Coordination of large collections of “uncertain” switched systems

Necmiye Ozay, EECS
University of Michigan, Ann Arbor

Workshop “Roberto Tempo” on Uncertain Dynamical Systems
Banyuls sur mer
July 7, 2017

Research partly funded by

Petter Nilsson

Johanna Mathieu

Contents ☺

- What is in this talk?
  - Multi-agent systems
  - Positive systems
  - Switched systems
  - Optimization (relaxations)
  - "Uncertainty"
  - Graph theory
- What is not in this talk?
  - Stability
  - Frequency domain

Motivation and applications

- Large-scale, complex, distributed sensing, actuation and control systems:
  - **Smart grid, Smart buildings, Aircraft systems, Automotive, Robotics, Manufacturing & Automation, Security & Surveillance**

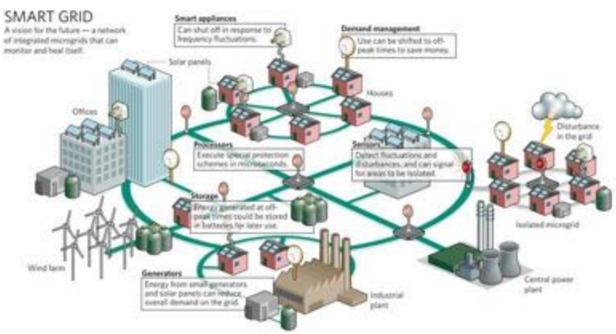

Observations:

- A very large number of (discrete & continuous) states and decision variables
- Complex requirements → need controllers too comp designed/analyzed b

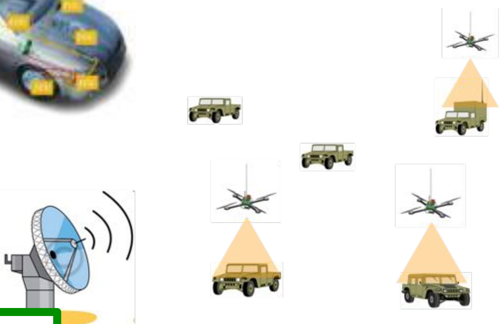

Scalable tools for data analysis, **control design** and verification (theory and software) are lagging!!!

Formal methods in control

- Models for:
  - the system (usually hybrid/ switched ODEs, with continuous/ discrete inputs, disturbances and parametric uncertainty)
  - the environment (faults, external events)
- Formalized assumptions and requirements
  - linear temporal logic and its extensions
- Methods for verification and synthesis
  - algorithms that can process formal models and requirements to do analysis and control synthesis

