

Programming Collective Adaptive Systems by relying on Attribute-based Communication

Rocco De Nicola

Joint work with

M. Loreti, Y. Abd Alrahman and Tan Duong

IFIP W.G. 2.2 - Wien - September 2019

IMT- School for Advanced Studies - Lucca

Languages for supporting the engineering of different classes of modern distributed systems

- ▶ network-aware programming
- ▶ service-oriented computing
- ▶ autonomic computing.

Programming collective adaptive systems

AbC: A calculus for Attribute based Programming

- ▶ Syntax
- ▶ Semantics
- ▶ Implementations
- ▶ Verification

Why a Languages based Approach

Languages

Languages play a key role in the engineering of systems

- ▶ Systems must be specified as naturally as possible
- ▶ Distinctive aspects of the domain need to be first-class citizens
- ▶ Intuitive/concise specifications are possible and encodings can be avoided

Models

Models strictly related to languages are at least as important for effective analysis

- ▶ high-level abstract models guarantee feasible investigations
- ▶ the scrutiny of results (e.g., counterexample) based on system features, rather than on their low-level representation, guarantees better feedbacks.

Major challenge

The big challenge for language designers is to devise appropriate abstractions and linguistic primitives to deal with the specificities of the systems under consideration while relying on an appropriate semantic model.

A possible approach

Combined use of formal methods with model-driven software engineering. Key ingredients are

1. A specification language equipped with a formal semantics
2. A programming framework with associated runtime environment
3. A number of verification techniques and associated tools

Our Contributions: A timeline

Service-oriented computing

- services composition
- heterogeneous components
- code reuse
- interoperability



2006-2009 SCC - COWS - CASPIS

Collective adaptive systems progr.

- large number of components
- decentralised control
- unpredictable environment
- emergent behaviour



2016 AbC


1998 Klaim



Network-aware programming

- awareness of the network infrastructure
- asynchronous interactions
- open-ended non-determ. environment
- computation mobility

2012 SCEL




Autonomic computing

- reduced maintenance cost
- no human intervention
- continuous monitoring
- adaptation

Collective Adaptive Systems



We are surrounded by examples of collective systems, in the **natural world**

- ▶ Bees, Fishes, Birds, ...



... and in the **man-made world**

- ▶ Traffic, Epidemics, Robots, ...



Many components

From a computer science perspective, collective adaptive systems can be viewed as consisting of a **large number of interacting entities**.

Local behaviour

Each entity may have its **own properties, objectives and actions** and at the system level the entities combine to create the **collective - emergent - behaviour**.

Mutual Influence

The behaviour of the system is dependent on that of the individual entities and the behaviour of the individuals will be influenced by the state of the overall system.

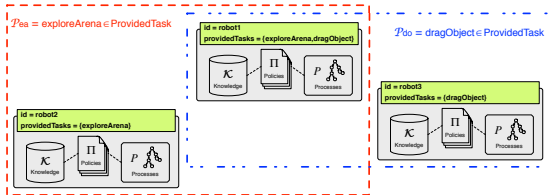
No Central Control

CAS need to operate **without centralised control** or direction. When conditions within the system change it may not be feasible to have human intervention to adjust behaviour appropriately and systems must **autonomously adapt**.

The SCEL language: ensembles

Ensembles formation

- ▶ **Attributes** are used by the system components to dynamically organize themselves into **ensembles**
- ▶ **Predicates** P over attributes are used by components to specify the targets of communication actions.




- ▶ Ensembles are determined by the predicates validated by each component
- ▶ There is no coordinator, hence no bottleneck or critical point of failure
- ▶ A component might be part of more than one ensemble

The SCEL language

Introduced to deal with the challenges posed by the design of **ensembles** of autonomic components

An autonomic component in SCEL:



- ▶ **Knowledge repositories** where components store and retrieve information about their **working environment** and to use it for determining and adapting their behaviour
- ▶ **Policies** regulating the inter- and intra-components interaction
- ▶ **Interfaces** consisting of a collection of **attributes**, like provided functionalities, spatial coordinates, group memberships, trust level, response time, ...

