Strategic Argumentation in Rigorous Persuasion Dialogue

Joseph Devereux and Chris Reed

School of Computing, University of Dundee, Dundee, DD1 4HN, UK
{josephdevereux,chris}@computing.dundee.ac.uk

Abstract. Philosophical dialogue games have been used widely as models for protocols in multi-agent systems to improve flexibility, expressiveness, robustness and efficiency. Many dialogue games, however, are effectively based on propositional logic, which is not always sufficiently expressive for artificial reasoning. In particular, they do not allow for a strong connection between computational models of dialogic argument and mature mathematical models of abstract argument structures, which support a range of sophisticated agent reasoning systems. In this paper we describe how an existing dialogue game — Walton & Krabbe's RPD_0 — may be extended and improved by using Dung Argumentation Frameworks in place of propositional logic. We call this new dialogue game RPD_{GD}, and describe some of its advantages over RPD_0, chiefly (i) that it allows the proponent to win by exploiting not just defects in the opponent's reasoning or inconsistency in its knowledge base, but also the incompleteness of its knowledge; and (ii) that it thus provides much wider scope for strategic sophistication. Specifically, we present a technique for dialogic strategy based on two key observations: first, that there are minimal amounts of information that one agent requires in order to successfully persuade (or know that it cannot persuade) another; and second, that specific structural manipulation of argumentation frameworks can be engineered to yield that persuasive effect if it exists. The former we present as the concept *Minimum Sufficient Contextual Knowledge*, *MSCK*, and the latter we define as a process of *fortification*. By determining *MSCK* and then developing an appropriate set of fortifications, agents can achieve strategic advantage in an open, heterogeneous system. In even very simple situations, we demonstrate that such strategy can mean the difference between winning and losing a given encounter.

1 Introduction

Dialogue games have been attracting more and more interest in multi-agent systems during recent years [1], as the philosophical questions which motivated their development overlap broadly with practical questions which arise in inter-agent communication. Dialogue games are concerned generally with enforcing standards of dialogic argumentation, and are typically defined in terms of unambiguous rules to which a player's adherence is externally verifiable. The standards in some cases refer merely to logical coherence, but there are also games with more restrictive standards, designed to exclude dialogic behaviour that is generally unhelpful (but not necessarily illogical), or tailored to a particular sort of goal-directed dialogue (such as persuasion dialogue). Thus dialogue games are a promising approach to the regulation of interactive behaviour, either in general or in narrower goal-oriented contexts.

Here we adapt Walton & Krabbe's *RPD₀* [2] to define the game RPD_{GD}. Whereas RPD_0 is based on propositional logic, RPD_{GD} is based on Dung Argumentation Frameworks (*DAF*s) [3], and one of our main concerns is to show how *DAF*s allow much wider scope for strategic sophistication relative to propositional logic. We regard a *strategy* as a player's approach to making decisions required for gameplay. *Strategic sophistication* exists (variously) in all strategies except the simple strategy, whereby all decisions are made by random choice.

We proceed (i) to summarise RPD_0; (ii) to describe RPD_{GD} and its advantages over RPD_0 in general; (iii) to demonstrate in particular the greater scope offered by RPD_{GD} for strategic sophistication, (iv) to compare related work; and (v) to emphasise some general conclusions.

2 RPD_0

RPD_0 is a protocol for *Rigorous Persuasion Dialogue*, intended to capture dialogues in which relevance is enforced by close restrictions on roles, commitment-withdrawal and choice of moves. Such dialogues involve a *proponent* and an *opponent* (of a proposal), and their rigour consists mainly in restrictions imposed on the opponent, whose role is purely responsive. The opponent has simply to challenge or concede the proponent's assertions, incurring irrevocable commitments as it does so. The proponent, while also restricted in its moves, always has the initiative, and rigorous persuasion dialogue is essentially a matter of the proponent extracting commitments from the opponent. The opponent's commitments are very important, because the opponent may lose not only by conceding the proposal, but also by lapsing into inconsistency — by challenging (or conceding the complement of) a proposition to which it is already committed. Because RPD_0 is based on propositional logic, inconsistency implies everything, including the proposal.

The opponent is committed to whatever it concedes, and may also incur commitment through its challenges. If the opponent challenges only part of the proposal, the proponent may ascertain exactly where disagreement lies, and perhaps thereby induce a concession. Alternatively the proponent may question the opponent on its commitments, and thereby induce challenges or concessions. Otherwise the proponent may quiz the opponent on altogether new theses. Such *free questions* must be accompanied by assertions of the corresponding theses, and thus must be either conceded or challenged by the opponent. Only a limited number of free questions may be asked.

The limitation of free questions encourages strategic sophistication, because when the quota is exhausted, the proponent has only the opponent's commitments and challenges to work with, and thus scope for further dialogue is constrained. Beyond this, however, RPD_0 offers very little scope for strategic sophistication, because it is founded on propositional logic. Thus for each of the proponent's assertions, the opponent may determine how the assertion relates to its beliefs, and thereby avoid challenging the assertion (if it is implied) or conceding it (if it is inconsistent). Such consistency-enforcement is uncommon in the human dialogues for which RPD_0 was designed, but could be the norm in a computational society. Furthermore, even if the opponent could not perfectly manage its knowledge base (KB), the proponent's scope for strategic sophistication would not be very much greater. The proponent might then follow strategies to *quicken* the process of either (i) finding existing inconsistency in the opponent's KB or (ii) 'tricking' the opponent into inconsistency by exploiting its defective reasoning. However, without limitation on the number of free questions, such strategies would affect merely the timing of the result. Though this could, in realistic implementations, alone justify the development of strategic reasoning; we are here interested in a broader strategic purview.

3 $\mathbf{RPD_{GD}}$

One way of amending RPD_0 to enhance the scope for strategic sophistication is to replace the underlying logic with a nonmonotonic formalism. RPD_{GD} shows how this could be done, as its underlying logic is provided by Dung Argumentation Frameworks (*DAF*s) interpreted according to the grounded semantics [3]. Dung's is one of the best-known argument-based nonmonotonic formalisms, while grounded semantics is one of the simplest of its semantics.

A *DAF* is a pair $\langle AR, attacks \rangle$, where AR is a set of abstract *arguments*, and *attacks* is a binary relation on AR defining attacks between arguments. Arguments may thus form chains linked by the attacks relation, and for any such chain $a_n \rightharpoonup a_{n-1} \ldots a_0$, odd-numbered and even-numbered arguments *indirectly attack* and *indirectly defend* a_0 respectively.

