

A Logical Framework for Self-Optimizing Networked Cyber-Physical Systems

Mark-Oliver Stehr, Minyoung Kim, Carolyn Talcott

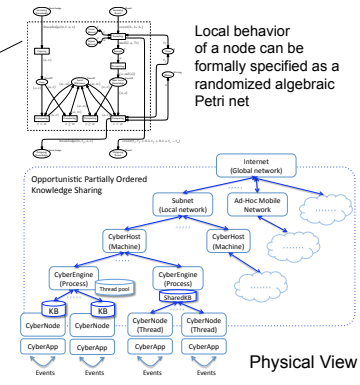
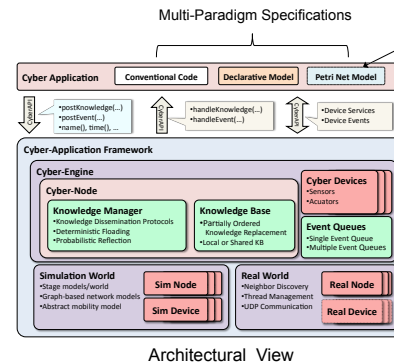
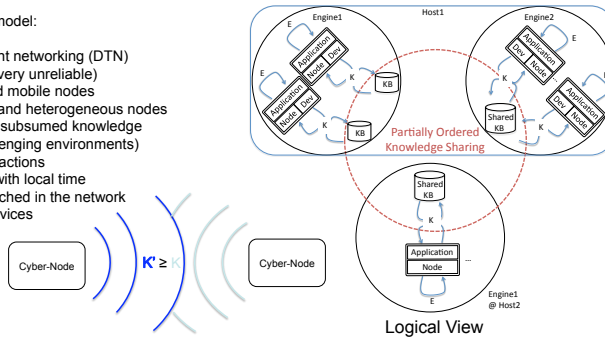
<http://ncps.csl.sri.com>

Students: Jinwoo Kim (Seoul National Univ.), Vincent Wang (UPenn)



An Application Framework for Networked CPS

- Based on new loosely-coupled distributed computing model: **Partially Ordered Knowledge Sharing**
- Inspired by our earlier work on delay-/disruption-tolerant networking (DTN)
- Minimal assumptions on network connectivity (can be very unreliable)
- Works with dynamic topologies, network partitions, and mobile nodes
- Designed for heterogeneous networking technologies and heterogeneous nodes
- Partial order allows the network to replace obsolete or subsumed knowledge
- Global consistency is not enforced (impossible in challenging environments)
- Avoids strong non-implementable primitives, e.g. transactions
- Locally each cyber-node uses an event-based model with local time
- Events are local, but knowledge can be shared and cached in the network
- Each cyber-node can have attached cyber-physical devices
- Framework supports
 - model-based simulation/analysis mode
 - real-world deployment/execution mode
- Applications cannot distinguish between simulation and reality



Distributed Declarative Control

Logical Theory for Distributed Surveillance by a Team of Mobile Robots in an Instrumented Space

Forward Clauses:
 F1: $Noise(T, A) \Rightarrow Trigger(T, A)$.
 F2: $Motion(T, A) \Rightarrow Trigger(T, A)$.
 F3: $Adjacent(A, B) \Rightarrow Adjacent(B, A)$.

Backward Clauses:

