

Interaction & Coordination in Distributed Systems

Distributed Systems
Sistemi Distribuiti

Andrea Omicini
andrea.omicini@unibo.it

Dipartimento di Informatica – Scienza e Ingegneria (DISI)
ALMA MATER STUDIORUM – Università di Bologna a Cesena

Academic Year 2013/2014



Outline

- 1 Distributed Systems Engineering & Interaction
- 2 Interaction & Coordination
- 3 Enabling vs. Governing Interaction
- 4 Classes of Coordination Models

Outline

- 1 Distributed Systems Engineering & Interaction
- 2 Interaction & Coordination
- 3 Enabling vs. Governing Interaction
- 4 Classes of Coordination Models

Scenarios for Concurrent / Distributed Systems

Issues

- Concurrency / Parallelism
 - Multiple independent activities / loci of control
 - Active simultaneously
 - Processes, threads, actors, active objects, agents. . .
- Distribution
 - Activities running on different and heterogeneous execution contexts (machines, devices, . . .)
- “Social” Interaction
 - Dependencies among activities
 - Collective goals involving activities coordination / cooperation
- “Environmental” Interaction
 - Interaction with external resources
 - Interaction within the time-space fabric

Basic Engineering Principles

Principles

- Abstraction
 - Problems should be faced / represented at the most suitable level of abstraction
 - Resulting “abstractions” should be expressive enough to capture the most relevant problems
 - Conceptual integrity
- Locality & encapsulation
 - Design abstractions should embody the solutions corresponding to the domain entities they represent
- Run-time vs. design-time abstractions
 - Incremental change / evolution
 - On-line engineering [Fredriksson and Gustavsson, 2004]
 - (Cognitive) Self-organising systems [Omicini, 2012]

Which Components?

Open systems

- No hypothesis on the component's life & behaviour


Distributed systems

- No hypothesis on the component's location & motion

Heterogeneous systems

- No hypothesis on the component's nature & structure

The Space of Interaction



Components of an Interactive System

What is a component of an interactive system?

- A computational abstraction characterised by
 - an independent computational activity
 - I/O capabilities
- Two independent dimensions
 - elaboration / *computation*
 - *interaction*

(Non) Algorithmic Computation I

Elaboration / computation

- Turing Machine (TM)
 - gets an input, elaborates it, throws an output
 - no interaction during computation
- Black-box algorithms
- Church's Thesis and computable functions
 - in short, a function is *algorithmically computable* iff can be computed by a TM
 - so, all computable functions are computable by a TM

(Non) Algorithmic Computation II

The power of interaction [Wegner and Goldin, 2003]

Real computational systems are not rational agents that take inputs, compute logically, and produce outputs. . . It is hard to draw the line at what is intelligence and what is environmental interaction. In a sense, it does not really matter which is which, as all intelligent systems must be situated in some world or other if they are to be useful entities. [Brooks, 1991]

. . . a theory of concurrency and interaction requires a new conceptual framework, not just a refinement of what we find natural for sequential [algorithmic] computing. [Milner, 1993]

(Non) Algorithmic Computation III

Beyond Turing Machines

- Turing's *choice machines* and *unorganised machines* [Wegner and Goldin, 2003]
- Wegner's Interaction Machines [Goldin et al., 2006]
- Examples: AGV, Chess oracle [Wegner, 1997]

Basics of Interaction

Component model

A simple component exhibits

Computation Inner behaviour of a component

Interaction Observable behaviour of a component as *input* and *output*

Coupling across component's boundaries

- Control?
- Information
- Time & Space – internal / computational vs. external / physical

Information-driven interaction

Output shows part of its state outside

Input bounds a portion of its own state to the outside

Compositionality vs. Non-compositionality

Compositionality

- Sequential composition $P1; P2$
- $behaviour(P1; P2) = behaviour(P1) + behaviour(P2)$

Non-compositionality

- Interactive composition $P1|P2$
- $behaviour(P1|P2) = behaviour(P1) + behaviour(P2) + \mathbf{interaction}(P1, P2)$
- Interactive composition is more than the sum of its parts

Non-compositionality

Issues

- Compositionality vs. formalisability
 - A notion of formal model is required for stating any compositional property
 - However, formalisability does not require compositionality, and does not imply predictability
 - *Partial formalisability* may allow for proof of properties, and for partial predictability
- Emergent behaviours
 - Fully-predictable / formalisable systems do not allow by definition for emergent behaviours
- Formalisability vs. expressiveness
 - Less / more formalisable systems are (respectively) more / less expressive in terms of potential behaviours

