# Static and Dynamic Analysis of Timed Distributed Traces

Parasara Sridhar Duggirala, Taylor T Johnson, Adam Zimmerman and Sayan Mitra

Coordinated Science Laboratory

University of Illinois at Urbana Champaign

33rd IEEE Real-Time Systems Symposium, San Juan, Puerto Rico

# Detecting global properties from local data recordings

How to infer global properties from independently recorded local data logs or traces of a distributed cyber-physical system?

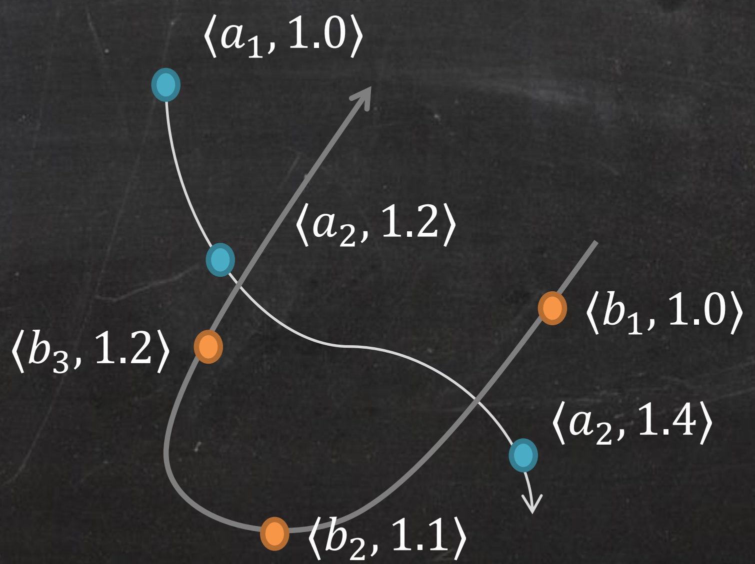

# Related Work

- Inferring global properties from independently recorded local traces in a distributed ~~cyber-physical~~ systems
  - [Babaoglu & M. Raynal 93], [Garg et al. 96-2005], [Cooper and K. Marzullo 91]
  - Lattice of "happens before" relation on events, exponential algorithms for traversing the lattice, polynomial for special classes of predicates

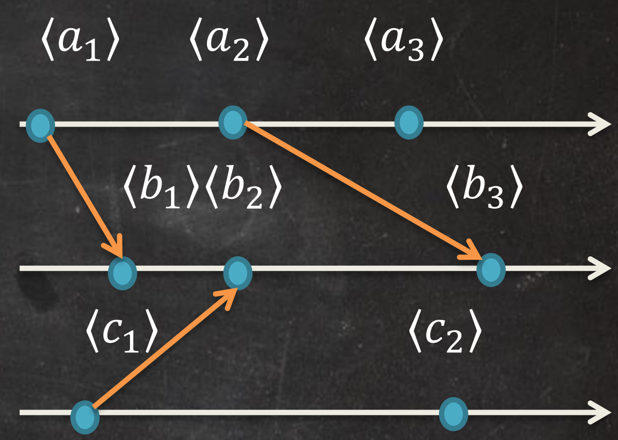


- CPS Challenges
  - Programs generate events
  - Clocks, albeit imperfectly synchronized
  - Discrete data but continuous evolution

# Traces and Consistent Executions

- An observation $\langle x_i, clk_i \rangle$ of age
  - $x_i$ : recorded state
  - $clk_i$: Timestamp from local clock
- A trace $\beta_i = \langle x_{i1}, clk_{i1} \rangle, \langle x_{i2}, c$
- System trace $\beta$ is a collection of

*Model **A*** comes from modeler or from static analysis.

- A: $x_i$ is Lipchitz
- A: Polynomial
- A: Hybrid automaton

1.0⟩

.4⟩

- **ConExec$(\beta, A)$**: Given a system model *A*, this is the set of behaviors of A that is consistent with $\beta$
- Property : predicates on states of individual agents
  - E.g., Did robots ever get closer than r ?

# Traces, Executions and Properties

Given a system trace $\beta$, a system model A, and a global property P does every consistent execution satisfy P at *all times* ?

