Formal ReSpecT in the A&A Perspective

Andrea Omicini[1]

Alma Mater Studiorum—*Università di Bologna,*
via Venezia 52, 47023 Cesena, Italy

Abstract

Coordination languages and models have found a new course in the context of MAS (multiagent systems). By re-interpreting results in terms of agent-oriented abstractions, new conceptual spaces are found, which extend the reach of coordination techniques far beyond their original scope. This is for instance the case of coordination media, when recasted in terms of coordination artifacts in the MAS context.

In this paper, we take the well-established ReSpecT language for programming tuple centre behaviour, and adopt the A&A (agents and artifacts) meta-model as a perspective to reinterpret, revise, extend and complete it. A formal model of the so-called A&A ReSpecT language is presented, along with an example illustrating its use for MAS coordination.

Key words: Tuple-based Coordination, Artifacts for MAS, A&A, Tuple Centres, ReSpecT.

---

1 Introduction

In the last decade, the field of coordination models and languages has produced a wide range of results on the general issue of governing interaction in complex systems: such results are today finding their natural exploitation in hot areas of computational system research, such Web Service Orchestration, WfMS (workflow management systems), and MAS (multiagent systems). It is then seemingly appropriate to take well-established results from the coordination field, and go beyond their mere inter-disciplinary application—gearing instead toward a full trans-disciplinary approach. This means essentially that findings by coordination researchers should first be taken and used to address the issues of interaction management in complex systems, as they emerge from other research areas (inter-disciplinarity); then, they should be recast according to the new conceptual framework, suitably revised and extended along to the

---

[1] Email: `andrea.omicini@unibo.it`

This paper is electronically published in
Electronic Notes in Theoretical Computer Science
URL: `www.elsevier.nl/locate/entcs`

new lines of interpretation, and in such a new form brought back to where they came from (trans-disciplinarity).

Coordination is today acknowledged as one of the key issues in the modelling and engineering of complex systems: as such, it has been the subject of numerous investigations in areas like Sociology, Economics and Organisational Theory [16]. There, coordination is generally conceived as a means to integrate a multiplicity of diverse activities or processes in such a way that the resulting ensemble exhibits some desired / required features. The design of coordination mechanisms is particularly challenging in the field of MAS, as they are usually embedded in highly dynamic environments, and neither the number nor the behaviour of agents are possibly known at design time.

However, conceptual foundations of the MAS area are still under impetuous development, pushed by deep and heterogeneous inputs from distributed computing, programming languages, software engineering, simulation, artificial intelligence, and other related areas that are today converging toward agent-orientedness [30]. Among the most promising approaches, the A&A meta-model [19] re-interprets MAS in terms of two fundamental abstractions: *agents* and *artifacts*. Agents are the active entities encapsulating control, which are in charge of the goals/tasks that altogether build up the whole MAS behaviour. Artifacts are instead the passive, reactive entities in charge of the services and functions that make individual agents work together in a MAS, and that shape agent environment according to the MAS needs. Altogether, the A&A meta-model has a deep impact on the way in which MAS are engineered [10], programmed [26], and simulated [11]. Along this line, coordination artifacts can be conceived as a generalisation of coordination media, as specialised artifacts encapsulating coordination services for MAS [19].

Bringing back the A&A meta-model to the coordination fields suggests a number of interesting considerations. For instance, features of artifacts (like inspectability, forgeability, linkability, etc.) can be used to build up a framework for classifying coordination media, to understand and compare them along the coordination literature [18]. More generally, coordination models and languages can be suitably reinterpreted within the A&A conceptual framework, and revised and extended accordingly. In this paper, we follow the latter line of thought.

In particular, we take the ReSpecT language for programming the behaviour of tuple centres [15] and its formal model [14], and discuss its reformulation in the A&A framework. Section 2 shortly discusses the A&A meta-model, and the main features of artifacts. Section 3 briefly recalls the essentials of the TuCSoN model and of the ReSpecT language. Section 4 introduces the new, revised ReSpecT syntax, along with an example of A&A ReSpecT coordination of agents. Section 5 presents the formal semantics of A&A ReSpecT, obtained by largely revising and extending the original one [14] along the A&A main lines. After Section 6 discusses relations with previous work and related literature, Section 7 provides for final remarks and future lines of work.

2 The A&A Meta-Model for MAS

Our approach to MAS coordination is grounded on the A&A (agents and artifacts) meta-model, which adopts artifacts—along with agents—as the basic MAS building blocks to program and, more generally, to engineer complex software systems [19].

In the A&A meta-model, agents are the basic abstractions to represent active, task-/goal-oriented components, designed to pro-actively carry on one or more activities toward the achievement of some kind of objective, requiring different levels of skills and reasoning capabilities. On the other hand, artifacts are the basic abstractions to represent passive, function-oriented building blocks, which are constructed and used by agents, either individually or cooperatively, during their working activities.

Taking human society as a metaphor, agents play the role of humans, while artifacts coincide with the objects and tools (called artifacts in the human society, too) used by humans as the means to support their work and achieve their goals, or as the target of their activities. The role of artifacts in the context of human activities—in particular social activities—is one of the most important and investigated points of many theories—Activity Theory (AT) [9] and Distributed Cognition in particular—well-known and used in fields such as CSCW and HCI [27,12].

According to AT, any activity carried on by one or more components of a systems—individually or cooperatively—cannot be conceived or understood without considering the tools or artifacts mediating the actions and interactions of the components. Artifacts on the one side mediate the interaction between individual components and their environment (including the other components); on the other side they embody the portion of the environment that can be designed and controlled to support components' activities. Moreover, as an observable part of the environment, artifacts can be monitored along with the development of the activities to evaluate overall system performance and keep track of system history. In other words, mediating artifacts become first-class entities for both the analysis and synthesis of individual as well as cooperative working activities inside complex systems. Such a vision is also promoted by Distributed Cognition [8], a branch of cognitive science that proposes that human cognition and knowledge representations, rather than being confined within the boundaries of an individual, is distributed across individuals, tools and artifacts in the environment.

The complexity of activities of the social systems accounted for by AT and Distributed Cognition can be found nowadays in MAS. This is why we consider the inter-disciplinary study of such conceptual frameworks as fundamental for the analysis and synthesis of social activities inside MAS, and in particular of the artifacts mediating such activities [24]. Examples range from coordination abstractions such as tuple centres [15], to pheromone infrastructure [22] in the context of stigmergy coordination, to the Institution abstraction in electronic-institution approaches [5], to cite some.

Unlike agents, artifacts are not meant to be autonomous or exhibit a pro-active behaviour, neither to have social capabilities. Among the main properties that are useful according to artifact purpose and nature [18], one could list: *(i) inspectability* and *controllability*, i.e. the capability of observing and controlling artifact structure, state and behaviour at runtime, and of supporting their on-line management, in terms of diagnosing, debugging, testing; *(ii) malleability*, i.e. the capability of artifact function to be changed / adapted at runtime (on-the-fly) according to new requirements or unpredictable events occurring in the open environment, *(iii) linkability*, i.e. the capability of linking together at runtime distinct artifacts as a form of composition, as a means to scale up with complexity of the function to provide, and also to support dynamic reuse, *(iv) situation*, i.e. the property of being immersed in the MAS environment, and to be reactive to environment events and changes. It is worth to be remarked that most of these artifact features are not agent features: typically, agents are not inspectable, do not provide means for malleability, do not provide operations for their change, and do not compose with each other through operational links. On the other hand, agents are typically told to be situated: however, how this is realised, in particularly how pro-activity and re-activity features could be reconciled, is not an easy matter. Instead, once artifacts are situated, agent situatedness could be recasted in terms of their interaction with artifacts.

Coordination artifacts [19] are a primary example of artifacts for MAS, as artifacts designed to provide agents and MAS with specific coordination functionalities and services [29]. In human societies, coordination artifacts are as common as traffic lights, street signs, post-its on whiteboards; in computational systems, things like blackboards, event-services, shared message boxes, could be easily be seen as coordination artifacts. In the context of MAS, coordination artifacts are used to both enable and govern forms of *mediated interaction*—i.e., where agents do not communicate directly but through a medium—, which is essential to support forms of communication that are uncoupled along both the time and space dimensions.

So, the overall view of MAS adopting the A&A perspective is given by agents distributed across the networks that inter-operate and coordinate both by communicating via some kind of ACL (agent communication language)—such as FIPA ACL [6]—and by sharing and (co-)using different kind of artifacts. Generally speaking, the A&A meta-model recasts the space of interaction within MAS. So, the components of a MAS interact in three different ways: agents *speak* with agents; agents *use* artifacts; artifacts *link* with artifacts.

