Fluid Model Checking

Fluid Approximation for Checking Logic Properties in Markov Population Models

Luca Bortolussi[1,2]

[1]Dipartimento di Matematica e Geoscienze
Università degli studi di Trieste
[2]CNR/ISTI, Pisa
luca@dmi.units.it

Joint work with Jane Hillston and Roberta Lanciani

Bertinoro Summer School in Formal Methods
June 17-21, 2013

COLLECTIVE DYNAMICS

The behaviour of many systems can be interpreted as the result of the collective behaviour of a large number of interacting entities.

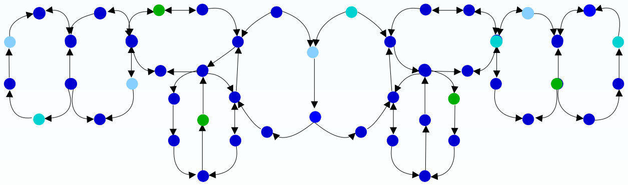

For such systems we are often as interested in the population level behaviour as we are in the behaviour of the individual entities.

COLLECTIVE BEHAVIOUR

In the natural world there are many instances of collective behaviour and its consequences:

COLLECTIVE BEHAVIOUR

In the natural world there are many instances of collective behaviour and its consequences:

COLLECTIVE BEHAVIOUR

This is also true in the man-made and engineered world:

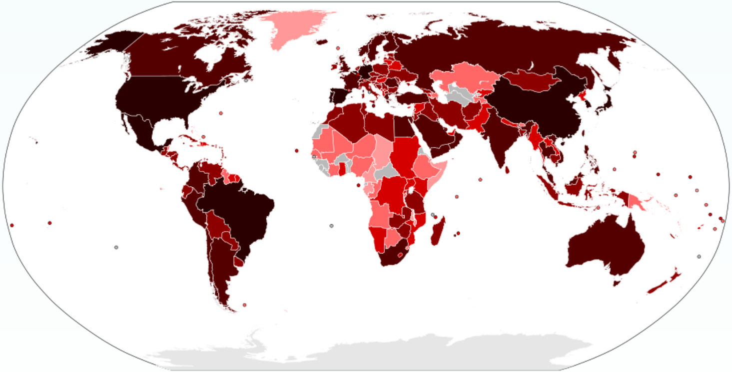

Spread of H1N1 virus in 2009

COLLECTIVE BEHAVIOUR

This is also true in the man-made and engineered world:

Love Parade, Germany 2006

COLLECTIVE BEHAVIOUR

This is also true in the man-made and engineered world:

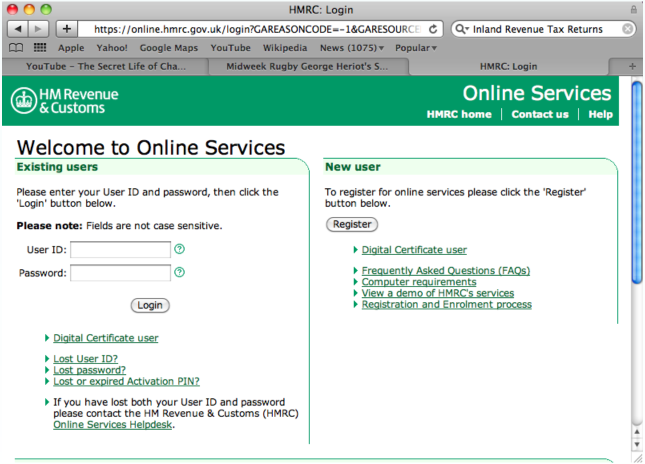

Self assessment tax returns 31st January each year

SOLVING DISCRETE STATE MODELS

With compositional modelling approaches we have a CTMC with global states determined by the local states of all the participating components.

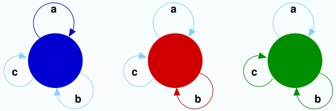

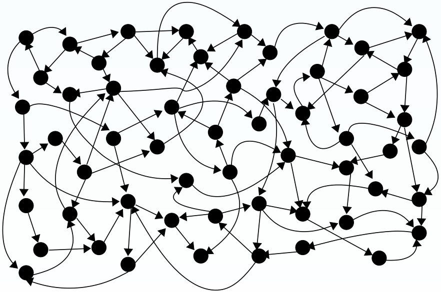

SOLVING DISCRETE STATE MODELS

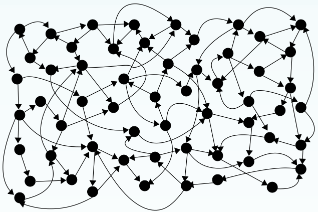

When the size of the state space is not too large they are amenable to NUMERICAL SOLUTION (linear algebra) to determine a STEADY STATE or TRANSIENT PROBABILITY DISTRIBUTION.

$$Q = \begin{pmatrix} q_{1,1} & q_{1,2} & \cdots & q_{1,N} \\ q_{2,1} & q_{2,2} & \cdots & q_{2,N} \\ \vdots & \vdots & & \vdots \\ q_{N,1} & q_{N,2} & \cdots & q_{N,N} \end{pmatrix}$$

$$\pi(t) = (\pi_1(t), \pi_2(t), \ldots, \pi_N(t))$$

SOLVING DISCRETE STATE MODELS

Alternatively they may be studied using STOCHASTIC SIMULATION. Each run generates a single trajectory through the state space. Many runs are needed in order to obtain average behaviours.

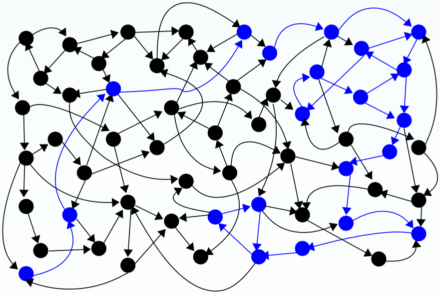

STATE SPACE EXPLOSION

As the size of the state space becomes large it becomes infeasible to carry out numerical solution and extremely time-consuming to conduct stochastic simulation.

In these cases we would like to take advantage of the MEAN FIELD or FLUID APPROXIMATION techniques.

Use CONTINUOUS STATE VARIABLES to approximate the discrete state space.

Use ORDINARY DIFFERENTIAL EQUATIONS to represent the evolution of those variables over time.

Appropriate for models in which there are large numbers of components of the same type, i.e. models of populations and situations of collective dynamics.

POPULATION MODELS - TIME SERIES ANALYSIS

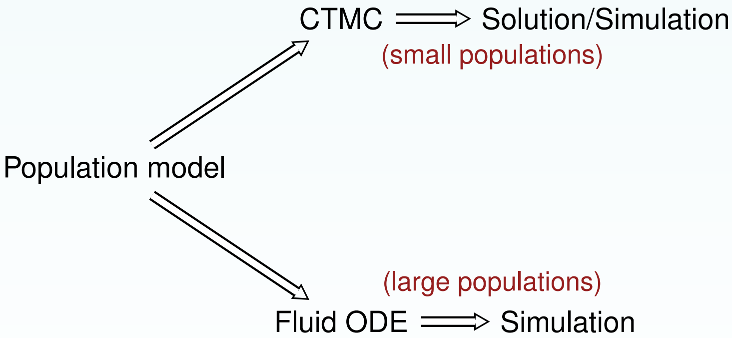

Fluid methods: approximate description of the collective (average) behaviour, estimate of certain passage times

- M. Tribastone, S. Gilmore, J. Hillston: Scalable Differential Analysis of Process Algebra Models. IEEE Trans. Softw Eng. 2012.
- R.A. Hayden, A. Stefanek, J.T. Bradley. *Fluid computation of passage-time distributions in large Markov models.* Theor. Comput. Sci. 2012.

POPULATION MODELS - MODEL CHECKING

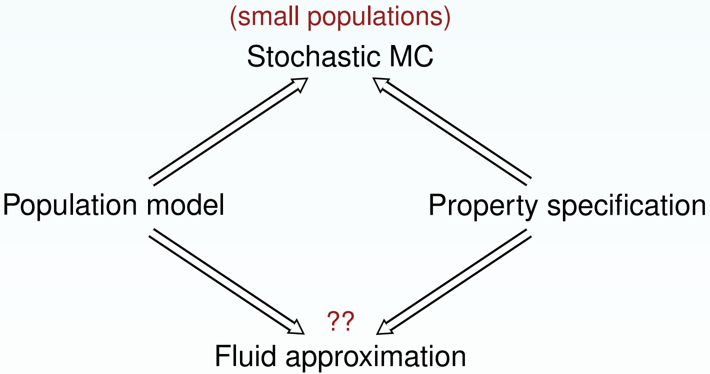

Understand how and to what extent fluid methods can be used to efficiently approximate stochastic model checking.

Goals

We will consider population models, composed of many interacting agents of one or more classes.

We will focus on questions related to the behaviour of individual agents for medium and large population size.

We will investigate:

- individual properties, concerned with the behaviour of a single or a few agents
- collective properties, concerned with the behaviour at the population level.

LECTURE PLAN

- Introduction to population CTMC and fluid approximation for collective and individual behaviour;
- Individual properties: model checking time-inhomogeneous CTMC, decidability, and correctness
- Collective properties: linear noise approximation (if there will be time — not in the book chapter).

- L. Bortolussi, J. Hillston, D. Latella, M. Massink.Continuous Approximation of Collective Systems Behaviour: a Tutorial. Performance Evaluation, 2013.
- L. Bortolussi, J. Hillston: Fluid Model Checking. CONCUR 2012.
- L. Bortolussi, J. Hillston: Model Checking Single Agent Behaviours by Fluid Approximation, submitted to Information and Computation.
- L. Bortolussi, R. Lanciani. Model Checking Markov Population Models by Central Limit Approximation. QEST 2013.

Outline

EXAMPLE: P2P NETWORK EPIDEMICS

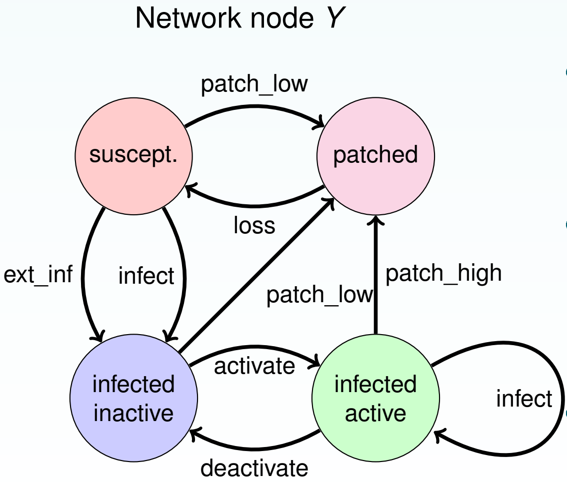

- A network is composed of N interconnected nodes
- Indistinguishable individual nodes ⇒ we only count of how many nodes are in each state
- Dynamics specified at the collective level

Population CTMC: individuals and collectives

Individuals

We have N individuals with state $Y_i^{(N)} \in S$, $S = \{1, 2, \ldots, n\}$ in the system (we can have multiple classes; the population is assumed constant for simplicity).

Collective Variables

$X_j^{(N)} = \sum_{i=1}^{N} \mathbf{1}\{Y_i^{(N)} = j\}$, and $\mathbf{X}^{(N)} = (X_1^{(N)}, \ldots, X_n^{(N)})$

Example: network epidemics

- Individual state space: $S = \{$susceptible (s),infected and inactive (d), infected and active (i), patched (p) $\}$
- Collective variables: $X_s^{(N)} = \sum_{j=1}^{n} \mathbf{1}\{Y_j^{(N)} = s\}$, $X_d^{(N)}$, $X_i^{(N)}$, $X_p^{(N)}$.

POPULATION CTMC: COLLECTIVE DYNAMICS

COLLECTIVE TRANSITIONS $\mathcal{T}^{(N)}$

$\tau \in \mathcal{T}^{(N)}$ describes a possible action/ event.

