

An Information Model for Computing Accountabilities

Matteo Baldoni^[0000-0002-9294-0408] (✉), Cristina Baroglio^[0000-0002-2070-0616], Katherine M. May, Roberto Micalizio^[0000-0001-9336-0651], and Stefano Tedeschi^[0000-0002-9861-390X]

Università degli Studi di Torino — Dipartimento di Informatica
c.so Svizzera 185, I-10149 Torino (Italy)
firstname.lastname@unito.it

Abstract. We propose an information model that describes which data should be available, together with their relationships, in order to identify accountabilities in a group of interacting parties. The model is intended for use in multi-agent systems, and is expressed by means of Object-Role Modeling, due to the relational nature of the concepts involved.

Keywords: Accountability · Responsibility · ORM · Information Model.

1 Introduction

Many contexts, both in the human world and in software, are characterized by the distribution of activities through a group of interacting parties: each member in the group takes care of a part of the activity, and the desired overall result is achieved only when each member behaves properly, and properly interacts with the others. This happens both in human organizations, and in distributed systems like multi-agent systems (MAS). Such contexts often resort to the concept of *responsibility* to refer to the assignment of a task (or a duty) to a member in the group, e.g. [11, 17, 16, 14]. Despite the centrality of the notion of responsibility, and despite the many discussions and models proposed in philosophy, psychology, law, etc. (among which the well-known triangle model of responsibility by Schlenke *et al.* [30]), there is a lack of computational support to the specification and the use of responsibilities inside information systems (that are thought for the business world), as well as in MAS. Even widely used tools like RBAC (role-based access control) or widely adopted agent organization platforms (like JaCaMo [7]) do not provide an understanding of the concept.

A particularly relevant exception in this landscape is the ReMMo conceptual model by Christophe Feltus [17], which proposes a conceptualization of how responsibility is structured. Interestingly, this is also the first proposal in its field that brings into the picture the notion of accountability: ReMMo, in fact, defines responsibility as “a charge assigned to a unique actor to signify its accountabilities concerning a unique business task.” Accountability is considered a central concept in many fields that study human interaction. In sociology (e.g., [15, 25, 18]) it is well-known that social order, i.e. the functioning of a group of individuals, depends on often complex relationships between the parties, relationships that the parties accept, and that bring about expectations on each other’s behavior. From the seminal work in [18], many studies identify in *accountability* the key notion on top of which interaction is built. In political sciences, e.g. [1],

accountability is seen as a major driving force of individuals when it comes to decide about their own behavior. Psychologists provide evidence that accountability increases the salience of goals [27]. Ethnomethodologists postulate that social behavior is configured by relying on the same mechanisms through which it is explained, which indeed give meaning to social action, e.g. [18]. Management studies, e.g. [29], consider it a framework for managing expectations. So, ReMMO is surely an important proposal but still it leaves questions to be answered. In particular, it does not provide an information model that ties the concepts to the environment, capturing constraints on how data evolves. Such a layer is necessary both for attributing responsibilities and, given a situation of interest, for identifying those who answer for it (forward and backward responsibility, according to [26]).

This paper pursues three research aims. First, it refines the characterization of accountability traced in [2, 3]. In particular, relying on a wide literature on accountability (including but not limited to [29, 18, 1, 22, 32, 27, 9]) we identify the main concepts that come into play in the process of accountability determination, such as *expectation* and *control*, as well as the relationships between them. Second, the paper proposes an *information model for accountability*, that is, it describes which data need to be available to develop systems that, in any situation of interest arising in a group of interacting agents, allow the identification of responsibilities. The model, that we provide in Object-Role Modeling (ORM¹) due to the relational nature of the represented concepts, contributes to the development of systems that support the governance of a group of interacting parties. Third, the paper sets the ground for the development of approaches to MAS programming where interaction, coordination, and exception handling will rely on the sibling notions of responsibility and accountability.

The paper is organized as follows. Section 2 characterizes the concept of accountability we rely upon. Section 3 outlines the main requirements that a model of accountability should satisfy and, then, presents the accountability model. Section 4 compares the proposal to other proposals for accountability. Discussions end the paper.

