Activity Theory as a Framework for MAS Coordination

Alessandro Ricci[1], Andrea Omicini[2], and Enrico Denti[3]

[1] DEIS, Università di Bologna
Via Rasi e Spinelli 176 47023 Cesena, FC, Italy
aricci@deis.unibo.it
[2] DEIS, Università di Bologna
Via Rasi e Spinelli 176 47023 Cesena, FC, Italy
aomicini@deis.unibo.it
[3] DEIS, Università di Bologna
Viale Risorgimento 2 40136 Bologna, Italy
edenti@deis.unibo.it

Abstract. Approaches to the coordination of multiagent systems (MAS) have been recently classified as *subjective* – typically coming from the distributed artificial intelligence (DAI) –, and *objective* – coming from the community of Coordination Models and Languages. Subjective and objective approaches have a very different impact on the engineering of social aspects of MAS, in particular with respect to the ability of specifying and enacting social laws to achieve global coherent behaviours. In this work, we provide a conceptual framework – influenced by the research on *Activity Theory* – where both subjective and objective coordination play an essential role, each providing effective means for the same coordination/cooperative problems at different abstraction and operational levels: *co-construction/co-operation* level for subjective coordination, and *co-ordination* level for objective coordination. In particular, the work shows the benefits of supporting dynamic transitions between such levels, alternating co-operation stages – in which agents reason about coordination and cooperatively forge coordination *artifacts* (laws, constraints, norms) – and co-ordination stages – where the artifacts, embodied in proper coordination media, are exploited, so as to enact automated, consistent and prescriptive coordination.

1 Objective vs Subjective Coordination in MAS

Interaction and coordination are interdisciplinary issues, and therefore it is not surprising that their study and development in multiagent systems (MAS) is supported by approaches coming from heterogeneous contexts (refer to [20, 13, 36, 27, 28] for comprehensive surveys).

The major contribution to coordination in MAS comes from the distributed artificial intelligence (DAI) field. Generally, DAI approaches explicitly deal with *subjective coordination* [32], considering coordination as "the process by which an agent reasons about its local action and the (anticipated) actions of others to try and ensure the community acts in a coherent manner" [13]. Consequently, the coordination of the overall system is determined by both the mental attitudes (beliefs, goals) and the role of individuals, and by the subjective perception of the inter-dependencies among the members of a society [24]. Well-known examples of subjective coordination techniques in MAS are *multi-agent planning* [10], where agents build plans and commit to behave in accordance with them, and *negotiation* [4, 20], where agents communicate so as to reach a mutually accepted agreement on some matter. Typically, subjective coordination is exploited in MAS composed by intelligent (such as BDI) agents, provided with an high-level agent communication language (ACL), such as KQML or FIPA, whose formal semantics can become the means to realise flexible coordination [3].

Other substantial contributions to MAS coordination are rooted in concurrent (parallel and distributed) systems. Starting from the need to explicitly separate coordination from computation issues in system design and development [12], several *coordination models and languages* have been developed [29], and applied to MAS [28]. Generally, these approaches explicitly deal with *objective coordination* [32], promoting the separation between the individual perception of coordination and the global coordination issues, enabling the modelling and shaping of the interaction space independently of the interacting entities [24]. This kind of coordination is called *objective* because it prescinds from the subjective view of the coordinated agents. It is also called *uncoupled* [36], since coordination is no longer coupled with the computational issues of the coordinated entities. According to [32], objective coordination is mainly concerned with inter-agent issues such as *(i)* the description of how the MAS environment is organised, and *(ii)* the management of interactions between both agents and their environment, and agent themselves. Objective coordination models are heavily based of the concept of *mediated interaction*: agents are provided with specific abstractions that enable their actions (typically communications), and mediate the emerging interactions, caused by the dependencies inside the agent ensemble. In the classification generally adopted by the coordination community [6], this abstraction role is assumed by the *coordination medium*, which rules agent interactions by applying *coordination laws*, which represent social laws and system constraints. *Tuple centres* [23] and the related TuCSoN coordination infrastructure [25] are an example of a coordination approach promoting objective coordination.

The variety and complexity of the interaction scenarios that characterise multiagent systems and agent societies require approaches with the properties of both subjective and objective coordination. For instance, distributed workflow management systems (WfMS) and office automation environments require the automatism and prescriptiveness that are easily provided by objective approaches; instead, unstructured environments (from the coordination point of view) and market-oriented competitive scenarios mostly rely on the reasoning and interaction capabilities of individual agents. Moreover, relevant coordination scenarios, such as business process coordination and computer supported cooperative work (CSCW), require models providing the means for balancing subjectivity and objectivity, mediating between application-centric and human-centric tasks, unstructured process and highly structure process structure [8]. In order to face such a complexity, we consider useful from both a scientific and an engineering point of view to provide a conceptual framework exploiting both the subjective and the objective coordination approaches: this is important to ground models and infrastructures aiming at supporting both approaches in MAS coordination.