$\tau = (R_\tau, r_\tau^{(N)})$, where

- $r_\tau^{(N)} = r_\tau^{(N)}(\mathbf{X}^{(N)})$ is the rate function, giving the speed at which the event happens.
- R_τ is the multi-set of update rules, $R_\tau = \{i_1 \to j_1, \ldots, i_k \to j_k\}$. $m_{\tau,i\to j}$ is the multiplicity of $i \to j$ in R_τ

UPDATE VECTOR

With each transition τ, we associate an update vector $\mathbf{v}_\tau$, giving the net change in collective variables due to τ:

$$\mathbf{v}_{\tau,i} = \sum_{(i\to j)\in R_\tau} m_{\tau,i\to j}\mathbf{e}_j - \sum_{(i\to j)\in R_\tau} m_{\tau,i\to j}\mathbf{e}_i,$$

EXAMPLE: P2P NETWORK EPIDEMICS

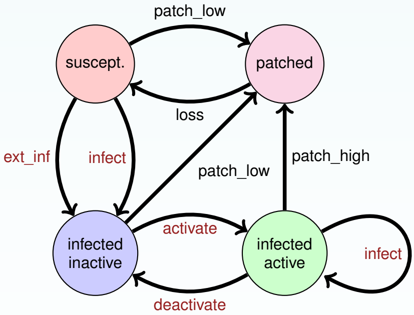

ext_inf:	$R_{\mathtt{ext\_inf}} = \{s \to d\}$,	$r^{(N)}_{\mathtt{ext\_inf}} = k_{ext} X_s$;
infect:	$R_{\mathtt{infect}} = \{s \to d, i \to i\}$,	$r^{(N)}_{\mathtt{infect}} = \frac{k_{inf}}{N} X_s X_i$;
activate:	$R_{\mathtt{activate}} = \{d \to i\}$,	$r^{(N)}_{\mathtt{activate}} = k_{act} X_d$;
deactivate:	$R_{\mathtt{deactivate}} = \{i \to d\}$,	$r^{(N)}_{\mathtt{deactivate}} = k_{deact} X_i$;

EXAMPLE: P2P NETWORK EPIDEMICS

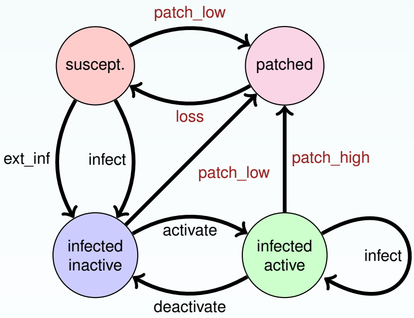

$$\texttt{patch\_s}: \quad R_{\texttt{patch\_s}} = \{s \to p\}, \quad r^{(N)}_{\texttt{patch\_s}} = k_{low} X_s;$$
$$\texttt{patch\_d}: \quad R_{\texttt{patch\_d}} = \{d \to p\}, \quad r^{(N)}_{\texttt{patch\_d}} = k_{low} X_d;$$
$$\texttt{patch\_i}: \quad R_{\texttt{patch\_i}} = \{i \to p\}, \quad r^{(N)}_{\texttt{patch\_i}} = k_{high} X_i;$$
$$\texttt{loss}: \quad R_{\texttt{loss}} = \{p \to s\}, \quad r^{(N)}_{\texttt{loss}} = k_l X_p;$$

Population CTMC

A population model is thus given by a tuple $\mathcal{X}^{(N)} = (\mathbf{X}^{(N)}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$, where

- $\mathbf{X}^{(N)}$ are the collective variables;
- $\mathcal{T}^{(N)}$ are the collective transitions;
- $\mathbf{x}_0^{(N)}$ is the initial state.

State Space

$\mathcal{S}^{(N)} = \{\mathbf{x} \in \mathbb{N}^n \mid \sum x_i = N\}$

CTMC infinitesimal generator $Q = (q_{\mathbf{x},\mathbf{x}'})$

$$q_{\mathbf{x},\mathbf{x}'} = \sum \{r_\tau(\mathbf{x}) \mid \tau \in \mathcal{T},\ \mathbf{x}' = \mathbf{x} + \mathbf{v}_\tau\}.$$

EXAMPLE: CLIENT SERVER INTERACTION

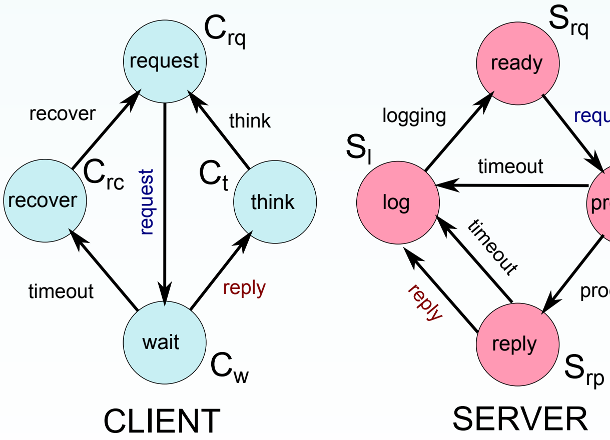

EXAMPLE: CLIENT SERVER INTERACTION

VARIABLES

- 4 variables for the client states: C_{rq}, C_w, C_{rc}, C_t.
- 4 variables for the server states: S_{rq}, S_p, S_{rp}, S_l.

TRANSITIONS

There are 7 transition in totals.

- request: $C_{rq} \to C_w, S_{rq} \to S_p;\ kr \cdot min(C_{rq}, S_{rq})$
- reply: $C_w \to C_t, S_{rp} \to S_l;\ min(k_w C_w, k_{rp} S_{rp})$
- timeout: $C_w \to C_{rc};\ k_{to} C_w$
- ...

Outline

FLUID APPROXIMATION

- It applies to population CTMC models with large population size N (studies the limit as $N \to \infty$)
- It applies to population densities (normalisation step), under suitable scaling of rate functions.
- It is a functional version of the law of large numbers: in any finite time horizon, the trajectories of the PCTMC converge to a deterministic trajectory, solution of the fluid ODE.

An intuition

As population increases, we observe more events each having a smaller impact on the population density vector.

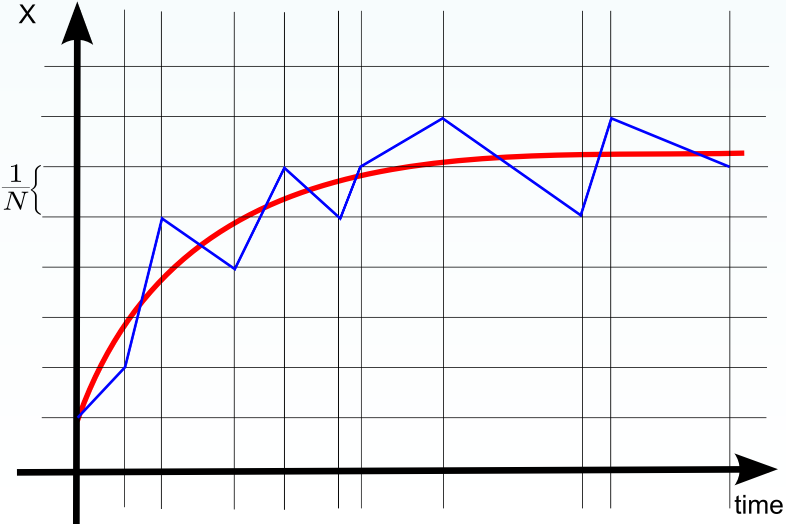

AN INTUITION

As population increases, we observe more events each having a smaller impact on the population density vector.

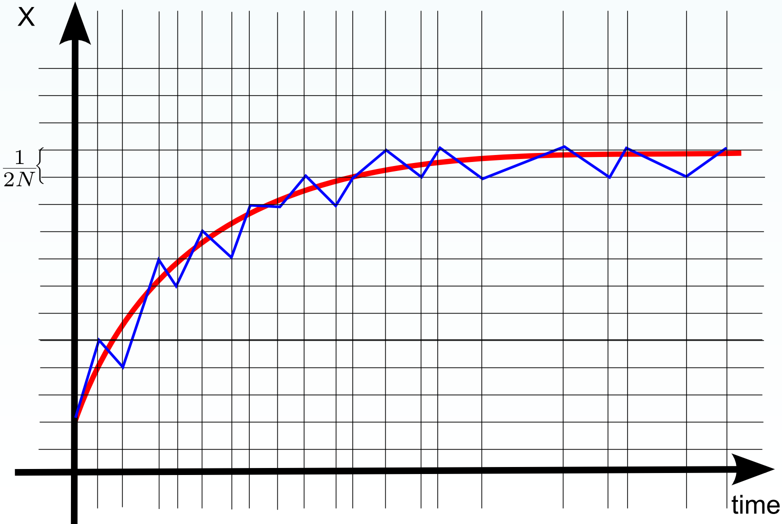

NORMALIZATION

The normalized model $\hat{\mathcal{X}}^{(N)} = (\hat{\mathbf{X}}, \hat{\mathcal{T}}^{(N)}, \hat{\mathbf{x}}_0^{(N)})$ associated with $\mathcal{X}^{(N)} = (\mathbf{X}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$ is defined by:

- Variables: $\hat{\mathbf{X}} = \frac{\mathbf{X}}{N}$
- Initial conditions: $\hat{\mathbf{x}}_0^{(N)} = \frac{\mathbf{x}_0^{(N)}}{N}$
- Normalized transition $\hat{\tau} = (R_\tau, \hat{r}_\tau^{(N)}(\hat{\mathbf{X}}))$ from $\tau \in \mathcal{T}^{(N)}$:
  - rate $\hat{r}_\tau^{(N)}\left(\frac{X}{N}\right) = r_\tau^{(N)}(\mathbf{X})$.
  - update vector $\frac{1}{N}\mathbf{v}_\tau$.

We assume to have a sequence of (normalised) models $\hat{\mathcal{X}}^{(N)}$, $N > 0$, that differ only in the total population size.

EXAMPLE

We will consider the normalised P2P network epidemics model, for an increasing number of netwkr nodes.

SCALING ASSUMPTIONS

- $E \subset \mathbb{R}^n$ is a open (or compact) set containing the state space of each $\hat{\mathbf{X}}^{(N)}(t)$ for each N. As here the population remains constant, it can be taken as the unit simplex in $\mathbb{R}^n$: $\{\mathbf{x} \in [0,1]^n \mid \sum_i x_i = 1\}$.
- $\frac{1}{N}\hat{r}_\tau^{(N)}$ is required to converge uniformly to a locally Lipschitz continuous and locally bounded function f_τ:

$$\sup_{\mathbf{x} \in E} \left\| \frac{1}{N}\hat{r}_\tau^{(N)}(\mathbf{x}) - f_\tau(\mathbf{x}) \right\| \to 0.$$

- If $\frac{1}{N}\hat{r}_\tau^{(N)} = f_\tau$ does not depend on N, the rate satisfies the density dependence condition.
- The following theorem works also under less restrictive assumptions (e.g. random increments with bounded variance and average).

DRIFT AND LIMIT VECTOR FIELD

DRIFT

The drift or mean increment at level N is

$$F^{(N)}(\mathbf{x}) = \sum_{\tau \in \mathcal{T}} \frac{\mathbf{v}_\tau}{N} \hat{r}_\tau^{(N)}(\mathbf{x})$$

By the scaling assumptions, $F^{(N)}$ converges uniformly to F, the limit vector field (locally bounded and Lipschitz continuous):

$$F(\mathbf{x}) = \sum_{\tau \in \mathcal{T}} \mathbf{v}_\tau f_\tau(\mathbf{x}).$$

THE FLUID ODE IS

$$\frac{d\mathbf{x}(t)}{dt} = F(\mathbf{x}(t))$$

CONVERGENCE TO THE FLUID ODE

THEOREM (KURTZ 1970)

If $\hat{\mathbf{x}}_0^{(N)} \to \hat{\mathbf{x}}_0 \in E$ in probability, then for any finite time horizon $T < \infty$, it holds that:

$$\mathbb{P}\left\{ \sup_{0 \leq t \leq T} ||\hat{\mathbf{X}}^{(N)}(t) - \mathbf{x}(t)|| > \varepsilon \right\} \to 0.$$

THE MOMENT CLOSURE POINT OF VIEW

Alternatively, the fluid ODE can be seen as a (first order) approximation of the ODE for the average of the PCTMC.