Yehia Abd Alrahman, Rocco De Nicola, Michele Loreti:

[On the Power of Attribute-Based Communication](#)

FORTE 2016 and Information and Computation 2019

...

...

Rocco De Nicola, Tan Duong, and Michele Loreti:

[ABEL - A DSL for programming with attribute-based communication](#)

Coordination 2019


AbC: a calculus distilled from SCEL

- ▶ Systems are represented as sets of parallel components, each equipped with a set of **attributes** whose values can be modified by internal actions.
- ▶ Communication actions (**send** and **receive**) are decorated with **predicates over attributes** that partners have to satisfy to make the interaction possible.
- ▶ Communication takes place in an **implicit multicast** fashion: partners are selected via predicates over the attributes exposed in their interfaces.
- ▶ Components are **unaware of the existence of each other** and receive messages only if they satisfy senders requirements.
- ▶ Components can offer **different views** of themselves and can communicate with different partners according to different criteria.
- ▶ Semantics for **output actions is non-blocking** while **input actions are blocking**: they can take place through synchronization with an available sent message.

AbC: basic ingredients.

An AbC system consists of a set components that contain

- ▶ a **behaviour** - a set of running processes
- ▶ an **environment** - a map from attributes names to values



Processes can:

- ▶ **send** a message to all the components satisfying a given predicate;
- ▶ **receive** a message from a component satisfying a given predicate;
- ▶ **change** the environment;
- ▶ **wait** until a given predicate is locally satisfied.

Components	$C ::= \Gamma ; P \mid C_1 \parallel C_2 \mid [C]^{<f} \mid [C]^{>f}$
Processes	$P ::= 0 \mid \Pi(\tilde{x}).U \mid (\tilde{E})@ \Pi.U \mid \langle \Pi \rangle P \mid P_1 + P_2 \mid P_1 P_2 \mid K(x_1, \dots, x_n)$
Updates	$U ::= [a := E]U \mid P$
Predicates	$\Pi ::= \text{tt} \mid \text{ff} \mid p_k(E_1, \dots, E_k) \mid \Pi_1 \wedge \Pi_2 \mid \Pi_1 \vee \Pi_2 \mid \neg \Pi$
Expressions	$E ::= v \mid x \mid a \mid \text{this}.a \mid o_k(E_1, \dots, E_k)$

A basic component, $\Gamma :_I P$, is a process P associated with an **attribute environment** Γ , and an **interface** I .

- ▶ The **attribute environment** $\Gamma : \mathcal{A} \rightarrow \mathcal{V}$ is a partial map from attribute identifiers with $a \in \mathcal{A}$ to values $v \in \mathcal{V}$ where $\mathcal{A} \cap \mathcal{V} = \emptyset$. A value could be a number, a name (string), a tuple, etc.
- ▶ The **interface** $I \subseteq \mathcal{A}$ consists of a *finite* set of **attributes names** that are exposed by a component to control the interactions with other components.
- ▶ Attributes in I are **public**, and to those in $dom(\Gamma) - I$ are **private**.

Two operators $[C]^{\triangleleft f}$ and $[C]^{\triangleleft f}$ are introduced to **restrict information flow**. Function f associates a predicate Π to each tuple of values $\tilde{v} \in \mathcal{V}^*$ and attribute environment Γ .

- ▶ $[C]^{\triangleleft f}$ is used to restrict the messages that component C can send.
 - ▶ When the message outgoing $[C]^{\triangleleft f}$, the target predicate is updated to consider also predicate $\Pi' = f(\Gamma, \tilde{v})$
 - ▶ Only components satisfying $\Pi \wedge \Pi'$ will receive the message.
 - ▶ To prevent a **specific secret** s from being spread outside C , one can use $f_s(\Gamma, \tilde{v}) = \text{tt}$ if $s \notin \tilde{v}$ and $f_s(\Gamma, \tilde{v}) = \text{ff}$ otherwise.
- ▶ $[C]^{\triangleleft f}$ is used to restrict the messages that component C can receive.
 - ▶ If a component with public attribute environment Γ sends a message \tilde{v} to components C satisfying Π , only those components in C that satisfy $\Pi \wedge f(\Gamma, \tilde{v})$ are eligible to receive the message.

