

Collective resource bounded reasoning in concurrent multi-agent systems

Valentin Goranko
Stockholm University

(based on joint work with **Nils Bulling**)



Workshop on Logics for Resource-Bounded Agents
ESSLLI'2015, Barcelona, August 14, 2015

Overview of the talk

- Collective agency and resource boundedness
- Concurrent game models with resource costs and action guards
- A running example: robot team on a mission
- QATL*: a quantitative extension of the logic ATL* and its use in collective resource bounded reasoning.
- Concluding remarks.

Introduction:

Collective agency and resource sharing

When acting towards achievement of their (qualitative) objectives, agents often act as teams.

As such, they share resources and their collective actions bring about collective updates or redistribution of these resources.

The cost / resource consumption may depend not only on the objective but also on the team of agents acting together.

Some examples:

- A family sharing household and budget
- Joint venture business. Industrial plants.
- Conference fees and organising expenses
- Petrol consumption per passenger in a car, sharing a taxi, paying for a dinner party, etc.

Collective agency and resource sharing: formal modelling and logical reasoning

Here we propose abstract models and logical systems for reasoning about **resource sharing collective agency**.

The models extend concurrent game models with resource update mechanism, associating with every action profile a table of resource costs for all agents.

Thus, resource consumption and updates are (generally) collective.

The logic extends ATL* with 'resource counters', one per agent. (Extension to multiple resources is quite straightforward.)

Thus, in the course of the play, resources change dynamically.

The available actions for a given agent at every given state depend on the current resource credit of that agent.

All these lead to combined quantitative-qualitative reasoning.

Arithmetic constraints over resources

A simple formal language for dealing with resource updates:

- $R_{\mathbb{A}} = \{r_{\mathbf{a}} \mid \mathbf{a} \in \mathbb{A}\}$:
set of special variables to refer to the accumulated resources;
- Given sets X and $A \subseteq \mathbb{A}$, the set $T(X, A)$ of **terms over X and A** is built from $X \cup R_A$ by applying addition.
- Terms are evaluated in domain of payoffs D (usually, \mathbb{Z} or \mathbb{R}).
- The set $AC(X, A)$ of **arithmetic constraints** over X and A :

$$\{t_1 * t_2 \mid * \in \{<, \leq, =, \geq, >\} \text{ and } t_1, t_2 \in T(X, A)\}$$

- **Arithmetic constraint formulae:**
 $ACF(X, A)$: the set of Boolean formulae over $AC(X, A)$.

Concurrent game models with resource guards and updates

A **CGM with guards and resources (CGMGR)** is a tuple

$$\mathfrak{M} = (\mathcal{S}, \text{resource}, \{g_{\mathbf{a}}\}_{\mathbf{a} \in \mathbb{A}},$$


where $\mathcal{S} = (\mathbb{A}, \text{St}, \{\text{Act}_{\mathbf{a}}\}_{\mathbf{a} \in \mathbb{A}}, \{\text{act}_{\mathbf{a}}\}_{\mathbf{a} \in \mathbb{A}}, \text{out}, \text{Prop}, L)$ is a CGM and:

- $\text{resource} : \mathbb{A} \times S \times \text{Act}_{\mathbb{A}} \rightarrow D$ is a **resource update function**.
- **accumulated resource** of a player \mathbf{a} at a state of a play: the sum of the initial resource credit and all \mathbf{a} 's resource updates incurred in the play so far.
- $g_{\mathbf{a}} : S \times \text{Act}_{\mathbf{a}} \rightarrow \text{ACF}(X, \{a\})$, for $\mathbf{a} \in \mathbb{A}$, is a **guard function** such that $g_{\mathbf{a}}(s, \alpha)$ is an ACF for each $s \in \text{St}$ and $\alpha \in \text{Act}_{\mathbf{a}}$.
 - ▷ The action α is available to \mathbf{a} at s iff the current accumulated resource of \mathbf{a} satisfies $g_{\mathbf{a}}(s, \alpha)$.

The guard must enable at least one action for \mathbf{a} at s .

Example: robots on a mission

Scenario: a team of 3 robots is on a mission. The team must accomplish a certain task, e.g., formalized as 'reaching state *goal*'.




The robots work on batteries which need to be charged in order to provide the robots with sufficient energy to be able to function.

We assume the robots' energy levels are non-negative integers.

Every action of a robot consumes some of its energy.

Collective actions of all robots may, additionally, increase or decrease the energy level of each of them.