A LOOK AT K. THEOREM PROOF FOR DENSITY DEPENDENT RATES

ODE SOLUTION, INTEGRAL FORM

$$\mathbf{x}(t) = \mathbf{x}(0) + \int_0^t F(\mathbf{x}(s))ds$$

PERTURBED ODE REPRESENTATION OF A CTMC

$$\hat{\mathbf{X}}^{(N)}(t) = \hat{\mathbf{X}}^{(N)}(0) + \int_0^t F(\hat{\mathbf{X}}^{(N)}(s))ds + \mathbf{M}^{(N)}(t)$$

$\mathbf{M}^{(N)}(t)$ is a stochastic process, in particular a martingale, and by applying some martingale inequality (e.g. Doob's), one has that

$$\varepsilon_N = \sup_{s \leq t} \|\mathbf{M}^{(N)}(s)\| \to 0 \text{ as } N \to \infty$$

The theorem then follows as for proving uniqueness of solutions for Lipschitz vector fields (Grönwall inequality).

EXAMPLE: P2P NETWORK EPIDEMICS NORMALISED MODEL

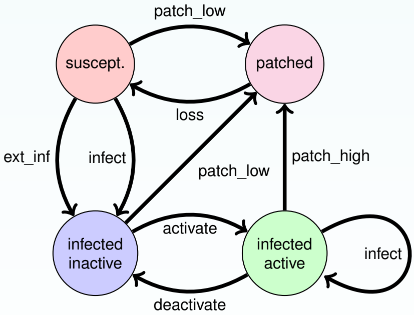

$$
\begin{aligned}
&\texttt{ext\_inf}: && \mathbf{v}_{\texttt{ext\_inf}} = \tfrac{1}{N}(-1,1,0,0), && \hat{r}^{(N)}_{\texttt{ext\_inf}} = Nk_{ext}\tfrac{X_s}{N} = Nk_{ext}\hat{X}_s;\\
&\texttt{infect}: && \mathbf{v}_{\texttt{infect}} = \tfrac{1}{N}(-1,1,0,0), && \hat{r}^{(N)}_{\texttt{infect}} = Nk_{inf}\tfrac{X_s}{N}\tfrac{X_i}{N} = Nk_{inf}\hat{X}_s\hat{X}_i;\\
&\texttt{activate}: && \mathbf{v}_{\texttt{act}} = \tfrac{1}{N}(0,-1,1,0), && \hat{r}^{(N)}_{\texttt{act}} = Nk_{act}\hat{X}_d;
\end{aligned}
$$

P2P NETWORK EPIDEMICS: FLUID EQUATIONS

$$\begin{cases} \dfrac{dx_s(t)}{dt} = -k_{ext}x_s - k_{inf}x_sx_i - k_{low}x_s + k_{loss}x_p \\ \dfrac{dx_d(t)}{dt} = k_{ext}x_s + k_{inf}x_sx_i - k_{act}x_d - k_{low}x_d + k_{deact}x_i \\ \dfrac{dx_i(t)}{dt} = k_{act}x_d - k_{deact}x_i - k_{high}x_i \\ \dfrac{dx_p(t)}{dt} = k_{low}x_s + k_{low}x_d + k_{high}x_i - k_{loss}x_p \end{cases}$$

P2P NETWORK EPIDEMICS: FLUID AT WORK

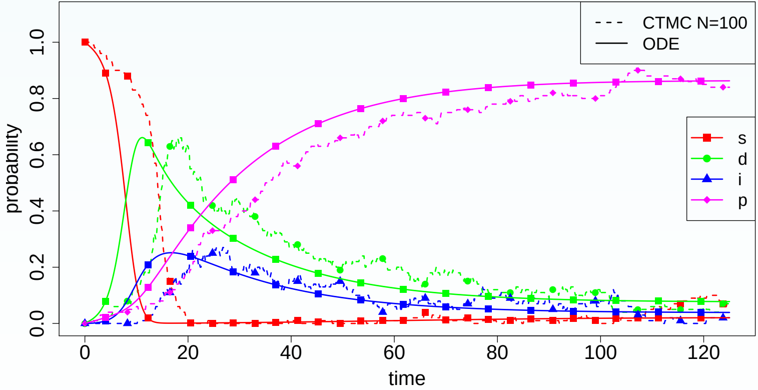

$N = 100$

P2P NETWORK EPIDEMICS: FLUID AT WORK

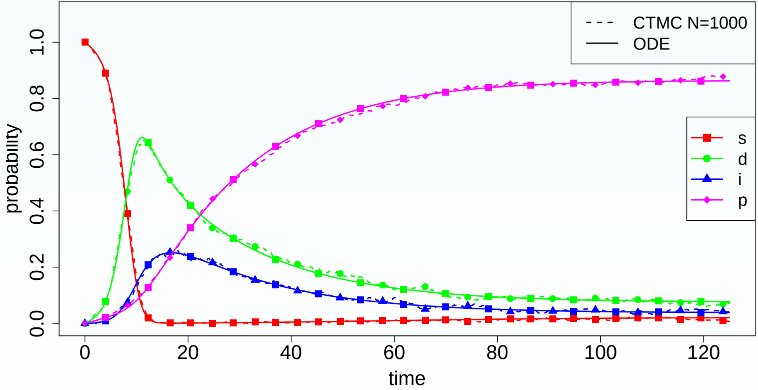

$N = 1000$

STEADY STATE BEHAVIOUR

- Kurtz theorem in general cannot be extended to convergence of the steady state.
- The problem is for instance with multi-stable fluid ODEs (more than one attracting equilibrium):
  in this case, in the long run the CTMC will always keep jumping between these different equilibria, although it will spend a long time in each attractor.

Kurtz theorem holds also for steady state distributions only if the fluid ODE has a unique globally attracting steady state.

- L. Bortolussi, J. Hillston, D. Latella, M. Massink. Continuous Approximation of Collective Systems Behaviour: a Tutorial. Performance Evaluation, 2013.

SINGLE AGENT ASYMPTOTIC BEHAVIOUR

- Focus on single individuals $Y_h^{(N)}$.
- Fix h and let $Z^{(N)} = Y_h^{(N)}$ be the single-agent stochastic process with state space S (not necessarily Markov).
- Let $Q^{(N)}(\mathbf{x})$ be defined by

$$\mathbb{P}\{Y_h^{(N)}(t+dt) = j \mid Y_h^{(N)}(t) = i, \hat{\mathbf{X}}^{(N)}(t) = \mathbf{x}\} = q_{i,j}^{(N)}(\mathbf{x})dt,$$

  with $Q^{(N)}(\mathbf{x}) \to Q(\mathbf{x})$.
- Let $z(t)$ be the time inhomogeneous-CTMC on S with infinitesimal generator $Q(t) = Q(\mathbf{x}(t))$, $\mathbf{x}(t)$ fluid limit.

THEOREM (FAST SIMULATION THEOREM)

For any $T < \infty$, $\mathbb{P}\{Z^{(N)}(t) \neq z(t), t \leq T\} \to 0$.

R. Darling, J. Norris. Differential equation approximations for Markov chains. *Probability Surveys*, 2008.

P2P NETWORK EPIDEMICS

SINGLE NODE

$Y^{(N)} \in \{s, d, i, p\}$

RATES OF $Z^{(N)}$

- `ext_inf`: $\frac{1}{X_s^{(N)}} r^{(N)}_{\texttt{ext\_inf}}(\mathbf{X}^{(N)}) = \frac{1}{X_s^{(N)}} k_{ext} X_s^{(N)} = k_{ext}$
- `infect`: $\frac{1}{X_s^{(N)}} r^{(N)}_{\texttt{infect}}(\mathbf{X}^{(N)}) = \frac{1}{N} k_{inf} X_i^{(N)} = k_{inf} \hat{X}_i^{(N)}$

RATES OF z

- `ext_inf`: k_{ext}
- `infect`: $k_{inf} x_i$

P2P NETWORK EPIDEMICS

The single agent infinitesimal generator is then $Q^{(N)}(\mathbf{x}) = Q(\mathbf{x})$, giving the following time dependent Q-matrix $Q(\mathbf{x}(t))$, where $\mathbf{x}(t)$ is the solution of the fluid equations.

$$\begin{pmatrix} -k_{ext} - k_{inf}x_i(t) - k_{low} & k_{ext} + k_{inf}x_i(t) & 0 & k_{low} \\ 0 & -k_{act} - k_{low} & k_{act} & k_{low} \\ 0 & k_{deact} & -k_{deact} - k_{high} & k_{high} \\ k_{loss} & 0 & 0 & -k_{loss} \end{pmatrix}$$

Transient probabilities for the fluid approximation of the single agent can be computed by solving the forward Kolmogorov equations

$$\frac{d\Pi(0,t)}{dt} = \Pi(0,t)Q(t).$$

P2P NETWORK EPIDEMICS: TRANSIENT PROBABILITIES

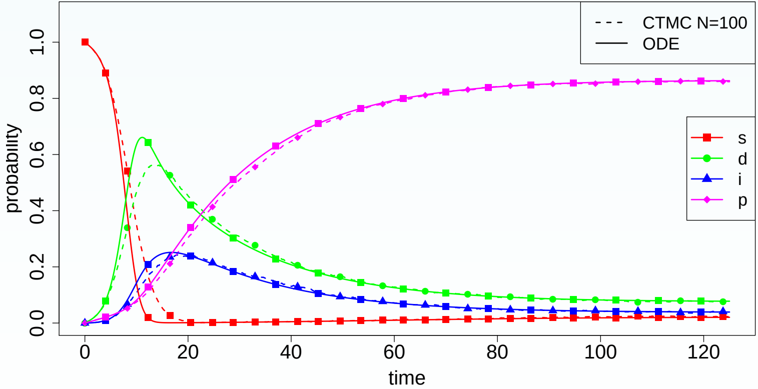

$N = 100$

P2P NETWORK EPIDEMICS: TRANSIENT PROBABILITIES

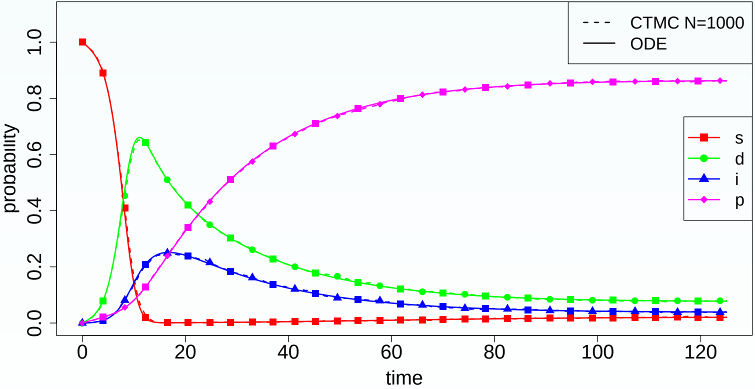

$N = 1000$

CLIENT SERVER EXAMPLE

SINGLE CLIENT

$Y^{(N)} \in \{rq, w, t, rc\}$

RATES OF $Z^{(N)}$

- request: $\frac{1}{C_{rq}^{(N)}} k_r \min(C_{rq}^{(N)}, S_{rq}^{(N)})$
- reply: $\frac{1}{C_{w}^{(N)}} \min(k_w C_w^{(N)}, k_{rp} S_{rp}^{(N)})$
- timeout: k_{to}; recover: k_{rc}

RATES OF z

- request: $k_r \min(1, \frac{s_{rq}(t)}{c_{rq}(t)})$
- reply: $\min(k_w, k_{rp} \frac{s_{rp}(t)}{c_w(t)})$
- timeout: k_{to}; recover: k_{rc}

CLIENT-SERVER: TRANSIENT PROBABILITIES

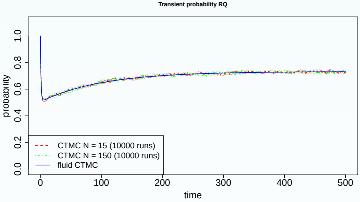

request

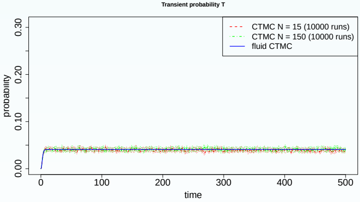

think

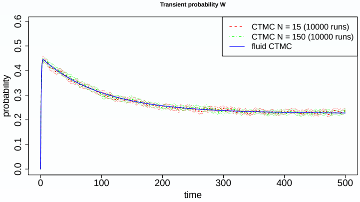

wait

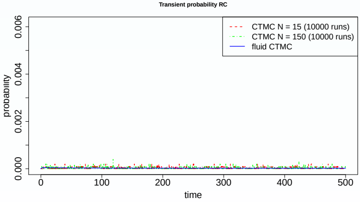

recover

Outline

INDIVIDUAL PROPERTIES

We are interested in the behaviour of a (random) individual.