Outline

- 1 Distributed Systems Engineering & Interaction
- 2 Interaction & Coordination**
- 3 Enabling vs. Governing Interaction
- 4 Classes of Coordination Models



Coordination in Distributed Programming

Coordination model as a glue

A coordination model is the glue that binds separate activities into an ensemble [Gelernter and Carriero, 1992]

Coordination model as an agent interaction framework


A coordination model provides a framework in which the interaction of active and independent entities called agents can be expressed [Ciancarini, 1996]

Issues for a coordination model

A coordination model should cover the issues of creation and destruction of agents, communication among agents, and spatial distribution of agents, as well as synchronization and distribution of their actions over time [Ciancarini, 1996]

What is Coordination?

Ruling the space of interaction



New Perspective on Computational Systems

Programming languages

- Interaction as an orthogonal dimension
- Languages for interaction / coordination

Software engineering

- Interaction as an independent design dimension
- Coordination patterns


Artificial intelligence

- Interaction as a new source for intelligence
- Social intelligence

Coordination: Sketching a Meta-model

The

- “fills” the interaction space
- enables / promotes / governs the admissible / desirable / required interactions among the interacting entities
- according to some *coordination laws*
 - enacted by the behaviour of the medium
 - defining the semantics of coordination



Coordination: A Meta-model [Ciancarini, 1996]

A constructive approach

Which are the components of a coordination system?

Coordination entities Entities whose mutual interaction is ruled by the model, also called the *coordinables*

Coordination media Abstractions enabling and ruling interaction among coordinables

Coordination laws Laws ruling the observable behaviour of coordination media and coordinables, and their interaction as well

Coordinables

Original definition [Ciancarini, 1996]

These are the entity types that are coordinated. These could be Unix-like processes, threads, concurrent objects and the like, and even users.

examples Processes, threads, objects, human users, agents, ...

focus Observable behaviour of the coordinables

question Are we anyhow concerned here with the internal machinery / functioning of the coordinable, in principle?

→ This issue will be clear when comparing Linda & TuCSoN agents

Coordination Media

Original definition [Ciancarini, 1996]

These are the media making communication among the agents possible. Moreover, a coordination medium can serve to aggregate agents that should be manipulated as a whole. Examples are classic media such as semaphores, monitors, or channels, or more complex media such as tuple spaces, blackboards, pipelines, and the like.

examples Semaphors, monitors, channels, tuple spaces, blackboards, pipes, ...

focus The core around which the components of the system are organised

question Which are the possible computational models for coordination media?

→ *This issue will be clear when comparing Linda tuple spaces & ReSpecT tuple centres*

Coordination Laws I

Original definition [Ciancarini, 1996]

A coordination model should dictate a number of laws to describe how agents coordinate themselves through the given coordination media and using a number of coordination primitives. Examples are laws that enact either synchronous or asynchronous behaviors or exploit explicit or implicit naming schemes for coordination entities.

Coordination Laws II

- Coordination laws rule the observable behaviour of coordination media and coordinables, as well as their interaction
 - a notion of (*admissible interaction*) *event* is required to define coordination laws
 - The interaction events are (also) expressed in terms of
 - the *communication language*, as the syntax used to express and exchange data structures
- examples* tuples, XML elements, FOL terms, (Java) objects, ...
- the *coordination language*, as the set of the admissible interaction primitives, along with their semantics
- examples* in/out/rd (Linda), send/receive (channels), push/pull (pipes), ...

Outline

- 1 Distributed Systems Engineering & Interaction
- 2 Interaction & Coordination
- 3 Enabling vs. Governing Interaction**
- 4 Classes of Coordination Models

Toward a Notion of Coordination Model

What do we ask to a coordination model?

- to provide high-level *abstractions* and powerful *mechanisms* for distributed system engineering
- to enable and promote the construction of *open, distributed, heterogeneous* systems
- to intrinsically *add properties* to systems independently of components
 - e.g. flexibility, control, intelligence, ...