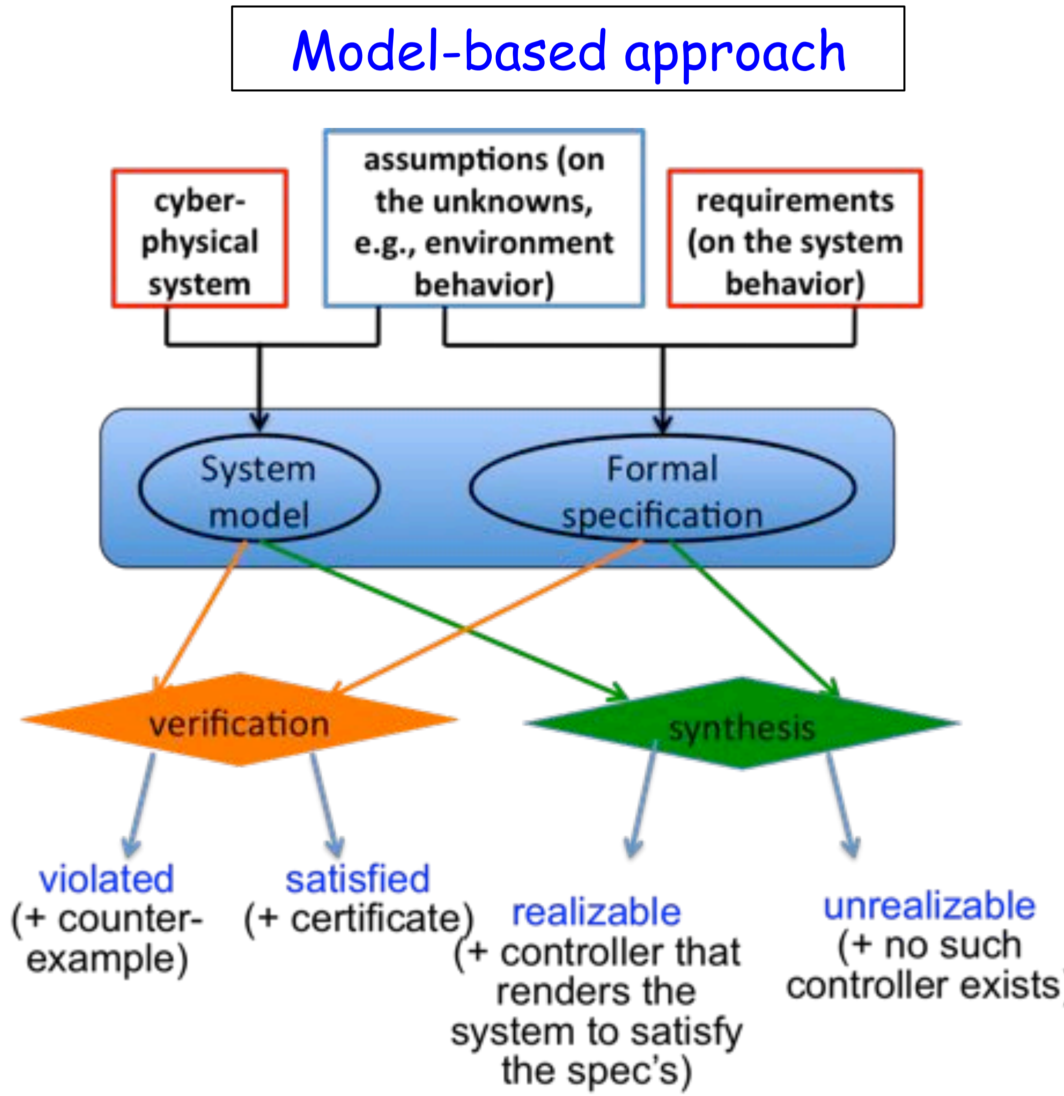

System models

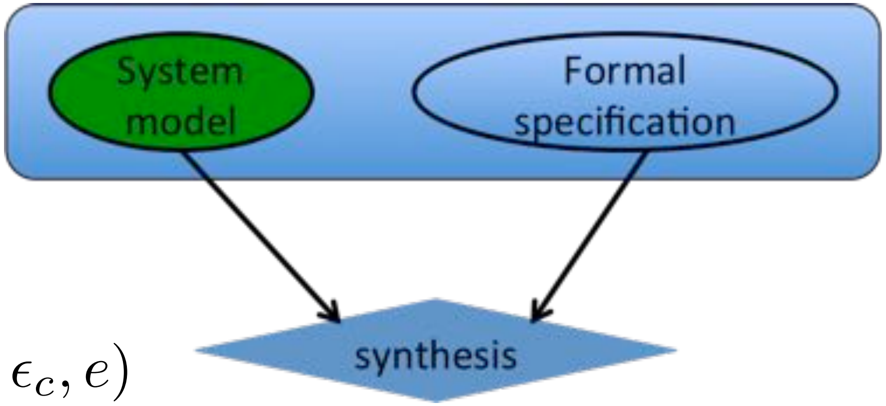

Differential equations (continuous-time):

$$\dot{x} = f(x, u_c, u_d, \epsilon_c, e)$$

Or, difference equations (discrete-time):

$$x(k+1) = f(x(k), u_c(k), u_d(k), \epsilon_c(k), e(k))$$

$$\begin{aligned} x \in \mathcal{X} &: \text{ state} \\ u_c \in \mathcal{U}_c &: \text{ continuous control input} \\ u_d \in \mathcal{U}_d &: \text{ discrete control input} \\ \epsilon_c \in \mathcal{D}_c &: \text{ disturbance input} \\ e \in \mathcal{D}_d &: \text{ discrete uncontrollable input} \end{aligned}$$

$$\boxed{\mathcal{X} \subset \mathbb{R}^N}$$

Some characteristics:

- Hard constraints (on input and states)
- Infinite horizon specifications
- Hybrid (either the system or the controller or both)
- Robust/reactive

State-of-the-art in formal methods in control (incomplete list!)

- Hard state/input constraints, hybrid dynamics, complex specifications (e.g., temporal logics)
  - Belta, Fainekos, Girard, Murray, Pappas, Tabuada, Tomlin …
- Applications (with "small" state-space dim.)
  - Robotics, building thermal management, adaptive cruise control, aircraft subsystems, traffic control
- "Medium"-scale systems
  - Monotonicity (Hafner & Del Vecchio 11, Coogan & Arcak 15)
  - Multi-scale abstractions for safety (Girard et al. 13)
  - Compositional synthesis (Nilsson & Ozay, Chen et al., Kim et al.), incremental abstractions (Nilsson & Ozay)

State-of-the-art in formal methods in control (incomplete list!)

- Hard state/input constraints, hybrid dynamics, complex specifications (e.g., temporal logics)
  - Belta, Fainekos, Girard, Murray, Pappas, Tabuada, Tomlin …
- Applications (with “small” state-space dim.)
  - Robotics, building thermal management, adapt
    aircraft subsystems, traffic control

Recurring theme:
structural properties

- “Medium”-scale systems
  - Monotonicity (Hafner & Del Vecchio 11, Coogan & Arcak 15)
  - Multi-scale abstractions for safety (Girard et al. 13)
  - Compositional synthesis (Nilsson & Ozay, Chen et al., Kim et al.), incremental abstractions (Nilsson & Ozay)
- “Large”-scale (but not synthesis)
  - Parametric verification of rectangular hybrid automata (Johnson & Mitra 12)
  - Abstractions of large collections of stochastic systems (Soudjani & Abate 15)

Large collections of systems

Example 1: Emergency response with a robotic swarm

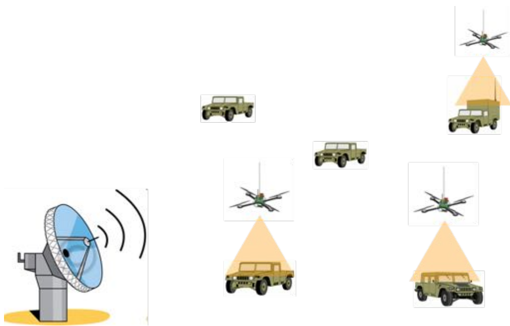

Creative commons public license

- Deploy a large collection of robots (e.g., quadrotors, ground vehicles) for search and rescue mission
- Plan trajectories by taking dynamic constraints into account
- Requirements:
    - Sufficiently many robots in certain areas at any given time
    - Not too many robots in certain regions (danger zones)
    - Collision avoidance
    - Charging/reporting constraints

Large collections of systems

Example 2: Coordination of thermostatically controlled loads (TCLs)

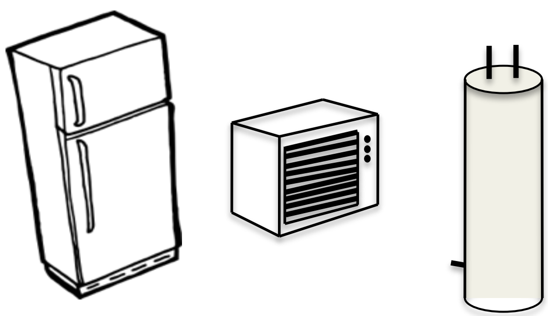

TCLs

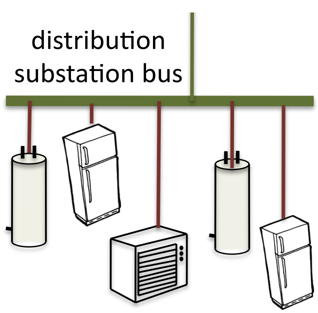

thermostatically
controlled loads
(TCLs)