Robots on a mission: agents and states



For every collective action: an ‘energy update table’ is associated, representing the net change – increase or decrease – of the energy level after that collective action is performed at the given state.


In this example the energy level of a robot may never go below 0.

Here are the detailed descriptions of the components of the model:

Agents: The 3 robots: **a, b, c.**

States: The ‘base station’ state (*base*) and the target state *goal*.

Robots on a mission: actions and transitions



Actions. The possible actions are:

R : 'recharge', N : 'do nothing', G : 'go to goal', B : 'return to base'.

All robots have the same functionalities and abilities to perform actions, and their actions have the same effect.

Each robot has the following actions possibly executable at the different states: $\{R, N, G\}$ at state *base* and $\{N, B\}$ at state *goal*.

Transitions. The transition function is specified in the figure.

NB: since the robots abilities are assumed symmetric, it suffices to specify the action profiles as multisets, not as tuples.

Robots on a mission: some constraints

- The team has one recharging device which can recharge at most 2 batteries at a time and produces a total of 2 energy units in one recharge step.

So if 1 or 2 robots recharge at the same time they receive a pro rata energy increase, but if all 3 robots try to recharge at the same time, the device does not charge any of them.

- Transition from one state to the other consumes a total of 3 energy units. If all 3 robots take the action which is needed for that transition (G for transition from *base* to *goal*, and B for transition from *goal* to *base*), then the energy cost of the transition is distributed equally amongst them.

If only 2 of them take that action, then each consumes 2 units and the extra unit is transferred to the 3rd robot.

- An attempt by a single robot to reach the other state fails and costs that robot 1 energy unit.

Robots on a mission: resource updates

Resource updates. Resource updates are given below as vectors with components that correspond to the order of the actions in the triple, not to the order of the agents who have performed them.

From state *base*:

Actions	Successor	Payoffs
<i>RRR</i>	<i>base</i>	(0,0,0)
<i>RRN</i>	<i>base</i>	(1,1,0)
<i>RRG</i>	<i>base</i>	(1,1,-1)
<i>RNN</i>	<i>base</i>	(2,0,0)
<i>RNG</i>	<i>base</i>	(2,0,-1)
<i>RGG</i>	<i>goal</i>	(3,-2,-2)
<i>NNN</i>	<i>base</i>	(0,0,0)
<i>NNG</i>	<i>base</i>	(0,0,-1)
<i>NGG</i>	<i>goal</i>	(1,-2,-2)
<i>GGG</i>	<i>goal</i>	(-1,-1,-1)

From state *goal*:

Actions	Successor	Payoffs
<i>NNN</i>	<i>goal</i>	(0,0,0)
<i>NNB</i>	<i>goal</i>	(0,0,-1)
<i>NBB</i>	<i>base</i>	(1,-2,-2)
<i>BBB</i>	<i>base</i>	(-1,-1,-1)

Robots on a mission: guards

At state *base*:

Action	Guard
<i>R</i>	$v \leq 2$
<i>N</i>	<i>true</i>
<i>G</i>	$v \geq 2$
<i>B</i>	<i>false</i>

At state *goal*:

Action	Guard
<i>R</i>	<i>false</i>
<i>N</i>	<i>true</i>
<i>G</i>	<i>false</i>
<i>B</i>	$v \geq 2$

Guards. The same for each robot. The variable v denotes the current resource of the respective robot. Some explanations:

- Action *B* is disabled at state *base* and actions *R* and *G* are disabled at state *goal*.
- No requirements for the 'do nothing' action *N*.
- *R* can only be attempted if the current energy level is ≤ 2 .
- For a robot to attempt a transition to the other state, that robot must have a minimal energy level 2.
- Any set of at least two robots can ensure transition from one state to the other, but no single robot can do that.

Configurations, plays and histories in a CGMGR

Configuration in $\mathfrak{M} = (\mathcal{S}, \text{resource}, \{g_a\}_{a \in \mathbb{A}}, \{d_a\}_{a \in \mathbb{A}})$:
a pair (s, \vec{u}) of a state s and a vector $\vec{u} = (u_1, \dots, u_k)$ of currently accumulated resources of the agents at that state.

The set of possible configurations: $\text{Con}(\mathfrak{M}) = \mathcal{S} \times \mathbb{D}^{|\mathbb{A}|}$.