2 Accountability and Responsibility

Responsibility, accountability, answerability, liability, causation, sanction are all related terms, sometimes used as synonyms, some other times used to capture different shades of meaning. A thorough ontological analysis of the term is not in the scope of this work. We will provide just a minimal characterization, sufficient to understand those facets of the notions, on which we rely. In essence, we understand *accountability* as representing the deceptively simple concept of one principal holding another to account for his/her actions, both “good” and “bad.” From this definition, we identify two integral pieces to accountability: 1) a relationship between two entities in which one feels a liability to account for his/her actions to the other; and 2) a process of accounting in which actions are declared, evaluated, and scored. The mechanism of accountability contains two sides that we call *positive* or *negative accountability*. Positive accountability means that a principal is expected to act in a certain way and will be held to account for

¹For an introduction to ORM, please check <http://www.orm.net/pdf/ORMwhitePaper.pdf>

that expectation's fulfillment. Negative accountability, on the other hand, means that a principal is expected to not impede social progress and negatively impact others. The model that we propose captures the information necessary to positive accountability.

Concerning *responsibility*, Feltus [17] sees it as a charge assigned to an agent, which is always linked at least to one accountability. This view is compatible with the triangle model [30]², according to which the term bears two main understandings, each of which investigated at depth by philosophers: one amounting to causation (who did it?), the other to answerability (who deserves positive or negative treatment because of the event?). Schlenke *et al.* explain how responsibility, as individuals perceive it, depends on the strength of three linkages, each of which involves two out of three basic elements, that are prescriptions, events, and identities. Prescriptions come from regulative knowledge and (broadly speaking) concern what should be done or avoided. Events simply occur in the environment. Identities include but are not limited to roles of the individual that are relevant in the context. The three linkages, thus, respectively capture whether and to which extent: a prescription is considered to concern an event, an event is considered relevant for an identity, a prescription is considered to concern an identity.

Accountability has distinctive traits which do not allow making it a special kind of responsibility. First of all it involves two agents, the one who gives the account and the one who takes the account [30, 22, 27, 9]. Indeed, following [30], accountability is a pyramid, that comes into being when an accredited public watches a responsibility triangle. The account taker can only be someone who has some kind of authority on the account giver [29, 12]. The origin of such an authority may be various; for instance, it may be due to a principal-agent relationship, or to a delegation. The account taker is sometimes called the *forum* [9]. Following ReMMO, accountability is in relationship with responsibility (it concerns a responsibility), and a responsibility may be subject to many accountabilities. Accountability may also involve a sanction, as a social consequence of the account giver's achievement or non-achievement of what expected, and of its providing or not providing an account.

All such considerations yield that, in order to properly tackle accountability in a computational system, it is necessary to identify those data that are specific to accountability and their relationships. We now summarize the key aspects of accountability, that we rely upon in drawing the model, reporting also the reference literature:

- a) *Accountability implies agency*. If a principal does not possess the qualities to act "autonomously, interactively and adaptively," i.e. with agency, there is no reason to speak of accountability, for the agent would be but a tool, and a tool cannot be held accountable [31].
- b) *Accountability requires but is not limited to causal significance*. The plain physical causation (also called scientific causation in [30]), that does not involve awareness or choice, does not even create a responsibility, let aside accountability. This view is supported also by [9, 10].
- c) *Accountability does not hinder autonomy*. Indeed, accountability makes sense because of autonomy in deliberation [1, 30, 32, 10].

² A psychological model by Schlenke *et al.* which is widely used in the context of human resource management.

- d) *Accountability requires observability.* In order to make correct judgments, a forum must be able to observe the necessary relevant information. However, in order to maintain modularity, a forum should not observe beyond its scope. For example, if a principal buys a product and the product is faulty, that principal holds the factory as a whole accountable. The factory, in turn, holds one of its members accountable for shoddy production. In other words, accountability determination is strictly related to a precise context. In each context, the forum must be able to observe events and/or actions strictly contained in its scope and decipher accountability accordingly. As context changes, accountability will change accordingly. For this reason, a mechanism to compose different contexts and decide accountability comprehensively is essential.
- e) *Accountability requires control.* Without control decisions cannot be enacted, the agent does not have an impact on the situation. It will be ineffectual. In [21], control is defined as the capability, possibly distributed among agents, of bringing about events. Due to our focus on positive accountability, i.e. on bringing about a situation of interest, we follow this proposal and interpret omissions (not acting) as non-achievements. [13] gives a slightly different definition of control as the ability of an agent to maintain the truth value of a given state of affairs. Alternatives where control amounts to interference or constraint can be devised but are related to negative accountability.
- f) *Accountability requires a mutually held expectation.* Accountability is a directed social relationship concerning a behavior, that serves the purposes of sense-making and coordination in a group of interacting parties, sharing an agreement on how things should be done [18]. The role of expectation is widely recognized [18, 32, 1]. Both parties must be aware of such a relationship for it to have value (the account-taker to know when it has the right to ask for an account, the account giver to know when towards whom it is liable in case of request).
- g) *Accountability is rights-driven.* One is held accountable by another who, in the context, has the claim-right to ask for the account. Particularly relevant on this aspect the understanding of accountability that is drawn in tort law [12], where the account-taker is the only recognized authority who can ask for an account, and the account-giver has a liability towards the account-taker (to explain when requested). Further analysis is carried out in [19].