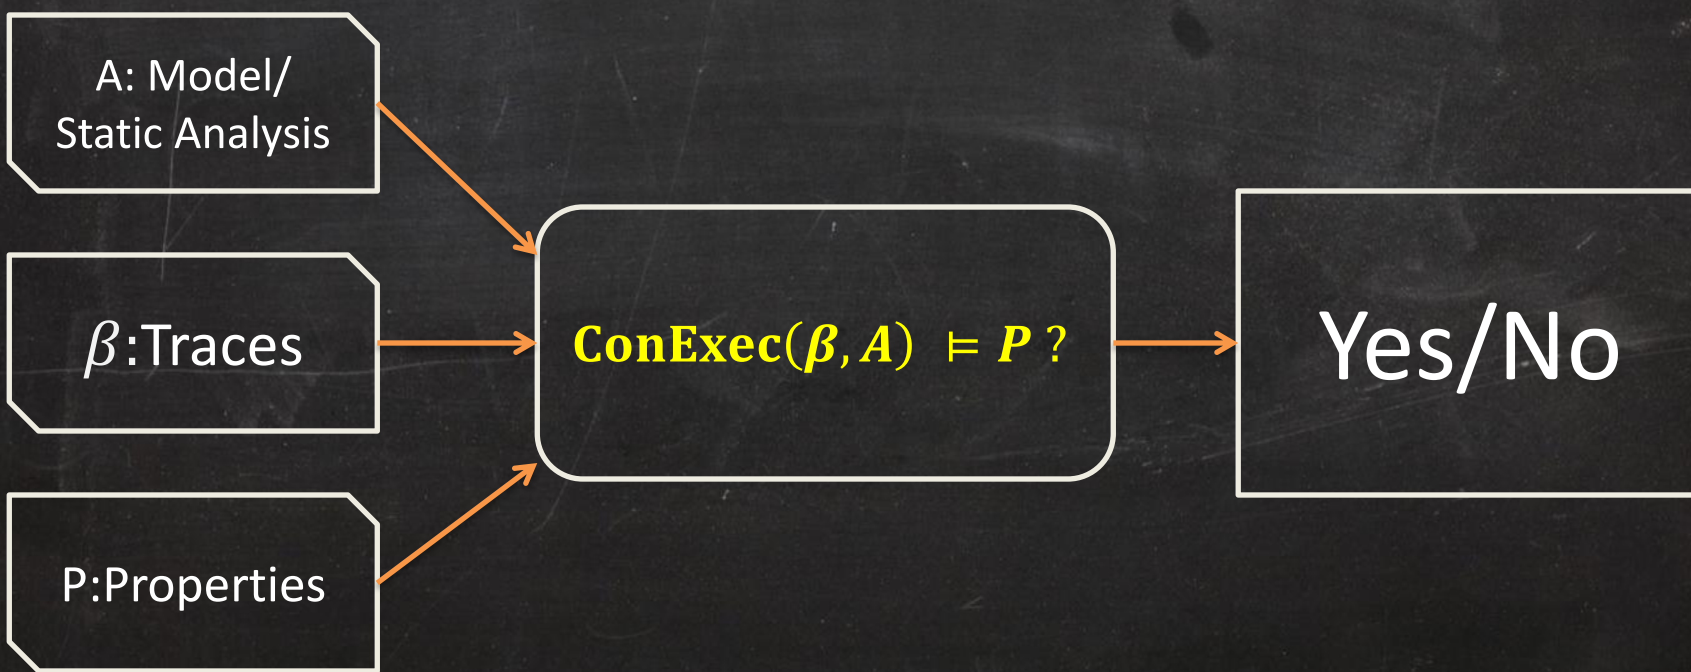

# Approach & outline

- Bounds on observation times from clocks and messages in

- Reachable states between observations from static analysis $\beta$

- Experiments with mobile robots on StarL

# Real-time bounds from local time stamps

- A trace $\beta$ is $\sigma$-synchronized if for every observation $\langle x, clk \rangle$, every consistent execution visits $x$ some time in $[clk_i - \sigma, clk_i + \sigma]$
- A trace $\beta$ is tightly $\sigma$-synchronized if (1) and for every time in the interval $[clk - \sigma, clk + \sigma]$ there is a consistent witness execution which visits $x$ at the time