*DAF*s may be used as simple KBs for agents, and the grounded semantics is a simple means of interpreting such KBs to identify beliefs. The grounded semantics divides a *DAF*'s arguments into the *grounded extension* and the rest. The grounded extension is defined by a recursive function, whereby the *DAF* is searched repeatedly for arguments whose attackers are all attacked by arguments identified in previous searches. Thus first all unattacked arguments are found; then all arguments defended against all attackers by those unattacked arguments; and so on until the search which finds no new arguments. Thus every *DAF* has one and no more than one grounded extension, and an agent's beliefs may unambiguously be defined as the grounded extension of its KB.

RPD_{GD} aims to remain as close as possible to the specification of RPD_0 whilst making full use of the underlying *DAF*-based logic. The main differences are necessary to accommodate the new logic, or have been included to take advantage of obvious new possibilities which it offers. The most important differences relate to negation, implication and questions. Negation is excluded as useless, as the arguments of a *DAF* are abstract entities, defined purely by the attacks relation.

Implication does not exist in *DAF*s either, but let us identify the *indirect defence* relation as a useful, partial equivalent — with the difference that indirect defence can be interpreted as a type of weak, defeasible, contextual implication, in contrast to the guaranteed universal implication of propositional logic. Questions are more diverse in $\mathrm{RPD_{GD}}$, as it includes the new *inquiry* locution-type, with which the proponent may simply inquire about the opponent's KB, without either referring to its commitments or simultaneously asserting the content of the question (cf. *bound* and *free* questions in $\mathrm{RPD_0}$). In $\mathrm{RPD_0}$ such inquiries would be less useful, because information acquired thereby could help the proponent only if the opponent's reasoning was defective, as will be demonstrated.

Let $\mathcal{L}$ be a simple *DAF*-language corresponding to the propositional-logical-language used in $\mathrm{RPD_0}$. $\mathcal{L}$'s atomic sentences are arguments, and conjunctions ($S \wedge S'$) and (inclusive) disjunctions ($S \vee S'$) of sentences are also sentences. Corresponding to the implication-sentences in $\mathrm{RPD_0}$ are the defence-implication sentences of the form $S \Rightarrow S'$, where '$\Rightarrow$' indicates 'indirectly defends'. Finally there are the attack sentences of the form $S \rightharpoonup S'$, where '$\rightharpoonup$' indicates 'attacks'.

$\mathrm{RPD_{GD}}$ may now be specified after the model of $\mathrm{RPD_0}$ (cf. [2, pp158-161]).

Locution Rules

1. Permitted Locutions are of the following types:
   (a) *Statement*: *A!* (A statement is either a ground sentence of $\mathcal{L}$ or $\emptyset$. The proponent only *asserts* statements (makes *assertions*), while the opponent only *concedes* statements (makes *concessions*). The $\emptyset$-statement expresses ignorance and may be used only by the opponent. A concession is made in response to the proponent, but the conceded sentence need not have been asserted by the proponent.)
   (b) *Challenge*: *A??* (A challenge to an assertion *in toto* — i.e. if A is a conjunction, a challenge to every conjoined sentence.)
   (c) *Challenge*: $(A \wedge B)??A?$ (A challenge to one half of a conjunction.)
   (d) *Challenge*: $(A \wedge B)??B?$ (A challenge to the other half of a conjunction.)
   (e) *Questions — Bound Questions*:
   Bound questions are bound in the sense that they refer to commitments of the opponent.
      i. $(A \wedge B)?A?$ (A question referring to one half of a conjunction.)
      ii. $(A \wedge B)?B?$ (A question referring to the other half of a conjunction.)
      iii. $(A \vee B)?$ (A question referring to a disjunction.)
      iv. $(A \Rightarrow B)?A!$ (A question referring to a defence-implication and accompanied by an assertion, with the assertion intended to elicit a concession or a challenge.)
   (f) *Questions — Free Questions*: $A?, A!$ (A question accompanied by an assertion; the opponent must either challenge or concede the assertion.)
   (g) *Questions — Inquiries* — *A???* (An unaccompanied question. A may be a ground or non-ground sentence in $\mathcal{L}$.)
   (h) *Final Remarks*
      i. *I give up!*

ii. *You said so yourself!*
iii. *Your position is absurd!*

2. Each move is the utterance of a single locution.

Commitment Rules

1. The proponent and the opponent have commitment stores—henceforth C_{P} and C_{O} respectively — whose members are sentences of $\mathcal{L}$. Every conceded statement (except the $\emptyset$-statement) is added to C_{O} and cannot be removed. Every asserted statement is added to C_{P} and is removed when challenged (regardless of whether the challenge was well-founded). Given the other rules, this means that C_{P} cannot contain more than one member.
2. Each challenge refers to the sole member of C_{P}.
3. Each bound question refers to an element of C_{O}.
4. If the opponent expresses ignorance ($\emptyset$) in response to an inquiry of the form $inquire(X \rightharpoonup a)$, a is added to C_{O}.

Structural Rules

1. The first move is a challenge of the proponent's proposal.
2. The proponent and the opponent move alternately.
3. In each of its moves, the proponent must either (i) defend a challenged assertion legally by the rules for challenge and defence; or (ii) question a member of C_{O} legally by the rules for question and answer; or (iii) ask a free question; or (iv) make an inquiry; or (v) make a final remark.
4. The proponent may not ask any question/inquiry such that one of its permitted answers (permitted by the rules for question and answer) is in C_{O}.
5. The proponent may not repeat an assertion unless C_{O} has expanded in the meantime.
6. The proponent may make the final remarks (i) '*You said so yourself!*' or (ii) '*Your position is absurd!*' if and only if, respectively (i) the opponent has challenged a member of C_{O}; or (ii) the opponent has conceded a sentence which it has previously challenged.
7. Unless it is '*I give up!*', each of the opponent's moves must refer to the proponent's preceding move. If the proponent's preceding move was an assertion, the opponent must challenge the assertion. If it was a bound question, free question or inquiry, the opponent must answer the question/inquiry legally by the rules for question and answer.
8. The opponent may utter '*I give up!*' in any move.

Win-and-Loss Rules

1. The player which utters '*I give up!*' loses the dialogue, and the other party wins.
2. The proponent wins if it utters (legally by the structural rules) either '*You said so yourself!*' or '*Your position is absurd!*'.

Table 1. Rules for challenge and defence.

Element of C_P challenged	Form of Challenge	Permitted Defence(s)
$A \vee B$	$(A \vee B)??$	$A! \mid B!$
$A \wedge B$	$(A \wedge B)??A?$	$A!$
$A \wedge B$	$(A \wedge B)??B?$	$B!$
A	$A??$	—

Table 2. Rules for question and answer.