A **process** P can be the:

- ▶ **inactive** process 0 ,
- ▶ **action-prefixed** process, $act.U$, where act is a communication action and U is a process possibly preceded by an **attribute update**,
- ▶ **self-aware** process $\langle \Pi \rangle P$, blocks the execution of P until predicate Π is satisfied within the attribute environment where the process is executing and triggers execution of P when the environment changes and $\Gamma \models \Pi$
- ▶ **nondeterministic choice** between two processes $P_1 + P_2$,
- ▶ **interleaving composition** of two processes $P_1 | P_2$, processes can only communicate indirectly through the attribute environment they share
- ▶ **parametrised process call** with a unique identifier K and a sequence of formal parameters (x_1, \dots, x_n) used in the process definition $K(x_1, \dots, x_n) \triangleq P$.

Using attributes

- ▶ **attribute-based output** $(\tilde{E})@P$ is used to send the evaluation of the sequence of expressions \tilde{E} to the components whose attributes satisfy the predicate P .
- ▶ **attribute-based input** $P(\tilde{x})$ is used to receive messages from any component whose attributes (and possibly transmitted values) satisfy the predicate P ; the sequence \tilde{x} acts as a placeholder for received values.
- ▶ **attribute update** $[a := E]$ is used to assign the result of the evaluation of E to the attribute identifier a . Updates are only possible after communication actions: they can be viewed as side effects of interactions. **Execution of a communication action and the following update(s) is atomic.**

Predicates can refer to **public** and **private** attributes of components.

$$(\text{"Req"}, 1, 3)@(i \geq \text{this.i})$$

can be used to send the message $(\text{"Req"}, 1, 3)$ to all components whose attribute i is not less than **this.i**.

Semantics rules: Potential Communications

$$\frac{\llbracket \tilde{E} \rrbracket_{\Gamma} = \tilde{v} \quad \{\Pi_1\}_{\Gamma} = \Pi}{\Gamma :_l (\tilde{E}) @ \Pi_1 . U \xrightarrow{\Gamma \downarrow l \triangleright \bar{\Pi}(\tilde{v})} \{\Gamma :_l U\}} \text{BRD}$$

Expressions in \tilde{E} are evaluated to \tilde{v} , and the *closure* Π of predicate Π_1 under Γ is computed then \tilde{v} , $\{\Pi_1\}_{\Gamma}$ and $\Gamma \downarrow l$. **Environment updates may be applied.**

$$\frac{\Gamma' \models \{\Pi_1[\tilde{v}/\tilde{x}]\}_{\Gamma_1} \quad \Gamma_1 \downarrow l \models \Pi}{\Gamma_1 :_l \Pi_1(\tilde{x}) . U \xrightarrow{\Gamma' \triangleright \Pi(\tilde{v})} \{\Gamma_1 :_l U[\tilde{v}/\tilde{x}]\}} \text{RCV}$$

A message can be received when $\Gamma_1 \downarrow l$ satisfies sender's predicate Π , and the environment of the sender Γ' satisfies the receiving predicate $\{\Pi_1[\tilde{v}/\tilde{x}]\}_{\Gamma_1}$. **Updates U under substitution $[\tilde{v}/\tilde{x}]$ may be applied.**

Atomicity of Communications and Updates

$$\{C\} = \begin{cases} \{\Gamma[a \mapsto \llbracket E \rrbracket] :_l U\} & C = \Gamma :_l [a := E] U \\ \Gamma :_l P & C = \Gamma :_l P \end{cases}$$

$$\begin{array}{c}
 C_1 \xrightarrow{\Gamma \triangleright \Pi(\tilde{v})} C'_1 \quad C_2 \xrightarrow{\Gamma \triangleright \Pi(\tilde{v})} C'_2 \\
 \hline
 C_1 \parallel C_2 \xrightarrow{\Gamma \triangleright \Pi(\tilde{v})} C'_1 \parallel C'_2
 \end{array}
 \text{ SYNC}
 \qquad
 \begin{array}{c}
 C_1 \xrightarrow{\Gamma \triangleright \bar{\Pi}(\tilde{v})} C'_1 \quad C_2 \xrightarrow{\Gamma \triangleright \Pi(\tilde{v})} C'_2 \\
 \hline
 C_1 \parallel C_2 \xrightarrow{\Gamma \triangleright \bar{\Pi}(\tilde{v})} C'_1 \parallel C'_2
 \end{array}
 \text{ COML}$$

