

Spatial Logics for Distributed Systems

Luca Cardelli

Microsoft Research

Oxford, Strachey Lecture, 2002-04-23

Reflecting joint work with
Luís Caires, Philippa Gardner, Andrew D. Gordon, Giorgio Ghelli.

Widely Distributed Systems

Concurrent systems that are *spatially* distributed

- Not in the same box.
- Not on the same LAN.
- Not inside the same firewall.
- Not always in the same place.

They have well-defined subsystems that:

- Fail independently.
- Recover independently.
- Hold secrets, mistrust each other.
- Move around.

The New Machine

The “machine” we now write programs for, is the whole Internet.

- New instruction sets (programming models):
 - Message-centric, asynchronous, often stateless.
Cannot rely on distributed consensus.
 - In striking contrast to shared-memory concurrency, and handshake-based (synchronous) concurrency.
- New type systems:
 - Traditional “strong” type systems have been (finally!) enthusiastically adopted as a foundation for security.
 - But entirely new type systems are needed for regulating communication, and to manage application-level security.
- New program logics:
 - Privacy/security concerns override everything else.
 - Need “location awareness”.

Talking About *Where*

Informal statements:

- Distribution: *Where* are things happening?
- Security: *Where* are things kept, and who can get there?
- Privacy: *Where* are things known, and where are they leaked?

We need a new way of reasoning (i.e. a new logic):

- Classical logic: *Whether* something is true.
- Intuitionistic logic: *How* something is true.
- Temporal logic: *When* something is true.
- *Spatial* logic: *Where* something is true.

Outline

We look, concretely, at specific logics for specific models:

- For trees.
- For graphs.
- For mobility.
- For communication.
- For privacy.

With some common, new-ish, techniques:

- Semantically: Modal logics for structured worlds.
- Syntactically: Many-world sequent calculi.

1: A Logic for Trees

Historically, this all began with a spatial logics for the Ambient Calculus (a process calculus based on trees of locations).

As a tutorial, we start by looking just at the trees.

Spatial interpretation: a formula holds at a particular (sub-)tree.

This has also its own applications.

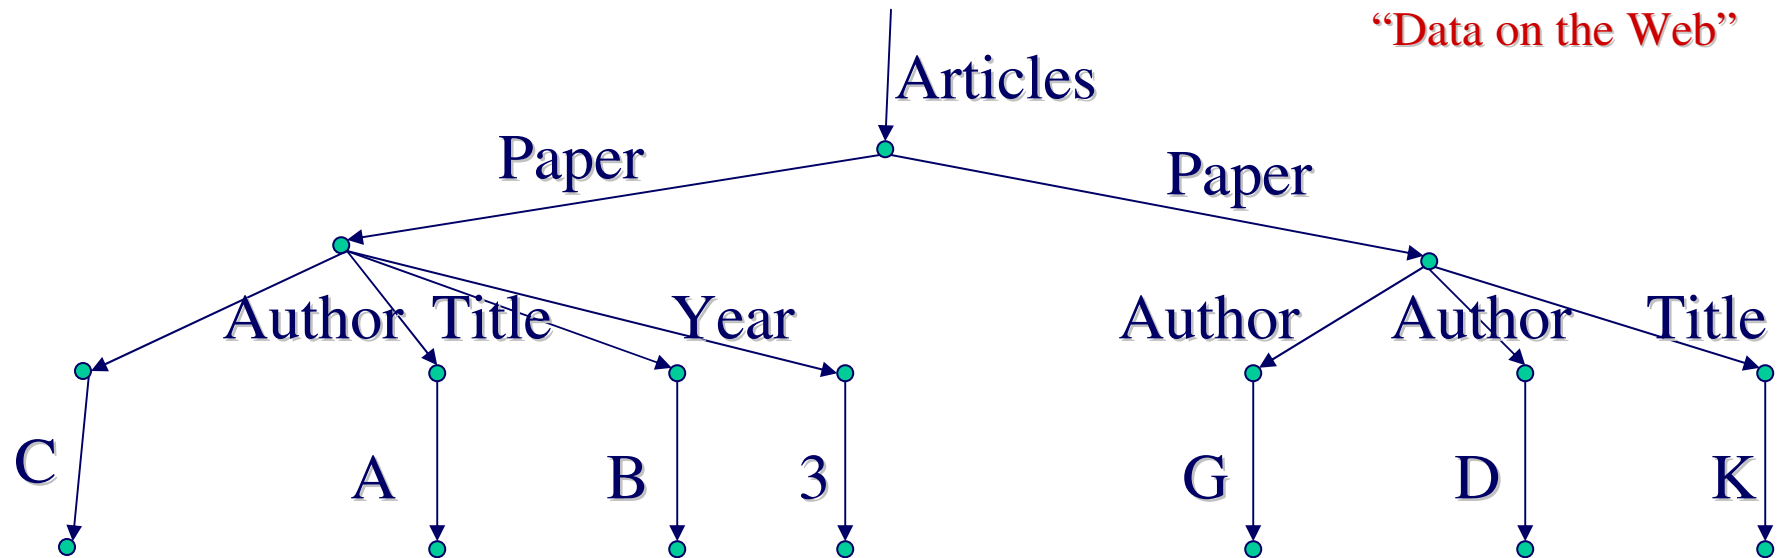
Cardelli: Describing Semistructured Data. SIGMOD Record.