3 Modeling Accountability with ORM

We have seen that accountability describes a relationship-centric approach, as it is fundamentally defined through relationships that permit one principal to take to account another. Without that relationship-permitting accounting, accountability would be reduced to a hybrid of traceability and blame-giving. To aid the modeling endeavor, use of ORM becomes a natural choice because of its built-in language advantage that places relationships at the center of its expressive power [20]. Throughout the process of modeling, we made use of both ORM and OWL (<https://www.w3.org/TR/owl2-syntax/>) to experiment with different tools and found that our instrument of choice, ORM, allowed us a more easily comprehensible description of accountability than OWL. We

found the centrality of the relationship, rather than the entity, to be particularly apt at expressing our desired concepts in ORM. In particular, the possibility to place constraints on role groups in relationships proved invaluable for describing how relationships interact with and depend on one another. We recall that in this paper we focus on positive accountability, that is when principals do not act as they should (rather than act as they shouldn't in the case of negative accountability). In other words, we are dealing with the *regulation of task completion* rather than task interference.

3.1 Accountability Requirements

For accountability to function, there must be a *base relationship* because of which a principal feels an obligation to account for his/her actions. The relationship entails a nuanced approach to unexpected outcomes or actions. For instance, a buyer may expect a seller to provide some goods and hold him/her to account should the seller not do so, but *not* in absence of payment. In other words, an unexpected action does not necessarily implicate wrongdoing, thanks to *mitigating circumstances* which circumscribe the scope of the accountability relationship – lack of payment is a mitigating circumstance. A straightforward case without mitigating circumstances would be a principal who acts with full understanding of his/her action effects and expected social role, that clearly caused the outcome/action in question, and could have chosen to act otherwise (i.e. has control over the given state of affairs). Thus, a forum would look for, among other qualities, *causation* as well as *autonomy* and *understanding* in action.

Due to the focus on positive accountability, a forum's interest lies in assessing possibility to act, that is, if a principal had complete potential and autonomy as author of the outcome. We can also say he/she effectively caused the outcome through inaction contrary to social expectations (e.g. a merchant who refuses to provide goods that were paid as agreed). A forum must also determine a principal's *situational knowledge* of his/her expected agreed-upon role (e.g. the merchant must be aware the client expects him/her to ship the purchased items). Therefore, a model of positive accountability must respect the following requirements:

- [R1] *Identify relationships of account giving between principals for certain outcomes;*
- [R2] *Account for mitigating circumstances;*
- [R3] *Establish a principal's qualities of:*
 - .1 *Agency*
 - .2 *Causal contribution to outcome;*
 - .3 *Possibility/opportunity to act;*
- [R4] *Allow for the passing of judgment.*

Requirement [R1] expresses the simplest: we must place in *relationship* two principals (who are not necessarily individuals but may, for instance, amount to organizations like a shipping company or an office) along with agreed *expected outcomes*. Note that for us accountability relationships are always the result of an act that was explicitly, deliberately, and voluntarily performed by the account giver, and that amount either to the creation of the accountability relationship by itself or to the acceptance some “rules of the game” – e.g. when enacting some role.

[R2] refines the base relationship through *context*. Context contains a series of conditions that stipulate when there are no mitigating circumstances. We reiterate that we speak in the absence of negative accountability, and can consequently discount interference from other system actors. Context, therefore, represents a kind of precondition to outcomes realization to be specified by the principal on who the expectations lie. The specialized context also protects the principal by disallowing unattainable obligations. The accountability relationship consequently takes the form as described in [10].

[R3] presents us with a more complex modeling problem. For agency, we adopt the definition given in [8], “The person is an autonomous, intentional, and planning agent who is capable of distinguishing right and wrong and good and bad.” A principal, arguably autonomous by design, satisfies the first part of the agency condition when s/he stipulates the ability to bring about an expectation from a context, that is, execute a plan, whether that be his/her own plan or a plan to hold another accountable. In order to know “good” from “bad,” a principal must have foreknowledge of his/her social expectations, whose completion for all intents and purposes are “good” and “right,” and, in absence of negative accountability, whose non-completion are the only “bad” and “wrong” actions.