- ▶ **SYNC** states that C_1 and C_2 can receive the same message.
- ▶ **COML** governs communication between components C_1 and C_2 .

$$\begin{array}{c}
 C \xrightarrow{\Gamma \triangleright \bar{\Pi}(\tilde{v})} C' \quad f(\Gamma, \tilde{v}) = \Pi' \\
 \hline
 [C] \triangleright^f \xrightarrow{\Gamma \triangleright \bar{\Pi} \wedge \Pi'(\tilde{v})} [C'] \triangleright^f
 \end{array}
 \text{ RESO}
 \qquad
 \begin{array}{c}
 C \xrightarrow{\Gamma \triangleright \Pi \wedge \Pi'(\tilde{v})} C' \quad f(\Gamma, \tilde{v}) = \Pi' \\
 \hline
 [C] \triangleleft^f \xrightarrow{\Gamma \triangleright \Pi(\tilde{v})} [C'] \triangleleft^f
 \end{array}
 \text{ RESI}$$

- ▶ **RESO**: if C evolves to C' via $\Gamma \triangleright \bar{\Pi}(\tilde{v})$ and $f(\Gamma, \tilde{v}) = \Pi'$ then $[C] \triangleright^f$ evolves via $\Gamma \triangleright \bar{\Pi} \wedge \Pi'(\tilde{v})$ to $[C'] \triangleright^f$.
- ▶ **RESI**: $[C] \triangleleft^f$ will receive \tilde{v} and evolve to $[C'] \triangleleft^f$ with a label $\Gamma \triangleright \Pi(\tilde{v})$ only when $C \xrightarrow{\Gamma \triangleright \Pi \wedge \Pi'(\tilde{v})} C'$ where $f(\Gamma, \tilde{v}) = \Pi'$.

Observable Barbs

Let $C \downarrow_{\Pi}$ mean that component C can send a message with a predicate $\Pi' \simeq \Pi$ (i.e., $C \xrightarrow{\nu \bar{x} \Pi' \check{v}}$ where $\Pi' \simeq \Pi$ and $\Pi' \neq \text{ff}$). We write $C \Downarrow_{\Pi}$ if $C \rightarrow^* C' \downarrow_{\Pi}$.

Barb Preservation

\mathcal{R} is barb-preserving iff for every $(C_1, C_2) \in \mathcal{R}$, $C_1 \downarrow_{\Pi}$ implies $C_2 \downarrow_{\Pi}$

Weak Reduction Barbed Congruence Relations

A Weak Reduction Barbed Relation is a symmetric relation \mathcal{R} over the set of AbC-components which is barb-preserving, reduction-closed, and context-closed.

Barbed Bisimilarity

Two components are weakly reduction barbed congruent, written $C_1 \cong C_2$, if $(C_1, C_2) \in \mathcal{R}$ for some weak reduction barbed congruent relation \mathcal{R} .

Weak Bisimulation

A symmetric binary relation \mathcal{R} over the set of AbC-components is a *weak bisimulation* if and only if for any $(C_1, C_2) \in \mathcal{R}$ and for any λ_1

$$C_1 \xrightarrow{\lambda_1} C'_1 \text{ implies } \exists \lambda_2 : \lambda_1 \simeq \lambda_2 \text{ such that } C_2 \xRightarrow{\hat{\lambda}_2} C'_2 \text{ and } (C'_1, C'_2) \in \mathcal{R}$$

Two components C_1 and C_2 are weakly bisimilar, written $C_1 \approx C_2$ if there exists a weak bisimulation \mathcal{R} relating them.

Theorem (Soundness)

$C_1 \approx C_2$ implies $C_1 \cong C_2$, for any two components C_1 and C_2 .