We will specify such a behaviour in Continuous Stochastic Logic (CSL). Other possibilities include DFA, DTA, LTL, MiTL.

P2P NETWORK EPIDEMICS EXAMPLE

- What is the probability of a node being infected within T units of time?
- Is the probability of a single node remaining infected for T units of time smaller than p_1?
- Is the probability of a node being patched before getting infected larger than p_2?
- What is the probability of being patched within time T_1, and then remaining uninfected with probability at least p_3 for T_2 units of time?

COLLECTIVE PROPERTIES

We will concentrate on collective properties of the form:

"What is the probability that a given fraction of individuals satisfies the local property ϕ (by time T)"?

P2P NETWORK EPIDEMICS EXAMPLE

- What is the probability of at most one tenth of nodes being infected within T units of time?
- Is the probability of at least one third of nodes remaining infected for T units of time smaller than p_1?
- Is the probability of at least half of nodes being patched before getting infected larger than p_2?

(TIME-BOUNDED) CONTINUOUS STOCHASTIC LOGIC

SYNTAX

$$\phi = a \mid \phi_1 \wedge \phi_2 \mid \neg\phi \mid \mathcal{P}_{\bowtie p}(\mathbf{X}^{[T_1,T_2]}\phi) \mid \mathcal{P}_{\bowtie p}(\phi_1 \mathbf{U}^{[T_1,T_2]}\phi_2)$$

- a is an atomic proposition;
- $\phi_1 \wedge \phi_2$ and $\neg\phi$ are the usual boolean connectives;
- $\mathcal{P}_{\bowtie p}(\mathbf{X}^{[T_1,T_2]}\phi)$ is the next state temporal modality.
- $\mathcal{P}_{\bowtie p}(\phi_1 \mathbf{U}^{[T_1,T_2]}\phi_2)$ is the until temporal modality.

DERIVED MODALITIES

EVENTUALLY: $F^{[0,T]}\phi \equiv true\,\mathbf{U}^{[0,T]}\phi$

ALWAYS: $G^{[0,T]}\phi \equiv \neg F^{[0,T]}\neg\phi$

CSL - RESTRICTIONS

SYNTAX

$$\phi = a \mid \phi_1 \wedge \phi_2 \mid \neg\phi \mid \mathcal{P}_{\bowtie p}(\mathbf{X}^{[T_1,T_2]}\phi) \mid \mathcal{P}_{\bowtie p}(\phi_1 \mathbf{U}^{[T_1,T_2]}\phi_2)$$

- We do not consider timed-unbounded operators: $0 \leq T_1, T_2 < \infty$;
- We do not consider steady state probabilities;
- We do not consider rewards.

Rewards can be easily added.

Time unbounded and steady state properties are more problematic: Kurtz theorem works only for time-bounded horizons.

CSL - NOTATION

We will interpret CSL formulae on a generic stochastic process $Z(t)$ on S, such that all relevant sets of paths (i.e. those satisfying until or next formulae) are measurable.

PATHS

A path σ of $Z(t)$ is a sequence

$$\sigma = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} \ldots,$$

with non null probability of jumping from s_i to s_{i+1}, for each i;

NOTATION

- $\sigma@t$ is the state of σ at time t;
- $\sigma[i]$ is the i-th state of σ;
- $t_\sigma[i]$ is the time of the i-th jump in σ;

CSL- SEMANTICS

STATE FORMULAE

- $s, t_0 \models a$ if and only if $a \in L(s)$;
- $s, t_0 \models \neg\phi$ if and only if $s, t_0 \not\models \phi$;
- $s, t_0 \models \phi_1 \wedge \phi_2$ if and only if $s, t_0 \models \phi_1$ and $s, t_0 \models \phi_2$;
- $s, t_0 \models P_{\bowtie p}(\psi)$ if and only if $\mathbb{P}\{\sigma \mid \sigma, t_0 \models \psi\} \bowtie p$.

PATH FORMULAE

- $\sigma, t_0 \models \mathbf{X}^{[T_1, T_2]}\phi$ if and only if $t_\sigma[1] \in [T_1, T_2]$ and $\sigma[1], t_0 + t_\sigma[1] \models \phi$.
- $\sigma, t_0 \models \phi_1 \mathbf{U}^{[T_1, T_2]}\phi_2$ if and only if $\exists \bar{t} \in [t_0 + T_1, t_0 + T_2]$ s.t. $\sigma@\bar{t}, \bar{t} \models \phi_2$ and $\forall t_0 \leq t < \bar{t}, \sigma@t, t \models \phi_1$.

EXAMPLE: P2P NETWORK INFECTION

- $\psi_1 = F^{[0,T]} a_{infected}$
  (a node is infected within T units of time);
- $\phi_1 = P_{<p_1}(G^{[0,T]} a_{infected})$
  (the probability of a single node remaining infected for T units of time is smaller than p_1);
- $\phi_2 = P_{>p_2}(\neg a_{infected} \mathbf{U}^{[0,T]} a_{patched})$
  (the probability of a node being patched before getting infected is larger than p_2);
- $\psi_2 = F^{[0,T_1]}(a_{patched} \wedge P_{\geq p_3}(G^{[0,T_2]} \neg a_{infected}))$
  (a node is patched within time T_1, and then remains not infected with probability at least p_3 for T_2 units of time).

THE IDEA

Approximate the behaviour of an agent Z in the system using the time-inhomogeneous Markov chain z.

Model check temporal logic formulae on z.

OUTLINE OF FOLLOWING TOPICS

- A model checking algorithm for CSL on time-inhomogeneous CTMC (ICTMC).
- Investigation of its decidability.
- Convergence results (asymptotic correctness for large N).

Outline

CSL MODEL CHECKING: BASIC IDEAS

- The model checking algorithm works by processing bottom up the parse tree of a formula.
- The intuition is that each state formula determines the set of states satisfying it. Once this set has been computed, one can treat the state formula as an atomic proposition.
- Dealing with atomic propositions and boolean connectives is easy: we just need to explain how to compute the satisfaction probability of path formulae.

CSL MODEL CHECKING: NEXT STATE OPERATOR

PATH PROBABILITY $\mathbf{X}^{[T_1,T_2]}\phi$

- We just need to evaluate the probability that, being in a state s, we jump within time $[T_1, T_2]$ to a state that satisfies ϕ.
- We know the set $\{s' \mid s' \models \phi\}$ by (inductive) hypothesis.
- We consider time-homogeneous CTMCs.
- The exit rate in state s is $q(s) = \sum_{s' \in \mathcal{S},\ s' \neq s} q(s, s')$.
- The rate at which we jump to a ϕ-state is $q_\phi(s) = \sum_{s' \models \phi,\ s' \neq s} q(s, s')$.

PROBABILITY DENSITY OF $\mathbf{X}\phi$

$$\frac{q_\phi(s)}{q(s)} q(s) \exp(-q(s)t) = q_\phi(s) \exp(-q(s)t)$$

CSL MODEL CHECKING: NEXT STATE OPERATOR

PROBABILITY DENSITY OF $\mathbf{X}\phi$

$$\frac{q_\phi(s)}{q(s)} q(s) \exp(-q(s)t) = q_\phi(s) \exp(-q(s)t)$$

PROBABILITY OF $\mathbf{X}^{[T_1,T_2]}\phi$

$$\begin{aligned} \mathbb{P}(s, \mathbf{X}^{[T_1,T_2]}\phi) &= \int_{T_1}^{T_2} q_\phi(s) \exp(-q(s)t) dt \\ &= \frac{q_\phi(s)}{q(s)} (\exp(-q(s)T_1) - \exp(-q(s)T_2)) \end{aligned}$$

- We then need to solve the inequality $\mathbb{P}(s, \mathbf{X}^{[T_1,T_2]}\phi) \bowtie p$ to decide if s satisfies $P_{\bowtie p}(\mathbf{X}^{[T_1,T_2]}\phi)$.
- This method requires the CTMC to be time-homogeneous

CSL MODEL CHECKING: UNTIL OPERATOR

- We start by considering the until path formula $\phi_1 \mathbf{U}^{[0,T]} \phi_2$.
- We need to compute the probability of all paths that remain in a ϕ_1-state before entering a ϕ_2 state before time T.
- The idea is that if we enter a $\neg\phi_1$-state, we should discard the path, while if we enter a ϕ_2-state, we are done.
- We can monitor these two events by "stopping" when they happen, making $\neg\phi_1$ and ϕ_2-states absorbing (i.e. removing outgoing transitions).

EXAMPLE

Consider the property $\texttt{notinfected}\,\mathbf{U}^{[0,T]}\texttt{patched}$. We need to make infected and patched states absorbing.

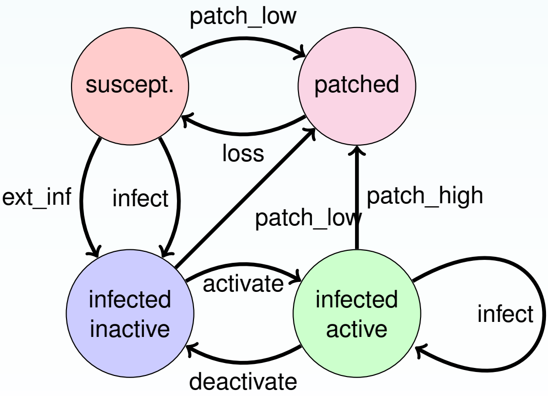

EXAMPLE

Consider the property $\texttt{notinfected}\,\mathbf{U}^{[0,T]}\texttt{patched}$.

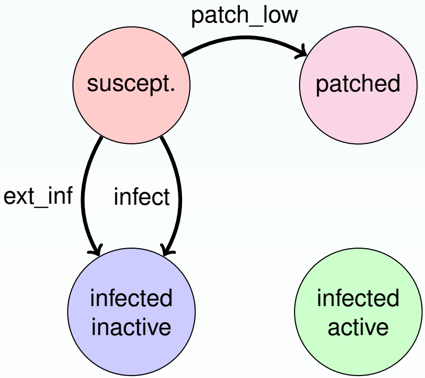

CSL MODEL CHECKING: UNTIL OPERATOR

Let Π be the probability matrix: $\Pi(0, T)[s, s']$ gives the probability of being in s' at time T, starting in s at time 0.

MODEL CHECKING ALGORITHM FOR $\phi_1 \mathbf{U}^{[0,T]} \phi_2$

1. Make $\neg\phi_1$ and ϕ_2 states absorbing
2. Compute the transient probability of the so modified CTMC at time T (using uniformisation or solving Kolmogorov equations): $\Pi_{\neg\phi_1 \vee \phi_2}(0, T)$,
3. The desired probability is

$$\mathbb{P}(\sigma \models \phi_1 \mathbf{U}^{[0,T]} \phi_2 \mid \sigma[0] = s) = \sum_{s' \models \phi_2} \Pi_{\neg\phi_1 \vee \phi_2}[s, s'](0, T)$$

CSL MODEL CHECKING: $\phi_1 \mathbf{U}^{[T_1,T_2]} \phi_2$

We split the problem in two parts:

1. Compute the probability of not entering a $\neg\phi_1$ in the first T_1 units of time, by making $\neg\phi_1$ states absorbing.
2. Compute the probability of the until formula $\phi_1 \mathbf{U}^{[0,T_2-T_1]} \phi_2$

MODEL CHECKING ALGORITHM FOR $\phi_1 \mathbf{U}^{[T_1,T_2]} \phi_2$

1. Compute $\Pi_{\neg\phi_1}(0, T_1)$ by transient analysis;
2. Compute $\Pi_{\neg\phi_1 \vee \phi_2}(0, T_2 - T_1)$ by transient analysis;
3. The desired probability $\mathbb{P}(\sigma \models \phi_1 \mathbf{U}^{[T_1,T_2]} \phi_2 \mid \sigma[0] = s)$ is

$$\sum_{s_1 \models \phi_1} \sum_{s_2 \models \phi_2} \Pi_{\neg\phi_1}(0, T_1)[s, s_1] \Pi_{\neg\phi_1 \vee \phi_2}[s_1, s_2](0, T_2 - T_1)$$

The method works only for time-homogeneous CTMCs.