Element of C_O questioned	Form of Question	Permitted Answer(s)
$A \vee B$	$(A \vee B)?$	$A! \mid B!$
$A \wedge B$	$(A \wedge B)?A?$	$A!$
$A \wedge B$	$(A \wedge B)?B?$	$B!$
$(A \Rightarrow B)$	$(A \Rightarrow B)?A!$	$B! \mid (C \wedge (C \rightharpoonup A))! \mid (D \wedge (D \rightharpoonup B))!$
—	$A?,A!$	$A! \mid A??$
—	$A???$ (A is ground)	$\emptyset! \mid A!$
—	$A???$ (A is non-ground)	$\emptyset! \mid A_i! \mid (A_1 \wedge \ldots \wedge A_n)!$ (each A_i is a ground instance of A)

3.1 Gameplay in $\mathrm{RPD_{SD}}$

Fundamental differences between $\mathrm{RPD_{GD}}$ and $\mathrm{RPD_0}$ with respect to gameplay correspond to differences between the underlying logics —

1. *DAF*s do not have *PL*'s *PL*-implication; hence not all of the commitment-inducing mechanisms of $\mathrm{RPD_0}$ are in $\mathrm{RPD_{GD}}$.
2. *DAF*s' weak, defensive implication is not in *PL*; correspondingly, $\mathrm{RPD_{GD}}$'s weak commitment-inducing mechanism has no exact equivalent in $\mathrm{RPD_0}$.
3. *DAF*s' non-monotonicity vs. *PL*'s monotonicity; correspondingly, $\mathrm{RPD_{GD}}$ has an *inquiry* locution-type (whereby the proponent may avoid blundering), whereas no such locution-type exists in $\mathrm{RPD_0}$, and would not be similarly useful if it did.

$\mathrm{RPD_{GD}}$ differs from $\mathrm{RPD_0}$ also in allowing the proponent unlimited free questions, but this is not essential. Limiting free questions would create extra scope for strategic sophistication, just as it does in $\mathrm{RPD_0}$. We allow unlimited free questions to highlight the extra possibilities for strategic sophistication created by the use of *DAF*s instead of *PL*. $\mathrm{RPD_{GD}}$ is consequently very asymmetric — not only does the proponent always have the initiative, but it is impossible for the proponent to blunder in such a way as to ensure that it loses when it 'should' have won — when the proposal is in the grounded extension of the union of the players' KBs.

Besides the particular differences 1–3 between gameplay under the two protocols, there are also two broader differences. The first is that, if the opponent is honest and open-minded, dialogues under $\mathrm{RPD_{GD}}$ are less likely to fall into irresolvable disagreement. If an agent has a propositional-logical KB, consistencymaintenance may demand that the agent unwaveringly rejects an assertion a which conflicts with a current belief $\neg a$, regardless of supporting arguments. Whereas a KB expressed by a DAF is much more open, because an argument may move in and out of acceptability as the KB expands. With an honest and open-minded opponent, a dialogue under RPD_{GD} would fall into irresolvable disagreement only when the proponent had uttered every argument in its KB attacking the opponent's position, and heard for every such argument a counter-argument independent of the argument(s) attacked.

The other broad difference between gameplay under the two protocols is that RPD_{GD} offers a greater variety of strategies to the proponent. Under RPD_0 the proponent might induce the opponent to expand its KB, and would never risk anything (except wasting free questions) in doing so. Whereas under RPD_{GD} such expansion may be riskier, as the proponent may thereby destroy some or all of its opportunities for victory. Such risks would exist where the proponent had only incomplete knowledge of the opponent's KB, and the opponent could construct new ways of defending its position with arguments and attacks asserted by the proponent. The point is important, because in RPD_{GD} expanding the opponent's KB might be an effective way (the only effective way, with a rational and perfectly-reasoning opponent) for the proponent to win.

Consider the example described by Fig. 1 and Table 3. The figure shows

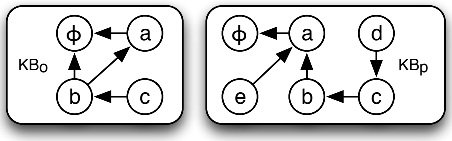

Fig. 1. The KBs of a game in which the proponent blunders.

the players' initial KBs, and the table shows their moves in the subsequent dialogue. The proponent might have won by citing $(e \wedge (e \rightharpoonup a))$, but instead cited $(b \wedge (b \rightharpoonup a))$ and the subsequent 'chain' of arguments. Thereby it caused the opponent to believe b, which in its KB attacks ϕ. The proponent could not subsequently attack b, and in the end abandoned the dialogue, with good reason — in whatever way the proponent had continued the dialogue, the opponent could have open-mindedly and honestly avoided conceding ϕ.

Thus the example suggests that a proponent may do better by following a particular strategy in expanding the opponent's KB, than by expanding it at random. The example also shows how such a strategy might be informed by the use of *inquiry* locutions. If the proponent had inquired into what attacked and defended a in KB_O, it would have discovered that its case for a must not reinstate b in KB_O.

Such reinstatement is possible only because the underlying logic is nonmonotonic, and is thus precluded in dialogues under RPD_0 (with its underlying PL),

Table 3. The moves of a game in which the proponent blunders.

Move	Opponent	Proponent
1	$challenge(\phi)$??	—
2	—	$question(b\ \wedge\ (b \rightharpoonup a))?,\ assert(b\ \wedge\ (b \rightharpoonup a))!$
3	$concede(b\ \wedge\ (b \rightharpoonup a))!$	—
4	—	$question(b \Rightarrow \phi)?,\ assert(b \Rightarrow \phi)!$
5	$concede(b \Rightarrow \phi)!$	—
6	—	$question(b \Rightarrow \phi)?\phi!$
7	$concede(c\ \wedge (c \rightharpoonup b))!$	—
8	—	$question(d\ \wedge\ (d \rightharpoonup c))?,\ assert(d\ \wedge\ (d \rightharpoonup c))!$
9	$concede(d\ \wedge\ (d \rightharpoonup c))!$	—
10	—	$question(d \Rightarrow \phi)?,\ assert(d \Rightarrow \phi)!$
11	$concede(d \Rightarrow \phi)!$	—
12	—	$question(d \Rightarrow \phi)?\phi!$
13	$concede(b\ \wedge (b \rightharpoonup \phi))!$	—
14	—	*I give up!*

except where the opponent's reasoning is defective. It follows that no correspondingly useful *inquiry* locution-type could be defined to extend RPD_0 — any instance of such a locution-type could be useful only if it was informed by insight into such defects, or (more indirectly) if it served to reveal such defects.

4 Strategy in RPD_{GD}

RPD_{GD} provides scope for three sorts of strategic sophistication. Moves may be chosen according to their perceived value in —

1. revealing existing inconsistency between the opponent's commitments and its beliefs;
2. exploiting defects in the opponent's reasoning, so as to create such inconsistency;
3. revealing blunders to be avoided in expanding the opponent's KB.