- Thermostatically controlled loads (e.g., refrigerators, air conditioners, water heaters) for demand response
- Thermal dynamics can be controlled via ON/OFF switches
- Requirements:
  - <u>Not too many</u> TCLs ON at the same time (to avoid line overload)
  - <u>Sufficiently many</u> ON all the time (to utilize renewable energy)
  - Local temperature constraints (never out of desired temperature range)

Mathieu, Koch, Callaway, IEEE Trans. on Power Systems

Common structural properties

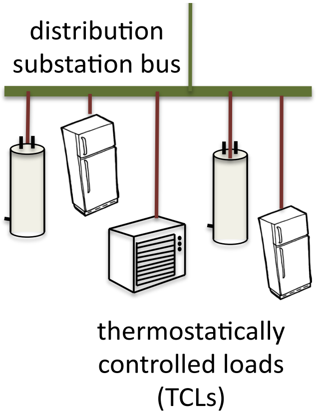

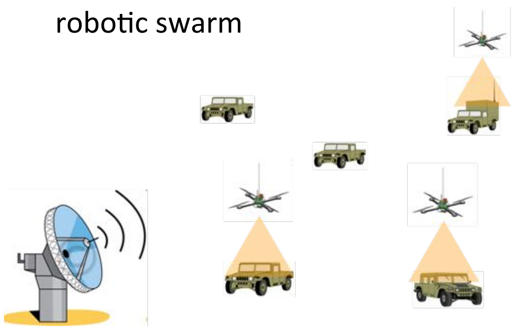

- Large number of systems, small number of classes
- Counting constraints: “how many in each mode?”, “how many in what region?”
- Identity of individual systems is not important

For simplicity, assume:

- dynamics are identical within each class
- (wlog) there is only one class

Mathematical formulation: TCLs

The temperature θ of a TCL has dynamics

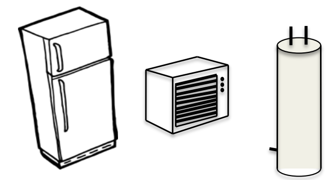

$$\dot{\theta}_i = \begin{cases} f_{on}(\theta_i), & \text{if TCL is on} \\ f_{off}(\theta_i), & \text{if TCL is off} \end{cases}$$

Suppose we have a collection of TCL's $\{\theta_i\}_{i \in [N]}$.

- Customers: Want TCL temperature to be close to a desired temperature θ_i^{des}, but small deviations are allowed.

$$\|\theta_i - \theta_i^{des}\| \le \Delta \quad (1)$$

- Utility company: Wants to control aggregate demand, i.e. the number of TCLs that are on

$$\sum_{i=1}^{N} \mathbb{1}_{\{\text{TCL } i \text{ is on}\}} \quad (2)$$

Goal: Find a switching (i.e., on/off) strategy that exploits the flexibility in (1) so that (2) can be controlled.

Mathematical formulation: General

- N identical switched system with M modes:

$$\dot{x}_i(t) = f_{\sigma_i(t)}(x_i(t)), \quad \sigma_i : \mathbb{R} \mapsto [M],$$

- Mode-specific unsafe sets: $\mathcal{U}_m$, $m \in [M]$
    - Equivalent to forced mode switches.
- Mode-counting bounds:

$$\underline{K}_m \leq \sum_{i=1}^{N} \mathbb{1}_m(\sigma_i(t)) \leq \overline{K}_m \tag{3}$$

Want to synthesize a switching strategy σ_i such that (3) satisfied over time.

Structural property: both the dynamics and the specification (counting constraints) are permutation invariant!

Solution overview

- Construct symbolic abstractions (i.e., a finite transition system) and aggregate dynamics and define “equivalent” problems on these structures
- (Analyze abstractions to understand fundamental limitations if any)
- An optimization-based solution approach
- Analysis of the solution approach

Solution overview

- Construct symbolic abstractions (i.e., a finite transition system)
  - ε-approximate bisimilar abstraction

$$\dot{x} = f(x, {\color{blue}u_1})$$

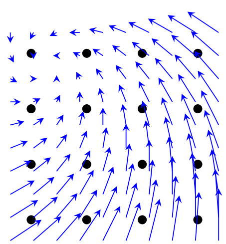

$$\dot{x} = f(x, {\color{red}u_2})$$

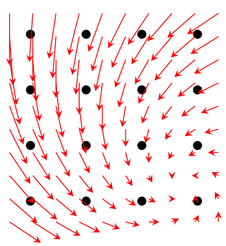

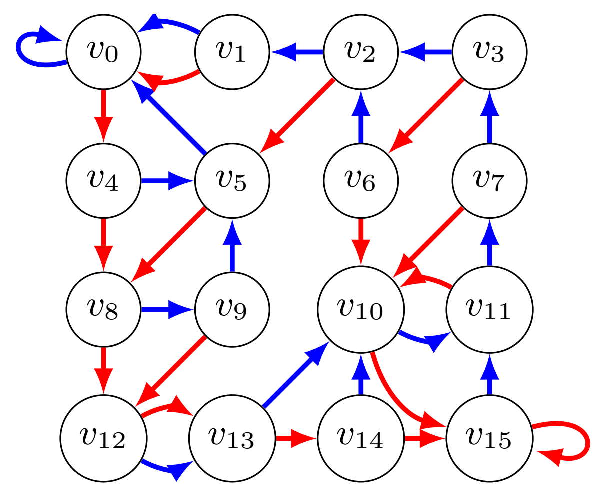

  - for each path on the finite transition system, there is a piecewise constant input that generates a trajectory such that time-sampled trajectory remains ε-close to the discrete states

Abstraction of individual dynamics

- Assume dynamics are δ-GAS with $\mathcal{KL}$ functions β_m

$$\|\phi_t^m(x) - \phi_t^m(y)\|_\infty \leq \beta_m\left(\|x-y\|_\infty, t\right). \tag{4}$$

Abstraction of individual dynamics

- Assume dynamics are δ-GAS with $\mathcal{KL}$ functions β_m

$$\|\phi_t^m(x) - \phi_t^m(y)\|_\infty \leq \beta_m\left(\|x - y\|_\infty, t\right). \tag{4}$$

- With discretization in time (τ) and space (η), an ϵ-approximate bisimilar model is obtained if $\beta_m(\epsilon, \tau) + \frac{\eta}{2} \leq \epsilon$.