Cardelli-Ghelli: A Query Language based on the Ambient Logic. ESOP'01.

Semistructured Data

(I.e.: XML after parsing)

Abiteboul, Buneman, Suciu:
“Data on the Web”



A tree (or graph), unordered (or ordered). With labels on the edges.

Invented for “flexible” data representation, for quasi-regular data like address books and bibliographies.

Adopted by the DB community as a solution to the “database merge” problem: merging databases from uncoordinated (web) sources.


Adopted by W3C as “web data”, then by everybody else.

Trees and their Descriptions

Trees



root



edge



join

Syntax for Trees (P,Q)

0 root

$n[P]$ edge

$P|Q$ join

Basic Descriptions (A,B)

0 there is only a root

$n[A]$ there is an edge n to a subtree

$A|B$ there are two joined trees

T there is anything

Example

Data

```
Cambridge[
  Eagle[
    chair[0] |
    chair[0]
  ]
]
```

Description

```
Cambridge[
  Eagle[
    chair[0] |
    T
  ] | T
]
```

data matches description

In Cambridge there is (nothing but) a pub called the Eagle that contains (nothing but) two empty chairs.

In Cambridge there is (at least) a pub called the Eagle that contains (at least) one empty chair.

Basic Tree Descriptions

$\mathbf{0}$ matches $\mathbf{0}$

$n[P]$ matches $n[\mathcal{A}]$ iff P matches \mathcal{A}

$P \mid Q$ matches $\mathcal{A} \mid \mathcal{B}$ iff P matches \mathcal{A} and Q matches \mathcal{B}

P matches \mathbf{T} always

Property: if P matches \mathcal{A} and $P \equiv Q$, then Q matches \mathcal{A}

where $P \equiv Q$ means that P and Q represent the same tree:

$$P_1 \mid P_2 \equiv P_2 \mid P_1$$

$$P_1 \mid (P_2 \mid P_3) \equiv (P_1 \mid P_2) \mid P_3$$

$$P \mid \mathbf{0} \equiv P$$

Propositional Descriptions

P	matches	$\neg A$	iff	P does not match	A
P	matches	$A \wedge B$	iff	P matches	A and P matches B
P	matches	$A \vee B$	iff	P matches	A or P matches B
P	matches	$A \Rightarrow B$	iff	if P matches	A then P matches B

Examples

“Vertical” implications about nesting

“Business Policy”

Borders[
 Starbucks[...] |
 Books[...]
]

Borders[**T**] \Rightarrow
Borders[*Starbucks*[**T**] | **T**]

If it's a **Borders**,
then it must contain a **Starbucks**

“Horizontal” implications about proximity

“Social Policy”

Smoker[...] |
NonSmoker[...] |
Smoker[...]

(*NonSmoker*[**T**] | **T**) \Rightarrow
(*Smoker*[**T**] | **T**)

If there is a **NonSmoker**,
then there must be a **Smoker** nearby

Other Descriptions

P matches $\forall x.A$ iff P matches A for any label n for x

P matches $\mu X.A$ iff P matches A where X is $\mu X.A$
(up to some well-formedness conditions)

Some definable descriptions:

$$\exists x.A \triangleq \neg \forall x. \neg A$$

$$\text{somewhere } A \triangleq \mu X. A \vee \exists y. y[X] \mid \mathbf{T}$$

Descriptions as Schemas

Descriptions are a “very rich type system”. We can comfortably represent various kinds of schemas.

Ex.: Xduce-like schemas (c.f. XML DTDs):

$\mathbf{0}$	the empty tree
$\mathcal{A} \mathcal{B}$	an \mathcal{A} next to a \mathcal{B}
$\mathcal{A} \vee \mathcal{B}$	either an \mathcal{A} or a \mathcal{B}
$n[\mathcal{A}]$	an edge n leading to an \mathcal{A}
\mathcal{A}^*	$\triangleq \mu X. \mathbf{0} \vee (\mathcal{A} X)$ the merge of zero or more \mathcal{A} s
\mathcal{A}^+	$\triangleq \mathcal{A} \mathcal{A}^*$ the merge of one or more \mathcal{A} s
$\mathcal{A}^?$	$\triangleq \mathbf{0} \vee \mathcal{A}$ zero or one \mathcal{A}

Descriptions as Queries

Yes/no: *Is there an empty chair at the Eagle?*

<i>Eagle</i> [<i>Eagle</i> [
<i>chair</i> [<i>John</i> [0]]	matches?	<i>chair</i> [0]	(Yes)
<i>chair</i> [<i>Mary</i> [0]]		T	
<i>chair</i> [0]]	
]			

With match variables \mathcal{X} : *Who is sitting at the Eagle?*

	<i>Eagle</i> [Yes: $\mathcal{X} = $ <i>John</i> [0]
matches?	<i>chair</i> [\mathcal{X}]	Yes: $\mathcal{X} = $ <i>Mary</i> [0]
	T	Yes: $\mathcal{X} = $ 0 oops!
]	

With match variables \mathcal{X} : *Who is really sitting at the Eagle?*

matches? *Eagle*[
chair[$\neg \mathbf{0} \wedge \mathcal{X}$] |
T
]

Yes: $\mathcal{X} = \textit{John}[\mathbf{0}]$
Yes: $\mathcal{X} = \textit{Mary}[\mathbf{0}]$

With *select-from*:

from Eagle[...]
match Eagle[*chair*[$\neg \mathbf{0} \wedge \mathcal{X}$] | **T**]
select person[\mathcal{X}]

Single result:
person[*John*[$\mathbf{0}$]] |
person[*Mary*[$\mathbf{0}$]]

2: A Logic for Graphs

Semistructured data is not actually trees, but either graphs, or graphs with tree backbones.