The interdisciplinary nature of interaction and coordination makes it possible to find relevant contributions in disciplines and theories outside computer science: social theories about organisation and coordination in human societies can be very effective, given the complexity of the interaction space involved in some MAS scenarios. An example is the theory of coordination developed by the Centre of Coordination Study (CCS) at MIT: from the analysis of heterogeneous contexts, from economy to computer science, coordination emerges as the process of *managing dependencies among activities* [17]. Accordingly, engineering coordination inside systems means first identifying the dependencies among the individual parts, then defining strategies to manage the identified dependencies (or interactions, when dependencies are among agent actions), and finally enacting the strategies. The concept of *dependency* is a very important element for both subjective and objective approaches: in the first case of subjective approaches, dependency appears as the basic relationship on which the construction of the whole agent social life is founded [5, 7], while objective coordination approaches are explicitly focused on the management of dependencies/interactions.

The coordination theory approach is clearly *bottom-up* (regulatory): dependencies are the starting point, then coordination activities are designed and developed on top of dependency analysis. However, the social issues that we aiming at engineering in MAS also account for a *top-down* (constructive) approach to coordination: the starting point here are the social goals and the properties that must be provided by the aggregation of agents as a system/society, then dependencies among the individual parts are determined/induced accordingly.

While the theory of coordination provides a good conceptual background for the bottom-up approach, we found *Activity Theory* [37] very effective to frame the role of objective and subjective coordination (and their relationship) in the top-down case. Initially developed to study dynamics in collective human work activity, Activity Theory has been recently introduced in computer science disciplines, such as CSCW and HCI (Human Computer Interaction). Both activity and dependency theories refer (either directly or indirectly) to the notion of mediated interaction: interaction media are studied, catalogued and used to manage dependencies in the theory of coordination, and media are forged as coordination *artifacts* in Activity Theory. The need to define abstractions and "social structures" mediating agent interaction is clearly visible also in the evolution path of some subjective coordination approaches in MAS: the notion of *social agency* [34] and *social laws* [33], the STEAM teamwork module in the TeamCore coordination architecture [35], and stigmergy coordination [30] are notable examples. In the last case, in particular, agent interactions are mediated and ruled by the environment, with its own laws and constraints, and coordination appears as an emergent phenomenon.

Given these premises, in this seminal work we provide a conceptual framework derived from Activity Theory where both subjective and objective approaches are exploited in the same coordination context, but at different conceptual and operational levels; the framework gives insights into the dynamics between the approaches/levels, showing the fundamental role of mediated interaction (exemplified by the coordination medium/artifact abstractions) for that purpose. Three hierarchical levels for analysing every social activity in MAS are identified: *co-construction*, *co-operation* and *co-ordination*. Accordingly, we show how subjective and objective approaches are both fundamental to support such levels: in particular subjective approaches for co-construction and co-operation, and objective ones for co-ordination. In this way, both high level cooperation protocols among intelligent agents – typically found in subjective approaches – and scripted coordination driven by laws embedded in coordination media – as typically found in objective coordination – turn out to be necessary to build MAS with autonomous agents behaving efficiently and coherently in a social/systemic context. Moreover, the dynamics between co-operation and co-ordination is discussed, providing the notion of coordination reflection and reification.

The remaining part of the work is organised as follows: Section 2 provides the basic elements of Activity Theory useful for this work. Section 3 comes back to coordination in MAS, and provides a conceptual framework – derived from previous section – in which co-ordination/co-operation activities are analysed, and the relationship with subjective and objective coordination discussed. Finally, Section 4 provides conclusions and future works.

2 Activity Theory

Activity Theory (AT) is a social psychological theory about the developmental transformation and dynamics in collective human work activity [37, 16, 2]. Recently, AT has been introduced in some fields of computer science – in particular in CSCW [15] and computer system design [18]. AT focuses on *human activities*, distinguished by their respective (physical and ideal) *objects*, that give them their specific directions, i.e. the *objectives* of the activities. Cooperation is understood as a *collaborative activity*, with one objective, but distributed onto several actors, each performing *actions* accordingly to the shared objective. Explicit norms and rules regulate the relationships among individual participants' work. Central to AT – as in the case of objective coordination – is the notion of mediated interaction: human activity is found to be always mediated by *coordination artifacts*, both physical and psychological, such as operating procedures, heuristics, scripts, individual and collective experiences, and languages. In this context, (scripted) *roles* can be understood as actors' expected behaviour with respect to *scripts*, that are the coordination artifacts embodying the shared objectives.

