R. van Eijk, S. Parsons, and L. Amgoud. A dialogue-game protocol for agent purchase negotiations. Report, Department of Computer Science, University of Liverpool, U.K., 2001.

[37] P. Mittelstaedt. *Quantum Logic*. D. Reidel, Dordrecht, The Netherlands, 1979.

[38] P. Noriega and C. Sierra. Towards layered dialogical agents. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III: Proc. 3rd Workshop on Agent Theories, Architectures, and Languages*, LNAI 1193, pages 173–188, Berlin, Germany, 1997. Springer.

[39] S. Parsons, C. Sierra, and N. R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.

[40] H. Prakken and T. F. Gordon. Rules of order for electronic group decision making — a formalization methodology. In J. A. Padget, editor, *Collaboration between Human and Artificial Societies: Coordination and Agent-Based Distributed Computing*, LNAI 1624, Berlin, Germany, 1999. Springer.

[41] C. Reed. Dialogue frames in agent communications. In Y. Demazeau, editor, *Proc. 3rd Intern. Conf. Multi-Agent Systems (ICMAS-98)*, pages 246–253. IEEE Press, 1998.

[42] W. Rehg, P. McBurney, and S. Parsons. Computer decision-support systems for public argumentation: Criteria for assessment. In H. V. Hansen, C. W. Tindale, J. A. Blair, and R. H. Johnson, editors, *Proc. 4th Biennial Conf. Ontario Soc. Study of Argumentation (OSSA-2001)*, Windsor, Ontario, Canada, 2001.

[43] O. Renn, T. Webler, and P. Wiedemann, editors. *Fairness and Competence in Citizen Participation: Evaluating Models for Environmental Discourse*. Kluwer, Dordrecht, The Netherlands, 1995.

[44] H. M. Robert. *Robert's Rules of Order. The Standard Guide to Parliamentary Procedure*. Bantam Books, New York, NY, USA, 1986.

[45] J. H. Roberts and G. L. Lilien. Explanatory and predictive models of consumer behavior. In J. Eliashberg and G. L. Lilien, editors, *Handbooks in Operations Research and Management Science: Volume 5: Marketing*, pages 27–82. North-Holland, Amsterdam, The Netherlands, 1993.

[46] F. Sadri, F. Toni, and P. Torroni. Logic agents, dialogues and negotiation: an abductive approach. In M. Schroeder and K. Stathis, editors, *Proc. Symp. Information Agents for E-Commerce, AISB-2001*, York, UK, 2001. AISB.

[47] T. W. Sandholm. Distributed rational decision making. In G. Weiss, editor, *Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence*, pages 201–258. MIT Press, Cambridge, MA, USA, 1999.

[48] M. P. Singh. A social semantics for agent communications languages. In F. Dignum, B. Chaib-draa, and H. Weigand, editors, *Workshop on Agent Communication Languages, IJCAI-99*, Berlin, Germany, 2000. Springer.

[49] D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany, 1995.

[50] H. Wohlrapp. A new light on non-deductive argumentation schemes. *Argumentation*, 12:341–350, 1998.

[51] M. Wooldridge. Semantic issues in the verification of agent communication languages. *Journal of Autonomous Agents and Multi-Agent Systems*, 3(1):9–31, 2000.