It is interesting to generalize to a logic for graph.


But this is actually *much* harder. Questions of expressiveness and, particularly, how to define query languages, are still open.

Likely, this is a subset of Second-Order Monadic logic, but it allows direct “local” reasoning about subgraphs, without explicit disjointness assumptions.

Cardelli-Gardner-Ghelli: A Spatial Logic for Querying Graphs. ICALP'02.

Graphs and their Descriptions

labeled graphs



empty

edge



join

(identify common nodes)

Syntax for Graphs (P,Q)

0 empty

$n(x,y)$ edge

$P \mid Q$ join

Basic Descriptions (A,B)

0 there is an empty graph

$n(x,y)$ there is an edge from x to y labeled n

$\mathcal{A} \mid \mathcal{B}$ there are two disjoint parts of a graph

T there is anything

Example: Paths



$$n(a,b) \mid n(b,a)$$

There is a path of length 3 (Yes!):

$$\exists w,x,y,z. (n(w,x) \mid \mathbf{T}) \wedge (n(x,y) \mid \mathbf{T}) \wedge (n(y,z) \mid \mathbf{T})$$

There is a non-repeating path of length 3 (No!):

$$\exists w,x,y,z. n(w,x) \mid n(x,y) \mid n(y,z) \mid \mathbf{T}$$

As shown in the second case, we can encode implicitly and compactly disjointness assumptions about subgraphs. This is due to the spatial nature of the logic: a formula holds at a subgraph; $\mathcal{A} \mid \mathcal{B}$ means that \mathcal{A} and \mathcal{B} hold at distinct subgraphs. (A “subgraph” here is a subset of edges, not of nodes.)

3: A Logic for Mobility



We now look at our original spatial logic: a logic for mobility.

Security specifications are a concern here: additional logical operators arise naturally, and turn out to be logical adjuncts.

Cardelli-Gordon: Anytime, Anywhere: Modal Logics for Mobile Ambients: POPL'00.

Mobility

Mobility is change of spatial structures over time.





$a[Q \mid c[out\ a.\ in\ b.\ P]]$

$\mid b[R]$

Mobility

Mobility is change of spatial structures over time.





$a[Q]$

$| c[in\ b.\ P] | b[R]$

Mobility

Mobility is change of spatial structures over time.