2.1 Co-operation and Co-ordination in AT

AT identifies three hierarchical levels defining the structure of collaborative activities: *co-ordinated*, *co-operative*, *co-constructive* [2, 11]:

- *co-ordinated* aspect of work captures the normal and routine flow of interaction. Participants follow their scripted roles, each focusing on the successful performance of their actions, implicitly or explicitly assigned to them; they share and act upon a common object, but their individual actions are only externally related to each other. Scripts coordinating participants' actions are not questioned or discussed, neither known/understood in all their complexity: in this stage actors act as "wheels in the organisational machinery" [15], and co-ordination ensures that an activity is working in harmony with surrounding activities.
- *co-operative* aspect of work concerns the mode of interactions in which actors focus on a common object and thus share the objective of the activity; unlike previous case, actors do not have actions or roles explicitly assigned to them: with regard to the common object, each actor has to balance his/her own actions with other agent actions, possibly influencing them to achieve the common task. So, in this case the object of the activity is stable and agreed upon: however the means for realising the activity is not yet defined.
- *co-constructive* aspect of work concerns interactions in which actors focus on re-conceptualising their own organisation and interaction in relation to their shared objects. Neither the object of work, nor the scripts are stable, and must be collectively constructed, i.e. *co*-constructed.

In the analysis of collaborative activities, AT emphasises that an activity cannot be said to exist on one level alone: co-ordination, co-operation, and co-construction are *analytical* distinctions of the same collaborative activity, and concur in different times and modes to its development.

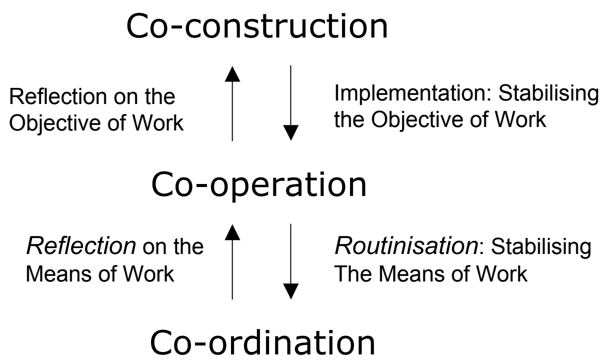

Fig. 1. Dynamics of the Cooperative Work as reported in [2]

2.2 Dynamic Transformation between Collaborative Levels

The notion of dynamic transformation between the hierarchical levels is central to AT. Transformations are strictly related to the stability of the means of work and of the object of work (see Fig. 1, taken from [2]): upward transformations are caused by reflection on the means for doing the work or on the object of the work itself, while downward transformations are caused by resolving conflicts and problems, and re-embodying the solution in the lower levels.

In the context of this article, *reflection* on the means of work – going from co-ordination to co-operation, and *routinisation* – going from co-operation to co-ordination – are the most important transitions. The former happens when the coordinated flow of work, relying on stable means of work such as scripts, rules, coordination artifacts in general, needs to be cooperatively re-established according to the object of work; the reasons can be either coordination breakdown, or a deliberate re-conceptualisation of the way the work is achieved currently. The latter works in the opposite directions, by re-establishing co-ordinated work where the means of collaboration are stabilised, and (mediating) artifacts are provided to be exploited by participants in the co-ordination stage.

The other transitions account for reflecting on the object of work – when the objective become unstable within the collaborating ensemble and transformation to the co-constructive level of collaboration is necessary –, and for implementation – when conflicts on a common objective are resolved, and communication/negotiation ensuring its commitment can begin.

3 Back to MAS: Reconciling Objective and Subjective Approaches

The conceptual framework provided by the AT is effective here to understand the (distinct) roles of objective and subjective coordination in MAS, and the relationships that occur between them. We reconsider the hierarchical structure of collaborative activities provided by the AT, and contextualise it according to the theory of coordination. Accordingly, in context of MAS coordination three main stages can be identified (see Fig. 2):

- **co-construction** – about the understanding and reasoning on what dependencies and interactions have to be faced and managed;
- **co-operation** – about planning what actions must be taken to manage the dependencies and interactions that have been identified in the previous (co-construction) stage;
- **co-ordination** – about enforcing/automating the activities to manage the interactions, planned in the co-operation stage.

