Reasoning about normative update

Natasha Alechina   Mehdi Dastani   Brian Logan

EUMAS 2013

What is normative update?

- Normative update of a system with a set of norms is the result of applying the set of norms to the system
- The question we are interested in is: how do norms change agents' behaviours?
- How to check the properties of the normative update?
- Some related work:
  - Ågotnes, van der Hoek, Wooldridge 2008 (logic of norm compliance)
  - Dastani, Grossi, Meyer 2011 (normative update with counts-as rules)
  - Knobbout and Dastani 2012 (acting under norm compliance)

Key contributions

- Previous work on verifying properties of normative updates has considered only a relatively simple view of norms, where some actions or states are designated as violations
- We look at conditional norms and reason both about regimented norms (behaviours violating norms are impossible) and non-regimented norms (violations are possible, but incur a sanction)
- If an (undesirable) state is achievable by the agent(s), how many sanctions do they have to incur in order to achieve it?

Conditional norms

- Assume we have disjoint sets of brute facts propositional atoms Π_b and normative facts propositional atoms Π_s; Π_s contains a distinguished atom $san_\perp$
- Let $cond$, ϕ, d be boolean combinations of propositional variables from Π_b and $san \in \Pi_s$
- A conditional obligation is represented by the tuple

$$(cond, O(\phi), d, san)$$

- A conditional prohibition is represented by the tuple

$$(cond, P(\phi), d, san)$$

- A norm set N is a set of conditional obligations and conditional prohibitions

Meaning of conditional norms

- Conditional norms are evaluated on runs of a transition system
- A conditional norm $n = (cond, Y(\phi), s, san)$, where Y is O or P, is detached in a state satisfying its condition $cond$
- A detached obligation $(cond, O(\phi), d, san)$ is obeyed if no state satisfying d is encountered before execution reaches a state satisfying ϕ, and violated if a state satisfying d is encountered before execution reaches a state satisfying ϕ
- A detached prohibition $(cond, P(\phi), d, san)$ is obeyed if no state satisfying ϕ is encountered before execution reaches a state satisfying d, and violated if a state satisfying ϕ is encountered before execution reaches a state satisfying d
- If a detached norm is violated in a state s, the sanction corresponding to the norm is applied in s

State violating a norm

- A state $\rho[i]$ violates a conditional obligation $(cond, O(\phi), d, san)$ on run ρ iff

  $$\rho, i \models d \wedge \neg\phi \wedge (((\neg\phi \wedge \neg d)\, Since\, (cond \wedge \neg\phi \wedge \neg d)) \vee cond)$$

  i.e., obligations are violated if ϕ does not become true in or before the deadline state

- $\rho[i]$ violates a conditional prohibition $(cond, P(\phi),\ d, san)$ iff

  $$\rho, i \models \phi \wedge \neg d \wedge (((\neg\phi \wedge \neg d)\, Since\, (cond \wedge \neg\phi \wedge \neg d)) \vee cond)$$

  i.e., prohibitions are violated in the first state where ϕ becomes true

Regimentation sanctions and resource sanctions

- Regimentation sanctions ensure that certain behaviours never occur
- If a norm labels a state with the distinguished sanction atom $san_\perp$, then the run containing this state is removed from the set of runs of the system by the normative update
- Resource sanctions treat sanctions essentially like fines or taxes
- Penalize rather than eliminate certain execution paths by reducing the resources of the agent

Normative update

- Let $M = (S, R, V)$ be a finite transition system with initial state s_o and N a finite set of conditional obligations and prohibitions
- A normative update of M with N, $M^N = (S^N, R^N, V^N)$, is a tree unravelling $T(M)$ of M where all norms from N are enforced on all runs:
  - in each tree node s', $V^N(s')$ contains sanction atoms for all norms violated in s'
  - paths which contain a state satisfying the distinguished sanction atom $san_\perp$ are removed from M^N

Example

Consider an obligation $(c, O(q), d, sanO)$ and a prohibition $(c', P(p), d', sanP)$

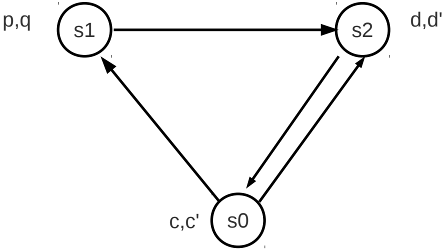

Example

Consider an obligation $(c, O(q), d, sanO)$ and a prohibition $(c', P(p), d', sanP)$

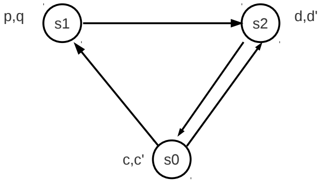

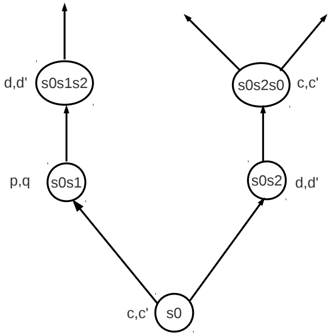

Example: normative update

Consider an obligation $(c, O(q), d, sanO)$ and a prohibition $(c', P(p), d', sanP)$