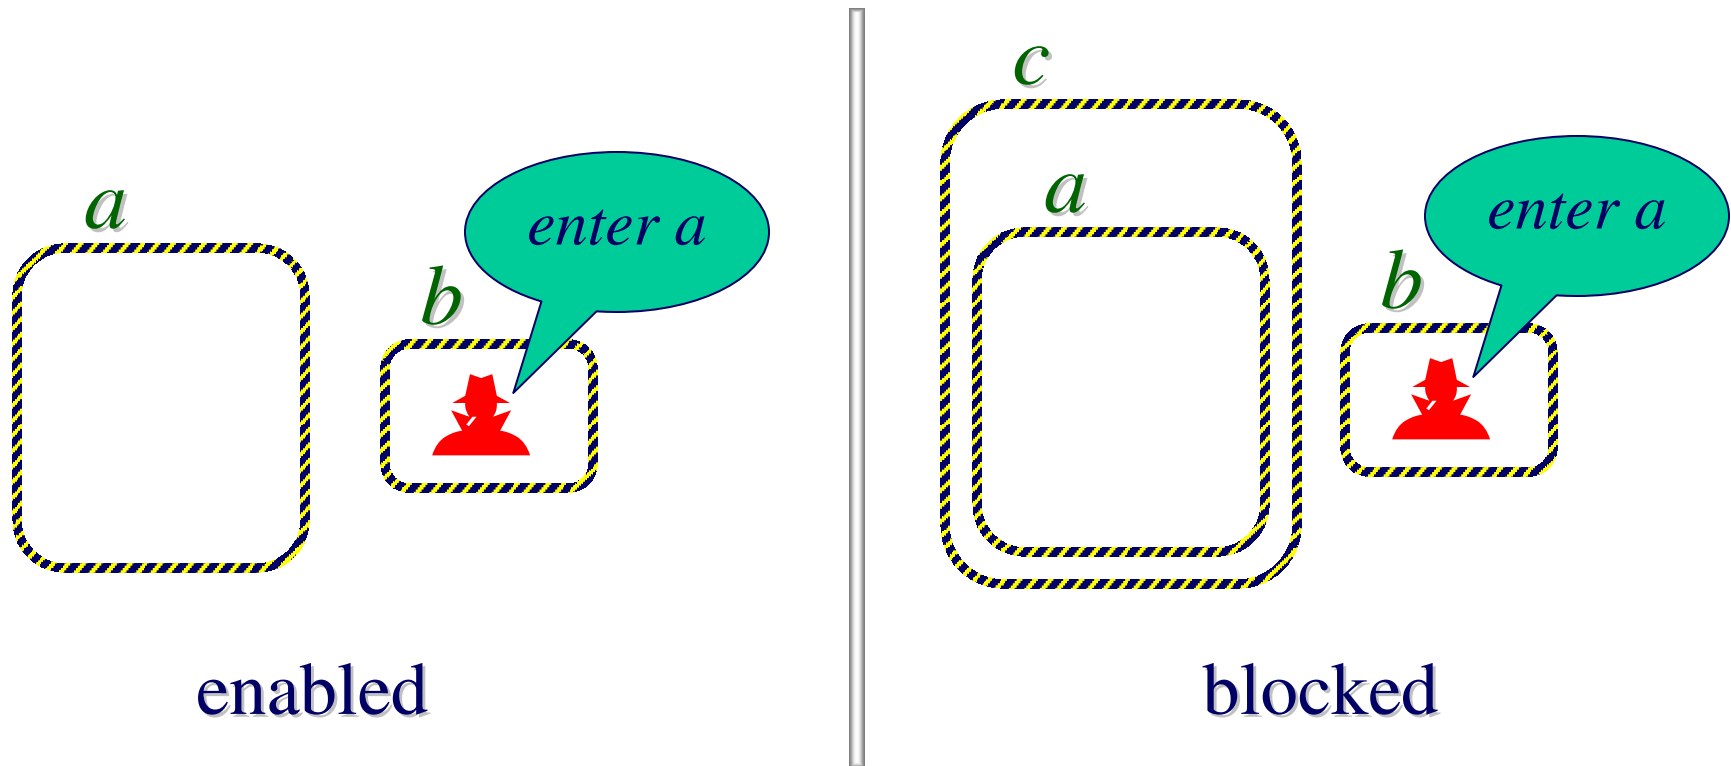
$a[Q]$

$| b[R | c[P]]$

Security

Security issues are reduced to the capability of entering and exiting locations (as well as creating and destroying locations).

- E.g.: Firewall-crossing protocols.
- Capabilities can be exercised only the the right places:



Spatial-style Protocol Specification

Right now, we have a spatial configuration, and later, we have another spatial configuration.

E.g.: Right now, the agent is outside the firewall, ...




$(agent[\mathbf{T}] \mid firewall[\mathbf{T}] \mid \mathbf{T})$

Spatial-style Protocol Specification

Right now, we have a spatial configuration, and later, we have another spatial configuration.

E.g.: Right now, the agent is outside the firewall, and later (after running an authentication protocol), the agent is inside the firewall.

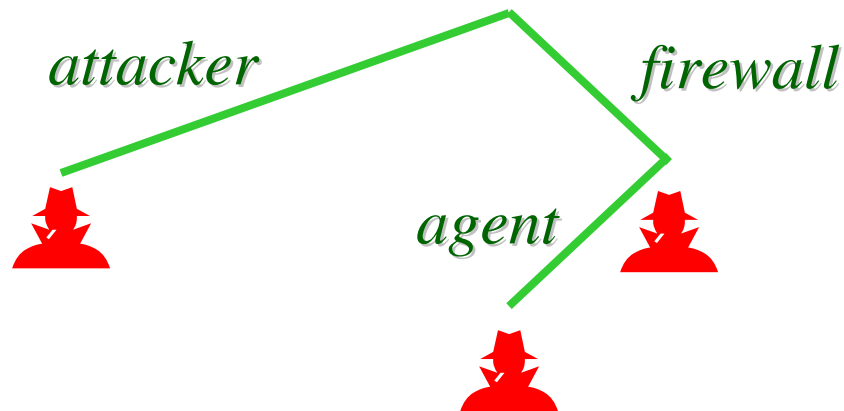


$$(agent[\mathbf{T}] \mid firewall[\mathbf{T}] \mid \mathbf{T}) \wedge \diamond (firewall[agent[\mathbf{T}] \mid \mathbf{T}] \mid \mathbf{T})$$

Spatial-style Protocol Specification

Right now, we have a spatial configuration, and later, we have another spatial configuration.

E.g.: Right now, the agent is outside the firewall, and later (after running an authentication protocol), the agent is inside the firewall. And this works in presence of any (reasonable) attacker.



Attack $\triangleright ((agent[\mathbf{T}] \mid firewall[\mathbf{T}] \mid \mathbf{T}) \wedge \diamond (firewall[agent[\mathbf{T}] \mid \mathbf{T}] \mid \mathbf{T}))$

Ambient Calculus (without Restriction)

$P \in \Pi ::=$	Processes		$M ::=$	Messages
$\mathbf{0}$	void	} <i>spatial</i>	n	name
$P \mid P'$	composition		$in\ M$	entry capability
$!P$	replication		$out\ M$	exit capability
$M[P]$	ambient		$open\ M$	open capability
$M.P$	capability	} <i>temporal</i>	ε	empty path
$(n).P$	input		$M.M'$	composite path
$\langle M \rangle$	output			

$$n[] \triangleq n[\mathbf{0}]$$

$$M \triangleq M.\mathbf{0} \quad (\text{where appropriate})$$

Reduction Semantics

A structural congruence relation $P \equiv Q$:

- On spatial expressions, $P \equiv Q$ iff P and Q denote the same tree.
- On full ambient expressions, $P \equiv Q$ if in addition the respective threads are “trivially equivalent”.
- Prominent in the definition of the logic.