Given these three levels, subjective approaches can be considered fundamental for the co-operation stage. Here (intelligent) agents are necessary to reason about what kind of coordination is required, what kind of coordination laws must be developed to manage interactions identified in co-construction stage. The smartness of the agents is useful to build cooperatively – by means of negotiation and high level (semantics driven) interaction protocols – effective artifacts (i.e. coordination laws defining coordination media behaviour) to be used in the co-ordination stage.

Instead, objective coordination is fundamental for the co-ordination stage, where the coordination laws and organisational rules must be enacted in the most automated, fluid, optimised manner. The coordination medium abstraction (and coordination laws defining its behaviour) represents and embodies effectively the AT concept of artifact (and related mediating tools), embedding and enacting in the co-ordination stage the social laws and interaction constraints established in the co-operation stage. The parallel between AT artifacts and coordination media help to understand the role of the media inside MAS: as the artifacts, coordination media first are used to *enable* the interactions among the agents, and then to *mediate* them in the most automated manner. As the artifacts, media become *the* place where the (partial) knowledge about the global behaviour of the MAS is traced and can be further inspected. As the artifacts, media become the source of the *social intelligence* that characterises concretely the systemic/synergistic vision (as opposite of compositional) of MAS [6]; in this context, coordination laws become the macro-levers that can be used to tune and adapt dynamically such a collective intelligence.

The explicit identification of a co-ordination stage distinct from the co-operation stage, with objective coordination models exploited in the first case and subjective models in the latter, is useful to understand the relationships between agent autonomy and social order, between the autonomy of single behaviours and the achievement of a globally coherent behaviour. Agent autonomy is preserved both in the co-operation stage – essential to build the social laws according to the aims and capability of the single agents –, and in the co-ordination stage – where agents are not constrained to know/understand all the coordination machinery (embedded in media) in order to participate to coordination. In particular, in the co-ordination stage agents can focus on the actions related to their specific goals, and delegate implicitly to the medium itself (by means of its role of social mediator) the enactment of the laws that characterise the global coordination activities. In this way, also agents without high level coordination skills can exploit the coordination service; indeed, this is a very important aspect in the engineering of (open) MAS: the support for effective coordination among agents with heterogeneous computational models, not necessarily intelligent or BDI, not necessarily speaking with an high level ACL. The same thing happens in AT collaborative scenarios: often humans exploit coordination artifacts without knowing their whole behaviour, focusing only on the actions that are expected from them by means of their role.

An other important coordination issue taken into account within this conceptual framework is the performance / optimisation of coordination activities. Suitably designed coordination media make it possible to minimise the interactions involved in the co-ordination stage: medium laws automate the coordination flow, minimising the need of negotiation among coordinated agents –

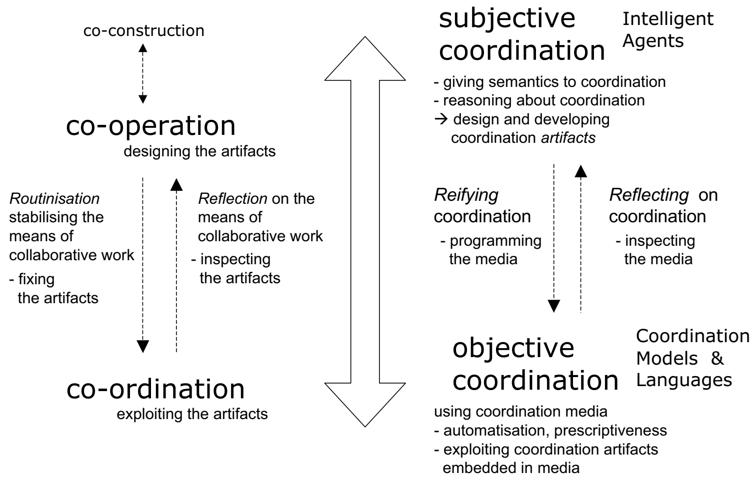

Fig. 2. Dynamics between Objective and Subjective Coordination

which typically characterises the co-operation stage. A good cooperation is then fundamental to produce effective laws that can affect the performance of the global behaviour of the MAS; of course, the impact depends also on the expressive power of the coordination medium, which constrains the possible social laws that can be specified and enforced. Also in this case, similarities with the human case are evident: the effectiveness of the routinised-work in the co-ordination stage depends on the characteristics of the artifacts used to mediate actors' interactions.