The strategic possibilities with respect to (1) and (2) correspond closely (but not exactly) to strategic possibilities offered by RPD_0, and will not be considered. We will concentrate instead on (3), presenting two concepts — *MSCK* and *fortification* — which are useful when considering how P should formulate its strategy. Henceforth we will refer to a proponent P with proposal ϕ, KB_{P} ($= \langle AR_{\text{P}}, attacks_{\text{P}} \rangle$) and $beliefs_{\text{P}}$, and an opponent O with KB_{O} ($= \langle AR_{\text{O}}, attacks_{\text{O}} \rangle$) and $beliefs_{\text{O}}$.

4.1 Minimum Sufficient Contextual Knowledge

Let us first consider what is most worth knowing about KB_{O} with respect to ϕ. Complete knowledge of $attacks_{\text{O}}$ includes complete knowledge of how ϕ is attacked and defended in KB_{O}. Thus it is sufficient to allow P to determine if and how it could devise a persuasive justification of ϕ. Everything in KB_{O} which remained hidden — those isolated arguments in KB_{O}, neither attackers nor attacked — would be useless[1]. Though complete knowledge of $attacks_{\mathrm{O}}$ is thus sufficient, it is not necessary — for instance, an attack in KB_{O} by ϕ against a non-attacking argument absent from KB_{P} would be irrelevant to P's task. Thus, while this first insight may be useful, it cannot be the only such insight.

Let us now take a more systematic approach. The value of the intuitive insight is that it excludes certain qualities of KB_{O} — those arguments which are both unattacked and unattacking, and the beliefs in $beliefs_{\mathrm{O}}$ corresponding to them — as irrelevant to P's task. All other arguments and beliefs are (respectively) included in and derivable from $attacks_{\mathrm{O}}$. More generally, what is relevant and irrelevant about KB_{O} in a persuasion dialogue depends on (i) KB_{O}, (ii) KB_{P}, and (iii) ϕ. Let us specify the range of qualities of any DAF-based KB, to include every conceivably relevant attribute, and to exclude every other attribute.

Definition 1. q *is a* quality *of a* DAF*-based KB* KB_X *iff (i) everything that* q *refers to is either (a) the presence in or absence from* KB_X *of an argument/attack or set of arguments/attacks; or (b) the presence in or absence from* $\mathrm{beliefs}_X$ *of a belief/set of beliefs; and (ii)* q *refers to at least one item described by (a) or (b).* $\mathrm{KB}_X^{\mathrm{Q}}$ *denotes the set of all qualities of* KB_X*.*

Thus $KB_{\mathrm{O}}^{\mathrm{Q}}$ collects expressions of various reference: to each $a \in AR_{\mathrm{O}}$ and each $att \in attacks_{\mathrm{O}}$; to each $b \in beliefs_{\mathrm{O}}$; to whether or not there exist arguments/attacks/beliefs fulfilling certain criteria in $KB_{\mathrm{O}}/beliefs_{\mathrm{O}}$; to which (if any) do so; and so on. It collects also all the converse qualities, expressing that $a' \notin AR_{\mathrm{O}}$; that $att' \notin attacks_{\mathrm{O}}$; and so on.

Let us now specify *Minimum Sufficient Contextual Knowledge* (*MSCK*) as follows.

Definition 2. *Let* $\mathrm{S} \subset \mathrm{KB}_{\mathrm{O}}^{\mathrm{Q}}$ *and let* E *be the set of all arguments (whether actually in* KB_O *or not)* explicitly *referred to in* S. S *is an instance of* MSCK *wrt* $\langle \mathrm{KB}_P, \mathrm{KB}_O, \phi \rangle$ *(more briefly,* ${}^{\mathrm{P}}_{\mathrm{O}}\mathrm{MSCK}_{\phi}$*) iff*

1. *knowledge of S permits* P *to know whether and how it can induce* O *to accept* ϕ*; and*
2. *there is no* S_2 *such that* $\mathrm{S}_2 \subset S$ *and* S_2 *is an instance of* ${}^{\mathrm{P}}_{\mathrm{O}}\mathrm{MSCK}_{\phi}$*; and*
3. *there is no other* $\mathrm{S}' \subset \mathrm{KB}_O^{\mathrm{Q}}$*, such that* $\mathrm{E}' \subset \mathrm{E}$*, where* E' *corresponds to* S' *as* E *corresponds to* S*, being the set of all arguments (whether actually in* KB_O *or not)* explicitly *referred to in* S'*.*

Let ${}^{\mathrm{P}}_{\mathrm{O}}\mathrm{msck}_{\phi}^{All} = \{{}^{P}_{O}\mathrm{msck}_{\phi}^{1} \ldots {}^{P}_{O}\mathrm{msck}_{\phi}^{n}\}$ *be the set of all instances of* ${}^{\mathrm{P}}_{\mathrm{O}}\mathrm{MSCK}_{\phi}$*.*

Thus condition (1) ensures sufficiency, and conditions (2) and (3) minimality. The intuition behind (3) is that an instance of *MSCK* should be minimal wrt explicit reference to arguments — as abstract as possible. Consider, for instance:

[1] If ϕ was itself one of these 'islands', it would be in O's beliefs. Thus, if P found that ϕ was not in the (weakly) connected parts of KB_{O}, it could safely conclude that O either believed or was unaware of ϕ; and thus would not challenge it.

$$\begin{aligned}
KB_{\mathrm{O}} &= \langle\{a, b, \phi\}, \{a \rightharpoonup \phi\}\rangle \\
KB_{\mathrm{P}} &= \langle\{a, c, \phi\}, \{c \rightharpoonup a\}\rangle \\
S &= \{attacks_O = \{(a, \phi)\}\} \\
S' &= \{(\{a, b, \phi\} = arguments_{\mathrm{O}}), (a \rightharpoonup \phi), \neg(b \rightharpoonup \phi)\}
\end{aligned}$$

S is an instance of ${}^{\mathrm{P}}_{\mathrm{O}}\mathrm{MSCK}_\phi$, but S' is not.

Counter-intuitively, acquiring an instance of ${}^{\mathrm{P}}_{\mathrm{O}}MSCK_\phi$ might require further knowledge of KB_{O}. But the concept is not thereby invalidated — *MSCK* is wholly separate from minimal knowledge required to acquire *MSCK*, and whether the latter is included in every instance of the former (as would be ideal) depends on the information-gathering methods available to P.