A reduction relation $P \rightarrow^* Q$:

- Defining the meaning of mobility and communication actions.
- Closed up to structural congruence:

$$P \equiv P', P' \rightarrow^* Q', Q' \equiv Q \Rightarrow P \rightarrow^* Q$$

- Any details about reduction are “hidden” in the logic within a temporal modality.

Logical Formulas

$\mathcal{A}, \mathcal{B} \in \Phi ::=$	Formulas	(η is a name n or a variable x)
F	false	
$\mathcal{A} \wedge \mathcal{B}$	conjunction	
$\mathcal{A} \Rightarrow \mathcal{B}$	implication (adjunct to \wedge)	
0	void	
$\mathcal{A} \mathcal{B}$	composition	
$\mathcal{A} \triangleright \mathcal{B}$	guarantee (adjunct to $ $)	
$\eta[\mathcal{A}]$	location	
$\mathcal{A} @ \eta$	placement (adjunct to $\eta[\]$)	
$\spadesuit \mathcal{A}$	somewhere	
$\diamond \mathcal{A}$	sometime	
$\forall x. \mathcal{A}$	quantification over names	

Satisfaction Relation

$P \models \mathbf{F}$	<i>never</i>
$P \models \mathcal{A} \wedge \mathcal{B}$	$\triangleq P \models \mathcal{A} \wedge P \models \mathcal{B}$
$P \models \mathcal{A} \Rightarrow \mathcal{B}$	$\triangleq P \models \mathcal{A} \Rightarrow P \models \mathcal{B}$
$P \models \mathbf{0}$	$\triangleq P \equiv \mathbf{0}$
$P \models \mathcal{A} \mid \mathcal{B}$	$\triangleq \exists P', P'' \in \Pi. P \equiv P' \mid P'' \wedge P' \models \mathcal{A} \wedge P'' \models \mathcal{B}$
$P \models \mathcal{A} \triangleright \mathcal{B}$	$\triangleq \forall P' \in \Pi. P' \models \mathcal{A} \Rightarrow P \mid P' \models \mathcal{B}$
$P \models n[\mathcal{A}]$	$\triangleq \exists P' \in \Pi. P \equiv n[P'] \wedge P' \models \mathcal{A}$
$P \models \mathcal{A} @ n$	$\triangleq n[P] \models \mathcal{A}$
$P \models \heartsuit \mathcal{A}$	$\triangleq \exists P' \in \Pi. P \downarrow^* P' \wedge P' \models \mathcal{A}$
$P \models \diamond \mathcal{A}$	$\triangleq \exists P' \in \Pi. P \rightarrow^* P' \wedge P' \models \mathcal{A}$
$P \models \forall x. \mathcal{A}$	$\triangleq \forall m \in \Lambda. P \models \mathcal{A}\{x \leftarrow m\}$

$P \downarrow P'$ iff $\exists n, P''. P \equiv n[P'] \mid P''$

\downarrow^* is the reflexive and transitive closure of \downarrow

Satisfaction for Basic Operators

- $\models 0$



$$\begin{array}{c} | \\ n \\ \triangle \\ P \end{array} \models n[\mathcal{A}] \quad \text{if} \quad \triangle P \models \mathcal{A}$$

$$\triangle P \mid Q \models \mathcal{A} \mid \mathcal{B} \quad \text{if} \quad \triangle P \models \mathcal{A} \quad \text{and} \quad \triangle Q \models \mathcal{B}$$

$$\triangle P \models \mathcal{A}@n \quad \text{if} \quad \begin{array}{c} | \\ n \\ \triangle \\ P \end{array} \models \mathcal{A}$$

$$\triangle P \models \mathcal{A} \triangleright \mathcal{B} \quad \text{if for all} \quad \triangle Q \models \mathcal{A} \quad \text{we have} \quad \triangle P \mid Q \models \mathcal{B}$$

Satisfaction for Somewhere/Sometime



N.B.: instead of $\diamond A$ and $\heartsuit A$ we can use a “temporal next step” operator $\gg A$, along with the existing “spatial next step” operator $n[A]$, together with μ -calculus style recursive formulas.

Basic Fact: satisfaction is invariant under structural congruence:

I.e.: $\{P \in \Pi \mid P \models A\}$ is closed under \equiv .

Hence, formulas describe only congruence-invariant properties.

Applications

Verifying security+mobility protocols.

- Still hard, but we have good techniques. [Gordon-Cardelli '99]

Modelchecking security+mobility assertions:

- If P is $!$ -free and \mathcal{A} is \triangleright -free, then $P \models \mathcal{A}$ is decidable. (PSPACE-complete [Cheratonik et al. '01].)
- If P is $!$ -free and \mathcal{A} is \forall -free, then $P \models \mathcal{A}$ is decidable. [Cardelli-Calcagno-Gordon '02].)
- These provides ways of mechanically checking (certain) assertions about (certain) mobile processes.

Expressing mobility/security policies of host sites.

- Conferring more flexibility than just sandboxing the agent.