This coordination framework is strictly related to MAS planning/scheduling and execution approaches, with two important points that distinguish the former from the latter at the foundation level: *(i)* the main focus is not on the individual agents and their computational/planning behaviour, but on the agent system and interactions; *(ii)* the coordination rules planned and scheduled in the co-operation stage do not "melt" into the agent computational behaviour in the co-ordination stage, but have a conceptual and physical body in the coordination media, outside agents, used by agents. It is worth to remark that Activity Theory has been exploited as a conceptual basis for *situated planning* [1], providing the connection between plans and the contextual condition for realising the plans in actual work.

3.1 Coordination Media are not Agents

Generally, in the context of MAS, agents are the only abstractions used for the system engineering – especially at the development and deployment stage. In particular, mediator services are provided by agents, too: *middle-agents* are a well-known example, providing coordinating activities (such as matchmaking, brokering) among agent providers and requesters (information, goods, or expertise), locating and connecting agents in open environments [14]. Accordingly, middle-agents can be understood as a suitable way to embody the AT artifacts at the co-ordination stage, and so acting as coordination media mediating agent interactions. However, as AT clearly distinguishes the ontological properties of the artifacts – and related mediating tools – and of the actors designing/developing (co-operation stage) and exploiting (co-ordination stage) the artifacts, the same approach can be adopted here for agents and coordination media. Among the main properties expected from a coordination medium there are:

- *inspectability* – the behaviour of a coordination medium should be inspectable, both from human and artificial agents. Moreover, the inspected coordination specification should be described in a declarative way, possibly with a formally defined semantics, to make the comprehension of coordination activities semantics for intelligent agents easier;
- *efficiency/specificity* – a coordination medium should be specialised in the management of interactions, in order to maximise performance in applying the coordination rules; moreover, a medium should be specialised to support the concurrent actions (communications) of multiple agents, possibly providing security, reliability and fault tolerance capabilities;
- *predictability* – the behaviour of a coordination medium should exactly reflect the coordination laws upon which it has been forged (autonomous, unpredictable behaviour is typically not desired); a formal semantics should be defined for the coordination model, in order to precisely define the effect of the coordination laws on the state of the medium and, more generally, on the agent interaction space;
- *flexibility* – a coordination medium should allow its behaviour to be forged dynamically.

Agents are generally not meant to provide such properties, since they are supposed to be autonomous, pro-active, situated entities, and typically speak by means of a general-purpose / high-level communication language [38].

These characteristics remark the conceptual and physical difference between the agent and the medium abstractions. In the TuCSoN model, for instance, coordination media are represented by programmable tuple spaces (or reactive logic-based associative blackboards, from the AI perspective) called tuple centres; agent interaction is enacted through the exchange of information via these spaces. The coordination laws that define the behaviour of tuple centres are expressed in the ReSpecT logic-based language [22, 23]: therefore, coordination artifacts are forged by suitably programming tuple centres, "engraving" the coordination media according to the need. As an example, the chunk of ReSpecT code in Fig. 3 enforces a simple coordination policy synchronising three activities possibly executed by three different agents.

In particular, ReSpecT coordination media provide the properties remarked above. As a general-purpose language for coordination, ReSpecT features a reaction-based computational model, and primitives specialised in the management of

```
reaction(out(task_end(first_task)),(
    in_r(task_end(first_task)),in_r(task_end(second_task)),out_r(new_task_to_do(third_task)))).
reaction(out(task_end(second_task)),(
    in_r(task_end(first_task)),in_r(task_end(second_task)),out_r(new_task_to_do(third_task)))).
```

Fig. 3. The coordination law enforced by the above ReSpecT code induces the execution of the task `third_task` after the completion of tasks `first_task` and `second_task`. An agent takes in charge a task execution by retrieving the tuple `new_task_to_do(`*TaskId*`)` from the tuple centre (by means of an `in` operation) and inserts the tuple `task_end(`*TaskId*`)` (by means of an `out` operation) when the task is done. So, according to the above code and the semantics of ReSpecT reactions [22], when an agent manifests the end of either `first_task` or `second_task`, the medium reacts and produces the tuple required for the execution of `third_task` *iff* both the tuples `task_end(first_task)` and `task_end(second_task)` are currently present in the tuple centre – that is, *iff* both tasks have been successfully brought to end.

interactions / communications. Such a specialisation impacts on both expressiveness – since higher-level coordination pattern/policy can be easily described by composing the primitive elements – and efficiency – since the primitive elements that affect the coordination performance are clearly identifiable and tunable. ReSpecT formal semantics [22] makes it possible to exactly predict the behaviour of the coordination medium given a ReSpecT specification and the current tuple centre state. Also, inspectability of the coordination laws is made possible by their explicit representations in terms of ReSpecT tuples, and their embodiment within tuple centres, where they can be accessed by both agents – through either meta-level coordination primitives or high-level abstractions like *agent coordination contexts* [21] – and engineers – through specific deployment tools like the Inspector [9]. Finally, the behaviour of a tuple centre can be modified dynamically, by changing at run-time its ReSpecT specification, thus providing the required flexibility.