We are here interested in *MSCK* itself, and especially those qualities of KB_{O} which are outside *every* ${}^{\mathrm{P}}_{\mathrm{O}}msck^i_\phi$. *MSCK* as currently defined does not provide any guidance on the matter. A more sophisticated definition would be costlier to consult, but would allow P to recognise areas of KB_{O} and $beliefs_{\mathrm{O}}$ which it need not consider. We may envisage a series of sets $\langle {}^{\mathrm{P}}_{\mathrm{O}}irrelevant^1_\phi \ldots {}^{\mathrm{P}}_{\mathrm{O}}irrelevant^n_\phi\rangle$, where each ${}^{\mathrm{P}}_{\mathrm{O}}irrelevant^i_\phi$ collects every element of $KB^{\mathrm{O}}_{\mathrm{O}}$ which fulfils criterion i and is thereby necessarily excluded from every instance of ${}^{\mathrm{P}}_{\mathrm{O}}MSCK_\phi$. By way of example, we identify three such criteria in Definition 3.

Definition 3. *Let* $\mathrm{KB}_\cup$ *be the graph union of* KB_P *and* KB_O*, and let* $\mathrm{paths}^{KB_\cup}_\phi$ *be the set of all acyclic paths in* $\mathrm{KB}_\cup$ *which start in a leaf node and terminate in* ϕ*.* Attacking *and* defending *paths contain respectively odd and even numbers of attacks, and* attacking *and* defending *arguments occur alternately in each path. A path's* elements *are its arguments and attacks. Let* islands*,* paths_{silly} *and* paths_{safe} *be as follows —*

1. *Let* islands *be the set of arguments neither attacking nor attacked in* KB_O*.*
2. *Let* paths_{silly} *be the set containing every* $\mathrm{path} \in \mathrm{paths}^{KB_\cup}_\phi$ *such that (2i) the final attack in* path *(i.e. the attack on* ϕ*) is not in* KB_O*; and (2ii) no element, possibly excepting* ϕ*, of* path *is an element of any other* $\mathrm{path}' \in \mathrm{paths}^{KB_\cup}_\phi$ *not fulfilling condition (2i).*
3. *Let* paths_{safe} *be the set containing every* $\mathrm{path} \in \mathrm{paths}^{KB_\cup}_\phi$ *such that (3i)* path *defends* ϕ*; and (3ii) all elements of* path *are in* KB_O*; and (3iii) no element, possibly excepting* ϕ*, in* path *is an element of any other* $\mathrm{path}' \in \mathrm{paths}^{KB_\cup}_\phi$ *not fulfilling both conditions (3i) and (3ii).*

Thus by recursion, for any path *multiple* others *may be considered when determining whether* $\mathrm{path} \in \mathrm{paths}_{silly}$*, and if any* $\mathrm{o} \in \mathrm{others}$ *does not fulfil condition (2i), neither* path *nor any of* others *is in* paths_{silly}*; otherwise all are. And similarly for* paths_{safe} *and conditions (3i) and (3ii).*

In the artificial example of Figure 2 the vast majority of KB_O is covered by the three concepts, and an instance of ${}^{\mathrm{P}}_{\mathrm{O}}MSCK_\phi$ may refer explicitly to no arguments except o, p and ϕ.

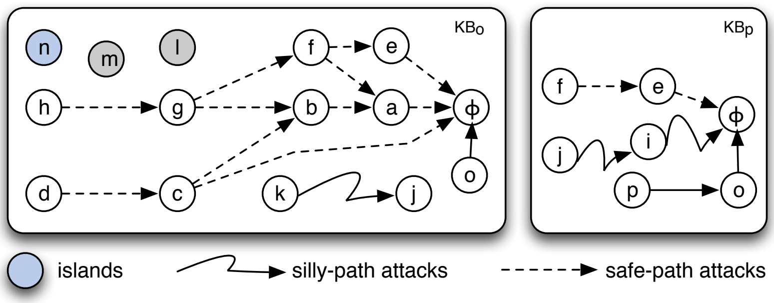

Fig. 2. Illustrating *islands*, $paths_{\text{silly}}$ and $paths_{\text{safe}}$.

Proposition 1. *Let* $\text{qualities}_{islands}$, qualities_{silly} *and* qualities_{safe} *contain all qualities* explicitly *referring to any* $i \in$ islands *and any element of any path in* paths_{silly} *and* paths_{safe}, *respectively. No* $\text{q} \in (\text{qualities}_{islands} \cup \text{qualities}_{silly} \cup \text{qualities}_{safe})$, *such that* q *refers explicitly to anything other than* ϕ, *may be in any instance of* ${}^{\text{P}}_{\text{O}}MSCK_{\phi}$.

Proof. Let $KB^{\text{sub}}_{\text{O}}$ be a weakly connected subgraph of KB_{O}, possibly containing ϕ. P would need to consider $KB^{\text{sub}}_{\text{O}}$ only if either (i) it contained an undefeated attacker of ϕ; or (ii) P could, by making a case for ϕ, inadvertently convert one of $KB^{\text{sub}}_{\text{O}}$'s arguments into an undefeated attacker of ϕ. Suppose (i) did not hold. (ii) would not hold either, if (a) there were no attacks in $KB^{\text{sub}}_{\text{O}}$. But even if there were attacks, there would be no danger if (b) $KB^{\text{sub}}_{\text{O}}$ contained no argument attacking ϕ, and P could never rationally present a case which resulted in the connection of $KB^{\text{sub}}_{\text{O}}$ to ϕ. And even if $KB^{\text{sub}}_{\text{O}}$ did contain an argument attacking ϕ, there would be no danger if (c) every such argument was defeated in $KB^{\text{sub}}_{\text{O}}$, and P was unaware of every argument in $KB^{\text{sub}}_{\text{O}}$, except as it existed in $KB^{\text{sub}}_{\text{O}}$.

Now let q be such that $q \in (\text{qualities}_{islands} \cup \text{qualities}_{silly} \cup \text{qualities}_{safe})$, and q refers explicitly to anything other than ϕ. If $q \in \text{qualities}_{islands}$, q must refer only to subgraphs of KB_{O} fulfilling (a); if $q \in \text{qualities}_{silly}$, q must refer only to subgraphs of KB_{O} fulfilling (b); and if $q \in \text{qualities}_{safe}$, q must refer only to subgraphs of KB_{O} fulfilling (c).

We have just shown that there are areas of KB_{O} irrelevant to ${}^{\text{P}}_{\text{O}}MSCK_{\phi}$, and that our three categories of qualities refer to three such areas. For an alternative proof, let us now, conversely, consider those areas of KB_{O} which are relevant to ${}^{\text{P}}_{\text{O}}MSCK_{\phi}$, and that our three categories of qualities make no reference to them.

Any ${}^{\text{P}}_{\text{O}}msck^{i}_{\phi}$ must include knowledge of every undefeated direct attacker *att* of ϕ in KB_{O}. Thus for every *att*, ${}^{\text{P}}_{\text{O}}msck^{i}_{\phi}$ might include $attacks(att,\phi)$ and also either $belief(att)$, or $(\{X \mid X \text{ is an undefeated attacker of } att\} = \emptyset)$, or some equivalent. Define q as before. It follows trivially from the definition of *islands*, $paths_{\text{silly}}$ and $paths_{\text{safe}}$, that q cannot refer to *att*, and thus that q does not appear at this stage of the construction of ${}^{\text{P}}_{\text{O}}msck^{i}_{\phi}$.