Just-in-time verification of code containing mobility instructions

- By either modelchecking or proof-carrying code.

4: New Logics for Concurrency

In the process of making spatial sense of $n[\mathcal{A}]$, we also had to make spatial sense of $\mathcal{A} \mid \mathcal{B}$. The latter is, in fact, the harder part. So, in retrospect, it makes sense to consider it on its own.

An outcome is spatial logics for CCS/CSP-like process calculi. Basic idea: take a Hennessy-Milner modal logic and add an $\mathcal{A} \mid \mathcal{B}$ operator. ([Dam] Very hard to reconcile with bisimulation.)

One can go further and investigate spatial logics for restriction, with a *hiding quantifier* $Hx.\mathcal{A}$ (e.g. for π -calculus). This is essential for security/privacy specifications.

([Caires] Very hard to reconcile with bisimulation.)

We can make all that work smoothly by taking a very *intensional* point of view. The logical formulas are not *up-to-bisimulation*: they are *up-to-structural-congruence*.

Spatial Properties: Identifiable Subsystems

A system is often composed of identifiable subsystems.

- “A message is sent from Alice to Bob.”
- “The protocol is split between two participants.”
- “The virus attacks the server.”

Such partitions of a system are (obviously) spatial properties. They correspond to a spatial arrangement of processes in different places.

- Process calculi are *very* good at expressing such arrangements operationally (*c.f.*, chemical semantics, structural congruence).
- To the point where a process is often used as a specification of another process! (We consider this as an anomaly.)
- We want something equally good at the specification, or logical, level.

Spatial Properties: Restricted Resources

A system often restricts the use of certain resources to certain subsystems.

- “A shared private key n is established between two processes.”
- “A fresh nonce n is generated locally and transmitted.”
- “The applet runs in a secret sandbox.”

Something is *hidden/secret/private* if it is present only in a limited subsystem. So these are spatial properties too.

- If something is secret, by assumption it cannot be known. Still, we want to talk about it in specifications.
- We can talk about a secret name only by using a *fresh* name for it (we cannot assume the secret name matches any known name).
- So freshness will be an important concept. Logics of freshness are very new.

Typical Spatial Formulas for Concurrency

Processes

Formulas

$\mathbf{0}$ (void)

$\mathbf{0}$ (nothing here)

$P \mid Q$ (composition)

$\mathcal{A} \mid \mathcal{B}$ (two things here)



$(\nu n)P$ (restriction)

$n\textcircled{\mathcal{A}}$ (hidden thing here)

$n\langle m \rangle$ (message)

$n\langle m \rangle$ (a message here)

$n\langle m \rangle \vDash n\langle m \rangle$



Things one can say

Single-threaded (or void):

$$\neg(\neg\mathbf{0} \mid \neg\mathbf{0})$$

Output: outputs a message m on n (and is/does nothing else):

$$n\langle m \rangle$$

In presence of a message m on n , sends a message n on m and stops:

$$n\langle m \rangle \triangleright \gg m\langle n \rangle$$

Contains a name free:

$$\textcircled{n} \triangleq \neg n \textcircled{\mathbf{T}}$$

$$P \models \neg n \textcircled{\mathbf{T}} \text{ iff } \neg P \equiv (\nu n)P' \text{ iff } n \in \text{fn}(P)$$

Satisfaction for Hidden and Public Names

$$\begin{array}{c} \text{cloud } n \\ \triangle P \end{array} \models Hx.A \quad \text{if } \exists m \notin fn(P, A) \quad P \{ \triangle n \leftarrow m \} \models A \{ x \leftarrow m \}$$

$$\triangle P \models \odot n \quad \text{if } n \in fn(P)$$

(Technically, $Hx.A$ and $\odot n$ are both defined from $n \circledast A$ and a Gabbay-Pitts freshness quantifier.)

Example: “Shared Secret” Postcondition

Consider a situation where “a hidden name x is shared by two locations n and m , and is not known outside those locations”.

$$\text{H}x.(n[\odot x] \mid m[\odot x])$$

What can we do with such a spec? We can fully expand the definitions and work it out in the process calculus:

- $P \models \text{H}x.(n[\odot x] \mid m[\odot x])$
 $\Leftrightarrow \exists r \in \Lambda. r \notin \text{fn}(P) \cup \{n, m\} \wedge \exists R', R'' \in \Pi. P \equiv (\nu r)(n[R'] \mid m[R''])$
 $\wedge r \in \text{fn}(R') \wedge r \in \text{fn}(R'')$
- E.g.: take $P = (\nu p) (n[p[]] \mid m[p[]])$.

Or we can work logically at the formula level, within a proof system.