Dealing with the government of interaction, coordination models and infrastructures like `TuCSoN` [21] represent the most natural technologies upon which the A&A approach can be put to test. Therefore, revising TuCSoN models and languages under the A&A viewpoint is seemingly appropriate. In particular, in this paper we recast the ReSpecT language for programming TuCSoN tuple centres: we reintepret, revise, extend and complete it so as to make it fit the A&A meta-model for MAS.

3 TuCSoN & ReSpecT

TuCSoN (Tuple Centres Spread over the Network[2]) is a general-purpose agent-oriented model and infrastructure for MAS coordination [21]. TuCSoN is based on a coordination model providing *tuple centres* as first-class abstractions to design and develop general-purpose coordination artifacts [15]. TuCSoN tuple centres are programmed through the ReSpecT logic-based specification language. In the remainder of this section, we first recall the essentials of tuple centre coordination in TuCSoN (Subsection 3.1); then, we resume the main features of the original ReSpecT language for programming the behaviour of TuCSoN tuple centres (Subsection 3.2).

3.1 The TuCSoN Tuple Centre Coordination Model

A tuple centre is a tuple space enhanced with the possibility to program its behaviour in response to interactions.

So, first of all, agents can operate on a TuCSoN tuple centre in the same way as on a Linda tuple space [7]: by exchanging *tuples* (which are ordered collection of knowledge chunks) through a simple set of coordination primitive. An agent can write a tuple in a tuple centre with an `out` operation; or read a tuple from a tuple centre with operations such as `in`, `rd`, `inp`, `rdp` specifying a *tuple template*—that is, an identifier for a set of tuples, according to some *tuple matching* mechanism. Reading tuples can be destructive (`in`, `inp` remove the matching tuple) or non-destructive (`rd`, `rdp` simply read the matching tuple), suspensive (`in`, `rd` wait until a matching tuple is found) or non-suspensive (`inp`, `rdp` immediately return either the matching tuple or a failure result)—but is anyway non-deterministic: when more than one tuple in a tuple centre are found that match a tuple template, one is non-deterministically chosen among them.

Accordingly, a tuple centre enjoys all the many features of a tuple space, which can be classified along three different dimensions: generative communication, associative access, and suspensive semantics. The main features of generative communication (where information generated has an independent life with respect to the generator) are the forms of uncoupling (space, time, name) based on mediated interaction: sender and receiver do not need to know each other, to coexist in the same space or at the same time in order to

---

[2] The TuCSoN technology is available as an open source project at the TuCSoN web site [28]

communicate (to exchange a tuple, in particular). Associative access (access based on structure and content of information exchanged, rather than on location or name) based on tuple matching promotes synchronisation based on tuple structure and content: thus, coordination is information-driven, and allows for knowledge-based coordination patterns. Finally, suspensive semantics promotes coordination patterns based on knowledge availability, and couples well with incomplete or partial knowledge.

Even more, while the basic tuple centre model is independent of the type of tuple [15], TuCSoN tuple centres adopt logic tuples—both tuples and tuple templates are essentially Prolog facts—and unification is used as the tuple-matching mechanism. So, for instance, an agent `ag1` performing operation `we?in(activity(ag1,CaseID))` on tuple centre `we` containing tuples `activity(ag1,c16)` and `activity(ag2,c22)` will be returned tuple `activity(ag1,c16)` (the one unifying with the template) removed from `we`. Since the overall content of a tuple centre is a multiset of logic facts, it has a twofold interpretation as either a collection of messages, or a (logic) theory of communication among agents—thus promoting in principle forms of reasoning about communication.

The TuCSoN infrastructure makes it possible to exploit tuple centres as coordination services distributed over the network [21]. In particular, TuCSoN overall coordination space is constituted by an open set of TuCSoN *nodes*, which correspond to Internet hosts or servers connected by the network. Each node can contain any number of tuple centres, each identified by a unique (inside the node) logic name (e.g. `message_board`). An agent can refer tuple centres either specifying their *full name*, that is, their logic name plus the address of the node hosting the tuple centre (e.g. `message_board@acme.org`), or their *local name*, for tuple centres located on the same host where the agent is situated. As a result, agents can exploit either the local or the global coordination space by adopting either the local or the full name.

Finally, a tuple centre is a programmable tuple space—thus adding programmability of the coordination medium as a new dimension of coordination. While the behaviour of a tuple space in response to communication events is fixed (so, the effects of coordination primitives is fixed), the behaviour of a tuple centre can be tailored to the application needs by defining a set of specification tuples, or reactions, which determine how a tuple centre should react to incoming / outgoing events.

While the basic tuple centre model is not bound to any specific language to define reactions [15], TuCSoN adopts the logic-based language ReSpecT (Reaction Specification Tuples) to program tuple centres.

3.2 ReSpecT *as a Core Coordination Language*

The original ReSpecT [14] is a logic-based language for the specification of the behaviour of tuple centre adopted by TuCSoN. As a behaviour specification language, ReSpecT:

- enables the definition of computations within a tuple centre, called *reactions*, and
- makes it possible to associate reactions to events occurring in a tuple centre.

So, ReSpecT has both a declarative and a procedural part. As a *specification language*, it allows events to be declaratively associated to reactions by means of specific logic tuples, called *specification tuples*, whose form is `reaction(E,R)`. In short, given a event Ev, a specification tuple `reaction(E,R)` associates a reaction $R\theta$ to Ev if $\theta = mgu(E,Ev)$.[3] As a reaction language, ReSpecT enables reactions to be procedurally defined in terms of sequences of logic reaction goals, each one either succeeding or failing. A reaction as a whole succeeds if all its reaction goals succeed, and fails otherwise. Each reaction is executed sequentially with a transactional semantics: so, a failed reaction has no effect on the state of a logic tuple centre.

All the reactions triggered by a communication event are executed before serving any other event: so, agents perceive the result of serving the communication event and executing all the associated reactions altogether as a single transition of the tuple centre state. As a result, the effect of a communication primitive on a logic tuple centre can be made as complex as needed by the coordination requirements of a MAS. Generally speaking, since ReSpecT has been shown to be Turing-equivalent [3], any computable coordination law can be in principle encapsulated into a ReSpecT tuple centre. This is why ReSpecT can be assumed as a general-purpose core language for coordination: a language that can then be used to represent and enact policies and rules of any sort for collaboration support systems.

Adopting the declarative interpretation of ReSpecT tuples, a TuCSoN tuple centre has then a twofold nature [14]: a theory of communication (the set of the ordinary tuples) and a theory of coordination (the set of the specification tuples)—which allows in principle intelligent agents to reason about the state of collaboration activities, and possibly affect their dynamics. Furthermore, the twofold interpretation of ReSpecT specification tuples (either declarative or procedural) allows knowledge and control to be represented uniformly (as Prolog-like facts) and encapsulated in a unique coordination artifact.

3.3 TuCSoN *&* ReSpecT *in the A&A Perspective*

In the A&A perspective, TuCSoN provides agents with a multiplicity of distributed artifacts (the tuple centres) containing both shared knowledge and the logic of coordination expressed in terms of logic tuples. ReSpecT tuple centres are inspectable artifacts (not controllable), and are malleable, since their behaviour can be affected at run-time by changing their behaviour specification. While the original ReSpecT specification did not encompass neither

---

[3] *mgu* is the most general unifier, as defined in logic programming.

linkability nor situatedness [15], two extensions were already introduced that moved along such directions. First, a first extension was proposed in [25], which introduced the first linkability primitive for tuple centre composition, that is, `out_tc`. Then, Timed ReSpecT was defined in [17], which first proposed the notions of timed artifact and timed tuple centre, and allowed for the specification of time-dependent coordination policies, encapsulated within Timed ReSpecT tuple centres.

4 Introducing A&A ReSpecT

4.1 Adopting the A&A Perspective

Adopting the A&A perspective promotes a more articulated view over the space of MAS interaction. First of all, a more general notion of event is required. Since artifacts are passive entities, the only real sources of events in a MAS are agents and the environment. So, whatever happens in a MAS has its "prime cause" either in an agent action, or in an environment phenomenon. However, artifacts are reactive, and link with each other—so, they can operate one each other. As a first consequence, the direct cause of any artifact event may also be some link operation from another artifact—not the prime cause, anyway. So, a general event descriptor should include both the original cause of an event, and the most direct one—thus allowing the event chain to be fully observed, and artifact coordinative behaviours to be properly defined.