# Real-time bounds from messages

- $L(x)$: Greatest lower bound on real-time for occurrence of $x$,
- $\mathrm{L(x)} = max\left(clk - \sigma, \max_{y \leftarrow x} U(y)\right)$
- $U(x)$: Least upper bound = $min\left(clk + \sigma, \mathrm{mix}_{x \leftarrow y} L(y)\right)$

**y** ← **x** in $\beta$ if and only if

1) same agent and **x** recorded after **y** or
2) **y** = send(m) and **x** =receive(m) or
3) **y** ← **w** and **w** ← **x**

# Inferring Global Predicates from Traces

- When did the observations occur?
- For observation $\langle x, clk \rangle$ we have (*tight*) *observation intervals* $[L(x), U(x)]$
- What happens between observations ?
- How to *over-approximate* the set of states reachable through $\text{ConExe}(\beta, A)$? Use static analysis

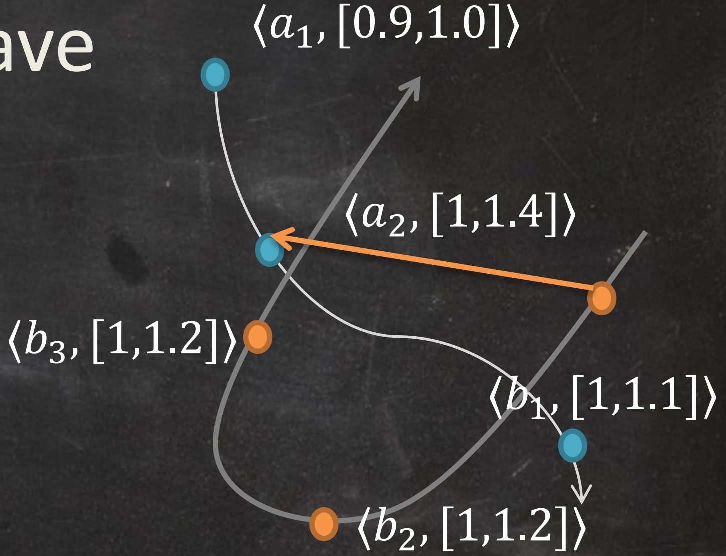

# Symbolic Over-approximation

$\mathbf{A}$: Model from static analysis

$Post(\mathbf{A}, x_j, t)$: Reach from $x_j$ in $t$ time

$Pre(\mathbf{A}, x_j, t)$: Reach to $x_j$ in $t$ time

$Reach(\{x_1, \ldots, x_m\}, t)$: Reachable through $x_1, \ldots, x_m$ at $t$

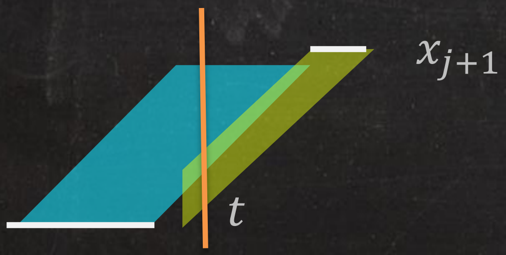


$\mathbf{A}$: $\dot{x} = Ax$

$Post(\mathbf{A}, x_j, t) = x_j e^{At}$

$\mathbf{A}$: $\dot{x} \in [a, b]$

$x_j + at \le Post(\mathbf{A}, x_j, t) \le x_j + bt$

$$Reach(\{x_1, \ldots, x_m\}, t) = \exists t_1 < \cdots < t_m: \bigwedge_{j=1}^{m} L(x_j) \le t_j \le U(x_j) \bigwedge_{j=1}^{m-1} [t_j \le t \le t_{j+1} \Rightarrow (Post(x_j, t - t_j) \wedge Pre(x_{j+1}, t_{j+1} - t))]$$