Ex: Immovable Object vs. Irresistible Force

$$Im \triangleq \mathbf{T} \triangleright \Box(obj\langle \rangle | \mathbf{T})$$

$$Ir \triangleq \mathbf{T} \triangleright \Box \Diamond \neg(obj\langle \rangle | \mathbf{T})$$

$$Im | Ir \vdash (\mathbf{T} \triangleright \Box(obj\langle \rangle | \mathbf{T})) | \mathbf{T}$$

$$\vdash \Box(obj\langle \rangle | \mathbf{T})$$

$$\vdash \Diamond \Box(obj\langle \rangle | \mathbf{T})$$

$$\mathcal{A} \vdash \mathbf{T}$$

$$(\mathcal{A} \triangleright \mathcal{B}) | \mathcal{A} \vdash \mathcal{B}$$

$$\mathcal{A} \vdash \Diamond \mathcal{A}$$

$$Im | Ir \vdash \mathbf{T} | (\mathbf{T} \triangleright \Box \Diamond \neg(obj\langle \rangle | \mathbf{T}))$$

$$\vdash \Box \Diamond \neg(obj\langle \rangle | \mathbf{T})$$

$$\vdash \neg \Diamond \Box(obj\langle \rangle | \mathbf{T})$$

$$\mathcal{A} \vdash \mathbf{T}$$

$$\Diamond \neg \mathcal{A} \vdash \neg \Box \mathcal{A}$$

$$\Box \neg \mathcal{A} \vdash \neg \Diamond \mathcal{A}$$

$$\text{Hence: } Im | Ir \vdash \mathbf{F}$$

$$\mathcal{A} \wedge \neg \mathcal{A} \vdash \mathbf{F}$$

5: A Common Formalization Style

Many-world sequents:

$$\langle S \rangle \Gamma \vdash \Delta$$

Validity: if all the constraints S_k and all the assumptions Γ_i are satisfied, then one of the conclusions Δ_j is satisfied

(Spatial) equivalence constraints
(denote structural congruence)

Indexes (denote processes)

$$\langle \dots u' \dot{=} v' \dots u'' \rightarrow^a v'' \dots \rangle \dots u : \mathcal{A} \dots \vdash \dots v : \mathcal{B} \dots$$

(Temporal) reduction constraints
(denote process reduction)

Formulas (denote properties)

Alex Simpson did this for temporal logic.

Caires-Cardelli: A Spatial Logic for Concurrency (Part II).

Recipe for Rules

- *Left rules, right rules.* Operate mainly on the $\Gamma \vdash \Delta$ part.
When operating on constraints $\langle S \rangle$:
 - Going up: One adds, the other checks constraints.
 - Going down: One removes, the other assumes constraints.They form cut elimination pairs.
- *World rules (optional).* Operate on the $\langle S \rangle$ part only.
Embody inversion lemmas: deep properties of process calculi.
(In temporal logic, they embody properties such as reflexivity and transitivity of the reachability relation.)
 - Going up: add deducible constraints.
 - Going down: remove redundant constraints.Commute easily with cuts.

Propositional Connectives

Identity, Cut, and Contraction:

(Id)

$$\frac{u \doteq_s u' \quad \mathcal{A} \equiv_s \mathcal{A}'}{\langle S \rangle \Gamma, u : \mathcal{A} \vdash u' : \mathcal{A}', \Delta}$$

(CL)

$$\frac{\langle S \rangle \Gamma, u : \mathcal{A}, u : \mathcal{A} \vdash \Delta}{\langle S \rangle \Gamma, u : \mathcal{A} \vdash \Delta}$$

(Cut)

$$\frac{\langle S \rangle \Gamma \vdash u : \mathcal{A}, \Delta \quad \langle S \rangle \Gamma, u : \mathcal{A} \vdash \Delta}{\langle S \rangle \Gamma \vdash \Delta}$$

(CR)

$$\frac{\langle S \rangle \Gamma \vdash u : \mathcal{A}, u : \mathcal{A}, \Delta}{\langle S \rangle \Gamma \vdash u : \mathcal{A}, \Delta}$$

Propositional Connectives:

(\wedge L)

$$\frac{\langle S \rangle \Gamma, u : \mathcal{A}, u : \mathcal{B} \vdash \Delta}{\langle S \rangle \Gamma, u : \mathcal{A} \wedge \mathcal{B} \vdash \Delta}$$

(\wedge R)

$$\frac{\langle S \rangle \Gamma \vdash u : \mathcal{A}, \Delta \quad \langle S \rangle \Gamma \vdash u : \mathcal{B}, \Delta}{\langle S \rangle \Gamma \vdash u : \mathcal{A} \wedge \mathcal{B}, \Delta}$$

(\Rightarrow L)

$$\frac{\langle S \rangle \Gamma \vdash u : \mathcal{A}, \Delta \quad \langle S \rangle \Gamma, u : \mathcal{B} \vdash \Delta}{\langle S \rangle \Gamma, u : \mathcal{A} \Rightarrow \mathcal{B} \vdash \Delta}$$

(\Rightarrow R)

$$\frac{\langle S \rangle \Gamma, u : \mathcal{A} \vdash u : \mathcal{B}, \Delta}{\langle S \rangle \Gamma \vdash u : \mathcal{A} \Rightarrow \mathcal{B}, \Delta}$$

(F L)

$$\frac{}{\langle S \rangle \Gamma, u : \mathbf{F} \vdash \Delta}$$

(F R)

$$\frac{\langle S \rangle \Gamma \vdash \Delta}{\langle S \rangle \Gamma \vdash u : \mathbf{F}, \Delta}$$

Spatial Connectives

Composition:

(0 L)

$$\frac{\langle S, u \dot{=} 0 \rangle \Gamma \vdash \Delta}{\langle S \rangle \Gamma, u : \mathbf{0} \vdash \Delta}$$