Likewise, for causal contribution, in absence of negative accountability and by declaring *control* over a *context* and *expectation*, a principal recognizes that one leads to another thanks to some intermediate action on her/his part. As a consequence, given context, the declared principal effectively *causes* either the outcome in question or its absence. Though a principal can subcontract out his/her plan, meaning another principal could be the “cause” of the first principals outcome, from a modular viewpoint the first retains his/her causal relevance for the outermost outcome as the “manager.”

The possibility/opportunity to act condition completes a concept already covered by the previously discussed requirements. Thanks to a principal’s *control* in a specified *context*, if that context comes about, that principal can *freely choose* to act or not act and bring about the desired *outcome*. Again, without negative accountability, through his/her control, that principal has the possibility and opportunity to act.

We choose to leave out the sanctioning piece, that is often yielded by accountability, to the requirement [R4], the passing of judgment, for future implementations of accountability. Our primary concern consists, indeed, in identifying the information that is needed to *support* the passing of judgement and the possible consequent sanction.

3.2 A Model Takes Shape

The accountability model is reported in Figure 1 as an ORM model. Concerning ORM, we report a little information for those who are not familiar with this notation; for details about ORM notation, please see [20]. In ORM all elements with the same name are the same entity/relationship. Entities/relationships can play identifiable (often named, e.g., *account-giver* or *subgoal*) roles inside relationships, that are explicitly depicted in the model. Part (or all) of the roles that make the relationship can be related to groups of roles inside another relationship by a constraint. For instance, in Figure 1 we use inclusion and exclusive or. Relationships represent facts. They are denoted by a reading that gives an intuition of the relationship, and can be characterized in many ways, for instance by graphically denoting reflexivity, a-/anti-/simmetry, commutativity (e.g. *contains* is asymmetric).

Central to the model lies the relationship of accountability (accountability requirement [R1]) which contains two principals, one (account-giver) who is accountable to the other (account-taker), for an achievement (... *is accountable to ... for ...*). An achievement is a pair that is made of a context and an outcome, meaning that the interest in the outcome raises in a specific context –thus, the context limits accountability for an outcome. An Achievement can be structured into subgoals. The asymmetric relation *contains* represents such structure. The accountability relationship is further constrained so that a principal who is accountable must be in control and that there must be a mutually held expectation on that principal to act. The idea of contextual control is then further specified by its two types: control that comes from one’s own efforts, and control that comes from one’s ability to hold another accountable.

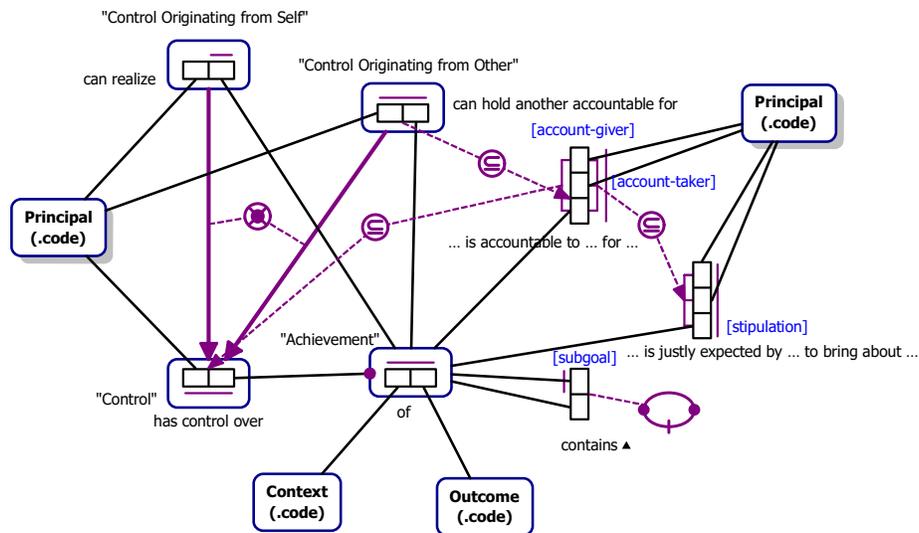


Fig. 1. ORM model for the accountability relationship.