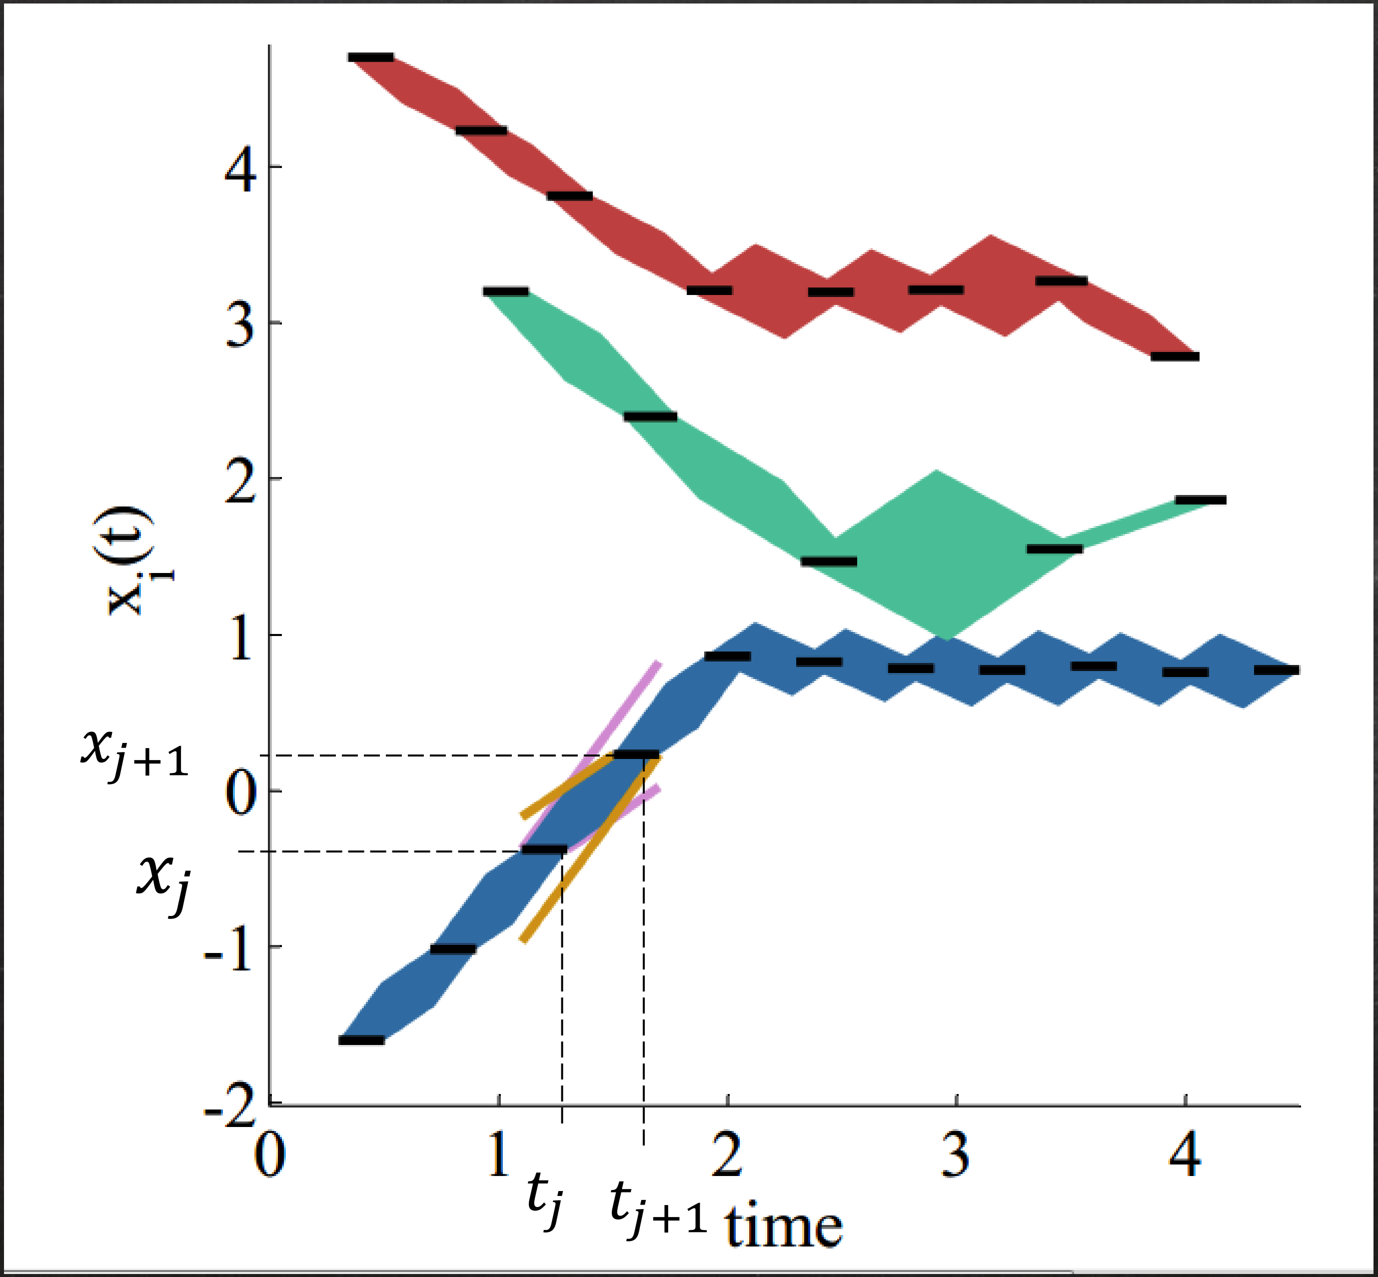

4
3
2
1
0
-1
-2
$x_i(t)$
$x_{j+1}$
$x_j$
0
1
2
3
4
$t_j$
$t_{j+1}$
time

# Soundness and Precision

- Fix a trace $\beta$ and a time t
  - (sound) At time t any consistent execution satisfies $Reach(\beta, t)$
  - (precision) If *Post()* and *Pre() are* exact, observation intervals are disjoint, and $\sigma$-synchronization is tight, then every state in $Reach(\beta, t)$ is visited by some consistent execution at time t
- Check property (separation, deadlock) over $Reach(\beta, t)$

# Experiments: Debugging robot apps!

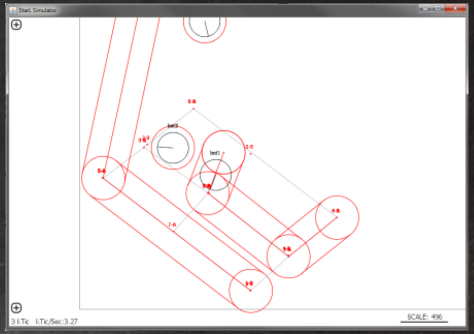

- StarL: API for distributed robotics
  - Primitive functions, e.g., mutual exclusion, leader election, motion control, …
  - Logs traces
  - Simulator

- Test bed: iRobot Create, Android Smartphone, Bluetooth, Vision-based indoor positioning
- Applications & properties
  - Waypoint following
  - GeoCast
  - Light painting