(0 R)

$$\frac{u \dot{=} 0}{\langle S \rangle \Gamma \vdash u : \mathbf{0}, \Delta}$$

(| L) *X, Y not free in the conclusion*

$$\frac{\langle S, u \dot{=} X | Y \rangle \Gamma, X : \mathcal{A}, Y : \mathcal{B} \vdash \Delta}{\langle S \rangle \Gamma, u : \mathcal{A} | \mathcal{B} \vdash \Delta}$$

(| R)

$$\frac{\langle S \rangle \Gamma \vdash v : \mathcal{A}, \Delta \quad \langle S \rangle \Gamma \vdash t : \mathcal{B}, \Delta \quad u \dot{=} v | t}{\langle S \rangle \Gamma \vdash u : \mathcal{A} | \mathcal{B}, \Delta}$$

Guarantee:

(▷ L)

$$\frac{\langle S \rangle \Gamma \vdash t : \mathcal{A}, \Delta \quad \langle S \rangle \Gamma, t | u : \mathcal{B} \vdash \Delta}{\langle S \rangle \Gamma, u : \mathcal{A} \triangleright \mathcal{B} \vdash \Delta}$$

(▷ R) *X not free in the conclusion*

$$\frac{\langle S \rangle \Gamma, X : \mathcal{A} \vdash v : \mathcal{B}, \Delta \quad v \dot{=} X | u}{\langle S \rangle \Gamma \vdash u : \mathcal{A} \triangleright \mathcal{B}, \Delta}$$

i.e.:

$$\frac{\mathcal{A} | \mathcal{B} \vdash C}{\mathcal{A} \vdash C \triangleright \mathcal{B}}$$

Additional World Structure:

(S10)

$$\frac{\langle S, u \dot{=} 0 \rangle \Gamma \vdash \Delta \quad ulv \dot{=} 0}{\langle S \rangle \Gamma \vdash \Delta}$$

(S11) *X, Y, U, V not free in the conclusion*

$$\frac{\langle S, u \dot{=} X | Y, v \dot{=} U | V, t \dot{=} X | U, w \dot{=} Y | V \rangle \Gamma \vdash \Delta \quad ulv \dot{=} t | w}{\langle S \rangle \Gamma \vdash \Delta}$$

Suppose $x | y = 0 \Rightarrow x = 0$. Then, if we can already deduce that $x | y \dot{=} 0$, we can eliminate a redundant assumption $x \dot{=} 0$.

Suppose $ulv = t | s \Rightarrow \exists x, y, z, w$ s.t. $u = x | y, v = z | w, t = x | z, s = y | w$. Then, if we can already deduce that $ulv \dot{=} t | w$, we can eliminate a redundant assumptions

$(\mathcal{A} \mid \mathcal{B}) \wedge \mathbf{0} \vdash \mathcal{A} \wedge \mathcal{B}$

6.2 $\langle S, u \doteq X \mid Y, u \doteq \mathbf{0}, X \doteq \mathbf{0} \rangle \Gamma, X : \mathcal{A}, Y : \mathcal{B} \vdash u : \mathcal{A}, \Delta$	(Id) since $u=X$
5.2 $\langle S, u \doteq X \mid Y, u \doteq \mathbf{0} \rangle \Gamma, X : \mathcal{A}, Y : \mathcal{B} \vdash u : \mathcal{A}, \Delta$	6.2, (S 0) since $X \mid Y = \mathbf{0}$
4.2 $\langle S, u \doteq X \mid Y \rangle \Gamma, X : \mathcal{A}, Y : \mathcal{B}, u : \mathbf{0} \vdash u : \mathcal{A}, \Delta$	5.2, (0 L)
3.2 $\langle S \rangle \Gamma, u : (\mathcal{A} \mid \mathcal{B}), u : \mathbf{0} \vdash u : \mathcal{A}, \Delta$	4.2, (L)
2.2 $\langle S \rangle \Gamma, u : (\mathcal{A} \mid \mathcal{B}) \wedge \mathbf{0} \vdash u : \mathcal{A}, \Delta$	3.2, (\wedge L)
...	
2.1 $\langle S \rangle \Gamma, u : (\mathcal{A} \mid \mathcal{B}) \wedge \mathbf{0} \vdash u : \mathcal{B}, \Delta$	Similarly
1 $\langle S \rangle \Gamma, u : (\mathcal{A} \mid \mathcal{B}) \wedge \mathbf{0} \vdash u : \mathcal{A} \wedge \mathcal{B}, \Delta$	2.1, 2.2, (\wedge R)

Conclusions

We set out to find logics for describing properties of distributed systems.
(After trying equational reasoning, traces, etc.)

Spatial logics exhibit the trade-offs of temporal logics: compact notation for implicit state, nice proof systems, reduced expressiveness.

Along the way, we discovered many other applications for the basic techniques. We believe there is something intriguing and new in the approach and its formalization.

With respect to traditional logics of concurrency, we are very *intensional*.
But another word for it is *precise*.

With Caires, we now have a logic and sequent calculus (with cut-elimination) for π -calculus, where we can express privacy properties.

Related work:

- With Calcagno and Godon: Model checking and validity checking.
- Sangiorgi: Spacetime bisimulation.
- O'Hearn and Pym: Logics for heaps.