As a further consequence, artifact operations should be available for exploitation to other artifacts, too: so, *usage* of artifacts by agents, and *linking* between artifacts should be as uniform as possible. So, primitives for artifact operations should be available for exploitation to both agents and artifacts—with no assumptions on the nature and behaviour of the invoker of a primitive.

As a meta-model for distributed computing, A&A also promote uncoupling of control: so, *(i)* linked artifacts should be fully uncoupled, *(ii)* agents should be left free to autonomously choose either synchronous or asynchronous primitives, while the behaviour of target artifacts remains unchanged and unaffected. As a result, every operation on an artifact should have a request / response structure: any invocation (request), once served, always implies a message of operation completed (response)—along with the result if needed—to be handled by the "operator" according to its nature: in case of an operation invoked by another artifact, in a completely asynchronous fashion—to ensure full uncoupling of artifact control; in case of an operation invoked by an agent, in either a synchronous or an asynchronous way according to the agent autonomous choice.

4.2 A&A ReSpecT*: The News*

Along the lines above, the original ReSpecT language has been revised and extended to follow the A&A perspective. The resulting core syntax of the

Table 1
Core syntax of A&A ReSpecT

```
      ⟨TCSpecification⟩ ::= {⟨SpecificationTuple⟩ .}
   ⟨SpecificationTuple⟩ ::= reaction( ⟨SimpleTCEvent⟩ , ⟨Guard⟩ , ⟨Reaction⟩ )
        ⟨SimpleTCEvent⟩ ::= ⟨SimpleTCOperation⟩ ( ⟨Tuple⟩ ) | time( ⟨Time⟩ )
                ⟨Guard⟩ ::= ⟨GuardPredicate⟩ | ( ⟨GuardPredicate⟩ {, ⟨GuardPredicate⟩} )
             ⟨Reaction⟩ ::= ⟨ReactionGoal⟩ | ( ⟨ReactionGoal⟩ {, ⟨ReactionGoal⟩} )
         ⟨ReactionGoal⟩ ::= ⟨TCOperation⟩ ( ⟨Tuple⟩ ) | ⟨EventObservation⟩ ( ⟨Tuple⟩ ) |
                            ⟨Computation⟩ | ( ⟨ReactionGoal⟩ ; ⟨ReactionGoal⟩ )
          ⟨TCOperation⟩ ::= ⟨SimpleTCOperation⟩ | ⟨TCLinkOperation⟩
      ⟨TCLinkOperation⟩ ::= ⟨TCIdentifier⟩ ? ⟨SimpleTCOperation⟩
    ⟨SimpleTCOperation⟩ ::= ⟨TCStateOperation⟩ | ⟨TCForgeOperation⟩
     ⟨TCStateOperation⟩ ::= in | inp | rd | rdp | out | no | get | set
     ⟨TCForgeOperation⟩ ::= ⟨TCStateOperation⟩_s
     ⟨EventObservation⟩ ::= ⟨EventView⟩_⟨EventInformation⟩
            ⟨EventView⟩ ::= current | event | start
     ⟨EventInformation⟩ ::= operation | tuple | source | target | time
       ⟨GuardPredicate⟩ ::= request | response | success | failure | endo | exo | intra | inter |
                            from_agent | to_agent | from_tc | to_tc | before( ⟨Time⟩ ) | after( ⟨Time⟩ )
                 ⟨Time⟩  is  a non-negative integer
                ⟨Tuple⟩  is  Prolog term
          ⟨Computation⟩  is  a Prolog-like goal performing arithmetic / logic computations
         ⟨TCIdentifier⟩ ::= ⟨TCName⟩ @ ⟨NetworkLocation⟩
               ⟨TCName⟩  is  a Prolog ground term
      ⟨NetworkLocation⟩  is  a Prolog string representing either an IP name or a DNS entry
```

newly-defined A&A ReSpecT is reported in Table 1.

The first apparent extension concerns the specification part of A&A ReSpecT: the `reaction` specification tuple has been extended to include a *guard specification*. Then, the behaviour of an A&A ReSpecT tuple centre is defined in terms of specification tuples of the form `reaction(E,G,R)`: such a tuple associates a reaction $R\theta$ to Ev if $\theta = mgu(E,Ev)$ and guard `G` is true. A guard is a sequence of guard predicates as defined by ⟨*GuardPredicate*⟩ in Table 1, whose semantics is defined in Table 5. A wide number of conditions over an event can now be checked before a reaction is triggered in a tuple centre: the event status, its source, its target, its time.

Along the same line, observation predicates have been generalised according to the new event model. Since an A&A ReSpecT event is defined according to the structure in Table 4, ⟨*EventObservation*⟩ predicates have now the form defined in Table 1: in particular, `event_` and `start_` predicates refer to the direct and "prime" cause of an event, respectively.

Another fundamental extension concerns uniformity of the operations upon tuple centres. Admissible primitives on a A&A ReSpecT tuple centres (⟨*TCStateOperation*⟩ in Table 1) can be invoked by an agent, but can also be used within reactions for a tuple centre to act on its state, or to act on another tuple centre state through a link operation. Also the semantics is essentially the same—with the only exception of the `in` and `rd` primitives, whose suspensive semantics is not preserved inside a reaction out of a link operation. Even more, the same class of predicates used for ordinary tuples can be used for specification tuples as well (⟨*TCForgeOperation*⟩ in Table 1), by simply adding the `_s` postfix—adding another dimension to uniformity.

Finally, A&A ReSpecT includes the first extension toward situatedness of coordination artifacts. In fact, following Timed ReSpecT [17], it includes time events (Table 4), timed reactions, as well as predicates to handle time (Table 1). More generally, further "situation" events could be envisioned, handling topology, or other environment issues: however, as shown in [17], handling time is one of the first, essential features for any real-world coordination model.

4.3 Distributed Dining Philosophers in A&A ReSpecT

In the classical Dining Philosopher problem, N philosopher agents share N chopsticks and a spaghetti bowl [4]. Each philosopher needs two chopsticks to eat, but each chopstick is shared by two adjacent philosophers: so, the two chopsticks have to be acquired atomically to avoid deadlock, and released atomically to ensure fairness.

In [14], we presented a ReSpecT-based implementation of the Dining Philosophers problem, where

- each philosopher agent acquires / releases his chopstick pairs as a tuple `chops(`i`,`j`)`: a philosopher willing to eat acquires the pair he needs by means of a single `in(chops(`i`,`j`))` operation from the `table` tuple centre, and releases it by means of a single `out(chops(`i`,`j`))` operation.
- individual chopsticks are represented as tuples of the kind `chop/1`: the result of philosopher's operations is the atomic removal / insertion of both `chop(`i`)` and `chop(`j`)` tuples from / in the `table` tuple centre.
- the `table` tuple centre works both as the knowledge repository for the table state—as a logic tuple space—and as the mediator between the two discrepant representations—as a programmable coordination artifact—through a suitable ReSpecT behaviour specification.

Here, we exploit some of the new features of A&A ReSpecT in order to implement a distributed version of the problem. The basic idea is to move the classical problem, which models multiple concurrent accesses to shared resources, to the distributed context, exploiting the intrinsic distribution promoted by the A&A meta-model in terms of agents and artifacts.

In the Distributed Dining Philosophers problem, N philosopher agents are supposed to be distributed around the network: each philosopher is assigned a *seat*, which is represented by a coordination artifact (a `seat(`i`,`j`)` tuple centre—meaning that `chops(`i`,`j`)` is the chopstick pair assigned to the

Table 2
Distributed Dining Philosophers: A&A ReSpecT code for `seat(i,j)` tuple centres.

```
reaction( out(wanna_eat), (request, from_agent), (          % (1)
    in(philosopher(thinking)), out(philosopher(waiting_to_eat)),
    current_target(seat(C1,C2)),
    table@node ? in(chops(C1,C2)) )
).
reaction( in(chops(C1,C2)), (response, inter, endo), (      % (2)
    in(philosopher(waiting_to_eat)), out(philosopher(eating)),
    out(chops(C1,C2)) )
).
reaction( out(wanna_think), (request, from_agent), (        % (3)
    in(philosopher(eating)), out(philosopher(waiting_to_think)),
    current_target(seat(C1,C2)), in(chops(C1,C2)),
    table@node ? out(chops(C1,C2)) )
).
reaction( out(chops(C1,C2)), (response, inter, endo), (     % (4)
    in(philosopher(waiting_to_think)), out(philosopher(thinking)) )
).
reaction( out(wanna_eat), (response, from_agent), in(wanna_eat) ).
reaction( out(wanna_think), (response, from_agent), in(wanna_think) ).
```

philosopher) located in the same TuCSoN node where the agent is. When a philosopher intends to eat / think, he just expresses his intention by emitting a tuple `wanna_eat` / `wanna_think` in his `seat(i,j)` tuple centre. In turn, each `seat(i,j)` tuple centre is in charge to handle its own agent intentions (to eat and to think), recording both the philosopher state (thinking, waiting to eat, eating, waiting to think) and the availability of chopsticks (`chops(i,j)` tuple), and interacting with the single `table` tuple centre (located in the `node` node), which holds and manages the `chop/1` tuples representing individual chopsticks on the table.