Let the set containing every undefeated direct attacker of ϕ in KB_O be S_{att}. Besides knowing every member of S_{att}, P must also know that S_{att} is complete. This is just a single quality of KB_{O}, and it cannot overlap with q, wrt explicit reference (apart from explict reference to ϕ).

Finally, P must know whether its KB contains a suitable set of arguments Def attacking every $att \in S_{att}$. This is the most complex stage, because any $arg \in Def$ might itself be defeated in KB_{O}; and any $S \subseteq Def$ might reinstate a defeated direct attacker of ϕ in KB_{O}. This introduces a great many complications, but none of them can involve q, because every complication must be to do with a path to ϕ *jointly* constructable by P and O. q cannot refer to any such path, because it refers only to *islands* or path(s) in $paths_{\mathrm{silly}}$ or $paths_{\mathrm{safe}}$. P could not rationally play any part in constructing any path in $paths_{\mathrm{silly}}$; and could not possibly play any part in constructing any path using any element of any $p_s \in paths_{\mathrm{safe}}$, such that the constructed path was not itself already in $paths_{\mathrm{safe}}$.

The concepts *islands*, $paths_{\mathrm{silly}}$ and $paths_{\mathrm{safe}}$ thus provide three criteria for determining whether a quality of KB_{O} is outside every instance of ${}^{\mathrm{P}}_{\mathrm{O}}MSCK_{\phi}$. How to determine all remaining such criteria, and how to compare their relative usefulness, and how to use multiple (perhaps overlapping) criteria efficiently are all questions for future work. For now we postulate that with more criteria, and hence a more sophisticated concept of *MSCK*, P would be able to define more precisely the nature of its task, and thereby become better informed on whether and how it might succeed. Furthermore, P might also be able to reason more efficiently in general about KB_{O}, by excluding from its reasoning all areas of KB_{O} which were necessarily outside any instance of ${}^{\mathrm{P}}_{\mathrm{O}}MSCK_{\phi}$.

4.2 Fortification

Let us now consider practical consequences of different approaches to *MSCK*. We first define *fortification* as a simple mechanism around which P may build a general approach to argumentation, in which more or less attention may be paid to *MSCK*. We then show how an agent which uses fortification might benefit by paying careful attention to *MSCK*.

Consider the subset S of $attacks_{\mathrm{O}}$ containing the attacks most obviously relevant to ϕ — for each $att \in S$, $att = (A \rightharpoonup \phi)$, where A is an undefeated argument in KB_{O}. Removing each $att \in S$ from KB_{O} produces a DAF containing ϕ in its grounded extension. However, KB_{O} is a monotonic KB, and thus $beliefs_{\mathrm{O}}$ can change only through the expansion of KB_{O}. Thus for every undefeated direct attacker a of ϕ in KB_{O}, P would need to know whether it could induce O to believe some b such that $b \rightharpoonup a$ held. To capture this idea, let us define the *fortification* function (FF) for DAFs and arguments as follows —

Definition 4. *Let $\mathcal{D}$ be a set of* DAF*s and $\mathcal{A}$ be a set of arguments.*
$\mathrm{FF} = (\mathcal{D} \times \mathcal{A}) \mapsto \mathcal{D}$

$$\begin{aligned} \textit{where, for every}\ \mathrm{DAF} &= \langle \mathrm{AR}, \mathrm{attacks} \rangle\ \textit{and} \\ \mathrm{DAF}_{\mathit{ffd}} &= \langle \mathrm{AR}_{\mathit{ffd}}, \mathrm{attacks}_{\mathit{ffd}} \rangle = \mathrm{FF}(\mathrm{DAF}, \phi)\ \textit{and} \\ \mathrm{DAF}_{\mathit{ffg}} &= \langle \mathrm{AR}_{\mathit{ffg}}, \mathrm{attacks}_{\mathit{ffg}} \rangle, \end{aligned}$$

1. $\mathrm{AR}_{\mathit{ffd}} = \mathrm{AR} \cup \mathrm{AR}_{\mathit{ffg}}$ *and* $\mathrm{attacks}_{\mathit{ffd}} = \mathrm{attacks} \cup \mathrm{attacks}_{\mathit{ffg}}$.
2. *For each* $\mathrm{attacks(X, Y)} \in \mathrm{attacks}_{\mathit{ffg}}$, Y *is an undefeated direct or indirect attacker of* ϕ *in* DAF.
3. *For each* $\mathrm{X} \in \mathrm{AR}_{\mathit{ffg}}$, *there exists some* $\mathrm{attacks(X, Y)} \in \mathrm{attacks}_{\mathit{ffg}}$.
4. *For each undefeated direct attacker* Z *of* ϕ *in* DAF, *Z is defeated in* $\mathrm{DAF}_{\mathit{ffd}}$.

$\mathrm{DAF}_{\mathit{ffd}}$ *is the* ϕ-fortified *version of* DAF, *while* $\mathrm{DAF}_{\mathit{ffg}}$ *is the* ϕ-fortifying *framework for* DAF.

So DAF_{ffd} is DAF expanded to defeat all direct attackers of ϕ. DAF_{ffg} is not necessarily minimal, but cannot be in any respect irrelevant, as every element must be involved in defending ϕ. *FF* is thus intended to reflect a general, abstract way in which P might expand KB_{O} if it had complete knowledge of ϕ's undefeated direct attackers in KB_{O}.

Suppose that P induced the expansion of KB_{O} into a ϕ-fortified version. Even so, O would not necessarily believe ϕ, as the examples in Figure 3 show. In DAF ϕ has one undefeated attacker — a — and both DAF' and DAF'' are sufficient to defeat a. But in so doing, DAF' converts b into an undefeated attacker of ϕ.

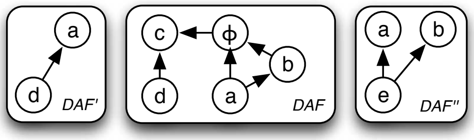

Fig. 3. Illustrating fortification — DAF' and DAF'' are both *ϕ-fortifying* frameworks for DAF, but while ϕ is acceptable in $DAF \cup DAF''$, it is not acceptable in $DAF \cup DAF'$

.