Pola, Girard, Tabuada 08

Abstraction of individual dynamics

- Assume dynamics are δ-GAS with $\mathcal{KL}$ functions β_m

$$\|\phi_t^m(x) - \phi_t^m(y)\|_\infty \leq \beta_m\left(\|x-y\|_\infty, t\right). \tag{4}$$

- With discretization in time (τ) and space (η), an ϵ-approximate bisimilar model is obtained if $\beta_m(\epsilon,\tau) + \frac{\eta}{2} \leq \epsilon$.
  - Mode 1 abstraction

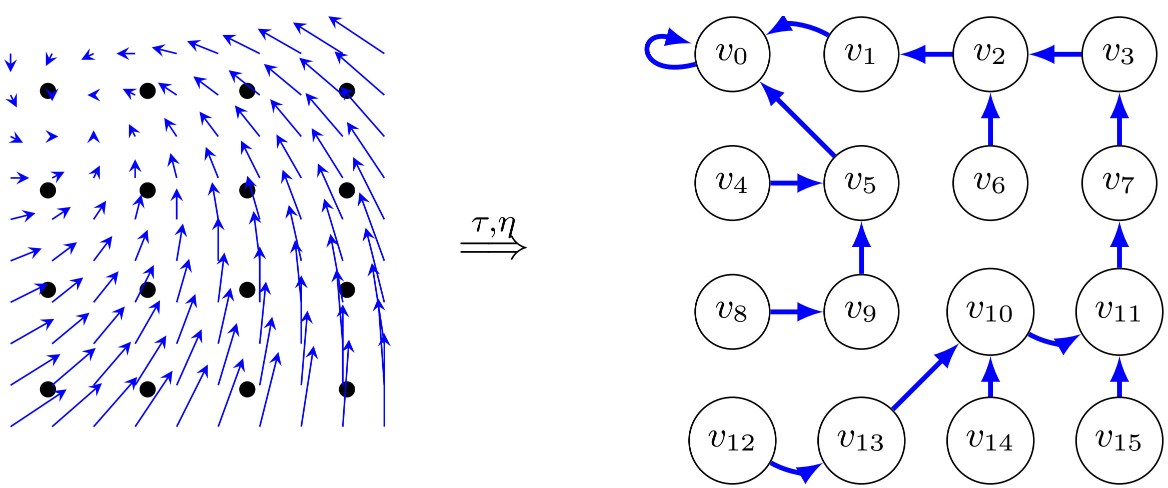

Abstraction of individual dynamics

- Assume dynamics are δ-GAS with $\mathcal{KL}$ functions β_m

$$\|\phi_t^m(x) - \phi_t^m(y)\|_\infty \leq \beta_m\left(\|x-y\|_\infty, t\right). \qquad (4)$$

- With discretization in time (τ) and space (η), an ϵ-approximate bisimilar model is obtained if $\beta_m(\epsilon,\tau) + \frac{\eta}{2} \leq \epsilon$.
  - Mode 2 abstraction

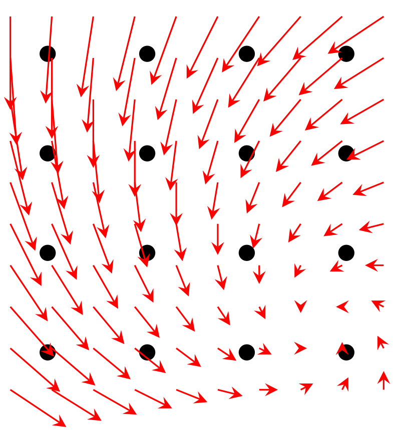

$\xRightarrow{\tau,\eta}$

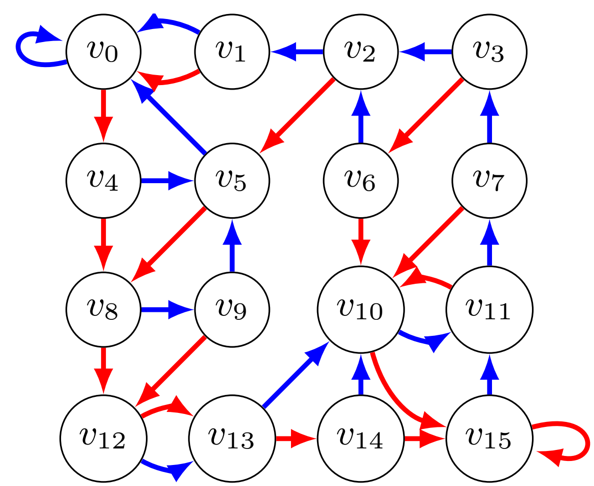

mode-transition graph $G = (V, E)$

Some observations

- For a homogeneous collection, each system will have an identical mode-transition graph
- Transition graphs are deterministic

mode-transition graph $G = (V, E)$

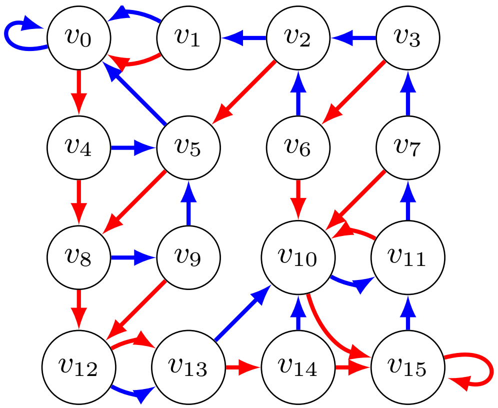

Some observations

- For a homogeneous collection, each system will have an identical mode-transition graph
- Transition graphs are deterministic
- Consider mild heterogeneity