In all, the Distributed Dining Philosophers problem requires $N+1$ coordination artifacts, connected in a star network with the `table` tuple centre in the middle, and the N `seat(i,j)` tuple centres around, mediating between the philosophers and the table. Connections between distributed tuple centres—required to maintain consistency of the global system behaviour—are based on linkability predicates introduced in A&A ReSpecT (⟨*TCLinkOperation*⟩ in Table 1).

In particular, the A&A ReSpecT code in Table 2 is the same for all the `seat(i,j)` tuple centres—the specific chopstick pair is recorded in the tuple centre name, and retrieved (reaction 1 in Table 2) via one of the predicates for event observation extended in A&A ReSpecT (⟨*EventObservation*⟩ in Table 1). Reactions 1 and 2 in Table 2 deal with philosopher's intentions, either

Table 3
Distributed Dining Philosophers: A&A **ReSpecT** code for the `table` tuple centre.

```
reaction( out(chops(C1,C2)), (response, from_tc), (          % (1)
    in(chops(C1,C2)), out(chop(C1)), out(chop(C2)) )
).
reaction( in(chops(C1,C2)), (request, from_tc), (            % (2)
    out(required(C1,C2)) )
).
reaction( in(chops(C1,C2)), (response, from_tc), (           % (3)
    in(required(C1,C2)) )
).
reaction( out(required(C1,C2)), (response, intra, endo), (  % (4)
    in(chop(C1)), in(chop(C2)),  out(chops(C1,C2)) )
).
reaction( out(chop(C)), (response, intra, endo), (           % (5)
    rd(required(C,C2))
    in(chop(C)), in(chop(C2)),  out(chops(C,C2)) )
).
reaction( out(chop(C)), (response, intra, endo), (           % (6)
    rd(required(C1,C))
    in(chop(C1)), in(chop(C)),  out(chops(C1,C)) )
).
```

retrieving or restoring the proper `chops(`i`,`j`)` tuple from / to the `table` tuple centre at node `node`. The semantics of linkability predicates (all operations have their completion, but are asynchronous) allows for a 4-state representation of philosopher agents: *thinking*, *waiting to eat*, *eating*, *waiting to think*—where the transition states (*waiting to eat / think*) can be easily handled by reacting to the completion of `in(chops(`i`,`j`))?table@node` and `out(chops(`i`,`j`))?table@node` operations (reactions 2 and 4 in Table 2).

Even though not discussed here, the availability of situation predicates like the time predicates in A&A ReSpecT would allow for more complex coordination patterns including fault tolerance schemes. For instance, it would be easy to associate timeouts to the different states of agent philosophers—which is of paramount importance in a distributed, non-reliable environment, where even a simple `out` operation could easily fail, making chopsticks disappear in the vacuum. The association of timed reactions to the philosopher's transition states could then permit the recovery from faulty situations, and also to allow the introduction of timed coordination policies—as the ones discussed in [17].

The code in Table 3 is the behaviour specification for the `table` tuple centre, and is more or less the translation in A&A ReSpecT of the code discussed in [14]. Worth to note, then, just the uniform syntax for all the tuple centre operations, as well as the introduction of the notion of guard (along with guard

Table 4
Events in A&A ReSpecT

$$\begin{aligned}
\langle \mathit{GeneralTCEvent} \rangle &::= \langle \mathit{StartCause} \rangle \, \textbf{,} \, \langle \mathit{Cause} \rangle \, \textbf{,} \, \langle \mathit{TCCycleResult} \rangle \\
\langle \mathit{StartCause} \rangle \, , \langle \mathit{Cause} \rangle &::= \langle \mathit{SimpleTCEvent} \rangle \, \textbf{,} \, \langle \mathit{Source} \rangle \, \textbf{,} \, \langle \mathit{Target} \rangle \, \textbf{,} \, \langle \mathit{Time} \rangle \\
\langle \mathit{Source} \rangle \, , \langle \mathit{Target} \rangle &::= \langle \mathit{AgentIdentifier} \rangle \ | \ \langle \mathit{TCIdentifier} \rangle \\
\langle \mathit{AgentIdentifier} \rangle &::= \langle \mathit{AgentName} \rangle \, \texttt{@} \, \langle \mathit{NetworkLocation} \rangle \\
\langle \mathit{AgentName} \rangle &\text{ is a Prolog ground term} \\
\langle \mathit{TCCycleResult} \rangle &::= \langle \mathit{Tuple} \rangle
\end{aligned}$$

predicates) that makes A&A ReSpecT reactions more general and expressive.

5 A&A ReSpecT: The Semantics

According to the framework defined in [13], a coordination medium is suitable for an operational characterisation in terms of an interactive transition system, where the state of communication is the system state, some transitions are triggered by interaction events, and some transitions generate output events. So, in order to formally denote the behaviour of a coordination artifact like a A&A ReSpecT tuple centre, we should first define its notion of *admissible tuple centre event*, then define its behaviour in terms of a transition system.

Definition 5.1 [A&A ReSpecT Event] An *admissible tuple centre event* for A&A ReSpecT (A&A ReSpecT event in short) is defined according to the structure in Table 4. Such a structure also defines implicitly the way in which an A&A ReSpecT event is denoted: if ϵ is an A&A ReSpecT event, then $\epsilon.\mathit{Cause}.\mathit{Source}$ denotes the entity whose activity directly caused the event, $\epsilon.\mathit{TCCycleResult}$ denotes the result of the operation, and so on.

5.1 *Semantics of A&A* ReSpecT *Reactions*

An A&A ReSpecT tuple centre is basically a logic tuple space enhanced with a behaviour specification that defines how the tuple centre reacts to events. Then, once A&A ReSpecT events have been defined, the reaction model can be given, in terms of the reactions triggered by an A&A ReSpecT event ϵ.

Definition 5.2 [A&A ReSpecT Triggered Reaction Multiset] Given a tuple centre c and its behaviour specification Σ, if ϵ is an A&A ReSpecT event, then the multiset of the ϵ triggered reactions is defined as

$$\mathsf{Z}_{\Sigma}(\epsilon) ::= \biguplus_{\texttt{reaction(e,G,R)}\in\Sigma} (\epsilon, \texttt{R}\theta) \ | \ \theta = \mathit{Unify}(\epsilon, \texttt{e}) \neq \bot, \mathit{Guard}(\epsilon, \texttt{G})$$

There,

$$\mathit{Unify}(\epsilon, \texttt{e}) ::= \mathit{mgu}(\texttt{e}, \epsilon.\mathit{Event}.\mathit{SimpleTCEvent})$$

while $\mathit{Guard}(\epsilon, \texttt{G})$ is defined according to Table 5.

Table 5
Guard Predicates in A&A ReSpecT

Guard atom	True if
$Guard(\epsilon, (g, G))$	$Guard(\epsilon, g) \wedge Guard(\epsilon, G)$
$Guard(\epsilon, \texttt{endo})$	$\epsilon.Cause.Source = c$
$Guard(\epsilon, \texttt{exo})$	$\epsilon.Cause.Source \neq c$
$Guard(\epsilon, \texttt{intra})$	$\epsilon.Cause.Target = c$
$Guard(\epsilon, \texttt{inter})$	$\epsilon.Cause.Target \neq c$
$Guard(\epsilon, \texttt{from\_agent})$	$\epsilon.Cause.Source$ *is an agent*
$Guard(\epsilon, \texttt{to\_agent})$	$\epsilon.Cause.Target$ *is an agent*
$Guard(\epsilon, \texttt{from\_tc})$	$\epsilon.Cause.Source$ *is a tuple centre*
$Guard(\epsilon, \texttt{to\_tc})$	$\epsilon.Cause.Target$ *is a tuple centre*
$Guard(\epsilon, \texttt{before}(t))$	$\epsilon.Cause.Time < t$
$Guard(\epsilon, \texttt{after}(t))$	$\epsilon.Cause.Time > t$
$Guard(\epsilon, \texttt{request})$	$\epsilon.TCCycleResult$ *is undefined*
$Guard(\epsilon, \texttt{response})$	$\epsilon.TCCycleResult$ *is defined*
$Guard(\epsilon, \texttt{success})$	$\epsilon.TCCycleResult \neq \bot$
$Guard(\epsilon, \texttt{failure})$	$\epsilon.TCCycleResult = \bot$

Hypotheses: c is the reacting tuple centre; ϵ is an admissible A&A ReSpecT event; g is an atomic guard predicate; G is a sequence of atomic guard predicates; t is a non-negative integer.