3.2 Reflecting and Reifying Coordination

The notion of dynamic transformation between levels in collaborative activities is central to AT. Accordingly, central to MAS coordination is the support for dynamic transformation from co-operation – that is the subjective coordination level – to co-ordination – that is the objective coordination level –, and viceversa. In particular, this is important in the context of *open* systems, where the environment is subject to change, and collective goals, norms, organisational rules must be adapted accordingly. This dynamism is captured by two basic transitions, the *reflection* and the *reification* of coordination, which must be supported dynamically, anytime required, during system execution. These transitions are strictly related to the transformations seen in the AT, and account for (see Fig. 2):

- **reification** – in this transition, coordination laws that have been designed and developed in the co-operation stage are reified in coordination media: intelligent agents forge / program the coordination media behaviour in order to reflect the social rules established in the co-operation stage, and to be used as artifacts in the co-ordination stage. It is worth noting that coordination media are meant to embed not only the rules promoting cooperation

among agents, but also the laws ruling interactions, useful to represent also norms and environment constraints, mediating possible agent competitive (not cooperative) behaviour.
– **reflection** – in this transition, the behaviour of the coordination media deployed in co-ordination stage is inspected. Agents can retrieve and understand the coordination laws underlying medium behaviour, in order to either evolve them according to changes in coordination policies or in environmental conditions, or learn how to exploit efficiently the artifacts.

By making it possible to balance dynamically subjectiveness and objectiveness, providing the tools to establish at runtime the distribution of the coordination burden between media and agents, objective coordination models endorse a relevant engineering impact. As a useful picture, we can imagine a "coordination engineering segment" whose extreme points are on the one side all coordination intelligence upon agents, and on the other side all coordination burden upon media. The main issue here is to provide the expressive tools to locate, and dynamically move, the coordination engineering of the system in any point inside the segment, thus determining a different morphology of coordination artifacts, providing a smooth, by-need transition from subjective to objective orientation (see Fig. 4). The position of the point depends on both the considered coordination scenario and the dynamics inside that scenario: the more automation/prescriptiveness is required and the social rules are well-defined (such as for workflow systems), the more the coordination media are charged; the more in the coordination context it is not possible (or feasible) to clearly identify collective rules/constraints, the more the individual agents are charged of the coordination burden. In order to support these capabilities, coordination models and technologies must provide the means (languages and tools) for coordination reflection (from objective to subjective transition), to inspect the coordination laws defining media behaviour, and coordination reification (from subjective to objective transition), programming the behaviour of the coordination media.

The dynamic support for balancing task automation with cooperation flexibly is among the most important requirements for state-of-the-art systems for workflow management, supply chain management, and CSCW [8, 19]. The ability to change the "engineering point" of coordination dynamically is specially required in open systems, where the environment can unpredictably change, the goals of the system can evolve, and coordination laws can be improved as a result of agent interaction. In the case of TuCSoN, this is achieved by means of the ReSpecT tuple centre model, and the tools provided by the infrastructure. The coordination laws that define the behaviour of the coordination media (tuple centres) expressed in the ReSpecT language can be inspected and changed dynamically by human and artificial agents by means of specific tools. We are verifying the effectiveness of this approach in scenarios such as pervasive computing and inter-organisational workflow management systems [31]. In the last context, for instance, tuple centres have been used to embody the role of workflow engines, and workflow rules have been described as ReSpecT coordination laws. Each workflow engine (tuple centre) acts as a coordination artifact provid-