mode-transition graph $G = (V, E)$

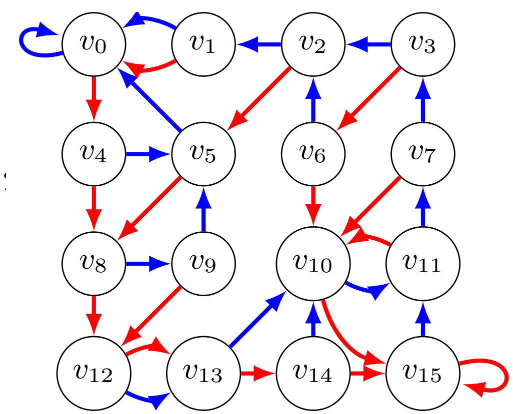

$$\dot{x}_i(t) = f_{\sigma_i(t)}(x_i(t), d_i(t)), \quad \sigma_i : \mathbb{R} \mapsto [M],$$

where $d_i \in \mathcal{D}$ (bounded parametric uncertainty or disturbance). If $f_m(x, d)$ is L_m-Lipschitz in x, and

$$||f_m(x, d) - f_m(x, 0)|| \leq \delta_m \quad \text{for all } \; d_i \in \mathcal{D},$$

then, with discretization in time (τ) and space (η), an ϵ-approximate bisimilar model is obtained if
$\beta_m(\epsilon, \tau) + \frac{\delta_m}{L_m}(e^{L_m \tau} - 1) + \frac{\eta}{2} \leq \epsilon$.

Aggregate dynamics on graph

Let $V = \{v_1, \ldots v_K\}$ denote the nodes of mode-transition graph $G = (V, E)$. Introduce the states $w_k^{m_1}$ and $r_k^{m_1,m_2}$.

- w_i^m represents **number of systems in mode** m **at** v_k.
- $r_k^{m_1,m_2}$ represents **number of systems at** v_k **that switch from** m_1 **to** m_2**.**

- The dynamics become

$$(w_k^{m_1})^+ = \sum_{j \in \mathcal{N}_k^{m_1}} \left(w_j^{m_1} + \sum_{m_2} r_j^{m_2,m_1} - r_j^{m_1,m_2} \right),$$

- Constrained control actions:

$$0 \leq \sum_{m_2} r_k^{m_1,m_2} \leq w_k^{m_1},$$

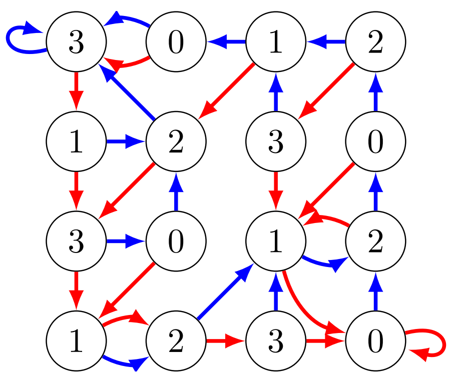

- Compact description: $\mathbf{w}^+ = A\mathbf{w} + B\mathbf{r}$

Nilsson, Ozay, HSCC 16, arxiv 17

Equivalent problem on aggregate dynamics

Theorem 1:

Consider aggregate dynamics $\Sigma_G : \mathbf{w}^+ = A\mathbf{w} + B\mathbf{r}$ with safety and mode-counting constraints:

$$w_k^m(t) = 0 \quad \forall k \in U_m, \tag{5}$$

$$\underline{K}_m \leq \sum_{i \in [N]} w_i^m(t) \leq \overline{K}_m. \tag{6}$$

Then,

- if $\exists$ sequence of control inputs $\mathbf{r}^\omega$ for Σ_G that enforce (5) and (6) with $U_m + B_\epsilon$, then $\exists$ a solution to the original problem.
- if $\nexists$ a sequence of control input $\mathbf{r}^\omega$ for Σ_G that enforces (5) and (6) with $U_m - B_\epsilon$, then no solution to the original problem.

We will focus on aggregate dynamics. We need infinite horizon strategies!

Solution strategy: from a given initial state, steer the system, while respecting the constraints, to a **nice state** from which a periodic input suffices.

Controllability-like conditions

Solution strategy: **from** a given **initial state**, **steer the system**, while respecting the constraints, **to** a **nice state** from which a periodic input suffices.

- Let's put the mode-counting constraints aside.
- Are there any fundamental limitations on what states can be reached from an initial condition?

$$\Sigma_G : \mathbf{w}^+ = A\mathbf{w} + B\mathbf{r}$$

with local safety and
input constraints

Controllability-like conditions

Solution strategy: from a given initial state, steer the system, while respecting the constraints, to a **nice state** from which a periodic input suffices.

$\Sigma_G : \mathbf{w}^+ = A\mathbf{w} + B\mathbf{r}$

with local safety and input constraints

- Let's put the mode-counting constraints aside.
- Are there any fundamental limitations on what states can be reached from an initial condition?

Definition: The period n of a strongly connected graph is the greatest common divisor of the lengths of its cycles.

Theorem 2: If the connected components of mode-transition graph has period n=1, any state is reachable from any other state (within the connected component). If n>1, then the reachable states live on a hyperplane arrangement with n hyperplanes.

Solution strategy

Solution strategy: from a given initial state, steer the system, while respecting the constraints, to a **nice state** from which a periodic input suffices.

- **Prefix:** for a fixed horizon T, given initial state, we will steer the state at time T to "**nice**" cycles
- **Suffix:** let individual systems circulate in the cycles

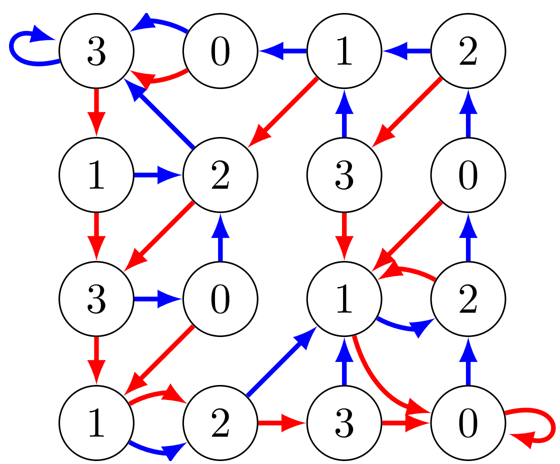

Cycle terminology

- Cycle $C = \{v_{c_1}, \ldots, v_{c_{|C|}}\}$ in G
- A *cycle assignment* for C is a function $\alpha : C \mapsto \mathbb{R}^+$.