Definition 5.3 [A&A ReSpecT Time-Triggered Reaction Multiset] Given a tuple centre c and its behaviour specification Σ, if nc is the local tuple centre time, then the multiset of the nc time-triggered reactions is defined as

$$\mathsf{Z}_\Sigma(nc) ::= \biguplus_{\texttt{reaction}(\texttt{time}(t),\texttt{G},\texttt{R})\in\mathsf{timed}(nc,\Sigma)} (\epsilon_t, \texttt{R}) \;\mid\; Guard(\epsilon_t, \texttt{G})$$

There,

$$\mathsf{timed}(nc, \Sigma) ::= \{\texttt{reaction}(\texttt{time}(t), \texttt{G}, \texttt{R}) \in \Sigma \;\mid\; t \leq nc\}$$

while $\epsilon_t ::= \langle \texttt{time}(t), c, c, t, \texttt{time}(t), c, c, t, t\rangle$, according to the event structure defined in Table 4.

So, given a tuple centre c at time nc with behaviour specification Σ, and an event ϵ, $\mathsf{Z}_\Sigma(\epsilon)$ denotes the multiset of triggered reactions caused by ϵ, while $\mathsf{Z}_\Sigma(nc)$ denotes the multiset of time-triggered reactions at time nc.

Once defined which reactions are triggered and when, the effects of reaction execution should be accounted for. This is encapsulated in the *reaction execution function.*

Definition 5.4 [Reaction Execution Function] Let R, R' be sequences of reaction goals, Tu, Tu' multi-sets of (ordinary) logic tuples, Σ, Σ' multi-sets of specification tuples, Re, Re' multi-sets of triggered reactions, and ϵ an A&A ReSpecT event, Out, Out' sequences of A&A ReSpecT events. A *reaction execution state* is then defined as a (labelled) quintuple $\langle R, Tu, \Sigma, Re, Out\rangle_\epsilon$,

Table 6
Operation predicate execution in A&A ReSpecT

Execution transition						
$\langle (r,R), Tu, \Sigma, Re, Out\rangle_\epsilon \longrightarrow_e \langle R', Tu', \Sigma', Re \cup \mathsf{Z}_\Sigma(\epsilon'), Out'\rangle_\epsilon$						
r	Tu'	Σ'	R'	$\epsilon'.Cause$	Out'	where
$op(\_T)?\mathtt{c}'$	Tu	Σ	R	$\langle op(\_T), c, c', nc\rangle$	$Out \cup \epsilon'$	
$\mathtt{out(T)}$	$Tu \cup \mathtt{T}$	Σ	R	$\langle \mathtt{out(T)}, c, c, nc\rangle$	Out	
$\mathtt{in(T)}$	Tu/tu	Σ	$R\theta$	$\langle \mathtt{in(T)}, c, c, nc\rangle$	Out	$\theta = mgu(tu, \mathtt{T})$
$\mathtt{inp(T)}$	Tu/tu	Σ	$R\theta$	$\langle \mathtt{inp(T)}, c, c, nc\rangle$	Out	$\theta = mgu(tu, \mathtt{T})$
$\mathtt{rd(T)}$	Tu	Σ	$R\theta$	$\langle \mathtt{rd(T)}, c, c, nc\rangle$	Out	$\theta = mgu(tu, \mathtt{T})$
$\mathtt{rdp(T)}$	Tu	Σ	$R\theta$	$\langle \mathtt{rdp(T)}, c, c, nc\rangle$	Out	$\theta = mgu(tu, \mathtt{T})$
$\mathtt{no(T)}$	Tu	Σ	R	$\langle \mathtt{no(T)}, c, c, nc\rangle$	Out	$\bot = mgu(tu, \mathtt{T})$
$\mathtt{set(TT)}$	$\mathtt{TT}$	Σ	R	$\langle \mathtt{set(TT)}, c, c, nc\rangle$	Out	
$\mathtt{get(TT)}$	Tu	Σ	$R\theta$	$\langle \mathtt{get(TT)}, c, c, nc\rangle$	Out	$\theta = mgu(Tu, \mathtt{TT})$
$\mathtt{out\_s(S)}$	Tu	$\Sigma \cup \mathtt{S}$	R	$\langle \mathtt{out\_s(S)}, c, c, nc\rangle$	Out	
$\mathtt{in\_s(S)}$	Tu	Σ/σ	$R\theta$	$\langle \mathtt{in\_s(S)}, c, c, nc\rangle$	Out	$\theta = mgu(\sigma, \mathtt{S})$
$\mathtt{inp\_s(S)}$	Tu	Σ/σ	$R\theta$	$\langle \mathtt{inp\_s(S)}, c, c, nc\rangle$	Out	$\theta = mgu(\sigma, \mathtt{S})$
$\mathtt{rd\_s(S)}$	Tu	Σ	$R\theta$	$\langle \mathtt{rd\_s(S)}, c, c, nc\rangle$	Out	$\theta = mgu(\sigma, \mathtt{S})$
$\mathtt{rdp\_s(S)}$	Tu	Σ	$R\theta$	$\langle \mathtt{rdp\_s(S)}, c, c, nc\rangle$	Out	$\theta = mgu(\sigma, \mathtt{S})$
$\mathtt{no\_s(S)}$	Tu	Σ	R	$\langle \mathtt{no\_s(S)}, c, c, nc\rangle$	Out	$\bot = mgu(\sigma, \mathtt{S})$
$\mathtt{set\_s(SS)}$	Tu	$\mathtt{SS}$	R	$\langle \mathtt{set\_s(SS)}, c, c, nc\rangle$	Out	
$\mathtt{get\_s(SS)}$	Tu	Σ	$R\theta$	$\langle \mathtt{get\_s(SS)}, c, c, nc\rangle$	Out	$\theta = mgu(\Sigma, \mathtt{SS})$

Hypotheses: ϵ, ϵ' are A&A ReSpecT events, such that $\epsilon'.StartCause = \epsilon.StartCause$ and $\epsilon'.TCCycleResult = \epsilon.TCCycleResult$; $tu \in Tu$ is a tuple; $\sigma \in \Sigma$ is a specification tuple; r is a reaction goal, R is a sequence of reaction goals; c, c' are tuple centres (c denotes the tuple centre currently in charge of the computation); nc is the time local to c when the execution takes place.

whereas a *reaction execution step* is a transition

$$\langle R, Tu, \Sigma, Re, Out\rangle_\epsilon \longrightarrow_e \langle R', Tu', \Sigma', Re', Out'\rangle_\epsilon$$

following the rules of Table 6 and Table 7. If a *reaction execution sequence* is a sequence of reaction execution steps, then

$$\langle R, Tu, \Sigma, Re, Out\rangle_\epsilon^*$$

denotes the *final state* of the reaction execution sequence whose initial state is $\langle R, Tu, \Sigma, Re, Out\rangle_\epsilon$, that is, the first state of the sequence for which no applicable rule exists in Table 6 and Table 7. Finally, if $\langle R, Tu, \Sigma, Re, Out\rangle_\epsilon^* = \langle R', Tu', \Sigma', Re', Out'\rangle_\epsilon$, then the *reaction execution function* E is defined as follows:

$$\mathsf{E}((\epsilon, R), Tu, \Sigma) ::= \begin{cases} (Tu', \Sigma', Re', Out') & \textit{if } \ R' = \emptyset \\ (Tu, \Sigma, \emptyset, \emptyset) & \textit{if } \ R' \neq \emptyset \end{cases}$$

To help intuition, at any step of a reaction execution sequence, R represents the reaction goals yet to be executed, Tu the current state of the space of ordinary tuples, Σ the current state of the space of specification tuples, Re the set of the reactions triggered by reaction goals already executed, Out the