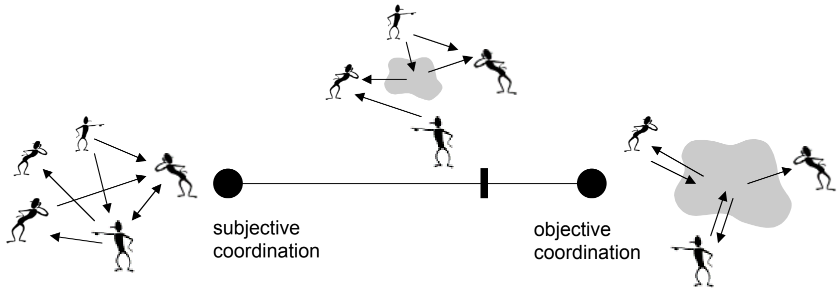

Fig. 4. The coordination engineering segment: all coordination in agents (subjective, left extreme), all coordination in media (objective, right extreme)

ing fluid coordination of the individual tasks executed autonomously by human and artificial agents. So, *(i)* workflow rules are inspectable by inspecting the coordination laws embedded in tuple centres (reflection stage); *(ii)* workflow rules are modifiable at runtime – as a consequence of exceptions, or changes in the business environment – by changing the coordination laws of tuple centres (reification stage).

4 Conclusions

Activity Theory provides a general framework for MAS coordination, enabling a precise understanding of the distinct roles of subjective and objective approaches, and of their mutual relationship as well. By stressing the role of coordination artifacts, AT points out that models for MAS coordination, to be effective, should support both the subjective and objective viewpoints, and explicitly enable the dynamic transitions between the two – allowing coordination laws cooperatively established by agents to be fixed as coordination artifacts in proper media, and, conversely, coordination media to be inspected and their behaviour understood by agents.

Coordination infrastructures and tools are necessary to support the models at runtime, providing coordination as a service: this accounts for *(i)* providing human and artificial agents with the means for dynamically understanding what coordination media/artifacts are available and how they can be exploited (protocols, permissions, etc.); *(ii)* providing agents with tools to dynamically create, access and exploit the coordination media/artifacts [9].

Explicitly accounting for subjective and objective coordination also gives the opportunity to flexibly balance the intelligence/burden of coordination between the agents and the coordination services provided by the infrastructure. This can be done according to the level of automation and prescriptiveness needed, according to the skills of the agents involved in the coordination, and to the expressiveness of the media embodying/acting-as the artifacts. In particular, reflection and reification transitions allow the balance to be tuned and adapted dynamically, according to the changes required by the application scenarios.

References

1. J. Bardram. Plans as situated action: An activity theory approach to workflow systems. In *Proceedings of the ECSCW 97 Conference*, Sept. 1997.
2. J. Bardram. Designing for the dynamics of cooperative work activities. In *Proc. of the 1998 ACA4 Conference on Computer Supported Cooperative Work*, 1998.
3. F. Bergenti and A. Ricci. Three approaches in coordination of MAS. In *Proc. of the 2002 ACM Symposium on Applied Computing (SAC'02)*. ACM, 2002. Track on Coordination Models, Languages and Applications.
4. S. Bussmann and J. Mueller. A negotiation framework for co-operating agents. In D. S., editor, *Proocedings of CKBS-SIG*, pages 1–17. Dake Center, University of Keele, 1992.
5. C. Castelfranchi, M. Miceli, and A. Cesta. Dependence relations among autonomous agents. In Y. Demazeau and E. Werner, editors, *Decentralized AI*, pages 215–231. Elsevier, 1992.
6. P. Ciancarini, A. Omicini, and F. Zambonelli. Multiagent system engineering: the coordination viewpoint. In N. R. Jennings and Y. Lespérance, editors, *Intelligent Agents VI — Agent Theories, Architectures, and Languages*, volume 1767 of *LNAI*, pages 250–259. Springer-Verlag, Feb. 2000.
7. R. Conte and C. Castelfranchi. *Cognitive and Social Action*. University College London, 1995.
8. U. Dayal, M. Hsu, and L. Rivka. Business process coordination: State of the art, trends and open issues. In *Proceedings of 27th VLDB Conference*, 2001.
9. E. Denti, A. Omicini, and A. Ricci. Coordination tools for the development of agent-based systems. *Applied Artificial Intelligence*, 16(9), Oct. 2002. Special Issue "From Agent Theory to Agent Implementation".
10. E. Durfee. Planning in distributed artificial intelligence. In G. M. P. O'Hare and N. R. Jennings, editors, *Foundations of Distributed Artificial Intelligence*. Wiley, 1996.
11. Y. Engeström, K. Brown, L. Christopher, and J. Gregory. Coordinatio, cooperation, and communication in the courts. In Y. Engeström and O. Vasquez, editors, *Mind, Culture, and Activity*. Cambridge University Press, 1997.
12. D. Gelernter and N. Carriero. Coordination languages and their significance. *Communications of the ACM*, 35(2):97–107, Feb. 1992.
13. N. Jennings. Coordination techniques for distributed artificial intelligence. In G. M. P. O'Hare and N. R. Jennings, editors, *Foundations of Distributed Artificial Intelligence*, pages 187–210. Wiley, 1996.
14. M. Klusch and K. Sycara. Brokering and matchmaking for coordination of agent societies: A survey. In Omicini et al. [26], chapter 8, pages 197–224.
15. K. Kuutti. The concept of activity as a basic unit of analysis for CSCW research. In *Proceedings of the Second European Conference on CSCW*, pages 249–264. Kluwer Academic Publisher, 1991.
16. A. Leontjev. *Activity, Consciousness, and Personality*. Prentice Hall, 1978.
17. T. Malone and K. Crowstone. The interdisciplinary study of coordination. *ACM Computing Surveys*, 26(1):87–119, 1994.
18. B. A. Nardi. *Context and Consciousness: Activity Theory and Human-Computer Interaction*. MIT Press, 1996.
19. G. Nutt. The evolution toward flexible workflow systems. *Distributed Systems Engineering*, 3(4):276–294, Dec. 1996.