Theorem (Completeness)

$C_1 \cong C_2$ implies $C_1 \approx C_2$, for any two components C_1 and C_2 .

Encoding other paradigms

A number of alternative communication paradigms can be easily modelled by relying on AbC primitives.

Explicit Message Passing

A $b\pi$ -calculus process P is rendered as an AbC component $\Gamma:P$ where $\Gamma = \emptyset$ and the communication channel is sent as a part of the transmitted values with the receiver checking its compatibility.

Group based Communications

The group name is encoded as an attribute in AbC. The constructs for joining or leaving a given group can be encoded as attribute updates.

Publish-Subscribe

A Publisher sends tagged messages for all subscribers by exposing from his environment only the current topic while subscribers check compatibility of messages according to their subscriptions.

Many challenges:

- ▶ Which kind of Middleware?
 - ▶ Centralized?
 - ▶ Distributed?
- ▶ Whom checks the predicates?
 - ▶ the sender?
 - ▶ the receiver?
 - ▶ a central entities?
- ▶ For the moment: four implementations
 - ▶ one in Java
 - ▶ two in Erlang
 - ▶ one in Go



- ▶ **AbaCus** - Java: a centralized broker, broadcast, missed performance evaluation [ISOLA'16]
- ▶ **AErlang** - Erlang: a centralized broker with different dispatching policies [COORD'17]
 - ▶ **Broadcast**: Receivers checks both sending and receiving predicates
 - ▶ **Push**: broker checks sending predicates, receivers check receiving predicates
 - ▶ **Pull**: broker checks receiving predicates, receivers check sending predicates
 - ▶ **Push-pull**: broker checks both sending and receiving predicates

dynamically handling messages, good performance, **deviated** semantics
- ▶ **GoAt** - Go: a set of broker connected in different shapes [ISOLA'18]
 - ▶ Semantics preserving implementation
 - ▶ Performance evaluation showed a tree-based structure performs best
 - ▶ However, deriving Goat code from AbC code is not immediate

ABEL - A programming framework for AbC

ABEL - Erlang) is a recent implementation of AbC combining previous experience [COORD'19]

- ▶ Providing Inter-coordinators (tree-based) and intra-coordinators interaction
- ▶ Supporting **total-ordering** and **relaxed ordering** of message delivery