- B1: $Interest(T, I, R) \Leftarrow Result(T, Tr, 0, I), Deliver(T, Tr, 1, I, R)$.
- B2: $Deliver(T, Tr, Nd, I, R) \Leftarrow Delivered(T, Tr, Nd, I, R)$.
- B3: $Delivered(T, Tr, Nd, I, R) \Leftarrow Position(Tp, R, A), Position(Tp, R', A'), R' \neq R, MoveTo(T, Tr, Nd, 0, \infty, R', A), Deliver(T, Tr, Nd, I, R)$.
- B4: $Result(T, Tr, Nd, I) \Leftarrow ComImage(T, Tr, Nd, I), I' = Extract(I)$.
- B5: $ComImage(T, Tr, Nd, I) \Leftarrow RawImage(T, Tr, Nd, I), I' = Compress(I)$.
- B6: $RawImage(T, Tr, Nd, I) \Leftarrow Trigger(T, A), T_1 \leq Tr, MoveTo(T, Tr, Nd, 0, Tr + \Delta_{ad}, R, A), TakeSnapshot(T, Tr, Nd, 0, Tr + \Delta_{ad}, A, I)$.
- B7: $TakeSnapshot(T, Tr, Nd, D, A, I) \Leftarrow Snapshot(T, Tr, Nd, Ts, A, I), T_1 \leq Ts, Ts \leq D$.
- B8: $MoveTo(T, Tr, Nd, W', D, R, B) \Leftarrow Position(T, Tr, Nd, I), I' \leq D$.
- B9: $MoveTo(T, Tr, Nd, W', D, R, B) \Leftarrow Adjacent(A, B), W' > -b_w, W = W' - 1, MoveTo(T, Tr, Nd, W, D, R, A), MoveTo(T, Tr, Nd, W', D, R, A, B)$.

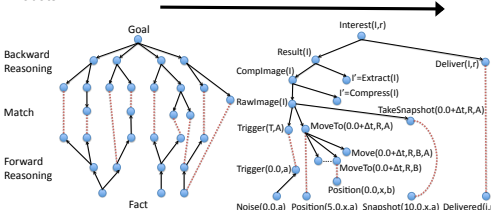
Replacement Ordering: (f denotes a fact and g a goal and x denotes either)
 O1: $f: Position(t_p, r, \dots) < f: Position(t_p, r, \dots)$ if $t_p < t_p$.
 O2: $x: X(t_1, \dots, t_n) < g: Interest(t_1, \dots, t_n)$ if $t_1 < t_1$.
 O3: $x: X(t_1, t_2, t_3, \dots) < f: Result(t_1, t_2, t_3, \dots)$ if $x \neq f: Result$.
 O4: $x: X(t_1, t_2, t_3, \dots) < f: Deliver(t_1, t_2, t_3, \dots)$ if $x \neq f: Deliver$.

Variables: T : time, D : snapshot deadline, A and B : area, R : robot, I : image or derived information, N : identifier, W : weight
 Constants: Δ_{ad} : relative snapshot deadline (max. delay from trigger event), b_w : bound for weight (diameter of the floor plan)

Challenges:

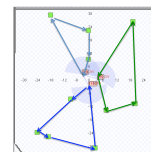
- Traditional logics are not designed for distributed reasoning
- Logics are traditionally closed systems, i.e. not interactive
- Here we consider the NCPS as a single asset
- Logical theory/specification is available to all nodes
- Nodes contribute resources according to their capabilities
- Knowledge = Facts + Goals is transparently shared
- Facts can represent observations
- Goals can represent control objectives
- Distributed logical framework
- Integrated forward and backward reasoning
- Partial order is essential part of the distributed logic
- Tested with abstract mobility model and Stage multi-robot simulator

Abstract Mobility Model with two Robots



Distributed Dynamic Optimization

- Distributed optimization fits well into the partial-ordered knowledge-sharing model
- Replacement order is defined by objective function (solution fitness)
- Case study: Multi TSP applied to sensor data collection by a team of mobile robots
- Algorithm: Distributed version of a quantum-evolutionary optimizer
- Optimizer performance is studied on Internet-wide testbed (Planet Lab)
- Same optimizer can operate as a cyber-node in an NCPS (e.g. in each robot)
- Optimization problem becomes dynamic (feedback through environment)
- Key features: robustness and fault-tolerance



Possible next steps:

- Generalization: use of a generic optimization framework
- Distributed multi-objective optimization (Pareto optimality)
- Combining optimization and declarative control
- Use of Weighted/Quantitative/Probabilistic Logic

Stage Multi-Robot Simulator

Robot Failure

Opportunistic Knowledge Sharing