Based on the identified requirements, we identify the following key nouns, which will take the shape of objects. *Principal* partially satisfies the agency requirement [R3] by representing an autonomous individual or organization who might potentially be thought of in legal terms as a persona juris. *Context* and *Outcome* both represent sets of facts that characterize some states of interest. They are always associated with one another in the model but the meaning of such association depends on further information. An outcome represents a condition to achieve, i.e. a set of facts that should be brought about and that can be verified. Context corresponds either to mitigating circumstances (accountability requirement [R2]), or to some preconditions under which a principal has the possibility to pursue an outcome, or the conditions under which an expectation that some outcome will be achieved is activated. The relationships between such objects, based on the requirements that are explained in Section 3.1, are as follows:

1. *Accountability: Principal is accountable to Principal for Achievement.* Represents the actual atomic accountability relationship in which the first Principal can be held to account to the second about Achievement. It has three roles and is constrained by the presence of some of those elements both in a relationship of expectation (outgoing arrow towards ... *is justly expected by ... to bring about ...*), and in one of control (outgoing arrow towards ... *has control over ...*). This fact type satisfies both [R1] (relationship) and [R2] (mitigating circumstances) requirements.
2. *Expectation: Principal is justly expected by Principal to bring about Stipulation.* Partially satisfies the foreknowledge requirement of the agency condition in order to distinguish right from wrong. By this relationship, the principal recognizes that if given the opportunity, realizing the achievement is good and not doing so is bad. With the word justly, we require that both principals are in agreement over the expectation.
3. *Control: Principal has Control over Achievement.* Control expresses contextual autonomy in that a Principal can effectively decide whether or not to realize an Achievement. Control is then divided into two types of control: Control from Self and Control from Other.
4. *Control Originating from Self: Principal can realize Achievement.* This type of control represents a modularization of knowledge. If a Principal could realize Achievement (that is, given Context it could realize Outcome), that Principal is effectively declaring knowledge and ability. We should also emphasize that, as this model defines positive accountability, so does control define a positive control. Control therefore speaks to intended consequences.
5. *Control Originating from Others: Principal can hold another accountable for Outcome.* Similarly we can interpret a kind of implicit knowledge declaration in which a Principal has control over an Achievement only if that Principal can hold another accountable for that Achievement (subset relationship between *Control Originating from Other* and account-taker and achievement in *Accountability*).

The model includes the following constraints between the above listed relationships. The subset constraint between relationships *Accountability* and *Expectation* in Figure 1 satisfies the “agency requirement” (Section 3.1, accountability requirement [R3.1]) by ensuring that accountability is only possible in the presence of a previously established expectation. A Principal *A*, in order to be accountable to Principal *B* for a given achievement, must be justly expected by Principal *B* to realize achievement. In other words, it is not enough for a Principal to exert autonomy, there must also be a socially established agreement of expectation. Simply put, in order to be accountable, a Principal must not only exert situational autonomy, but must also be expected to do so.

Subset constraint between account-giver and achievement in *Accountability*, on one side, with *Control*, on the other side, instead satisfies both the “causal” as well as the “possibility/opportunity” condition ([R3.2] and [R3.3]), as well as links mitigating circumstances to a principal-specified context. The constraint expresses in words that a Principal can be accountable for an Achievement only if that same Principal also has control over that Achievement, which means over an Outcome in same given Context. That is, in order to be accountable for an Outcome, a Principal must exhibit causal rele-

vance over, and have possibility to realize, that Outcome and, thus, be a crucial influence in the truth value result of Outcome in Context.

4 Comparison and Discussion

The 2002 December Report of the Auditor General of Canada [23] recognizes accountability as a critical element of representative democratic government, and proposes a well-known (and widely accepted) definition of accountability, that takes recent developments in public management and governance into account. Here, accountability is defined as “a relationship based on obligations to demonstrate, review, and take responsibility for performance, both the results achieved in light of agreed expectations and the means used.” All the elements of the definition are explained in Exhibit 9.1 [24]. The model we propose captures this understanding of accountability in a straightforward way. Of particular importance is the notion of *agreed expectation* stemming from *either a formal or informal agreement* as a key part of accountability. As underlined in our model, the creation of a just expectation relationship requires an agreement between the parties; an internal expectation, generated from one’s opinions, is not enough to hold another accountable. Another important point is that the key focus of accountability is on results accomplished or not accomplished, i.e. on outcomes, for which principals explicitly take responsibility. Exhibit 9.2 explains that performance should be clearly linked with each party’s capacity (e.g., skills) to deliver. This is what we capture by control. Exhibit 9.2 also explains that besides the accomplishments expected, also the operating constraints to be respected should be explicit, understood, and agreed upon. Operating constraints amount to those mitigating circumstances which, in our model, are explicitly captured by the context and agreed upon by the principals involved.