# Experiments 1: Diversity and Scaling

| N | x = 75 ms | 150 ms | 250 ms | 500 ms |
|---|---|---|---|---|
| 4 | 42 | 24 | 10 | 5 |
| 8 | 92 | 48 | 22 | 10 |
| 12 | 246 | 114 | 34 | 16 |
| 16 | 10 m | 4 m | 49 | 24 |
| 20 | 20 m | 8 m | 67 | 34 |

Always separation (d = 10 cm) for 5 mins @ x ms

| Property | N | Sat? | Ana Time (sec) | Mem (Mb) | Frmla size (Kb) |
|---|---|---|---|---|---|
| Always Separation (d = 25) | 4 | Yes | 1.5 | 3.07 | 3.9 |
| | 12 | **Yes** | 14 | 8.66 | 14.9 |
| | 20 | **Yes** | 81 | 18.6 | 31.6 |
| Always Separation (d = 10) | 4 | Yes | 1.5 | 3.07 | 3.9 |
| | 12 | **No** | 14 | 8.66 | 14.9 |
| | 20 | **No** | 81 | 18.6 | 31.6 |
| Always Georecv | 4 | Yes | 1 | 1.24 | 3.2 |
| | 12 | Yes | 1.7 | 3.67 | 9.5 |
| | 20 | Yes | 1.9 | 8.35 | 16 |

# Experiments 2: Impact of Precision of Static Analysis

- System model **precision**
  - VB: velocity bounds
  - OI: observation intervals