Let us now consider how ϕ could be excluded from the grounded extension of any $DAF_{\mathrm{ffd}} = (DAF \cup DAF_{\mathrm{ffg}})$. *Controversy* would necessarily be involved. An argument is *controversial* with respect to another if it both (directly or indirectly) attacks and indirectly defends that argument [3]. Because DAF_{ffg} defeats all undefeated attackers of ϕ in DAF, there must be controversy in either or both of (i) DAF and (ii) $(DAF \cup DAF_{\mathrm{ffg}})$. However, as there may be controversy not involving ϕ significantly, it is not a sufficient condition for ϕ's exclusion from the grounded extension of DAF_{ffd}. In effect *MSCK* and fortification contribute to a notion of dialectical relevance in the sense of [4].

We finish by returning to the example of Figure 1 and Table 3, and showing how P might have won if it had applied fortification with careful attention to *MSCK*. An alternative set of moves could be as shown in Table 4, where P follows the simple strategy of fortifying (whenever possible) for its proposal and anything which may serve in its defence, until either fortification is accomplished for all such arguments, or it finds something for which it cannot fortify. Thus at Move 2 P prepares to fortify for ϕ by inquiring about direct attackers of ϕ. O

Table 4. The moves of a game in which the proponent avoids blundering through careful use of fortification.

Move	Opponent	Proponent
1	$challenge(\phi)??$	—
2	—	$inquire(X \rightharpoonup \phi)???$
3	$concede((a \rightharpoonup \phi) \wedge (b \rightharpoonup \phi))!$	—
4	—	$inquire((X \rightharpoonup e) \vee (Y \rightharpoonup c))???$
5	$concede(\emptyset)!$	—
6	—	$question(e \wedge (e \rightharpoonup a) \wedge c \wedge (c \rightharpoonup b))?$, $assert(e \wedge (e \rightharpoonup a) \wedge c \wedge (c \rightharpoonup b))!$
7	$concede(e \wedge (e \rightharpoonup a) \wedge c \wedge (c \rightharpoonup b)!$	—
8	—	$question(e \Rightarrow \phi)?$, $assert(e \Rightarrow \phi)!$
9	$concede(e \Rightarrow \phi)!$	—
10	—	$question(e \Rightarrow \phi)?\phi!$
11	*I give up!*	—

mentions a and b (Move 3), thus informing P that fortification for ϕ depends on fortification for a direct attacker for each of a and b. P knows of only one direct attacker for each — e and c respectively — and thus it prepares to fortify for e and c by inquiring about direct attackers of each (Move 4). O mentions none (Move 5), so P can expect that it will succeed — as in fact it does. P's strategy here is not perfect (it unnecessarily argues against b), but it is good enough: it secures victory, unlike the careless, simple strategy illustrated earlier.

Let us finally consider *MSCK* in this case. The knowledge of KB_{O} gained and used by P in the example forms the set

$$\begin{aligned} S = \{\ &(\{X \mid attacks(X,\phi)\} = \{a,b\}), (\{X \mid attacks(X,e)\} = \emptyset), \\ &(\{X \mid attacks(X,c)\} = \emptyset)\}. \end{aligned}$$

S is not an instance of ${}^{\mathrm{P}}_{\mathrm{O}}MSCK_{\phi}$, as it unnecessarily refers explicitly to b and c. P's approach thus ensures only sufficiency; but had KB_O been different, it might have yielded an instance of ${}^{\mathrm{P}}_{\mathrm{O}}MSCK_{\phi}$. As things are, two actual instances of ${}^{\mathrm{P}}_{\mathrm{O}}MSCK_{\phi}$ are S' and S'' as follows —

$$\begin{aligned} S' &= \{\ (\{X \mid attacks(X,\phi),\ belief(X)\} = a), \\ &\quad\ (\{attacks(a,Y) \mid Y \neq \phi\} = \emptyset), \neg(e)\}. \\ S'' &= \{\ (\{X \mid attacks(X,\phi),\ belief(X)\} = a), (\{attacks(a,Y) \mid Y \neq \phi\} = \emptyset), \\ &\quad\ (\{attacks(Y,Z) \mid Z = e\} = \emptyset)\}. \end{aligned}$$

5 Related Work

RPD_0 has been little used in multi-agent systems, though its authors have otherwise been widely influential in the field[2]. It is used in [5, 6], but in neither case

[2] It is not clear why this should be, as some of RPD_0's main motivating problems — those of relevance and enforcing it in human dialogue — may arise in multi-agent dialogue too.

is much modified or very prominent. Therefore in this section we focus on work on DAFs and strategy in computational dialogue games.

Amgoud & Maudet [7] consider strategy using a DAF-based framework which incorporates propositional logic and preference orderings over arguments, and a dialogue game based on Mackenzie's DC [8]. Among the strategic matters considered is choice of argumentative content, but as their choice-process involves only the choosing agent's KB — and not the other agent's — their work overlaps little with ours.

In later work with Hameurlain [9], Amgoud considers dialogue strategy differently, as a bipartite decision problem — for each move, an agent must decide (i) which type of locution to use, and (ii) what content to use in the locution. These decisions are influenced respectively by (i) strategic goals and strategic beliefs and (ii) functional goals and basic beliefs, where strategic elements are purely to do with dialogic practice, while functional/basic elements are to do with everything else. Thus which locution-type to use is a strategic matter, while what content to use is a functional/basic matter. They use a framework for argumentation-based decision-making under uncertainty to show how this decision problem may be tackled, which takes into account the weights of beliefs and the priorities of goals.

Amgoud & Hameurlain's approach is relevant to the current paper, as the strategic vs. functional/basic distinction may be useful for strategy in dialogue games, where moves are typically defined by locution-type and locution content. However, they do not go very far beyond the definition of the decision problem and the high-level reasoning framework — they do not, for instance, illustrate how an agent could take into account the dialogue-history at either the strategic or functional/basic levels. Thus their work is currently of relevance only as providing a wider context for the ideas in this paper. Kakas et al. [10] present an alternative framework, which may also accommodate our ideas, especially as it accounts for both protocol and strategy.

Oren et al. [11] have looked at the question of confidentiality as a strategic concern in argumentation. This is relevant to $\mathrm{RPD_{GD}}$ — for example, in the dialogue described in Table 4, the proponent reveals more than is necessary, and in some contexts such superfluity may be unacceptable. An intuitive general approach would be to make fewer assertions and to ask more questions, but questions might reveal information too.

Black & Hunter [12] have considered strategy in enthymeme inquiry dialogue. Like ours, their approach uses dialogue games, and a player's strategy determines each of its moves from the range permitted by the protocol. However, in contrast to the adversarial, asymmetrical protocol presented here, their protocol is cooperative and symmetrical, with players aiming to jointly construct a (possibly enthymematic) argument acceptable to them both. In addition, their strategy is to be used by both players, and invariably leads to success.