Outline

CSL MODEL CHECKING FOR ICTMC

The fluid limit z of a single agent in a population model is a time-inhomogeneous CTMC.

IMPLICATIONS

- We cannot use the same algorithms sketched before, because we cannot always start transient computations from time 0.
- Non-nested properties can still be dealt with similarly, the difficulties arises with nested properties.

CSL MODEL CHECKING FOR ICTMC

Consider a ICTMC with state space S and rates $Q = Q(t)$.
Focus on a non-nested until formula of the type

$$\mathcal{P}_{\bowtie p}(\phi_1 U^{[0,T]} \phi_2)$$

which can be model checked as customary by solving the following reachability problem:

What is the probability of reaching a ϕ_2-state within time T without entering a $\neg\phi_1$-state?

SOLUTION

Make $\neg\phi_1 \vee \phi_2$-states absorbing, and compute the probability of reaching a goal state at time T (e.g., by solving the Kolmogorov equations or by uniformisation for ICTMC).

P2P NETWORK EPIDEMICS: THE MODEL

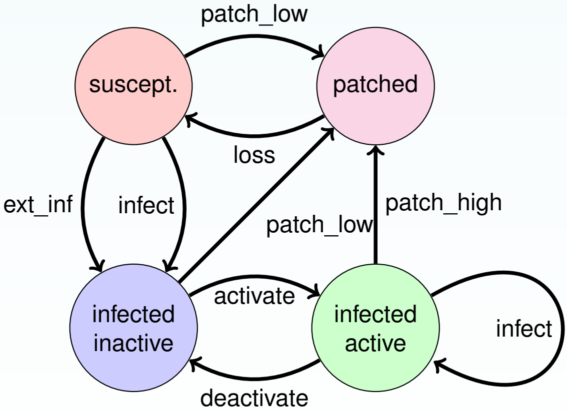

P2P NETWORK EPIDEMICS: $F^{[0,T]} a_{infected}$ FROM STATE s

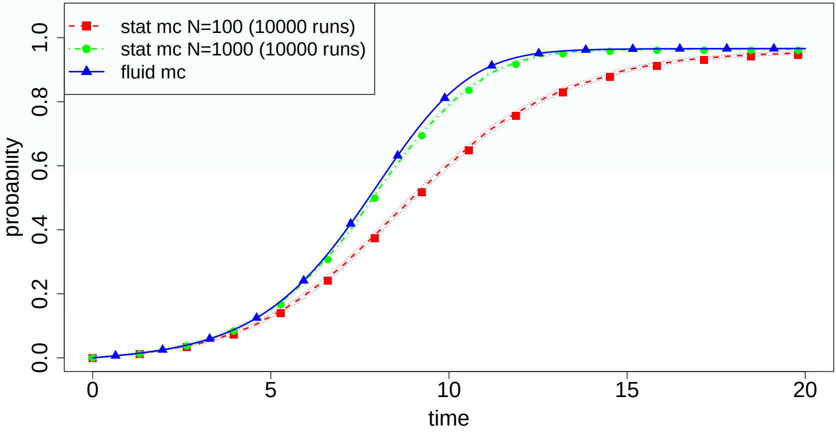

P2P NETWORK EPIDEMICS: $\neg a_{infected} \mathbf{U}^{[0,T]} a_{patched}$ FROM STATE s

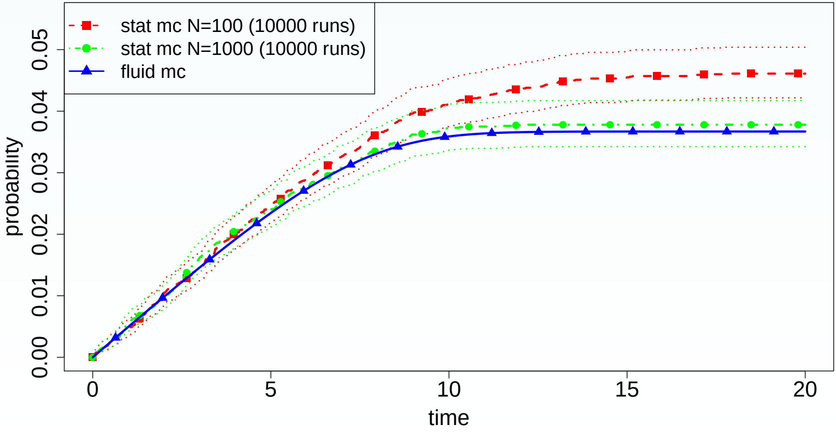

NEXT-STATE PROBABILITY

PROBABILITY OF $\mathbf{X}^{[T_1,T_2]}\phi$ STARTING AT TIME t_0

$$P_{next}(t_0)[s] = \int_{t_0+T_1}^{t_0+T_2} q_\phi(s,t) \cdot e^{-\Lambda(t_0,t)[s]} dt$$

where $\Lambda(t_0,t)[s] = \int_{t_0}^{t} -q_{s,s}(\tau) d\tau$ is the cumulative rate.

We can reduce the computation of the previous integral to the following initial value problem from $t_0 + T_1$ to $t_0 + T_2$.

$$\begin{cases} \dfrac{d}{dt}P(t) = q_{s,S_0}(t) \cdot e^{-L(t)} \\ \dfrac{d}{dt}L(t) = -q_{s,s}(t) \end{cases}$$

with $P(t_0 + T_1) = 0$ and $L(t_0 + T_1) = \Lambda(t_0, t_0 + T_1)$.

P2P NETWORK EPIDEMICS: $\mathbf{X}^{[0,T]} a_{infected}$ FROM STATE s

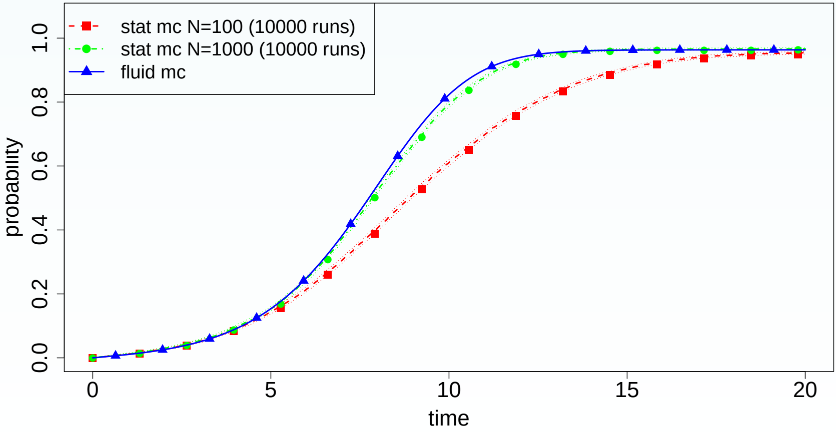

CLIENT-SERVER: THE MODEL

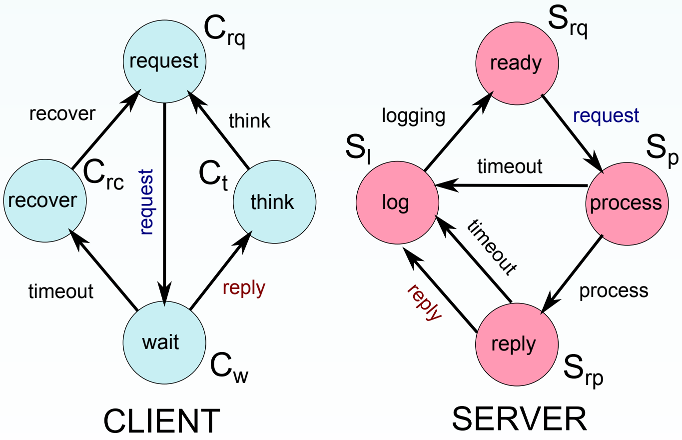

CLIENT-SERVER: $\mathcal{P}_{=?}(F^{\leq T} a_{timeout})$

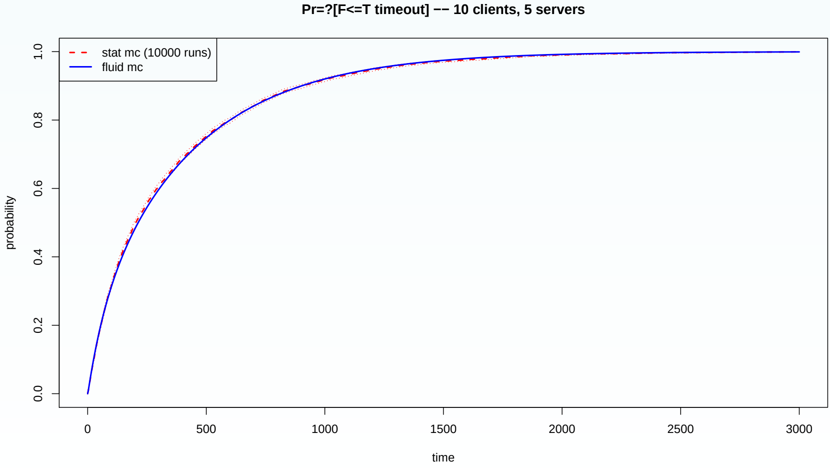

CLIENT-SERVER: $\mathcal{P}_{=?}(a_{request} \vee a_{wait} U^{\leq T} a_{timeout})$

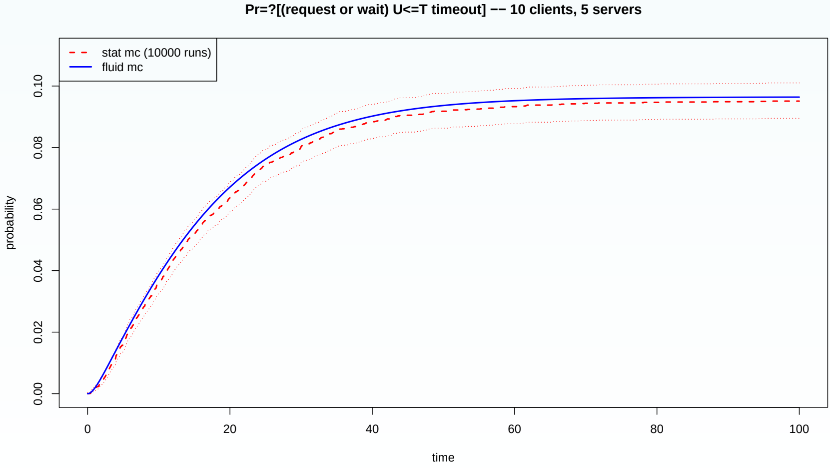

CLIENT-SERVER: COMPUTATIONAL COST

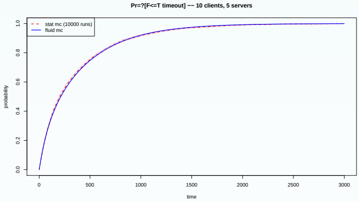

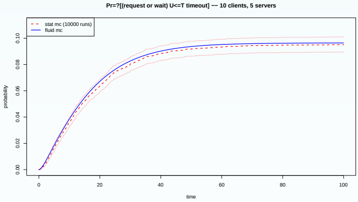

COMPUTATIONAL COST

- The cost of analysing the limit fluid system is independent of N.
- For the client server example (10 clients - 5 servers) it is $\sim$100 times faster than the simulation-based approach (which increases linearly with N).