Mode-counts on for a cycle assignment:

- Max-count $\overline{\Psi}^m(C, \alpha)$: maximal number of individual systems simultaneously in mode m when circulating α in C:
- Min-count $\underline{\Psi}^m(C, \alpha)$: minimal number of individual systems simultaneously in mode m when circulating α in C:

Illustration: cycles

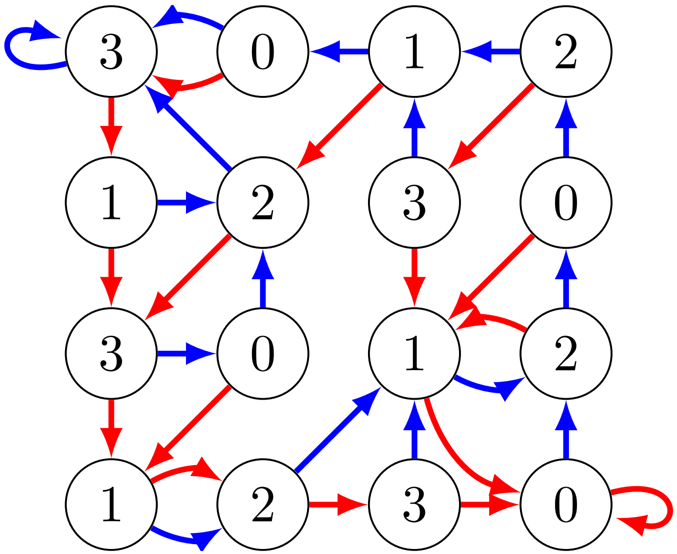

Illustration: cycles

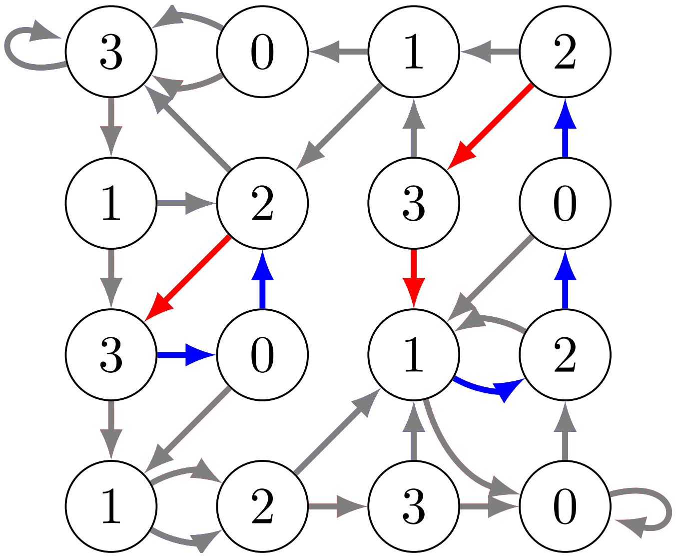

Illustration: cycles

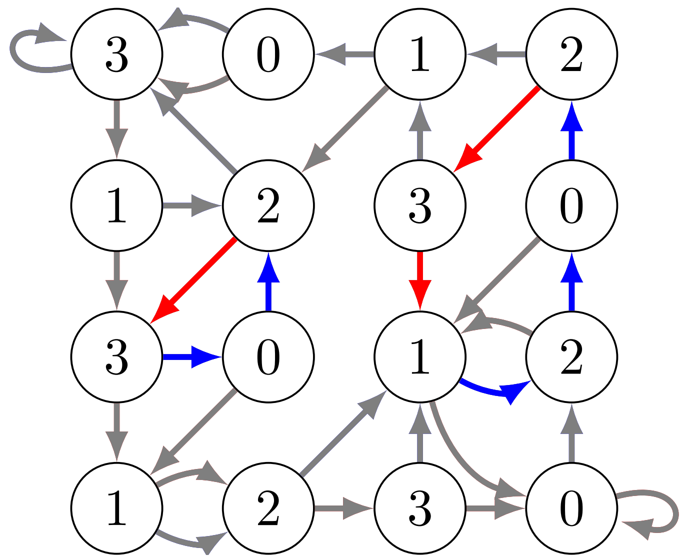

- Big cycle C_1, assignment $\alpha_1 = [1, 2, 0, 2, 3]$, gives red counts

$$\underline{\Psi}(C_1, \alpha_1) = 2, \quad \overline{\Psi}(C_1, \alpha_1) = 5$$

Illustration: cycles

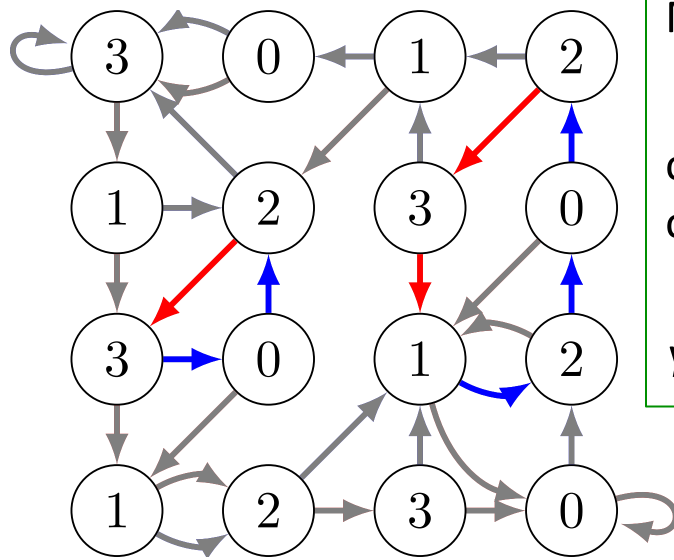

Mode-counting constraints

$$\underline{\Psi}^m(C,\alpha) \geq \underline{K}_m, \quad \overline{\Psi}^m(C,\alpha) \leq \overline{K}_m,$$

can be represented as linear constraints

$$\underline{K}_m \mathbf{1} \leq Y_C^m \alpha \leq \overline{K}_m \mathbf{1}$$

Y_c^m is a circulant matrix.