In [10], the concept of accountability requirement is defined as a directed relationship between two principals (account-giver and account-taker), an antecedent, and a consequent, which together constitute the mutual and conditional expectation between two parties. Should the antecedent become true, account-giver becomes accountable for the expected consequent with account-taker. An organization serves as a context for the accountability requirements. Principals enter the organization by playing some of its roles. Accountability requirements include commitments, authorizations, prohibitions, and empowerment. All the foreseen declinations of the accountability requirement only implicitly entail a notion of control.

[3, 2, 4] propose a protocol to ensure accountability as a design property in an organizational setting. The authors distill five principles for supporting accountability, which make use of the same concepts from the definition of multi-agent organizations (role, goal, power). A commitment-based protocol is, then, defined with the aim of guaranteeing the aforementioned principles, both in the construction of the organization and, through an enhanced monitoring functionality, while it operates. The ORM model we proposed here is more general: it does not explicitly include concepts like organization and role, but those concepts are still derivable. Principle one states that *all the collaborations subject to considerations of accountability among principals occur within a single scope called organization*. Naturally, in order to properly identify relationships of accountability, all agreements, including expectations and declarations of control,

must belong to the same scope to satisfy the subset constraints for accountability. The second and third principles concern the process of enrollment of an agent into an organization. Specifically, *a principal can enroll in an organization only by playing a role that is defined inside the organization and a principal willing to play a role in an organization must be aware of all the powers associated with such a role before adopting it*. As for [6], the ability to affect the institutional state is provided by the acquisition of some powers associated with the role itself at enactment time. The definition of a role delimits the range of actions an agent, who plays that role, can perform (i.e. its powers) inside the organization, thereby encoding the outcomes over which it has control. From the point of view of the ORM model, a role is essentially a set of outcomes and contexts that are grouped together with a purpose of organizing work. Playing a role, therefore, means having control (through powers) over a set of outcomes associated with it, and being aware of the fact that there could be a social expectation related to their realization. The fourth principle states that *a principal is accountable, towards the organization or another agent, for those goals s/he has explicitly accepted to bring about*. From the declaration of just (i.e. mutually agreed) expectation, in conjunction with the subset constraint from accountability, we can conclude that an agent must agree to a goal (through the expectation relationship) before that agent can be considered accountable. Finally, the fifth principle states that *a principal must have the leeway for putting before the organization the provisions s/he needs for achieving the goal to which s/he is committing. The organization has the capability of reasoning on the requested provisions and can accept or reject them*. By specifying control, a principal stipulates contextual conditions under which s/he can realize an outcome. Without a principal's declaration of control, s/he cannot be held accountable thanks to the subset constraint and will not be considered committed to an outcome.

5 Conclusions

In the proposed information model of accountability, just expectation is a key concept, because for an accountability relationship to hold, there must be a fact stating an active and agreed expectation between two principals about an outcome. The way in which this fact is determined is outside the scope of the proposal. For a discussion of expectations and their use in practical reasoning see, for instance [28].

The relationships, which in the model tie accountability, expectation, and control, are versatile and can be used in many ways. The availability of accountability facts influences the behavior of the whole group of interacting agents, both those who play the role of account-givers and those who play the role of account-takers. For instance, a principal, knowing that another principal is accountable for some achievement, and knowing that the account-giver is justly expected to bring about the achievement, will draw expectations on the behavior of that principal, and on the control exercised by that principal, and it will orient its own behavior consequently. Another principal, knowing to be accountable for an achievement, will likely take the outcome into account as a goal when the related context holds. It will be possible for the principal to pursue the achievement because, through one of the inclusion constraints, in order for an accountability fact to hold, there is a control fact stating the principal has control over the

achievement. In case control of the outcome is *from self*, it will perform some skill or power of its own. If, instead, its control is *from other*, the principal actually depends on accountability relationships involving some party. This recursive view of accountability, through control from other, accommodates hierarchies of authorities. For instance, a head of office is accountable over the office procedures not because s/he realizes them directly, but because of the accountability relationships his/her employees have towards him/her. These examples explain how accountability facts have an impact on future actions because in general they will increase awareness, thus accommodating the views in [1, 27]. Moreover, we discuss in [5], by relying on a practical example, how the proposed model overcomes some weaknesses in tackling goal distribution in business processes. In that context, the choice of a language, for expressing achievements, that allows for temporal expressions, introduces requirements on the way in which conditions are brought about, and shows how through the language it is possible to account also for the means used.