Table 7
Observation predicate execution in A&A ReSpecT

Execution transition	
$\langle (r, R), \mathit{Tu}, \Sigma, \mathit{Re}, \mathit{Out}\rangle_{\epsilon} \longrightarrow_{e} \langle R\theta, \mathit{Tu}, \Sigma, \mathit{Re}, \mathit{Out}\rangle_{\epsilon}$	
r	where
`event_operation(Obs)`	$\theta = \mathit{mgu}(\epsilon.\mathit{Cause}.\mathit{SimpleTCEvent}.\mathit{SimpleTCOperation}, \texttt{Obs})$
`event_tuple(Obs)`	$\theta = \mathit{mgu}(\epsilon.\mathit{Cause}.\mathit{SimpleTCEvent}.\mathit{Tuple}, \texttt{Obs})$
`event_source(Obs)`	$\theta = \mathit{mgu}(\epsilon.\mathit{Cause}.\mathit{Source}, \texttt{Obs})$
`event_target(Obs)`	$\theta = \mathit{mgu}(\epsilon.\mathit{Cause}.\mathit{Target}, \texttt{Obs})$
`event_time(Obs)`	$\theta = \mathit{mgu}(\epsilon.\mathit{Cause}.\mathit{Time}, \texttt{Obs})$
`start_operation(Obs)`	$\theta = \mathit{mgu}(\epsilon.\mathit{StartCause}.\mathit{SimpleTCEvent}.\mathit{SimpleTCOperation}, \texttt{Obs})$
`start_tuple(Obs)`	$\theta = \mathit{mgu}(\epsilon.\mathit{StartCause}.\mathit{SimpleTCEvent}.\mathit{Tuple}, \texttt{Obs})$
`start_source(Obs)`	$\theta = \mathit{mgu}(\epsilon.\mathit{StartCause}.\mathit{Source}, \texttt{Obs})$
`start_target(Obs)`	$\theta = \mathit{mgu}(\epsilon.\mathit{StartCause}.\mathit{Target}, \texttt{Obs})$
`start_time(Obs)`	$\theta = \mathit{mgu}(\epsilon.\mathit{StartCause}.\mathit{Time}, \texttt{Obs})$
`current_operation(Obs)`	$\theta = \mathit{mgu}(\texttt{current\_operation}, \texttt{Obs})$
`current_tuple(Obs)`	$\theta = \mathit{mgu}(\texttt{Obs}, \texttt{Obs}) = \{\}$
`current_source(Obs)`	$\theta = \mathit{mgu}(c, \texttt{Obs})$
`current_target(Obs)`	$\theta = \mathit{mgu}(c, \texttt{Obs})$
`current_time(Obs)`	$\theta = \mathit{mgu}(\mathit{nc}, \texttt{Obs})$
Hypotheses: ϵ is an A&A ReSpecT event; r is a reaction goal, R is a sequence of reaction goals; c denotes the tuple centre currently in charge of the computation; nc is the time local to c when the execution takes place.	

sequence of events to be emitted at the end of the execution, whereas ϵ is the event initially triggering reaction execution. Correspondingly, the execution of a triggered reaction (ϵ, R) in an A&A ReSpecT tuple centre whose tuple space is Tu and whose behaviour specification is Σ is represented by a sequence whose initial state is $\langle R, \mathit{Tu}, \Sigma, \emptyset, \emptyset\rangle_{\epsilon}$. The above definition of E also accounts for the success/failure transactional semantics of A&A ReSpecT reactions: if the sequence of the operations to be executed is empty, then reaction R triggered by event ϵ has been executed successfully, and a new ordinary-tuple multiset Tu', a new specification-tuple multiset Σ', along with the newly-triggered reaction set Re' and the sequence of events to be emitted Out' are provided for updating the tuple centre state. Otherwise, the old multisets Tu and Σ are returned, no new reactions are triggered, and no events to be emitted are added—so that no changes occur in the tuple centre state.

Transitions occur according to the rules of Table 6 and Table 7, where all the symbols retain their usual meanings. The final state of a sequence is reached whenever either no reaction goals are still to be executed, or there is no applicable rule available. Since each step actually deletes one goal from a reaction, and the number of reaction goals is finite for any reaction, each reaction is guaranteed to be executed in a finite number of steps.

5.2 Behaviour of A&A ReSpecT *Tuple Centres*

The state of a A&A ReSpecT tuple centre is expressed as a labelled quadruple ${}^{\mathit{InQ}}\langle \mathit{Tu}, \Sigma, \mathit{Re}, \mathit{Op}\rangle_{n}^{\mathit{OutQ}}$. There, Tu and Σ are the multisets of the ordinary and

Table 8
Service transition in A&A ReSpecT

<table>
<tr><th colspan="5">Service transition</th></tr>
<tr><th colspan="5">$^{InQ}\langle Tu, \Sigma, \emptyset, Op \cup \epsilon \rangle_{n}^{OutQ} \longrightarrow_{s} {}^{InQ}\langle Tu', \Sigma', \mathsf{Z}_{\Sigma}(\epsilon') \cup \mathsf{Z}_{\Sigma}(n), Op \rangle_{n'}^{OutQ,\epsilon'}$</th></tr>
<tr><th>$\epsilon.Cause.SimpleTCEvent$</th><th>$Tu'$</th><th>$\Sigma'$</th><th>$res$</th><th>where</th></tr>
<tr><td>$\mathtt{out(T)}$</td><td>$Tu \cup \mathtt{T}$</td><td>Σ</td><td>$\mathtt{T}$</td><td></td></tr>
<tr><td>$\mathtt{in(T)}, \mathtt{inp(T)}$</td><td>$Tu/tu$</td><td>$\Sigma$</td><td>$\mathtt{T}\theta$</td><td>$\theta = mgu(tu, \mathtt{T})$</td></tr>
<tr><td>$\mathtt{rd(T)}, \mathtt{rdp(T)}$</td><td>$Tu$</td><td>$\Sigma$</td><td>$\mathtt{T}\theta$</td><td>$\theta = mgu(tu, \mathtt{T})$</td></tr>
<tr><td>$\mathtt{inp(T)}, \mathtt{rdp(T)}$</td><td>$Tu$</td><td>$\Sigma$</td><td>$\perp$</td><td>$\perp = mgu(tu, \mathtt{T})$</td></tr>
<tr><td>$\mathtt{no(T)}$</td><td>Tu</td><td>Σ</td><td>$\mathtt{T}$</td><td>$\perp = mgu(tu, \mathtt{T})$</td></tr>
<tr><td>$\mathtt{no(T)}$</td><td>Tu</td><td>Σ</td><td>$\perp$</td><td>$\theta = mgu(tu, \mathtt{T})$</td></tr>
<tr><td>$\mathtt{set(TT)}$</td><td>$\mathtt{TT}$</td><td>Σ</td><td>$\mathtt{TT}$</td><td></td></tr>
<tr><td>$\mathtt{get(TT)}$</td><td>Tu</td><td>Σ</td><td>$\mathtt{TT}\theta$</td><td>$\theta = mgu(Tu, \mathtt{TT})$</td></tr>
<tr><td>$\mathtt{get(TT)}$</td><td>Tu</td><td>Σ</td><td>$\perp$</td><td>$\perp = mgu(Tu, \mathtt{TT})$</td></tr>
<tr><td>$\mathtt{out\_s(S)}$</td><td>Tu</td><td>$\Sigma \cup \mathtt{S}$</td><td>$\mathtt{S}$</td><td></td></tr>
<tr><td>$\mathtt{in\_s(S)}, \mathtt{inp\_s(S)}$</td><td>$Tu$</td><td>$\Sigma/\sigma$</td><td>$\mathtt{S}\theta$</td><td>$\theta = mgu(\sigma, \mathtt{S})$</td></tr>
<tr><td>$\mathtt{rd\_s(S)}, \mathtt{rdp\_s(S)}$</td><td>$Tu$</td><td>$\Sigma$</td><td>$\mathtt{S}\theta$</td><td>$\theta = mgu(\sigma, \mathtt{S})$</td></tr>
<tr><td>$\mathtt{inp\_s(S)}, \mathtt{rdp\_s(S)}$</td><td>$Tu$</td><td>$\Sigma$</td><td>$\perp$</td><td>$\perp = mgu(\sigma, \mathtt{S})$</td></tr>
<tr><td>$\mathtt{no\_s(S)}$</td><td>Tu</td><td>Σ</td><td>$\mathtt{S}$</td><td>$\perp = mgu(\sigma, \mathtt{S})$</td></tr>
<tr><td>$\mathtt{no\_s(S)}$</td><td>Tu</td><td>Σ</td><td>$\perp$</td><td>$\theta = mgu(\sigma, \mathtt{S})$</td></tr>
<tr><td>$\mathtt{set\_s(SS)}$</td><td>Tu</td><td>$\mathtt{SS}$</td><td>$\mathtt{SS}$</td><td></td></tr>
<tr><td>$\mathtt{get\_s(SS)}$</td><td>Tu</td><td>Σ</td><td>$\mathtt{SS}\theta$</td><td>$\theta = mgu(\Sigma, \mathtt{SS})$</td></tr>
<tr><td>$\mathtt{get\_s(SS)}$</td><td>Tu</td><td>Σ</td><td>$\perp$</td><td>$\perp = mgu(\Sigma, \mathtt{SS})$</td></tr>
<tr><td colspan="5">Hypotheses: ϵ, ϵ' are A&A ReSpecT events: $\epsilon \in \mathsf{sat}(Op, Tu, \Sigma)$, ϵ' is such that $\epsilon'.StartCause = \epsilon.StartCause$, $\epsilon'.Cause = \epsilon.Cause$ and $\epsilon'.TCCycleResult = res$; $tu \in Tu$ is a tuple; $\sigma \in \Sigma$ is a specification tuple.</td></tr>
</table>