Examples of Coordination Mechanisms I

Message passing

- communication among peers
- no abstractions apart from message
- no limitations
 - the notion of *protocol* could be added as a coordination abstraction
- no intrinsic model of coordination
- any pattern of coordination can be superimposed – again, protocols

Examples of Coordination Mechanisms II

Agent Communication Languages

- Goal: promote information exchange
- Examples: Arcol, KQML
- Standard: FIPA ACL
- Semantics: ontologies
- *Enabling communication*
 - ACLs *create* the space of inter-agent communication
 - they do not allow to *constrain* it
- No “real” coordination, again, if not with protocols

Examples of Coordination Mechanisms III

Service-Oriented Architectures

- Basic abstraction: service
- Basic pattern: Service request / response
- Several standards
- Very simple pattern of coordination

Examples of Coordination Mechanisms IV

Web Server

- Basic abstraction: resource (REST/ROA)
- Basic pattern: Resource request / representation / response
- Several standards
- Again, a very simple pattern of coordination
- Generally speaking, objects, HTTP, applets, JavaScript with AJAX, user interface
 - a multi-coordinated systems
 - “spaghetti-coordination”, no value added from composition
- How can we “fill” the space of interaction to add value to systems?
 - so, how do we get value from coordination?

Examples of Coordination Mechanisms V

Middleware

- Goal: to provide global properties across distributed systems
- Idea: fill the space of interaction with abstractions and shared features
 - interoperability, security, transactionality, ...
- Middleware can contain coordination abstractions
 - but, it can contain anything, so we need to look at *specific* middleware

Examples of Coordination Mechanisms VI

CORBA

- Goal: managing object interaction across a distributed systems in a transparent way
- Key features: ORB, IDL, CORBAServices. . .
- However, no model for coordination
 - just the client-servant pattern
- However, it can provide a shared support for any coordination abstraction or pattern

Enabling vs. Governing Interaction I

Enabling interaction

- ACL, middleware, mediators. . .
- enabling communication
- enabling components interoperation
- no models for coordination of components
 - no rules on what components should (not) say and do at any given moment, depending on what other components say and do, and on what happens inside and outside the system

Enabling vs. Governing Interaction II

Governing interaction

- ruling communication
- providing concepts, abstractions, models, mechanisms for meaningful component integration
- governing mutual component interaction, and environment-component interaction
- in general, a model that does
 - rule what components should (not) say and do at any given moment
 - depending on what other components say and do, and on what happens inside and outside the system

Outline

- 1 Distributed Systems Engineering & Interaction
- 2 Interaction & Coordination
- 3 Enabling vs. Governing Interaction
- 4 Classes of Coordination Models**

Two Classes for Coordination Models

Control-oriented vs. Data-oriented Models

- Control-driven vs. Data-driven Models
[Papadopoulos and Arbab, 1998]

Control-oriented Focus on the *acts* of communication

Data-oriented Focus on the *information* exchanged during communication

- Several surveys, no time enough here
- Are these really *classes*?
 - actually, better to take this as a criterion to observe coordination models, rather than to separate them

Control-oriented Models I

Processes as black boxes


- I/O ports
- events & signals on state

Coordinators...

- ...create coordinated processes as well as communication channels
- ...determine and change the topology of communication
- Hierarchies of coordinables / coordinators are possible

Control-oriented Models II

Coordinators as meta-level communication components



Control-oriented Models III

General features

- High flexibility, high control
- Separation between communication / coordination and computation / elaboration
- Examples
 - RAPIDE [Luckham et al., 1995]
 - Manifold [Arbab et al., 1993]
 - ConCoord [Holzbacher, 1996]
 - Reo [Arbab, 2004, Dastani et al., 2005]

A Classical Example: Manifold [Arbab et al., 1993]

Main features

- coordinators
- control-driven evolution
 - events without parameters
- stateful communication
- coordination via topology
- fine-grained coordination
- typical example: sort-merge

Control-oriented Models: Impact on Design

Which abstractions?

- Producer-consumer pattern
- Point-to-point communication
- Coordinator
- Coordination as configuration of topology

Which systems?

- Fine-grained granularity
- Fine-tuned control
- Good for small-scale, closed systems

An Evolutionary Pattern?