P2P Network Epidemics: computational cost

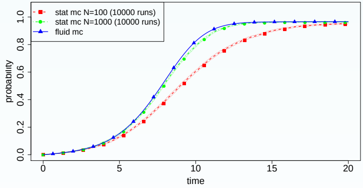

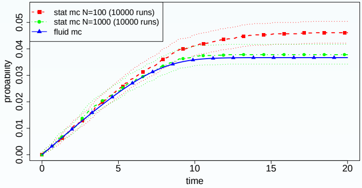

Computational cost

Checked property	Fluid MC	SMC ($N = 100$)	SMC ($N = 1000$)
Kolmogorov Equations	~ 0.1 s	~ 64 s	~ 101 s
$\mathbf{X}^{[0,T]} a_{infected}$	~ 0.06 s	~ 6 s	~ 24 s
$\neg a_{infected} \mathbf{U}^{[0,T]} a_{patched}$	~ 0.05 s	~ 5 s	~ 20 s

Outline

CSL MODEL CHECKING FOR ICTMC

Consider a ICTMC with state space S and rates $Q = Q(t)$.

$$\phi_1 \mathbf{U}^{[0,T]} \phi_2 \qquad \text{and} \qquad \mathbf{X}^{[T_1,T_2]} \phi$$

Time-homogeneity $\Rightarrow$ we can run each transient analysis/ integral computation from time $t_0 = 0$!

This is no more true in time-inhomogeneous CTMCs, as the probability of a path formula depends on the time at which we evaluate it.
Problems arise when we consider nested until formulae.

The truth value of ϕ in a state s depends on the time t at which we evaluate it.

TIME-DEPENDENT PROBABILITY OF $\mathbf{X}^{[T_1,T_2]}\phi$

PROBABILITY OF $\mathbf{X}^{[T_1,T_2]}\phi$ STARTING AT TIME t_0

$$P_{next}(t_0)[s] = \int_{t_0+T_1}^{t_0+T_2} q_\phi(s,t) \cdot e^{-\Lambda(t_0,t)[s]} dt$$

where $\Lambda(t_0,t)[s] = \int_{t_0}^{t} -q_{s,s}(\tau)d\tau$ is the cumulative rate.

INTUITION

- Compute $\frac{d}{dt_0} P_{next}(t_0)[s]$
- Construct an ODE for $P_{next}(t_0)$ and solve the i.v. problem.

CHECKING $P_{\bowtie p}(\mathbf{X}^{[T_1,T_2]}\phi)$

- Compute the path probability $P_{next}(t_0)[s]$ of $\mathbf{X}^{[T_1,T_2]}\phi$ as a function of t_0
- Solve the inequality $P_{next}(t_0)[s] \bowtie p$

P2P NETWORK EPIDEMICS: $\mathbf{X}^{[0,10]} a_{infected}$

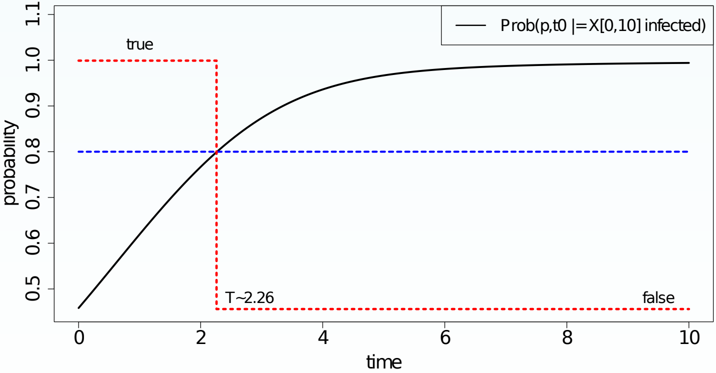

t_0 varying (Red line: $P_{\geq 0.8}(\mathbf{X}^{[0,10]} a_{infected})$)

TIME-DEPENDENT REACHABILITY PROBABILITY

Focus on $\mathcal{P}_{\bowtie p}(\phi_1 \mathbf{U}^{[0,T]} \phi_2)$. Assume that the truth of ϕ_1 and ϕ_2 does not depend on time.

Let $\Pi(t_1, t_2) = (\pi_{s_i,s_j}(t_1, t_2))_{i,j}$ be the probability matrix giving the probability of being in state s_j at time t_2, given that we are in state s_i at time t_1.

We consider $\Pi = \Pi_{\neg\phi_1 \vee \phi_2}$, the probability matrix of the CTMC in which $\neg\phi_1 \vee \phi_2$ states are made absorbing.

FORWARD AND BACKWARD KOLMOGOROV EQUATIONS

The device to compute the time dependent probability of an until formula $\phi_1 \mathbf{U}^{[0,T]} \phi_2$ are the Kolmogorov equations for ICTMCs.

FORWARD KOLMOGOROV EQUATION

$$\frac{d}{dt}\Pi(s,t) = \Pi(s,t)Q(t)$$

BACKWARD KOLMOGOROV EQUATION

$$\frac{d}{ds}\Pi(s,t) = -Q(s)\Pi(s,t)$$

COMPUTING $\Pi(t, t+T)$, FOR FIXED T

We just need to combine the two backward and forward equations by chain rule.

TIME-DEPENDENT REACHABILITY PROBABILITY

1. COMPUTE $\Pi(t, t+T)$, FOR $t \in [0, T_f]$

$\Pi(t, t+T)$, as a function of t, with initial conditions $\Pi(0, T)$, satisfies:

$$\frac{\mathrm{d}\Pi(t, t+T)}{\mathrm{d}t} = \Pi(t, t+T)Q(t+T) - Q(t)\Pi(t, t+T)$$

2. ADD PROBABILITY FOR GOAL STATES

$P_{\phi_1 \mathbf{U}^{[0,T]} \phi_2}(s, t)$ is equal to $\sum_{s' \models \phi_2} \Pi_{\neg\phi_1 \vee \phi_2}(t, t+T)[s, s']$.

3. COMPARE WITH THRESHOLD p

The truth value $\mathbf{T}(\phi, s, t)$ of formula ϕ in state s at time t is obtained by solving the inequality $P_{\phi_1 \mathbf{U}^{[0,T]} \phi_2}(s, t) \bowtie p$.

We need to find the zeros of the function $P_{\phi_1 \mathbf{U}^{[0,T]} \phi_2}(s, t) - p$.

P2P NETWORK EPIDEMICS: $G^{[0,10]}\neg a_{infected}$

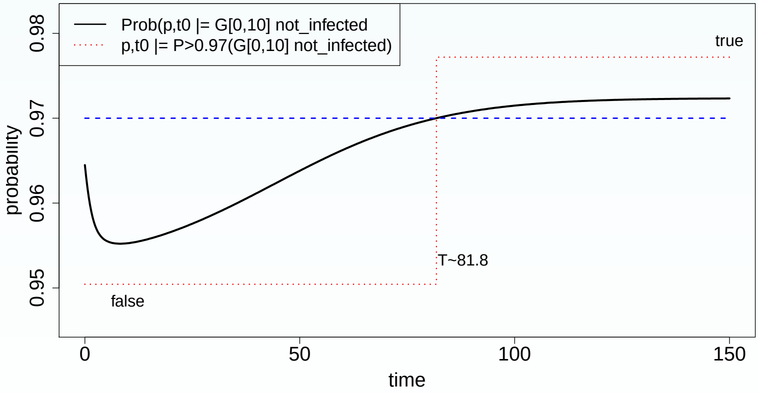

from state p (patched)

CLIENT SERVER: $\mathcal{P}_{=?}F^{\leq 50}a_{timeout}$ AS A FUNCTION OF t_0

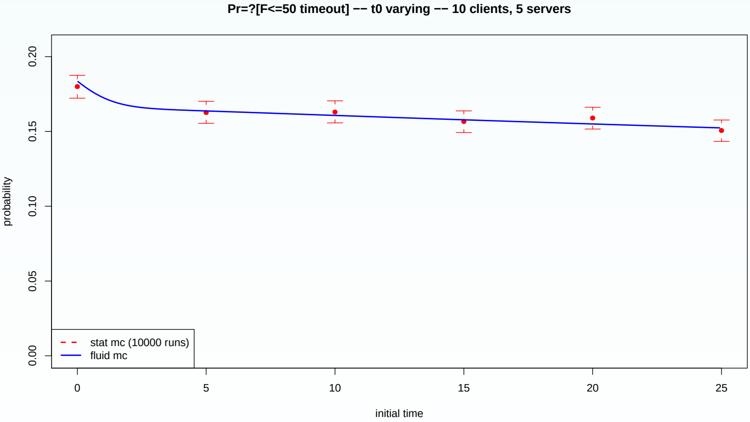

CLIENT-SERVER: $P_{<0.167}(F^{\leq 50}\, timeout)$

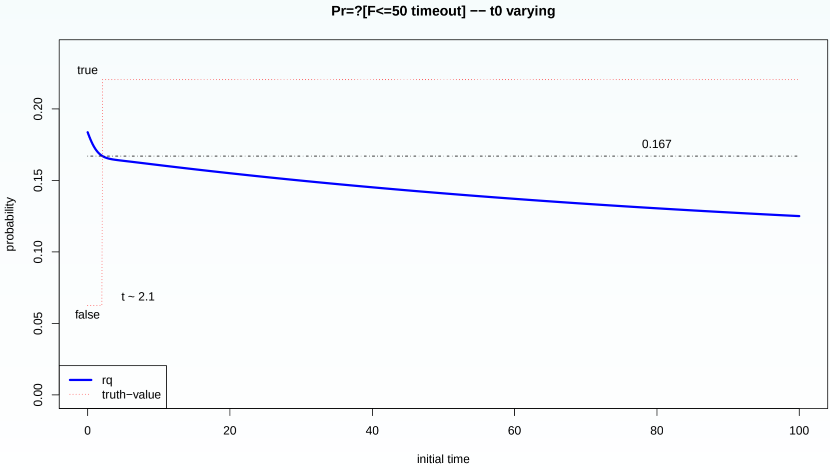

$P_{<0.167}(F^{\leq 50}\, timeout)$ from state *rq* of client.

COMPUTING THE TIME-DEPENDENT TRUTH IN PRACTICE

The equation $\frac{\mathrm{d}\Pi(t,t+T)}{\mathrm{d}t} = \Pi(t,t+T)Q(t+T) - Q(t)\Pi(t,t+T)$ is utterly stiff. Its integration error blows up even for the most accurate Matlab/Octave solvers.

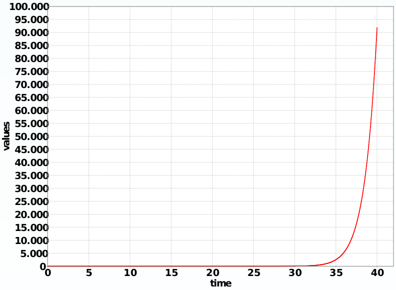

COMPUTING THE TIME-DEPENDENT TRUTH IN PRACTICE

The equation $\frac{d\Pi(t,t+T)}{dt} = \Pi(t, t + T)Q(t + T) - Q(t)\Pi(t, t + T)$ is utterly stiff. Its integration error blows up even for the most accurate Matlab/Octave solvers.

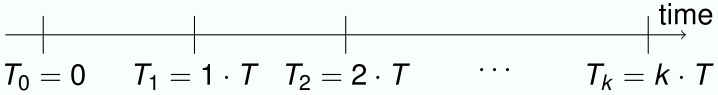

Practically, we can exploit the semigroup property

$$\Pi(t, t + T) = \Pi(t, T_j)\Pi(T_j, t + T)$$

and solve backward and forward equations separately, looping over j.

Outline

TIME-DEPENDENT TRUTH

- When computing the truth value of an until formula, we obtain a time dependent value $\mathbf{T}(\phi, s, t)$ in each state.
- When we consider nested temporal operators, we need to take this into account.
- The problem is that in this case the TOPOLOGY OF GOAL AND UNSAFE STATES in the CTMC can CHANGE IN TIME.