Taking a different perspective, the model tells which information is to be collected in order to support the identification of the responsibilities that are involved in the realization of a state of interest. Here, we give to the term responsibility the already discussed meaning provided by [17]. Recall that in that model each responsibility has associated a number of accountabilities. Suppose that the information conforms to our model and that its creation and evolution is collected by some monitoring framework. When a situation of interest arises (in particular, when some achievement is not satisfied), the availability of accountabilities makes it possible to identify the involved principals; not only those trivially involved in the last step that brought to the condition under scrutiny, but also those who, with their choices and actions, drove the execution towards its end. For instance, suppose the sellers in a shop share a view of the available items, and that it is their responsibility to timely inform the other sellers when some item is sold out. Seller Bob forgets to inform seller Alice that the last TV set was sold. Alice is, then, asked by a client about the availability of TV sets. She answers “yes” but when at some point of the sale she looks for the item, it turns out there is none and the sale fails. Accountability will help to identify Bob as the one who should answer. In this simple case only one is responsible and nothing can be done to fix the situation. In other cases it will, instead, be possible to activate compensation processes or to refine the executed processes, based on collected evidence. In fact, the principals, that are connected by the accountability and control relationships, are the executors of a network of related activities, and each will possibly contribute to the understanding of the situation under scrutiny. Note that if we think of each principal as corresponding to a different process, such an understanding will be built across processes, by exploiting the links that are captured by the accountability relationships.

On the other hand, the absence of a just-expectation fact will help to conclude the principal is not liable. For instance, consider a company where help-desk requests are a responsibility of some principal. A system administrator has access to and can operate on what is recorded in the database –thus also on help-desk requests, that are there stored– but is not expected to answer to such requests, so if a request is not attended the administrator will not be held to account. Similarly, a principal who is justly expected to satisfy some achievement, but who has no control over that achievement will not be

asked to account for the achievement –for instance, a new help-desk operator who does not have the proper access rights yet.

The availability of knowledge about the internal structure of an organization would support the realization of mechanisms to reason about accountability at different levels of granularities, as hoped for in Section 2. In other words, the model allows for modularity. So, for instance, a whole organization, acting as a principal, can be accountable for an outcome. That same organization, however, can have internal accountability relationships to add further nuance. Along these lines, it would also be interesting to combine the accountability model with a representation of *compensations*, e.g. [33], that should be executed when an outcome is not achieved. Related to these aims, further nuance could be added by taking into account role power differences in collective accountability. For instance, we might say that the director of the organization should assume more accountability than individual parts because that individual pulls more weight and has more decision power in the organization. We could then assign accountability weights to the various roles in the organization with values that reflect power differences in order to keep accountability proportional to one’s pull in the organization.

References

1. Anderson, P.A.: Justifications and precedents as constraints in foreign policy decision-making. *American Journal of Political Science* 25(4) (1981),)
2. Baldoni, M., Baroglio, C., May, K.M., Micalizio, R., Tedeschi, S.: ADOPT JaCaMo: Accountability-driven organization programming technique for JaCaMo. In: PRIMA 2017 - 20th Int. Conf., Proceedings. LNCS, vol. 10621. Springer (2017)
3. Baldoni, M., Baroglio, C., May, K.M., Micalizio, R., Tedeschi, S.: Supporting organizational accountability inside multiagent systems. In: Proc. of AI*IA 2017. LNCS, vol. 10640, pp. 403–417. Springer (2017)
4. Baldoni, M., Baroglio, C., May, K.M., Micalizio, R., Tedeschi, S.: Computational Accountability in MAS Organizations with ADOPT. *Journal of Applied Sciences* 8(4) (March 2018), special issue “Multi-Agent Systems”
5. Baldoni, M., Baroglio, C., Micalizio, R.: Goal Distribution in Business Process Models. In: Ghedini, C., Magnini, B., Passerini, A., Traverso, P. (eds.) Proc. of 17th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2018). Springer, Trento, Italy (2018), in this same volume.
6. Boella, G., van der Torre, L.W.N.: The ontological properties of social roles in multi-agent systems: definitional dependence, powers and roles playing roles. *Artif. Intell. Law* 15(3), 201–221 (2007)
7. Boissier, O., Bordini, R.H., Hübner, J.F., Ricci, A., Santi, A.: Multi-agent oriented programming with JaCaMo. *Sci. Comput. Program.* 78(6), 747–761 (2013), <http://dx.doi.org/10.1016/j.scico.2011.10.004>
8. Braham, M., van Hees, M.: An anatomy of moral responsibility. *Mind* 121(483) (2012)
9. Burgemeestre, B., Hulstijn, J.: *Handbook of Ethics, Values, and Technological Design*, chap. Designing for Accountability and Transparency. Springer (2015)
10. Chopra, A.K., Singh, M.P.: The thing itself speaks: Accountability as a foundation for requirements in sociotechnical systems. In: IEEE 7th Int. Workshop RELAW. IEEE Computer Society (2014), <http://dx.doi.org/10.1109/RELAW.2014.6893477>
11. Conte, R., Paolucci, M.: Responsibility for societies of agents. *Journal of Artificial Societies and Social Simulation* 7(4) (2004)