Partial configuration transition function:

$$\widehat{\text{out}} : \text{Con}(\mathfrak{M}) \times \text{Act}_{\mathbb{A}} \dashrightarrow \text{Con}(\mathfrak{M})$$

where $\widehat{\text{out}}((s, \vec{u}), \vec{\alpha}) = (s', \vec{u}')$ iff $\text{out}(s, \vec{\alpha}) = s'$ and:


- (i) the value u_a assigned to r_a satisfies $g_a(s, \alpha_a)$ for each $a \in \mathbb{A}$
- (ii) $u'_a = u_a + \text{resource}_a(s, \vec{\alpha})$ for each $a \in \mathbb{A}$

The **configuration graph on \mathfrak{M} with an initial configuration (s_0, \vec{u}_0)** consists of all configurations in \mathfrak{M} reachable from (s_0, \vec{u}_0) by $\widehat{\text{out}}$.

A **play** in \mathfrak{M} : an infinite sequence $\pi = c_0 \vec{\alpha}_0, c_1 \vec{\alpha}_1, \dots$ from $(\text{Con}(\mathfrak{M}) \times \text{Act})^\omega$ such that $c_n \in \widehat{\text{out}}(c_{n-1}, \vec{\alpha}_{n-1})$ for all $n > 0$.

A **history**: any finite initial sequence of a play in $\text{Plays}_{\mathfrak{M}}$.

Some configurations and plays in the robots example



Initial configuration: $(base, (0, 0, 0))$.


1. The robots do not coordinate and keep trying to recharge forever. The mission fails:

$(base; 0, 0, 0)(RRR), (base; 0, 0, 0)(RRR), (base; 0, 0, 0)(RRR), \dots$

2. Now the robots coordinate on recharging, 2 at a time, until they each reach energy levels at least 3. Then they all take action G and the team reaches state $goal$ and then succeeds to return to

$(base, 0, 0, 0)(RRN), (base, 1, 1, 0)(NRR), (base, 1, 2, 1)(RNR), (base, 2, 2, 2)(RRN),$
 $(base, 3, 3, 2)(NNR), (base, 3, 3, 4)(GGG)(goal, 2, 2, 3)(BBB), (base, 1, 1, 2) \dots$

More configurations and plays in the robots example



3. Again the robots coordinate on recharging, but after the first recharge Robot **a** goes out of order. Thereafter **a** does nothing while the other two robots try to accomplish the mission by each recharging as much as possible and then both taking action *G*. The team reaches state *goal* but cannot return to *base* and remains stuck at state *goal* forever, for one of the two functioning robots does not have enough energy to apply *B*:

$(base, 0, 0, 0)(RRN)$, $(base, 1, 1, 0)(NRR)$, $(base, 1, 2, 1)(NRR)$, $(base, 1, 3, 2)(NRR)$,
 $(base, 1, 3, 4)(NGG)$, $(goal, 2, 1, 2)(NNB)$, $(goal, 2, 1, 1)(NNN)$, ...

4. As above, but now **b** and **c** apply a cleverer plan and succeed together to reach *goal* and then return to *base*:

$(base, 0, 0, 0)(RRN)$, $(base, 1, 1, 0)(NRR)$, $(base, 1, 2, 1)(NRR)$, $(base, 1, 3, 2)(NGR)$,
 $(base, 1, 2, 4)(NRN)$, $(base, 1, 4, 4)(NGG)$, $(goal, 2, 2, 2)(NBB)$, $(base, 3, 0, 0) \dots$

The logic of qualitative strategic abilities ATL*

Alternating-time Temporal Logic ATL* involves:

- *Coalitional strategic path operators*: $\langle\langle A \rangle\rangle$ for any coalition of agents A . We will write $\langle\langle i \rangle\rangle$ instead of $\langle\langle \{i\} \rangle\rangle$.
- *Temporal operators*: \mathcal{X} (next time), \mathcal{G} (forever), \mathcal{U} (until)

Formulae:

$$\varphi := p \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \langle\langle A \rangle\rangle\varphi \mid \mathcal{X}\varphi \mid \mathcal{G}\varphi \mid \varphi_1 \mathcal{U} \varphi_2$$

Semantics: in concurrent game models.

Extends the semantics for LTL with the clause:

$\langle\langle A \rangle\rangle\varphi$: “The coalition A has a collective strategy to guarantee the satisfaction of the goal φ ” on every play enabled by that strategy.

The Quantitative ATL*: syntax and semantics

State formulae $\varphi ::= p \mid \text{ac} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle\gamma$

Path formulae: $\gamma ::= \varphi \mid \neg\gamma \mid \gamma \wedge \gamma \mid \mathcal{X}\gamma \mid \mathcal{G}\gamma \mid \gamma\mathcal{U}\gamma$

where $A \subseteq \mathbb{A}$, $\text{ac} \in \text{AC}$, and $p \in \text{Prop}$.