- **Lower precision** model ($\pm 20ms$) produces **more conservative** answers than the higher precision models ($\pm 5ms$)

| | VB = $\pm 0$ cm/s | VB = $\pm 20$ cm/s | VB = $\pm 20$ cm/s |
|---|---|---|---|
| | Separation (d=10 cm) | | |
| OI $= \pm 5ms$ | yes | yes | no |
| OI $= \pm 10ms$ | yes | no | no |
| OI $= \pm 20ms$ | no | no | no |
| | Georeceive | | |
| delay = 0ms | yes | yes | yes |
| delay = 20ms | yes | yes | no |
| delay = 50ms | no | no | no |

# Summary and Future Directions

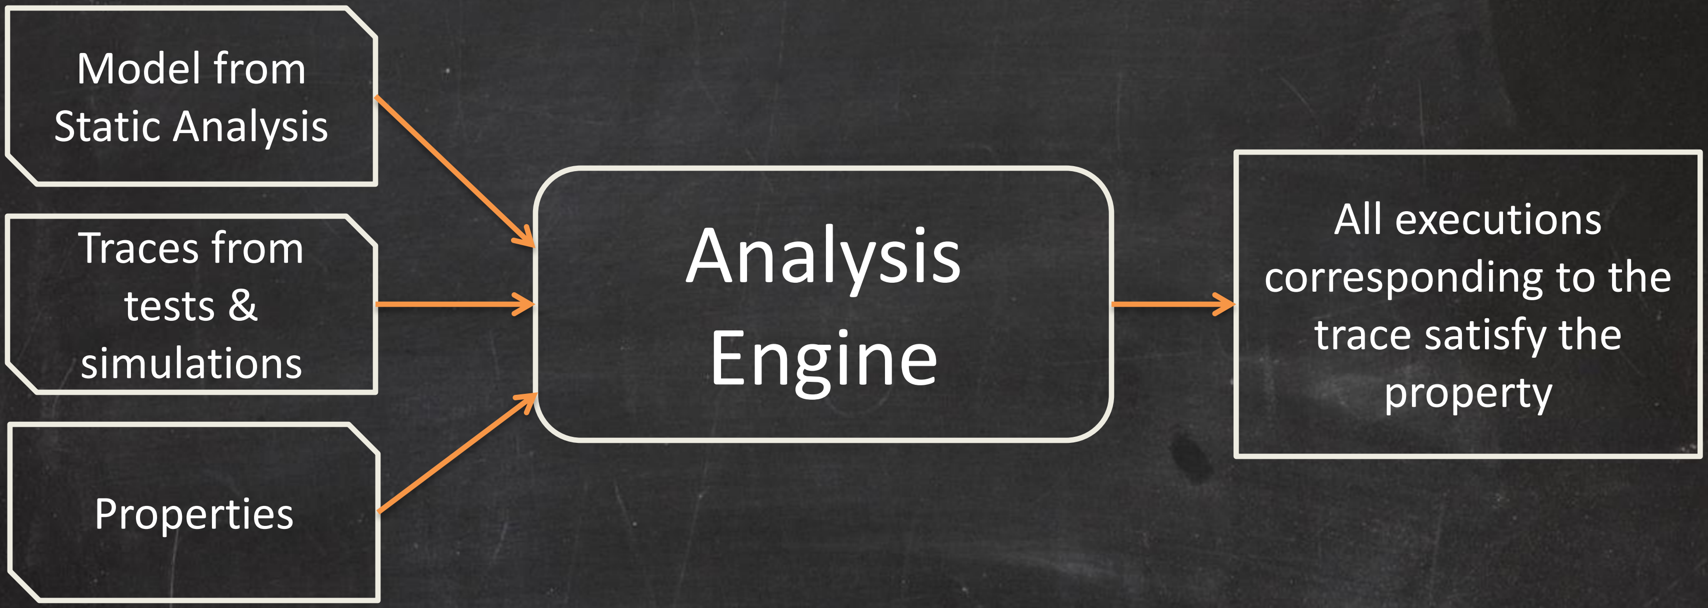


## Future directions:

- Investigate static vs. dynamic analysis trade-off
- Close the loop from the output of the Engine to the generation of traces

We gratefully acknowledge the support of NSF and AFOSR for this research

# Questions?

# Conclusions and Future Work

- Sound algorithm for analyzing traces of distributed real time systems
- Completeness in some cases
- Future work:
  1) Accuracy vs Sampling
  2) Generalization of Completeness results
  3) Application to other domains like Power-Grid systems or UAV navigation