12. Darwall, S.: *Morality, Authority, and Law: Essays in Second- Personal Ethics I*, chap. *Civil Recourse as Mutual Accountability*. Oxford University Press (2013)
13. Dastani, M., Lorini, E., Meyer, J.C., Pankov, A.: *Other-condemning anger = blaming accountable agents for unattainable desires*. In: *Proc. of AAMAS*. ACM (2017)
14. Dastani, M., Yazdanpanah, V.: *Distant group responsibility in multi-agent systems*. In: *PRIMA 2016: Principles and Practice of Multi-Agent Systems* (2016)
15. Durkheim, E.: *De la division du travail social*. PUF (1893)
16. Eshleman, A.: *Moral responsibility*. *The Stanford Encyclopedia of Philosophy* (2014)
17. Feltus, C.: *Aligning Access Rights to Governance Needs with the Responsibility MetaModel (ReMMo) in the Frame of Enterprise Architecture*. Ph.D. thesis, University of Namur, Belgium (March 2014)
18. Garfinkel, H.: *Studies in ethnomethodology*. Prentice-Hall Inc., Englewood Cliffs, New Jersey (1967)
19. Grant, R.W., Keohane, R.O.: *Accountability and Abuses of Power in World Politics*. *The American Political Science Review* 99(1) (2005)
20. Halpin, T., Morgan, T.: *Information Modeling and Relational Databases*. Morgan Kaufmann Publishers (2008)
21. Marengo, E., Baldoni, M., Baroglio, C., Chopra, A., Patti, V., Singh, M.: *Commitments with regulations: reasoning about safety and control in REGULA*, AAMAS '11, vol. 2, pp. 467–474. IFAAMAS (2011)
22. Nissenbaum, H.: *Accountability in a computerized society*. *Science and Engineering Ethics* 2(1), 25–42 (1996)
23. Office of the Auditor General of Canada: *2002 December Report of the Auditor General of Canada: Chapter 9* (2002), http://www.oag-bvg.gc.ca/internet/English/parl_oag_200212_09_e_12403.html
24. Office of the Auditor General of Canada: *Exhibit 9.1 The elements of accountability* (2002), http://www.oag-bvg.gc.ca/internet/English/att_20021209xe01_e_12282.html
25. Parsons, T.: *The Structure of Social Action*. Collier-Macmillan, London (1968)
26. van de Poel, I.: *The Relation Between Forward-Looking and Backward-Looking Responsibility*, pp. 37–52. Springer Netherlands (2011)
27. Quinn, A., Schlenker, B.R.: *Can accountability produce independence? goals as determinants of the impact of accountability on conformity*. *Personality and Social Psychology Bulletin* 28(4), 472–483 (April 2002)
28. Ranathunga, S., Cranefield, S., Purvis, M.K.: *Integrating expectation monitoring into BDI agents*. In: *Programming Multi-Agent Systems - 9th Int. Workshop, ProMAS Revised Selected Papers*. LNCS, vol. 7217, pp. 74–91. Springer (2011), https://doi.org/10.1007/978-3-642-31915-0_5
29. Romzek, B.S., Dubnick, M.J.: *Accountability in the Public Sector: Lessons from the Challenger Tragedy*. *Public Administration Review* 47(3) (1987)
30. Schlenker, B.R., Britt, T.W., Pennington, J., Rodolfo, M., Doherty, K.: *The triangle model of responsibility*. *Psychological Review* 101(4), 632–652 (October 1994)
31. Simon, J.: *The Online Manifesto: Being human in a hyperconnected era*, chap. *Distributed Epistemic Responsibility in a Hyperconnected Era*. Springer Open (2015)
32. Suchman, L.: *Discourse, Tools, and Reasoning: Essays on Situated Cognition*, chap. *Centers of Coordination: A Case and Some Themes*. Springer-Verlag, Berlin (1997)
33. Unruh, A., Bailey, J., Ramamohanarao, K.: *A framework for goal-based semantic compensation in agent systems*. In: *Safety and Security in Multiagent Systems - Research Results from 2004-2006*, LNCS, vol. 4324, pp. 130–146. Springer (2009)