Given: \mathfrak{M} be a CGMGR, c a configuration, φ state formula, $\gamma, \gamma_1, \gamma_2$ path formulae, \mathcal{S}^p and \mathcal{S}^o two classes of strategies.

$\mathfrak{M}, c \models p$ iff $p \in L(c^s)$;

$\mathfrak{M}, c \models \text{ac}$ iff $c^u \models \text{ac}$,

$\mathfrak{M}, c \models \langle\langle A \rangle\rangle\gamma$ iff there is a \mathcal{S}^p -strategy s_A such that for all \mathcal{S}^o -strategies $s_{\mathbb{A} \setminus A}$: $\mathfrak{M}, \text{outcome_play}^{\mathfrak{M}}(c, (s_A, s_{\mathbb{A} \setminus A})) \models \gamma$.

$\mathfrak{M}, \pi \models \varphi$ iff $\mathfrak{M}, \pi[0] \models \varphi$,

$\mathfrak{M}, \pi \models \mathcal{X}\gamma$ iff $\mathfrak{M}, \pi[1] \models \gamma$,

$\mathfrak{M}, \pi \models \mathcal{G}\gamma$ iff $\mathfrak{M}, \pi[i] \models \gamma$ for all $i \in \mathbb{N}$,

$\mathfrak{M}, \pi \models \gamma_1\mathcal{U}\gamma_2$ iff there is $j \in \mathbb{N}_0$ such that $\mathfrak{M}, \pi[j] \models \gamma_2$ and $\mathfrak{M}, \pi[i] \models \gamma_1$ for all $0 \leq i < j$.

Ultimately, we define $\mathfrak{M}, c \models \varphi$ iff $\mathfrak{M}, c, 0 \models \varphi$.

Expressing properties in QATL*: examples

Suppose the objective of the team of robots on mission, starting from state *base* where each robot has energy level 0, is to eventually reach the state *goal* and then return to the base station.

Below, 'base' is an atomic proposition true only at state *base* and, 'goal' is an atomic proposition true only at state *goal*.

The following QATL*-formulae are true at $(base, 0, 0, 0)$:

- $\langle\langle\langle\mathcal{G}\rangle\rangle\rangle (r_a \geq 0 \wedge r_b \geq 0 \wedge r_c \geq 0)$
- $\neg\langle\langle\langle\mathbf{a}\rangle\rangle\rangle \mathcal{F} \text{ goal} \wedge \neg\langle\langle\langle\mathbf{b}\rangle\rangle\rangle \mathcal{F} \text{ goal} \wedge \neg\langle\langle\langle\mathbf{c}\rangle\rangle\rangle \mathcal{F} \text{ goal}.$
- $\langle\langle\langle\mathbf{b}, \mathbf{c}\rangle\rangle\rangle \mathcal{F} (\text{goal} \wedge \langle\langle\langle\mathbf{a}, \mathbf{b}, \mathbf{c}\rangle\rangle\rangle (r_a > 0 \wedge r_b > 0 \wedge r_c > 0) \mathcal{U} \text{base}).$
- $\langle\langle\langle\mathbf{b}, \mathbf{c}\rangle\rangle\rangle \mathcal{F} (\text{goal} \wedge \langle\langle\langle\mathbf{b}, \mathbf{c}\rangle\rangle\rangle (r_a > 0) \mathcal{U} (\text{base} \wedge r_a > 0)).$
- $\neg\langle\langle\langle\mathbf{b}, \mathbf{c}\rangle\rangle\rangle \mathcal{F} (\text{goal} \wedge \langle\langle\langle\mathbf{b}, \mathbf{c}\rangle\rangle\rangle \mathcal{F} (\text{base} \wedge (r_b > 0 \vee r_c > 0))).$

Model checking in QATL*

Model checking in QATL* on finite models is generally undecidable, even under very weak assumptions.

Still, there are several practically important decidable cases where the configuration space remains finite, e.g.:

- When resources are not created, but only consumed or re-distributed, but cannot become negative.
- When the possible accumulated amount of resource per agent is bounded above and below.

There are some non-trivial decidable cases with infinite configuration spaces, too.

Concluding remarks

Collective resource bounded reasoning in concurrent multi-agent systems is natural and important.

QATL* is a natural and expressive logic for such reasoning.

Model checking is generally undecidable, but there are practically important decidable cases.

Many open problems and directions for further work.

One of them: allow collective guards on the resources of coalitions.

Good case studies are wanted, too.