specification tuples in the tuple centre, respectively; *Op* is the multiset of the operations served waiting for a response; *InQ* and *OutQ* are the incoming and outgoing event queues, respectively; finally, n is the local tuple centre time.[4] *InQ* is a queue that is automatically extended whenever incoming events affect a tuple centre—so no special transitions are required for incoming events. Dually, *OutQ* is automatically emptied by emitting the outgoing events, with no need again of special transitions.

The operational behaviour of an A&A ReSpecT tuple centre whose state is $^{InQ}\langle Tu, \Sigma, Re, Op \rangle_{n}^{OutQ}$ can now be modelled in terms of a transition system with four kind of different transitions—below, in order of decreasing priority:

reaction When $Re \neq \emptyset$, triggered reactions in *Re* are executed through a *reaction* transition ($\longrightarrow_r$).

time When $Re = \emptyset$ and $\mathsf{timed}(n, \Sigma) \neq \emptyset$, timed reactions can trigger new reactions through a *time* transition ($\longrightarrow_t$).

service When $Re = \mathsf{timed}(n, \Sigma) = \emptyset$, and $\mathsf{sat}(Op, Tu, \Sigma) \neq \emptyset$, operations waiting for a response can be served through a *service* transition ($\longrightarrow_s$).[5]

---

[4] Whenever not needed by the context, *InQ*, *OutQ* and n could be dropped from the representation of a tuple centre state.

[5] $\mathsf{sat}(Op, Tu, \Sigma) \subseteq Op$ is the subset of the *Op* operations waiting for a response that can be actually served given the current state of the tuple centre.

Table 9
Log transition in A&A ReSpecT

Log transition		
$^{\epsilon,InQ}\langle Tu, \Sigma, \emptyset, Op\rangle_n^{OutQ} \longrightarrow_l {}^{InQ}\langle Tu, \Sigma, \mathsf{Z}_\Sigma(\epsilon) \cup \mathsf{Z}_\Sigma(n), Op'\rangle_{n'}^{OutQ}$		
$\epsilon.Cause.SimpleTCEvent$	$\epsilon.TCCycleResult$	Op'
$op(\_)$	undefined	$Op \cup \epsilon$
$op(\_)$	defined	Op
Hypotheses: ϵ is an A&A ReSpecT event.		

log When $Re = \mathsf{timed}(n, \Sigma) = \mathsf{sat}(Op, Tu, \Sigma) = \emptyset$, and $InQ \neq \emptyset$, operations queued in InQ can be "logged" by a tuple centre through a *log* transition ($\longrightarrow_l$)

Reaction transition works as follows:

$$^{InQ}\langle Tu, \Sigma, Re \cup re, Op\rangle_n^{OutQ} \longrightarrow_r {}^{InQ}\langle Tu', \Sigma', Re \cup Re', Op\rangle_{n'}^{OutQ, Out'}$$

where $\mathsf{E}(re, Tu, \Sigma) = (Tu', \Sigma', Re', Out')$—thus computing according to the semantics of reaction predicates presented in the previous subsection.

Time transition takes instead the following form:

$$^{InQ}\langle Tu, \Sigma, \emptyset, Op\rangle_n^{OutQ} \longrightarrow_t {}^{InQ}\langle Tu, \Sigma/\mathsf{timed}(n, \Sigma), \mathsf{Z}_\Sigma(n), Op\rangle_{n'}^{OutQ}$$

where past timed reactions ($\mathsf{timed}(n, \Sigma)$) are evaluated and then discarded ($\Sigma/\mathsf{timed}(n, \Sigma)$), and possibly generate some time-triggered reactions ($\mathsf{Z}_\Sigma(n)$).

The more articulated service and log transitions are regulated according to Table 8 and Table 9, respectively.

6 Related Works

With respect to the original formulation [15], A&A ReSpecT, as presented in this paper, is largely changed and extended. First, the main things left unchanged are *(i)* the basic reaction model, with the two-level atomicity of reaction execution (system level and agent level), and *(i)* the logic-based syntax, with event descriptions unifying with reaction heads and the reaction bodies built as sequences of Prolog-like atoms.

Some syntax modifications have addressed known limitations in the original ReSpecT. The introduction of guards, for instance, has improved on the expressiveness of reactions, allowing programmers to minimise the number of unnecessarily-triggered reactions; a number of expressive guard predicates are introduced in A&A ReSpecT for this purpose. Also, the syntax of primitives is now uniform: the same primitive invocations can be used by agents on a tuple centre, and by a programmer within a reaction, either to access and change the internal tuple centre state, or to operate on other tuple centres; even the primivites for accessing and changing a tuple centre behaviour specification have essentially the same form. Even more, semantics of any agent invocation is now in a sense uniform: all the coordination primitives that an agent can invoke on a tuple centre have the same request / response behaviour (not only `ins` and `rds`, but also `outs`), and reactions can then be designed having in mind the same conceptual structure for any primitive involved.

More generally, the adoption of A&A as the underlying meta-model has led to a number of new features in the novel A&A ReSpecT. First of all, an extended event model is defined for A&A ReSpecT, which encompasses the A&A meta-model and its general event model. Then, linkability of artifacts is developed and extended up to its maximum reach: any tuple centre operation in a reaction can now be invoked to be executed either within the tuple centre itself, or on any other tuple centre.[6] Finally, situatedness of artifacts is recognised here as a general issue, which encompasses timed artifacts and their computational model. With respect to the original formulation of Timed ReSpecT [17], the model is then made more general, and also fully formalised within the overall A&A ReSpecT framework.

Linkability in its most general acceptation is not strictly a new idea in the field of coordination models and languages. The most prominent example is Reo [1], where channel composition is one of the most important and relevant features. Also, Reo has been recently experimented explicitly in the MAS field [2]. However, Linda-based approaches better cope with agent autonomy, since coordination is not forced upon the agents participating to the workflow, but is instead provided them as a service [29].

In the context of Linda-based models, to the best of our knowledge, only Lime [23] could exhibit some sort of mechanism for tuple space composition. However, such a mechanism is essentially implicit, and does not allow for the explicit control allowed instead by A&A ReSpecT linkability primitives.

7 Conclusions & Future Work

In this paper, we adopt the A&A (agents & artifacts) meta-model for MAS, and recast the ReSpecT language for programming the behaviour of tuple centres, and its formal model as well, according to the A&A perspective. The resulting model and language, called A&A ReSpecT, is introduced: the new syntax is defined, an example (the Distributed Dining Philosophers) is discussed, and the formal semantics of A&A ReSpecT is provided.

While implementation of the new A&A ReSpecT is underway, along with the new version of the TuCSoN infrastructure for MAS coordination, in the future we plan to experiment with A&A ReSpecT in the many domains where the original ReSpecT is already used—from e-learning to workflow management systems and case-handling, from simulation to self-organising systems. Meanwhile, we mean to further explore the issue of *situatedness* of coordina-

---

[6] The first element of linkability in ReSpecT was introduced in [25], in the limited form of an `out_tc` predicate, and used in [20] for distributed workflow.

tion artifacts, by extending the ability of tuple centres to react to environment events. In the version of A&A ReSpecT presented here, in fact, only time events are accounted for and treated: more general forms of environment events (like topological ones, for instance) should be instead made available and properly manageable.