- Big cycle C_1, assignment $\alpha_1 = [1, 2, 0, 2, 3]$, gives red counts

$$\underline{\Psi}(C_1, \alpha_1) = 2, \quad \overline{\Psi}(C_1, \alpha_1) = 5$$

- Small cycle C_2, assignment $\alpha_2 = [3, 0, 2]$, gives red counts

$$\underline{\Psi}(C_2, \alpha_2) = 0, \quad \overline{\Psi}(C_2, \alpha_2) = 3$$

Solution via linear programming

For cycles $C_1, \ldots, C_J$, required mode-counts K_m, horizon T

find $\alpha_1, \ldots, \alpha_J$ cycle assignments,

$\mathbf{r}(0), \ldots, \mathbf{r}(T-1)$,

$\mathbf{w}(0), \ldots, \mathbf{w}(T)$,

s.t. $\underline{K}_m \le \sum w_i^m(t) \le \overline{K}_m, \quad 0 \le t \le T-1$ mode-counting during prefix

$\underline{K}_m$

de-counting during suffix

$\sum_j$

$\Lambda($

undary conditions between efix and suffix

Feasibility problem with linear constraints:

- integrality constraints on the inputs (ILP)
- relaxing integrality (LP)

Number of constraints and variables are independent of the number of systems N!

$\mathbf{w}(t+1) = A\mathbf{w}(t) + B\mathbf{r}(t), \quad t = 0, \ldots, T-1,$ system dynamics

$\Lambda(\mathbf{w}(0)) = \lambda_0,$

$\sum_{m_2} r_j^{m_1, m_2} = w_j^{m_1}$ for all $j \in \bigcup_{i \in U_{m_1}} \mathcal{N}_i^{m_1},$

local safety constraints

$r_j^{m_2, m_1} = 0$ for all $m_2 \in [M], j \in U_{m_1},$

control constraints.

Analysis

- **Integer solutions (ILP)**
  - **Completeness of prefix-suffix solutions:** There exists a finite T and some maximal cycle length L such that ILP with all cycles with length less than L provides a complete solution to the original problem
  - From any feasible ILP solution, we can extract a solution to the original problem
- **Non-integer solutions (LP):**
  - Enough to consider simple cycles
  - Gives certificates for non-existence of solutions
- **Rounding a non-integer solution:**
  - A non-integer solution over the cycles can be rounded to an integer feasible solution with mode counting loss at most

$$\underline{\Psi}^m(C, \alpha_{int}) \leq \underline{\Psi}^m(C, \alpha_{avg}) + \frac{|C|}{4}$$

Intuition behind cycles: TCLs

$$\dot{\theta}_i = -a(\theta_i - \theta_a) - bP_m$$

θ :room temperature

θ_a :ambient temperature

$P_m = 0$ when OFF

$P_m = 5.6$ when ON

local safety
$\theta_i \in [21.5, 23.5]$

For an individual system if only local ON/OFF control is used (no demand response for extra switching), the temperature evolves as follows:

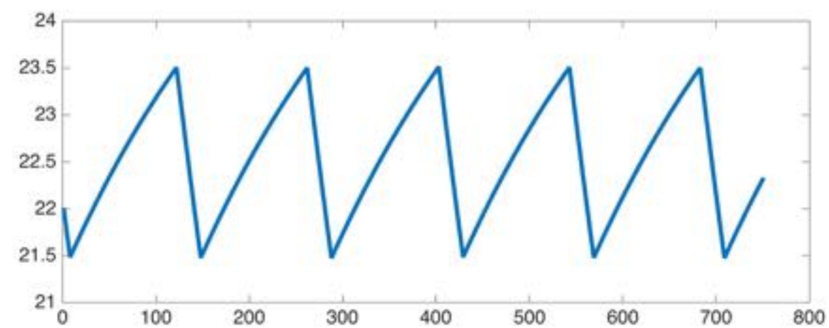

Parameters from Mathieu, Koch, Callaway, IEEE Trans. on Power Systems, 2013

Intuition behind cycles: TCLs

$$\dot{\theta}_i = -a(\theta_i - \theta_a) - bP_m$$

θ :room temperature

θ_a :ambient temperature

$P_m = 0$ when OFF

$P_m = 5.6$ when ON

local safety

$\theta_i \in [21.5, 23.5]$

For an individual system if only local ON/OFF control is used (no demand response for extra switching), the temperature evolves as follows:

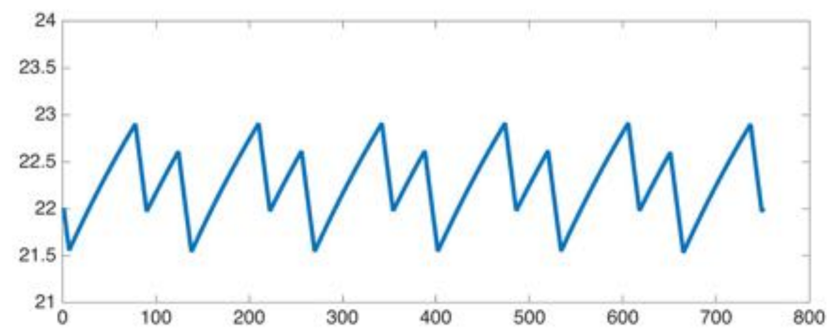

Roughly, cycles are defining new "bands" within the dead-band allowed by the local safety constraints. That is, we are changing the duty cycle.