TIME DEPENDENT TRUTH: $F^{\leq T}\phi$

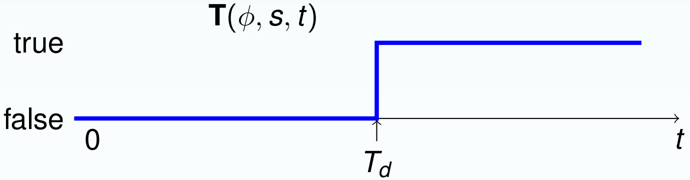

At discontinuity times, changes in topology introduce discontinuities in the probability values.

BUT...

Discontinuities happen at specific and FIXED time instants. We can solve Kolmogorov equations piecewise!

k DISCONTINUITIES $T_1, \ldots, T_k$ IN $[t, t+T]$

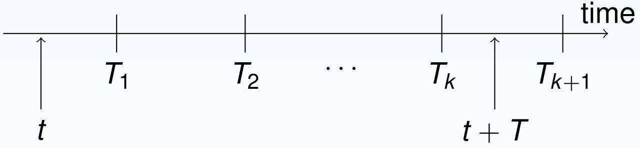

THE GENERIC CK EQUATION

$\Pi(t, t+T) = \Pi_1(t, T_1)\zeta(T_1)\Pi_2(T_1, T_2)\zeta(T_2)\cdots\zeta(T_k)\Pi_{k+1}(T_k, t+T).$
$\zeta(T_j)$ apply the proper bookkeeping operations to deal with changes in the topology of absorbing states.

- We can compute $\Pi(t, t+T)$ by an ODE obtained by derivation and application of chain rule.
- In advancing time, when we hit a discontinuity point (from below or above), the structure of the previous equation changes: integration has to be stopped and restarted.

THE ALGORITHM (SKETCHED)

Proceed bottom-up on the parse tree of a formula.
Case $\mathbf{T}(\mathcal{P}_{\bowtie p}(\phi_1 U^{[0,T]} \phi_2), t)$:

- Compute $\mathbf{T}(\phi_1, t)$ and $\mathbf{T}(\phi_2, t)$
- Let $T_1, \ldots, T_m$ be all the discontinuity points of $\mathbf{T}(\phi_1, t)$ and $\mathbf{T}(\phi_2, t)$ up to a final time T_f.
- Compute $\Pi(T_i, T_i + 1)$ for each i
- Compute $\Pi(0, T)$ using generalized CK equations
- Integrate $\frac{d}{dt}\Pi(t, t + T)$ up to T_f.
- Return $\mathbf{T}(\mathcal{P}_{\bowtie p}(\phi_1 U^{[0,T]} \phi_2), t) = \Pi(t, t + T) \bowtie p$.

The use of Kolmogorov equations is feasible if the state space is small (few dozens of states).

This is usually the case for single agent mean field models.

P2P NETWORK EPIDEMICS: $F^{[0,T]}(a_{patched} \wedge P_{\geq 0.97}(G^{[0,10]} \neg a_{infected}))$

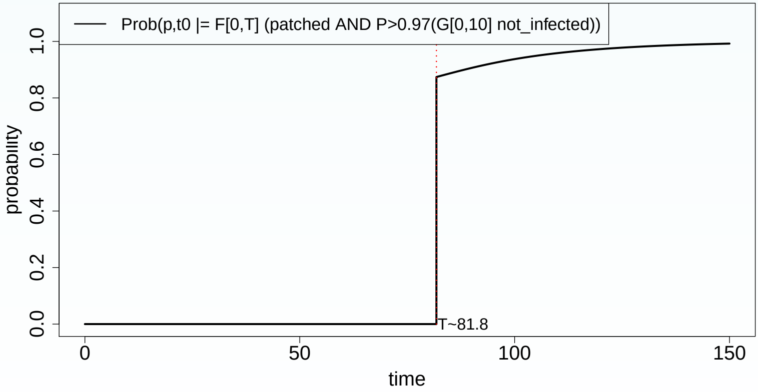

from state p (patched)

CLIENT-SERVER: $F^{\leq T}(P_{<0.167}(F^{\leq 50}\mathit{timeout}))$

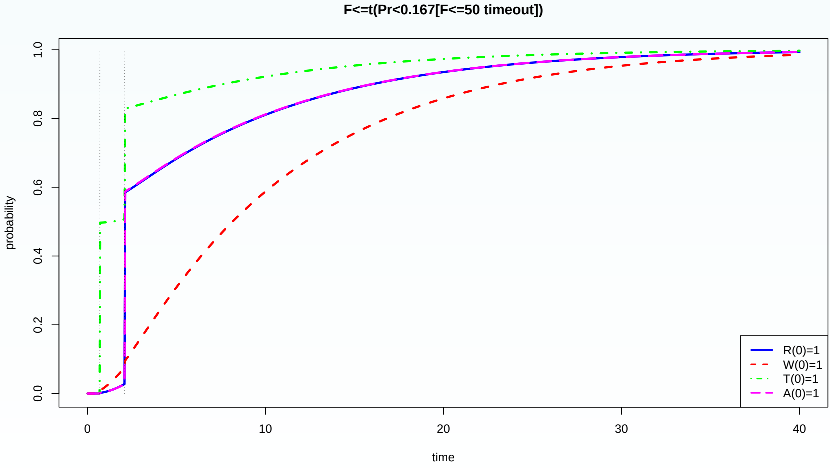

Outline

DECIDABILITY

DECIDABILITY

- We use algorithms to solve ODEs with error guarantee (interval analysis).
- We need to find zeros of function $P(s,t) - p$ (root finding), and guarantee their number to be finite (restrict to piecewise-real analytic functions).
- To answer the CSL query for main until formulae, we need to know if $P(s,0) \bowtie p$ (zero test).
- It is not known if root finding and zero test are decidable.

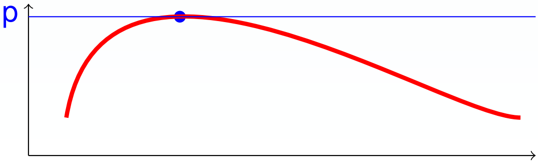

DECIDABILITY

DECIDABILITY

- We use algorithms to solve ODEs with error guarantee (interval analysis).
- We need to find zeros of function $P(s, t) - p$ (root finding), and guarantee their number to be finite (restrict to piecewise-real analytic functions).
- To answer the CSL query for main until formulae, we need to know if $P(s, 0) \bowtie p$ (zero test).
- It is not known if root finding and zero test are decidable.

THEOREM (QUASI-DECIDABILITY)

Let $\phi = \phi(\mathbf{p})$ be a CSL formula, with constants $\mathbf{p} = (p_1, \ldots, p_k) \in [0, 1]^k$ appearing in until formulae. The CSL model checking for ICTMC problem is decidable for $\mathbf{p} \in E$, where E is an open subset of $[0, 1]^k$, of measure 1.

CONVERGENCE OF CSL TRUTH

- We considered also convergence of CSL properties: are properties that are true in $z(t)$ ultimately true in $Z^{(N)}(t)$?
- Convergence suffers from similar issues as decidability (e.g., non-simple zeros , $P(s, 0) = p$).

THEOREM (ASYMPTOTIC CORRECTNESS)

Let $\phi = \phi(\mathbf{p})$ be a CSL formula, with constants $\mathbf{p} = (p_1, \ldots, p_k) \in [0, 1]^k$ appearing in until formulae. Then, for $\mathbf{p} \in E$, an open subset of $[0, 1]^k$ of measure 1, there exists N_0 such that $\forall N \geq N_0$

$$s, 0 \models_{Z^{(N)}} \phi \Leftrightarrow s, 0 \models_{z} \phi.$$

Outline

FROM LOCAL TO GLOBAL

We restrict the set of properties we consider to non-nested CSL path formulae ψ.

LOCAL PROPERTY

What is the probability that a given agent Z satisfies ψ?

$$\mathbb{P}\{Z^{(N)} \models \psi\} = ?$$

GLOBAL PROPERTY

What is the probability that a fraction α of agents satisfy ψ?

$$\mathbb{P}\left\{ \sum_j \mathbf{1}\{Z_j^{(N)} \models \psi\} \bowtie N\alpha \right\} = ?$$

FROM LOCAL TO GLOBAL

Consider the client-server model, and the local property:

$$\psi = (a_{request} \vee a_{wait}) U^{\leq T} a_{timeout}$$

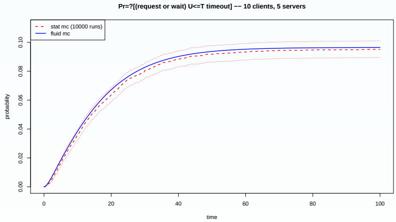

$\mathbb{P}\{Z^{(N)} \models \psi\}$ can be approximated by $\mathbb{P}\{z \models \psi\}$, using the fluid method presented above.

But how can we compute $\mathbb{P}\left\{ \sum_j \mathbf{1}\{Z_j^{(N)} \models \psi\} \geq N\alpha \right\}$?

FROM LOCAL TO GLOBAL: DECOUPLING OF AGENTS

One consequence of the fluid approximation theorem is that, in the limit, individual agents become independent. Hence

$$\mathbb{P}\{Z_1^{(N)} \models \psi, Z_2^{(N)} \models \psi\} \approx \mathbb{P}\{Z_1^{(N)} \models \psi\}\mathbb{P}\{Z_2^{(N)} \models \psi\}$$

BINOMIAL APPROXIMATION

$$\sum_j \mathbf{1}\{Z_j^{(N)} \models \psi\} \sim Bin(N, \mathbb{P}\{z \models \psi\})$$

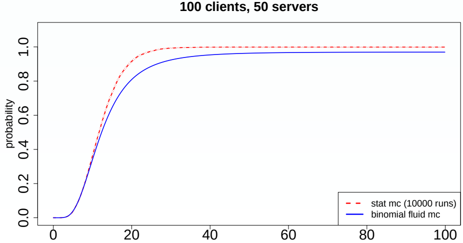

We ignore correlations between agents for finite N!

Outline

CENTRAL LIMIT APPROXIMATION

Master equation:

$$\frac{\partial P(\bar{\mathbf{X}}^{(N)},t)}{\partial t} = \sum_{\tau \in \mathcal{T}} \left(f_\tau^{(N)}(\bar{\mathbf{X}}^{(N)} - \bar{\mathbf{v}}_\tau) P(\bar{\mathbf{X}}^{(N)} - \bar{\mathbf{v}}_\tau, t) - f_\tau^{(N)}(\bar{\mathbf{X}}^{(N)}) P(\bar{\mathbf{X}}^{(N)}, t) \right)$$

If we approximate populations continuously and assume

$$\bar{\mathbf{X}}^{(N)}(t) = \mathbf{x}(t) + N^{-\frac{1}{2}}\zeta(t)$$

then the master equation can be approximated at zeroth order in N by a Fokker-Planck equation:

$$\frac{\partial \Pi(\zeta(t),t)}{\partial t} = -\sum_{s,h} \frac{\partial}{\partial \Phi_s} F_h(\mathbf{x}(t)) \left(\frac{\partial}{\partial \zeta_h} \zeta_s \Pi\left(\zeta(t),t\right) \right) +$$

$$+\sum_{\ell,r} \frac{1}{2} G_{\ell r}(\mathbf{x}(t)) \left(\frac{\partial^2}{\partial \zeta_\ell \zeta_r} \Pi\left(\zeta(t),t\right) \right); \quad G(\mathbf{x}) = \sum_{\tau \in \mathcal{T}} \mathbf{v}_\tau \mathbf{v}_\tau^T f_\tau(\mathbf{x}).$$

CENTRAL LIMIT APPROXIMATION

The solution $\Pi(\zeta, t)$ of the Fokker-Planck equation is a

Gaussian distribution

- mean $\mathbb{E}[\zeta(t)]$ such that

$$\begin{cases} \partial_t \mathbb{E}[\zeta(t)] = \mathbf{J}_F(\mathbf{x}(t))\mathbb{E}[\zeta(t)] \\ \mathbb{E}[\zeta(0)] = 0 \end{cases}$$

- covariance matrix $\mathrm{Cov}[\zeta(t)]$ such that

$$\begin{cases} \partial_t \mathrm{Cov}[\zeta(t)] = \mathbf{J}_F(\mathbf{x}(t))\mathrm{Cov}[\zeta(t)] + \mathrm{Cov}[\zeta(t)]\mathbf{J}_F^T(\mathbf{x}(t)) + G(\mathbf{x}(t)) \\ \mathrm{Cov}[\zeta(0)] = 0 \end{cases}$$

Hence $\mathbf{X}^{(N)}(t) \sim Norm\left(N \cdot \mathbf{x}(t), \sqrt{N \cdot \mathrm{Cov}[\zeta(t)]}\right)$.