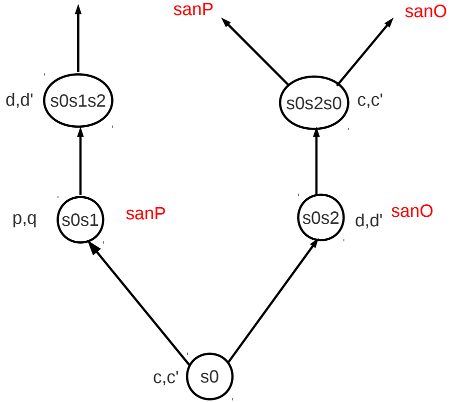

Language of *CTLS*

- We propose a new logic, *CTLS* , for reasoning about normative updates of single-agent systems
- *CTLS* is *CTL* with Sanction bounds
- Path quantifiers have the form $E^{\leq Z}$, where Z is a multiset of sanction bounds (where multiplicity of a sanction *san* is infinity, we represent this as $\infty * san$)
- $E^{\leq Z}$ means 'there exists a path of sanction cost at most Z '

$$p \in \Pi_b \cup \Pi_s \mid \neg\phi \mid \phi \wedge \phi \mid E^{\leq Z} X\ \phi \mid E^{\leq Z} \phi\ \mathit{Until}\ \phi \mid E^{\leq Z} G\ \phi$$

Semantics of *CTLS*

The truth of *CTLS* formulas is defined relative to a tree model T (intuitively, a normative update) and a state $s \in T$:

- $T, s \models E^{\leq Z} X\phi$ iff there exists a fullpath ρ' with $\rho'[0] = s$, such that $T, \rho'[1] \models \phi$ and $sanctions(\rho') \leq Z$
- $T, s \models E^{\leq Z}\phi \; Until \; \psi$ iff there exists a fullpath ρ' with $\rho'[0] = s$, such that for some $n \geq 0$, $T, \rho'[n] \models \psi$ and for every i, $i < n$, $T, \rho'[i] \models \phi$ and $sanctions(\rho') \leq Z$
- $T, s \models E^{\leq Z} G\ \phi$ iff there exists a fullpath ρ' with $\rho'[0] = s$, such that for every i, $T, \rho'[i] \models \phi$ and $sanctions(\rho') \leq Z$

Model-checking problem for normative update in *CTLS*

- The model-checking problem for a normative update in *CTLS* takes as inputs
  - a finite transition system $M = (S, R, V)$,
  - a state $s_0 \in S$,
  - a finite set of conditional norms N, and
  - a formula ϕ of *CTLS*
- It returns true if $M^N, s_0 \models \phi$, and false otherwise

Complexity of *CTLS* normative update checking

- The model-checking problem for a normative update in *CTLS* is in PSPACE (proof uses guessing and checking a polynomially representable path)
- It is PSPACE-hard by reduction of QSAT problem (proof idea adapted from Bulling and Jamroga's (IJCAI 2011) proof of PSPACE-hardness of CTL^{+})

ATLS

- We also consider normative update of a multi-agent system (concurrent game structure)
- Properties of a normative update of a MAS can be expressed in *ATLS* (*ATL* with Sanction bounds)

$$p \in \Pi_b \cup \Pi_s \mid \neg\phi \mid \phi \wedge \psi \mid \langle\langle C \rangle\rangle^{\leq Z} \mathcal{X}\phi \mid \langle\langle C \rangle\rangle^{\leq Z} G\phi \mid \langle\langle C \rangle\rangle^{\leq Z} \phi \mathcal{U} \psi$$

- $\langle\langle C \rangle\rangle^{\leq Z}\gamma$ means 'the group of agents C has a strategy, all executions of which incur at most Z sanctions and satisfy the formula γ, whatever the other agents in $\mathcal{A} \setminus C$ do'

ATLS semantics

- Inspired by Resource-Bounded ATL (Alechina, Logan, Nguyen & Rakib, IJCAI 2009), with sanction costs of strategies defined in terms of sanction costs of paths
- Normative update defined as for single agent case (assume all norms apply to individual agents)
- $M^N, s \models \langle\langle C \rangle\rangle^{\leq Z} \gamma$ iff there exists a strategy F_C of sanction cost at most Z in s such that for all $\rho \in out(s, F_C)$, $M^N, \rho \models \gamma$

Complexity for *ATLS* normative update checking

- The model-checking problem for a normative update in *ATLS* is in PSPACE
- It is PSPACE-hard from PSPACE-hardness of *CTLS*

Summary

- The model-checking problems for normative updates of both single and multi-agent systems is PSPACE-complete
- Future work: define normative update for group norms