References

[1] Farhad Arbab. Reo: A channel-based coordination model for component composition. *Mathematical Structures in Computer Science*, 14:329–366, 2004.

[2] Mehdi Dastani and OSGi – Open Services Gateway Initiatives Consortium. *Coordination and Composition of Multi-Agent Systems.* Invited talk, 1st International Workshop on Coordination and Organisation (CoOrg 2005), COORDINATION 2005, Namur, Belgium, 23 April 2005.

[3] Enrico Denti, Antonio Natali, and Andrea Omicini. On the expressive power of a language for programming coordination media. In *1998 ACM Symposium on Applied Computing (SAC'98)*, pages 169–177, Atlanta, GA, USA, 27 February–1 March 1998. ACM. Special Track on Coordination Models, Languages and Applications.

[4] Edsger W. Dijkstra. *Co-operating sequential processes.* Academic Press, London, 1965.

[5] Marc Esteva, Bruno Rosell, Juan Antonio Rodríguez-Aguilar, and Josep Lluís Arcos. Ameli: An agent-based middleware for electronic institutions. In Nicholas R. Jennings, Carles Sierra, Liz Sonenberg, and Milind Tambe, editors, *3rd international Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, volume 1, pages 236–243, New York, USA, 19–23July 2004. ACM.

[6] Foundation for Intelligent Physical Agents. FIPA home page. http://www.fipa.org.

[7] David Gelernter. Generative communication in Linda. *ACM Transactions on Programming Languages and Systems*, 7(1):80–112, 1985.

[8] David Kirsh. Distributed cognition, coordination and environment design. In *European Conference on Cognitive Science*, pages 1–11, 1999.

[9] Aleksie Nikolaevich Leontjev. *Activity, Consciousness, and Personality.* Prentice Hall, 1978.

[10] Ambra Molesini, Andrea Omicini, Enrico Denti, and Alessandro Ricci. SODA: A roadmap to artefacts. In Oğuz Dikenelli, Marie-Pierre Gleizes, and Alessandro Ricci, editors, *Engineering Societies in the Agents World VI*, volume 3963 of *LNAI*, pages 49–62. Springer, June 2006. 6th International Workshop (ESAW 2005), Kuşadası, Aydın, Turkey, 26–28 October 2005. Revised, Selected & Invited Papers.

[11] Sara Montagna, Alessandro Ricci, and Andrea Omicini. Agents & Artifacts for Systems Biology: Toward a framework based on TuCSoN. In Alessandro Genco, Antonio Gentile, and Salvatore Sorce, editors, *Industrial Simulation Conference 2006 (ISC 2006)*, pages 25–32, Palermo, Italy, 5–7 June 2006. EUROSIS (The European Simulation Society) & ETI (The European Technology Institute).

[12] Bonnie Nardi, editor. *Context and Consciousness: Activity Theory and Human-Computer Interaction.* MIT Press, 1996.

[13] Andrea Omicini. On the semantics of tuple-based coordination models. In *1999 ACM Symposium on Applied Computing (SAC'99)*, pages 175–182, San Antonio, TX, USA, 28 February– 2March 1999. ACM. Special Track on Coordination Models, Languages and Applications.

[14] Andrea Omicini and Enrico Denti. Formal ReSpecT. *Electronic Notes in Theoretical Computer Science*, 48:179–196, June 2001. Declarative Programming – Selected Papers from AGP 2000, La Habana, Cuba, 4–6 December2000.

[15] Andrea Omicini and Enrico Denti. From tuple spaces to tuple centres. *Science of Computer Programming*, 41(3):277–294, November 2001.

[16] Andrea Omicini, Sascha Ossowski, and Alessandro Ricci. Coordination infrastructures in the engineering of multiagent systems. In Federico Bergenti, Marie-Pierre Gleizes, and Franco Zambonelli, editors, *Methodologies and Software Engineering for Agent Systems: The Agent-Oriented Software Engineering Handbook*, volume 11 of *Multiagent Systems, Artificial Societies, and Simulated Organizations*, chapter 14, pages 273–296. Kluwer Academic Publishers, June 2004.

[17] Andrea Omicini, Alessandro Ricci, and Mirko Viroli. Time-aware coordination in ReSpecT. In Jean-Marie Jacquet and Gian Pietro Picco, editors, *Coordination Models and Languages*, volume 3454 of *LNCS*, pages 268–282. Springer-Verlag, April 2005. 7th International Conference (COORDINATION 2005), Namur, Belgium, 20–23 April 2005. Proceedings.

[18] Andrea Omicini, Alessandro Ricci, and Mirko Viroli. *Agens Faber*: Toward a theory of artefacts for MAS. *Electronic Notes in Theoretical Computer Sciences*, 150(3):21–36, 29 May 2006. 1st International Workshop "Coordination and Organization" (CoOrg 2005), COORDINATION 2005, Namur, Belgium, 22 April 2005. Proceedings.

[19] Andrea Omicini, Alessandro Ricci, and Mirko Viroli. Coordination artifacts as first-class abstractions for MAS engineering: State of the research. In Alessandro F. Garcia, Ricardo Choren, Carlos Lucena, Paolo Giorgini, Tom Holvoet, and Alexander Romanovsky, editors, *Software Engineering for Multi-Agent Systems IV: Research Issues and Practical Applications*, volume 3914 of *LNAI*, pages 71–90. Springer, April 2006. Invited Paper.

[20] Andrea Omicini, Alessandro Ricci, and Nicola Zaghini. Distributed workflow upon linkable coordination artifacts. In Paolo Ciancarini and Herbert Wiklicky,

editors, *Coordination Models and Languages*, volume 4038 of *LNCS*, pages 228–246. Springer, June 2006. 8th International Conference (COORDINATION 2006), Bologna, Italy, 14–16 June 2006. Proceedings.

[21] Andrea Omicini and Franco Zambonelli. Coordination for Internet application development. *Autonomous Agents and Multi-Agent Systems*, 2(3):251–269, September 1999. Special Issue: Coordination Mechanisms for Web Agents.

[22] H. Van Dyke Parunak, Sven Brueckner, and John Sauter. Digital pheromone mechanisms for coordination of unmanned vehicles. In *1st International Joint Conference on Autonomous Agents and Multiagent Systems AAMAS'02*, pages 449–450. ACM Press, 2002.

[23] Gian Pietro Picco, Amy L. Murphy, and Gruia-Catalin Roman. LIME: Linda Meets Mobility. In David Garlan, editor, *21st International Conference on Software Engineering (ICSE'99)*, pages 368–377, Los Angeles, CA, USA, May 1999. ACM Press.

[24] Alessandro Ricci, Andrea Omicini, Enrico Denti, and Marco Cadoli. Activity Theory as a framework for MAS coordination. In Paolo Petta, Robert Tolksdorf, and Franco Zambonelli, editors, *Engineering Societies in the Agents World III*, volume 2577 of *LNCS*, pages 96–110. Springer-Verlag, April 2003. 3rd International Workshop (ESAW 2002), Madrid, Spain, 16–17 September 2002. Revised Papers.

[25] Alessandro Ricci, Andrea Omicini, and Mirko Viroli. Extending ReSpecT for multiple coordination flows. In Hamid R. Arabnia, editor, *International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'02)*, volume III, pages 1407–1413, Las Vegas, NV, USA, 24–27 July 2002. CSREA Press.

[26] Alessandro Ricci, Mirko Viroli, and Andrea Omicini. Programming MAS with artifacts. In Rafael P. Bordini, Mehdi Dastani, Jürgen Dix, and Amal El Fallah Seghrouchni, editors, *Programming Multi-Agent Systems*, volume 3862 of *LNAI*, pages 206–221. Springer, March 2006. 3rd International Workshop (PROMAS 2005), AAMAS 2005, Utrecht, The Netherlands, 26 July 2005. Revised and Invited Papers.

[27] Tarja Susi and Tom Ziemke. Social cognition, artefacts, and stigmergy: A comparative analysis of theoretical frameworks for the understanding of artefact-mediated collaborative activity. *Cognitive Systems Research*, 2(4):273–290, December 2001.

[28] TuCSoN home page. http://tucson.alice.unibo.it.

[29] Mirko Viroli and Andrea Omicini. Coordination as a service. *Fundamenta Informaticae*, 73(4):507–534, 2006. Special Issue: Best papers of FOCLASA 2002.

[30] Franco Zambonelli and Andrea Omicini. Challenges and research directions in agent-oriented software engineering. *Autonomous Agents and Multi-Agent*

Systems, 9(3):253–283, November 2004. Special Issue: Challenges for Agent-Based Computing.