Outline

COMPUTING GLOBAL PROPERTIES

1. Modify the local agent model by creating unsafe and goal copies of its states.

Client-server model, local property $\phi = (a_{request} \vee a_{wait}) U^{\leq T} a_{timeout}$:

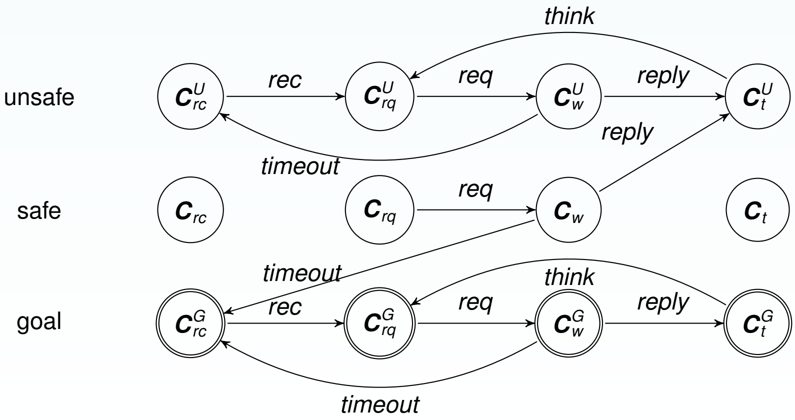

COMPUTING GLOBAL PROPERTIES

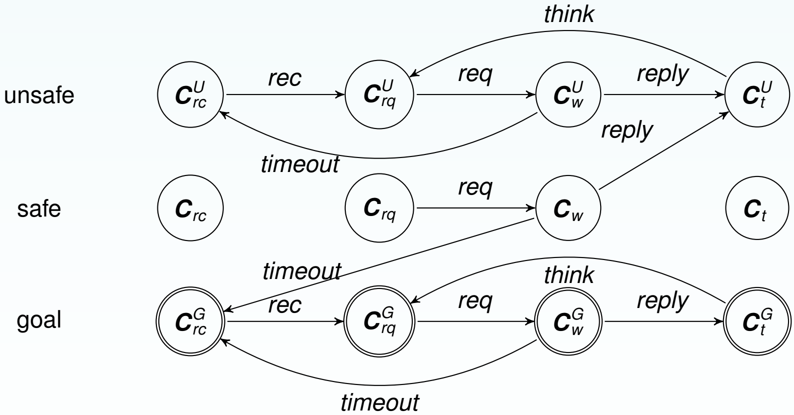

2. From the modified local model, construct a population model. Add a new variable G_ϕ, counting how many agents are in a goal state.
3. Apply central limit approximation to this new model.
4. Compute $\mathbb{P}\{G_\phi^{(N)} \geq \alpha N\}$ by $G_\phi^{(N)} \sim Norm\left(Ng_\phi(t), \sqrt{N\mathrm{Var}[\zeta_{g_\phi}(t)]}\right)$

CLIENT-SERVER - $\mathbb{P}\{G^{(N)}_{a_{request} \vee a_{wait} U^{\leq T} a_{timeout}} \geq N\theta\}$

$N = 150$, $\theta = 0.05$

CLIENT-SERVER - $\mathbb{P}\{G^{(N)}_{a_{request} \vee a_{wait} U^{\leq T} a_{timeout}} \geq N\theta\}$

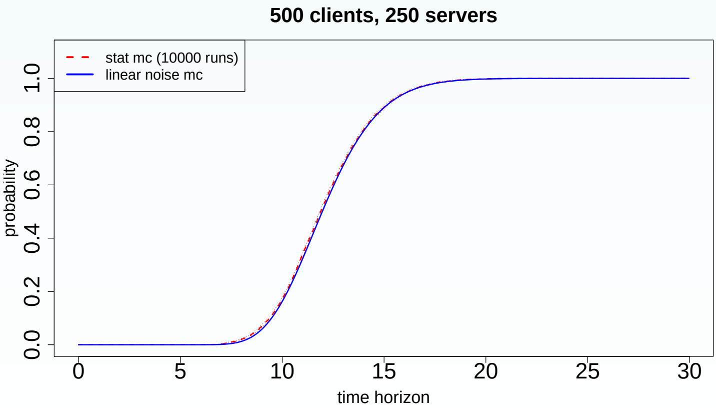

$N = 750$, $\theta = 0.05$

Client-Server - $\mathbb{P}\{G^{(N)}_{a_{request} \vee a_{wait} U^{\leq T} a_{timeout}} \geq N\theta\}$

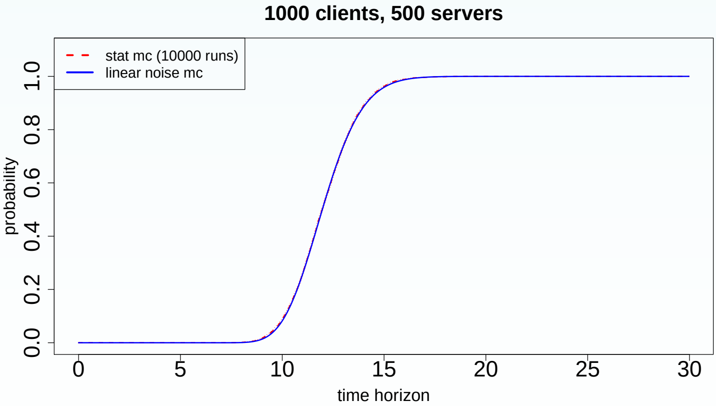

$N = 1500$, $\theta = 0.05$

CLIENT-SERVER - $\mathbb{P}\{G^{(N)}_{a_{request} \vee a_{wait} U^{\leq T} a_{timeout}} \geq N\theta\}$

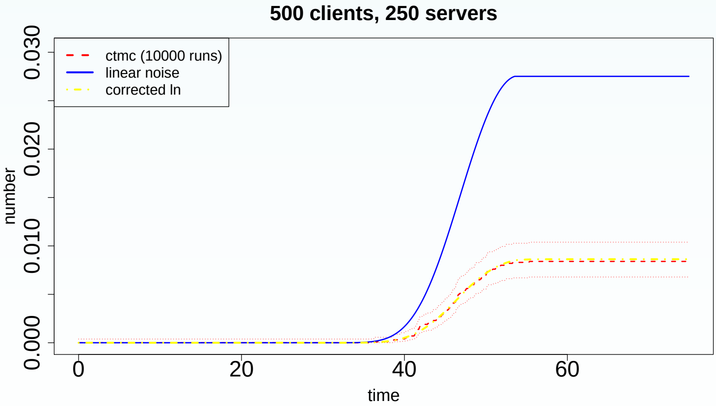

$N = 1500$, $\theta = 0.2$

CLIENT-SERVER - $\mathbb{P}\{G^{(N)}_{a_{request} \vee a_{wait} U^{\leq T} a_{timeout}} \geq N\theta\}$

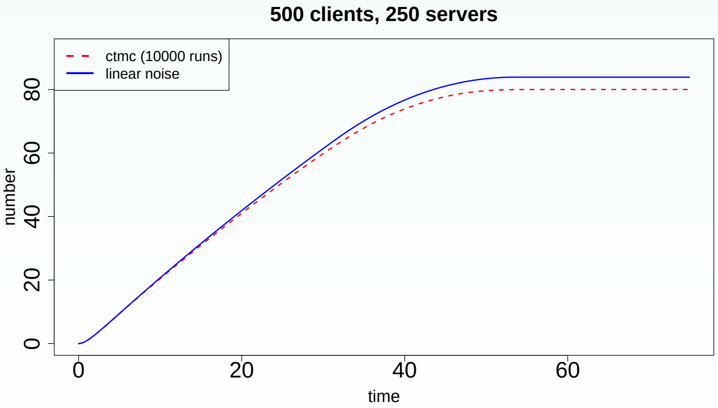

$N = 1500$, average value of Ng_ϕ and $G^{(N)}_\phi$.

CLIENT-SERVER - $\mathbb{P}\{G^{(N)}_{a_{request} \vee a_{wait} U^{\leq T} a_{timeout}} \geq N\theta\}$

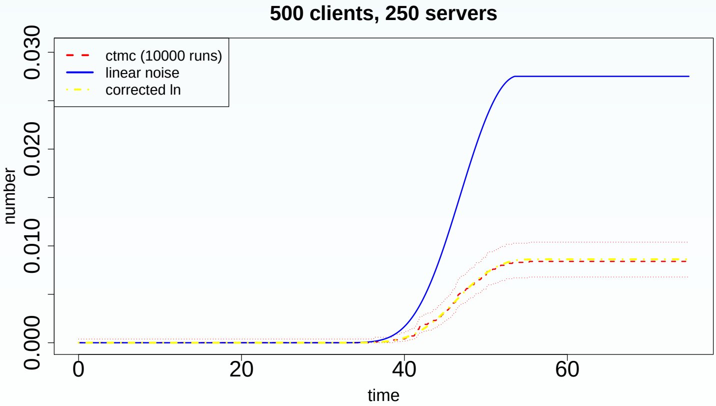

$N = 1500$, $\theta = 0.2$, corrected central limit

Outline

CONCLUSIONS

- We discussed an application of mean field theory to model check properties of medium and large population models.
- We considered first single agent properties, focussing on CSL and providing a method to model check CSL formulae versus time-inhomogeneous CTMC.
- We provided convergence results that guarantee quasi-consistence of the method.
- We then extended (non-nested) single agent properties to population level, using the central limit approximation.
- For collective properties, we have also considered a richer class of path properties specified by (restricted) DTA .

FUTURE WORK

- Use error bounds for mean field convergence to provide a (very rough) estimate of the error.
- Include rewards, and time-unbounded/ steady state, when possible.
- Working implementation.
- Consider other logics on single agents (e.g. MTL, LTL).
- Consider different properties for collective probabilities, specified by timed automata or LTL (in a local to global perspective and in a global perspective).
- Understand accuracy of central limit theorem.

BIBLIOGRAPHY

COURSE TOPICS

- L. Bortolussi, J. Hillston, D. Latella, M. Massink.Continuous Approximation of Collective Systems Behaviour: a Tutorial. Performance Evaluation, 2013.
- L. Bortolussi, J. Hillston: Fluid Model Checking. CONCUR 2012.
- L. Bortolussi, J. Hillston: Model Checking Single Agent Behaviours by Fluid Approximation, submitted to Information and Computation.
- L. Bortolussi, R. Lanciani. Model Checking Markov Population Models by Central Limit Approximation. QEST 2013.

BIBLIOGRAPHY

RELATED WORK

- M. Tribastone, S. Gilmore, J. Hillston: Scalable Differential Analysis of Process Algebra Models. IEEE Trans. Softw Eng. 2012.
- R.A. Hayden, A. Stefanek, J.T. Bradley. *Fluid computation of passage-time distributions in large Markov models.* Theor. Comput. Sci. 2012.
- R. A. Hayden and J. T. Bradley and A. Clark: Performance Specification and Evaluation with Unified Stochastic Probes and Fluid Analysis, IEEE Trans. Software Eng., 2013.
- R. Darling, J. Norris. Differential equation approximations for Markov chains. *Probability Surveys*, 2008.

BIBLIOGRAPHY

MODEL CHECKING ICTMC

- J.P. Katoen, A. Mereacre. Model Checking HML on Piecewise-Constant Inhomogeneous Markov Chains. FORMATS 2008.
- T. Chen, T. Han, J.P. Katoen, A. Mereacre: LTL Model Checking of Time-Inhomogeneous Markov Chains. ATVA 2009.
- T. Chen, T. Han, J.P. Katoen, A. Mereacre: Model Checking of Continuous-Time Markov Chains Against Timed Automata Specifications. Logical Methods in Computer Science 7, 2011.

THE END!

Thanks for the attention

Questions?