Parameters from Mathieu, Koch, Callaway, IEEE Trans. on Power Systems, 2013

Results on TCLs

N = 10000 units

10000-D state-space with 2^{10000} modes!

$$\dot{\theta}_i = -a(\theta_i - \theta_a) - bP_m$$

θ :room temperature

θ_a :ambient temperature

$P_m = 0$ when OFF

$P_m = 5.6$ when ON

local safety

$\theta_i \in [21.5, 23.5]$

Two different runs with different mode-counting constraints (also stricter constraints at the suffix)

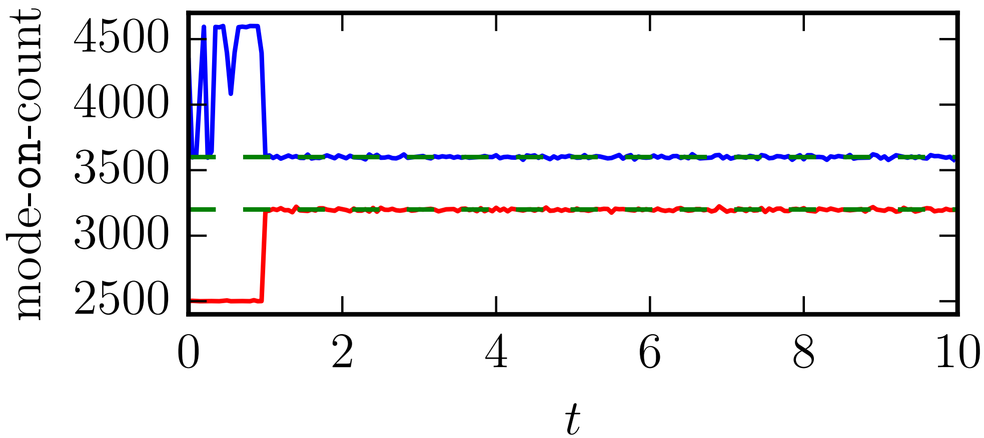

Lower mode-count:

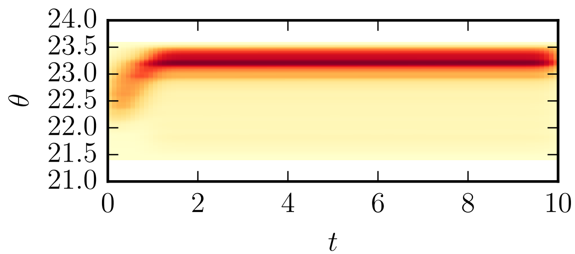

Higher mode-count:

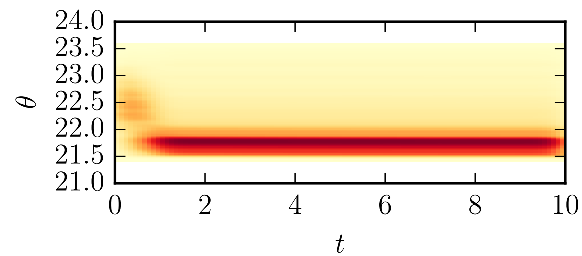

Parameters from Mathieu, Koch, Callaway, IEEE Trans. on Power Systems, 2013

Beyond mode counting

- Counting the agents in a region of state-space
- Time-evolution of counting constraints (counting LTL)

$$\varphi ::= \mathit{True} \mid cp \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi \mid \bigcirc\varphi \mid \varphi_1 \,\mathcal{U}\, \varphi_2,$$

- cp = [atom prop., count]
- Possible to encode asynchrony as well

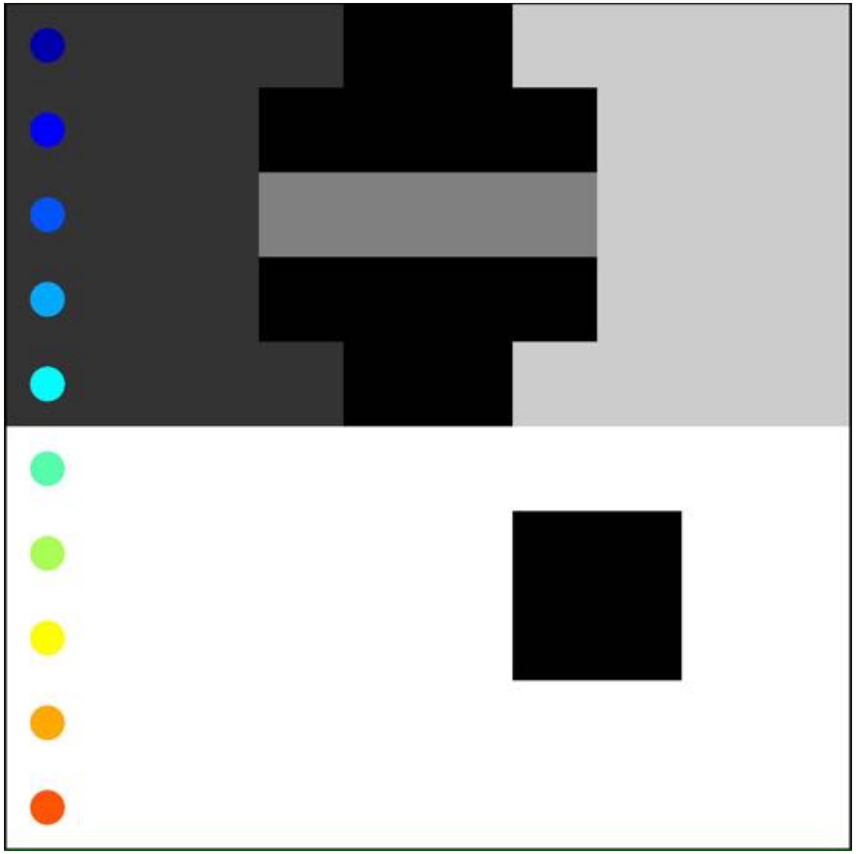

With Yunus Emre Sahin & Petter Nilsson
ICCPS17

Summary: structure for scalability

- A control synthesis method for large collections of systems with **counting constraints**
  - exploits the symmetry (permutation invariance) in the dynamics and in specifications
  - works across scales (10 to 10K or more systems)
  - with potential applications in different domains
  - extensions to asynchrony, counting temporal logic

Preprints and more information available @ http://web.eecs.umich.edu/~necmiye/

Summary: structure for scalability

- A control synthesis method for large collections of systems with **counting constraints**
  - exploits the symmetry (permutation invariance) in the dynamics and in specifications
  - works across scales (10 to 10K or more systems)
  - with potential applications in different domains
  - extensions to asynchrony, counting temporal logic
- Current work
  - partial information
  - non-deterministic abstractions (for not incrementally stable systems), asynchronous switching
  - tighter rounding bounds between LP and ILP
  - other types of symmetries that can be exploited

Preprints and more information available @ http://web.eecs.umich.edu/~necmiye/