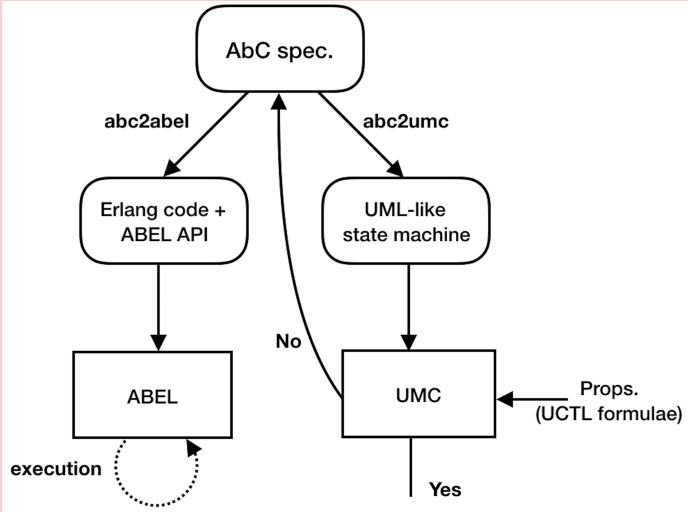


ABEL - A programming framework for AbC

ABEL API offers a one-to-one correspondence with AbC constructs

$C ::=$	new_component (<i>comp_name</i> , <i>Env</i> , <i>l</i>)	Create
	start_component (<i>C</i> , <i>BRef</i>)	Start
$BDef ::=$	<i>proc</i> (<i>C</i> , $\langle vars \rangle$) \rightarrow <i>Com</i> .	Definition
$BRef ::=$	fun ($\langle vars \rangle$) \rightarrow <i>proc</i> (<i>C</i> , $\langle vars \rangle$) end	Reference
	nil	
$Act ::=$	{ $\langle g \rangle$, <i>m</i> , <i>s</i> , $\langle u \rangle$ }	Output
	{ $\langle g \rangle$, <i>r</i> , $\langle u \rangle$ }	Input
$Com ::=$	prefix (<i>C</i> , { <i>Act</i> , <i>BRef</i> })	Prefix
	choice (<i>C</i> , [{ <i>Act</i> , <i>BRef</i> }]])	Choice
	parallel (<i>C</i> , [<i>BRef</i>])	Parallel
	call (<i>C</i> , <i>BRef</i>)	Call

A model-driven approach to AbC programming



An example: Stable Marriage with Attributes

- ▶ Match men and women based on their preferences on partner's attributes
 - ▶ attributes: agents characteristics
 - ▶ preferences: interested values of partners attributes
 - ▶ An examples with 2 attributes and 2 preferences

Id	Wealth	Body	Preferences
m1	rich	strong	eyes=amber \wedge hair=red
m2	rich	weak	eyes=green \wedge hair=dark
m3	poor	strong	eyes=green \wedge hair=red
m4	poor	weak	eyes=amber \wedge hair=red

Id	Eyes	Hair	Preferences
w1	amber	dark	wealth=poor \wedge body=weak
w2	amber	dark	wealth=rich \wedge body=strong
w3	green	red	wealth=rich \wedge body=strong
w4	green	dark	wealth=rich \wedge body=weak

- ▶ Man: iteratively proposes while gradually relaxing expectations (predicates)
- ▶ Woman: performs “select and swaps”

Two types of components M and W :

$$\blacktriangleright M_i \triangleq \Gamma_{mi} : \{id, w, b, pe, ph, \dots\} P_M$$

$$\blacktriangleright W_j \triangleq \Gamma_{wj} : \{id, e, h, pw, pb, \dots\} P_W$$

Specification for M : $P_M \triangleq Q \mid P \mid A \mid R$

$$\begin{aligned} \text{(Proposing)} \quad P &\triangleq \langle partner = 0 \wedge send = 1 \wedge \dots \rangle \\ &\quad (\text{'propose'}, this.id, \widetilde{msg}) @ (\Pi). [send := 0] P + \dots \end{aligned}$$

$$\begin{aligned} \text{(Positive answer)} \quad A &\triangleq (x = \text{'yes'})(x, y). (H(y) \mid A) \\ H(y) &\triangleq (\langle partner = 0 \rangle (\text{'confirm'}) @ (id = y). [partner := y] 0 \\ &\quad + \langle partner > 0 \rangle (\text{'toolate'}) @ (id = y). 0 \end{aligned}$$

$$\begin{aligned} \text{(Rejection answer)} \quad R &\triangleq (x = \text{'split'})(x, y). [send := 1, partner := 0, \dots] R \\ &\quad + (x = \text{'no'})(x, y). [send := 1, \dots] R \end{aligned}$$

Writing AbC in ABEL

$$R \triangleq (x = \text{'split'})(x, y).[\dots, \text{send} := 1, \text{partner} := 0]R \\ + (x = \text{'no'})(x, y).[\dots, \text{send} := 1]R$$

r(C) \rightarrow

... defining actions and continuations

Ref = fun() \rightarrow r(C) end,

choice(C, [{Act1**, Ref}], [{**Act2**, Ref}])).**

RP1 = fun(L, M, R) \rightarrow

size(M) == 2 andalso msg(1, M) == 'split'

end

U1 = [{send, 1}, {partner, 0}, ...]

Act1 = {RP1, U1}

We verified for all input spaces of problems of size of 2

- ▶ **Termination** - True
- ▶ **Soundness** of outcomes:
 - ▶ completeness - True
 - ▶ symmetry - True
 - ▶ uniqueness - False
- ▶ **Liveness properties**:
 - ▶ If a woman sends 'yes' she will eventually receive a 'toolate' or 'confirm' message - True
 - ▶ If a man receives a 'split', he will eventually send a new proposal - False (he may immediately receive another 'yes', and settle down)

Thank you!