Bentahar et al. [13] present a model for adversarial persuasion dialogue in which both players may attempt to persuade the other, and in which the moves of each player are almost fully specified. Their protocol is novel in several respects, but like ours is draws on dialogue games, being defined in terms of commitments, entry and exit conditions, and a dynamics determining the structure of dialogues. However, it differs from ours in leaving very little room for strategy. Agents have no choice regarding the locution-type used in their moves, except when required to decide whether to accept a commitment for which the other player has no justifying argument. On such occasions a player may *accept* or *refuse*, and chooses by consulting a complex trustworthiness model. It is on this aspect of strategy that the authors concentrate — no attention is paid to how, for instance, a player should choose which argument to use to attack the other player's previous commitment.

Dunne & McBurney [14] consider 'optimal utterances' in dialogue, but the dialogues they consider occur in fixed dialogue contexts, which are completely known to both players. Thus in the scope of the dialogue the players know one another's KBs, which contrasts with our scenario. Furthermore, the optimality they consider relates to dialogue length, rather than to whether one player persuades another.

We may finally compare our work with game theoretical approaches. Procaccia & Rosenschein [15], Riveret et al. [16] and Roth et al. [17] consider games with complete information, and thus their results are of limited relevance to this paper. However, Rahwan & Larson [18] use wholly ignorant agents in considering how their game-theoretical mechanism for DAF-based sceptical (grounded) argumentation may be strategy-proof. They prove that the the mechanism is strategy-proof only under the fairly restrictive condition that each player's set of arguments is free of direct/indirect conflict in the graph formed by the union of all players' arguments — otherwise an agent might benefit from revealing only a subset of its arguments. This last point is reflected in the examples of Figure 1 and Tables 3 and 4 — the proponent fails when it is careless about revealing its DAF, and succeeds when it takes care.

The game theoretical approach has thus produced interesting results. Game theoretical approaches tend, however, to have a different focus — namely on designing protocols and mechanisms for communication, and determining whether or not these protocols are strategy-proof [18]. The approach rests on an assumption that the best way to tackle the challenges of rich communication is to design closed systems that require adherence to highly restrictive protocols which preclude strategy in communication. Our approach, however, in contrast, is to maintain the flexibility, expressiveness, and sophistication of protocols that are not strategy-proof, and then invest in the reasoning required to construct communicative strategies that can be robust and effective in such an open environment.

6 Conclusions and Future Work

The work described here forms a part of a programme of research into strategic argumentation. Though that programme is as yet young, the exploration of RPD_{GD} presented here supports two significant conclusions. First, abstract argumentation frameworks allow much greater scope for strategic sophistication in inter-agent dialogue games than is available in the original versions of those games founded upon propositional logics. While RPD_0 is a useful protocol for investigating some forms of "tightened up" human argumentation, the proponent's strategy is limited largely to ensuring that free questions are not wasted. In games of $\mathrm{RPD}_{\mathrm{GD}}$, however, the proponent could blunder disastrously with any assertion, and thus strategy is far more prominent.

The second conclusion is that, while strategic reasoning is likely to be complex in the general case, it is possible to achieve significant strategic advances from even relatively simple foundational concepts such as *MSCK* and fortification. Fortification, in particular, is here under-developed, and exploring its role in more sophisticated strategies is an immediate avenue for further research.

The area of strategic agent argumentation is attracting rapidly increasing interest because of its centrality in building agents that can successfully compete in domains characterised by open, heterogeneous systems with complex market designs. This paper has shown how strategic techniques based on the intrinsic structure of argumentation frameworks can offer significant advantages in such challenging domains.

References

1. Prakken, H.: Formal systems for persuasion dialogue. The Knowledge Engineering Review **21** (2006) 163–88
2. Walton, D.N., Krabbe, E.C.W.: Commitment in Dialogue. Basic Concepts of Interpersonal Reasoning. SUNY Press, Albany (1996)
3. Dung, P.M.: On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. Artificial Intelligence **77**(2) (1995) 321–358
4. Walton, D.: Dialectical Relevance in Persuasion Dialogue. Informal Logic **19**(2&3) (1999) 119–143
5. Dignum, F., Dunin-Keplicz, B., Verbrugge, R.: Agent theory for team formation by dialogue. In: Intelligent Agents VII, Agent Theories Architectures and Languages. LNAI 1986, Springer (2001) 150–166
6. Dignum, F., Dunin-Keplicz, B., Verbrugge, R.: Creating collective intention through dialogue. Logic Journal of IGPL **9**(2) (2001) 289–304
7. Amgoud, L., Maudet, N.: Strategical considerations for argumentative agents (preliminary report). In: Proceedings of the 9th International Workshop on Non-Monotonic Reasoning (NMR). (2002) 409–417
8. Mackenzie, J.D.: Question-begging in non-cumulative systems. Journal of Philosophical Logic **8** (1979) 117–133
9. Amgoud, L., Hameurlain, N.: An Argumentation-Based Approach for Dialog Move Selection. In: Proceedings of ArgMAS 2006: Revised Selected and Invited Papers, Springer (2006)
10. Kakas, A., Maudet, N., Moraitis, P.: Modular representation of agent interaction rules through argumentation. Journal of Autonomous Agents and Multiagent Systems **11**(2) (2005) 189–206 Special Issue on Argumentation in Multi-Agent Systems.

11. Oren, N., Norman, T., Preece, A.: Loose lips sink ships: a heuristic for argumentation. In: Proceedings of the Third International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2006). (2006) 121–134
12. Black, E., Hunter, A.: Using Enthymemes in an Inquiry Dialogue System. In: Proceedings of AAMAS'08. (2008) 437–444
13. Bentahar, J., Moulin, B., Chaib-draa, B.: Specifying and implementing a persuasion dialogue game using commitments and arguments. In: Proceedings of the ArgMAS'04, Springer (2004) 130–148
14. Dunne, P., McBurney, P.: Concepts of optimal utterance in dialogue: selection and complexity. In Dignum, F., ed.: Advances in Agent Communication. Lecture Notes in Artificial Intelligence. Berlin, Germany: Springer Verlag (2003)
15. Procaccia, A., Rosenschein, J.: Extensive-form argumentation games. In: Proceedings of the Third European Workshop on Multi-Agent Systems (EUMAS-05), Brussels, Belgium. (2005) 312–322
16. Riveret, R., Prakken, H., Rotolo, A., Sartor, G.: Heuristics in argumentation: a game-theoretical investigation. In: Computational Models of Argument. Proceedings of COMMA-08. (2008)
17. Roth, B., Riveret, R., Rotolo, A., Governatori, G.: Strategic argumentation: a game theoretical investigation. In: Proceedings of the 11th International Conference on Artificial intelligence and law, ACM Press New York, NY, USA (2007) 81–90
18. Rahwan, I., Larson, K.: Mechanism design for abstract argumentation. In: Proceedings of AAMAS'08. (2008) 1031–1038