20. H. Nwana, L. Lee, and N. R. Jennings. Co-ordination in multi-agent systems. In H. Nwana and N. Azarmi, editors, *Software Agents and Soft Computing*, volume 1198 of *LNAI*. Springer-Verlag, 1997.
21. A. Omicini. Towards a notion of agent coordination context. In D. Marinescu and C. Lee, editors, *Process Coordination and Ubiquitous Computing*, pages 187–200. CRC Press, 2002.
22. A. Omicini and E. Denti. Formal ReSpecT. In A. Dovier, M. C. Meo, and A. Omicini, editors, *Declarative Programming – Selected Papers from AGP'00*, volume 48 of *Electronic Notes in Theoretical Computer Science*, pages 179–196. Elsevier Science B. V., 2001.
23. A. Omicini and E. Denti. From tuple spaces to tuple centres. *Science of Computer Programming*, 41(3):277–294, Nov. 2001.
24. A. Omicini and G. A. Papadopoulos. Why coordination models and languages in AI? *Applied Artificial Intelligence*, 15(1), January 2001. Special Issue: Coordination Models and Languages in AI.
25. A. Omicini and F. Zambonelli. Coordination for Internet application development. *Autonomous Agents and Multi-Agent Systems*, 2(3):251–269, Sept. 1999. Special Issue: Coordination Mechanisms for Web Agents.
26. A. Omicini, F. Zambonelli, M. Klusch, and R. Tolksdorf, editors. *Coordination of Internet Agents: Models, Technologies, and Applications*. Springer-Verlag, Mar. 2001.
27. S. Ossowski. *Co-ordination in Artificial Agent Societies: social structures and its implications for autonomous problem solving agents*, volume 1535 of *LNAI*. Springer-Verlag, 1998.
28. G. A. Papadopoulos. Models and technologies for the coordination of Internet agents: A survey. In Omicini et al. [26], chapter 2, pages 25–56.
29. G. A. Papadopoulos and F. Arbab. Coordination models and languages. *Advances in Computers*, 46:329–400, 1998.
30. V. D. Parunak. 'Go To The Ant': Engineering principles from natural agent systems. *Annals of Operations Research*, 75:69–101, 1997.
31. A. Ricci, E. Denti, and A. Omicini. Agent coordination infrastructures for virtual enterprises and workflow management. *Special Issue of International Journal of Cooperative Information Systems*, 2002. to appear.
32. M. Schumacher. *Objective Coordination in Multi-Agent System Engineering – Design and Implementation*, volume 2039 of *LNAI*. Springer-Verlag, Apr. 2001.
33. J. Shoham and M. Tennenholtz. Social laws for artificial agent societies: Off-line design. *Artificial Intelligence*, 73, 1995.
34. M. Singh. Agent communication languages: Rethinking the principles. *IEEE Computer*, 31(12), 1998.
35. M. Tambe and D. V. Pynadath. Towards heterogeneous agent teams. In M. Luck, V. Mrik, O. Stepankova, and R. Trappl, editors, *Multi-Agent Systems and Applications*, volume 2086 of *LNAI*, pages 1–16. Springer-Verlag, 2001.
36. R. Tolksdorf. Models of coordination. In A. Omicini, R. Tolksdorf, and F. Zambonelli, editors, *Engineering Societies in the Agents World*, volume 1972 of *LNAI*, pages 78–92. Springer-Verlag, Dec. 2000.
37. L. S. Vygotskij. *Mind and Society*. Harvard University Press, 1978.
38. M. J. Wooldridge and N. R. Jennings. Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.