Paradigms of sequential programming

- Imperative programming with “goto”
- Structured programming (procedure-oriented)
- Object-oriented programming (data-oriented)

Paradigms of coordination programming

- Message-passing coordination
- Control-oriented coordination
- Data-oriented coordination

Data-oriented Models I

Communication channel

- Shared memory abstraction
- Stateful channel

Processes


- Emitting / receiving data / information

Coordination

- Access / change / synchronise on shared data

Data-oriented Models II

Shared dataspace: constraint on communication



Data-oriented Models

General features

- Expressive communication abstraction
- information-based design
- Possible spatio-temporal uncoupling
- No control means no flexibility??
- Examples
 - Gamma / Chemical coordination
 - Linda & friends / tuple-based coordination

Summing Up

Coordination for distributed system engineering

- Engineering the space of interaction among components

Coordination as governing interaction

- Enabling vs. governing

Classes and features of coordination models

- Control-oriented vs. data-oriented models

References I



Arbab, F. (2004).

Reo: A channel-based coordination model for component composition.

Mathematical Structures in Computer Science, 14:329–366.



Arbab, F., Herman, I., and Spilling, P. (1993).

An overview of MANIFOLD and its implementation.

Concurrency: Practice and Experience, 5(1):23–70.



Brooks, R. A. (1991).

Intelligence without reason.

In Mylopoulos, J. and Reiter, R., editors, *12th International Joint Conference on Artificial Intelligence (IJCAI 1991)*, volume 1, pages 569–595, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

References II



Ciancarini, P. (1996).

Coordination models and languages as software integrators.

ACM Computing Surveys, 28(2):300–302.



Dastani, M., Arbab, F., and de Boer, F. S. (2005).

Coordination and composition in multi-agent systems.

In Dignum, F., Dignum, V., Koenig, S., Kraus, S., Singh, M. P., and Wooldridge, M. J., editors, *4rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005)*, pages 439–446, Utrecht, The Netherlands. ACM.

References III

- 
- Fredriksson, M. and Gustavsson, R. (2004).
Online engineering and open computational systems.
In Bergenti, F., Gleizes, M.-P., and Zambonelli, F., editors,
*Methodologies and Software Engineering for Agent Systems: The
Agent-Oriented Software Engineering Handbook*, volume 11 of
Multiagent Systems, Artificial Societies, and Simulated Organization,
pages 377–388. Kluwer Academic Publishers.
- 
- Gelernter, D. and Carriero, N. (1992).
Coordination languages and their significance.
Communications of the ACM, 35(2):97–107.
- 
- Goldin, D. Q., Smolka, S. A., and Wegner, P., editors (2006).
Interactive Computation: The New Paradigm.
Springer.

References IV



Holzbacher, A.-A. (1996).

A software environment for concurrent coordinated programming.
In Ciancarini, P. and Hankin, C., editors, *Coordination Languages and Models*, volume 1061 of *LNCS*, pages 249–266. Springer-Verlag.
1st International Conference (COORDINATION '96) Cesena, Italy, April 15–17, 1996.



Luckham, D. C., Kenney, J. J., Augustin, L. M., Vera, J., Bryan, D., and Mann, W. (1995).

Specification and analysis of system architecture using Rapide.
IEEE Transactions on Software Engineering, 21(4):336–354.



Milner, R. (1993).

Elements of interaction: Turing award lecture.
Communications of the ACM, 36(1):78–89.

References V



Omicini, A. (2012).

Agents writing on walls: Cognitive stigmergy and beyond.

In Paglieri, F., Tummolini, L., Falcone, R., and Miceli, M., editors, *The Goals of Cognition. Essays in Honor of Cristiano Castelfranchi*, volume 20 of *Tributes*, chapter 29, pages 543–556. College Publications, London.



Papadopoulos, G. A. and Arbab, F. (1998).

Coordination models and languages.

In Zelkowitz, M. V., editor, *The Engineering of Large Systems*, volume 46 of *Advances in Computers*, pages 329–400. Academic Press.



Wegner, P. (1997).

Why interaction is more powerful than algorithms.

Communications of the ACM, 40(5):80–91.

References VI

-  Wegner, P. and Goldin, D. (2003).
Computation beyond Turing machines.
Communications of the ACM, 46(4):100–102.



Interaction & Coordination in Distributed Systems

Distributed Systems
Sistemi Distribuiti

Andrea Omicini
andrea.omicini@unibo.it

Dipartimento di Informatica – Scienza e Ingegneria (DISI)
ALMA MATER STUDIORUM – Università di Bologna a Cesena

Academic Year 2013/2014

