Private Consensus: A Building Block for Secure Distributed Control

Zhenqi Huang, Sayan Mitra, and Geir Dullerud
Information Trust Institute
University of Illinois at Urbana Champaign

July 12th 2012

A Building Block for Distributed Control

Load balancing

Sensor fusion

Filtering

Flocking

Distributed coordination

…

Iterative Consensus Problem

- N agents
- Initial states $v(0) = [\theta_1, \ldots, \theta_N]$
- Agents interact & update states
  - $v(t+1) = P(t)\, v(t)$
  - $P(t)$: Matrix capturing communication and update rule
- Goal 1: **Converge**

  $$\lim_{t\to\infty} v(t) = v(0)$$

  - Average of the initial values
  - Equal spacing

A Use Case

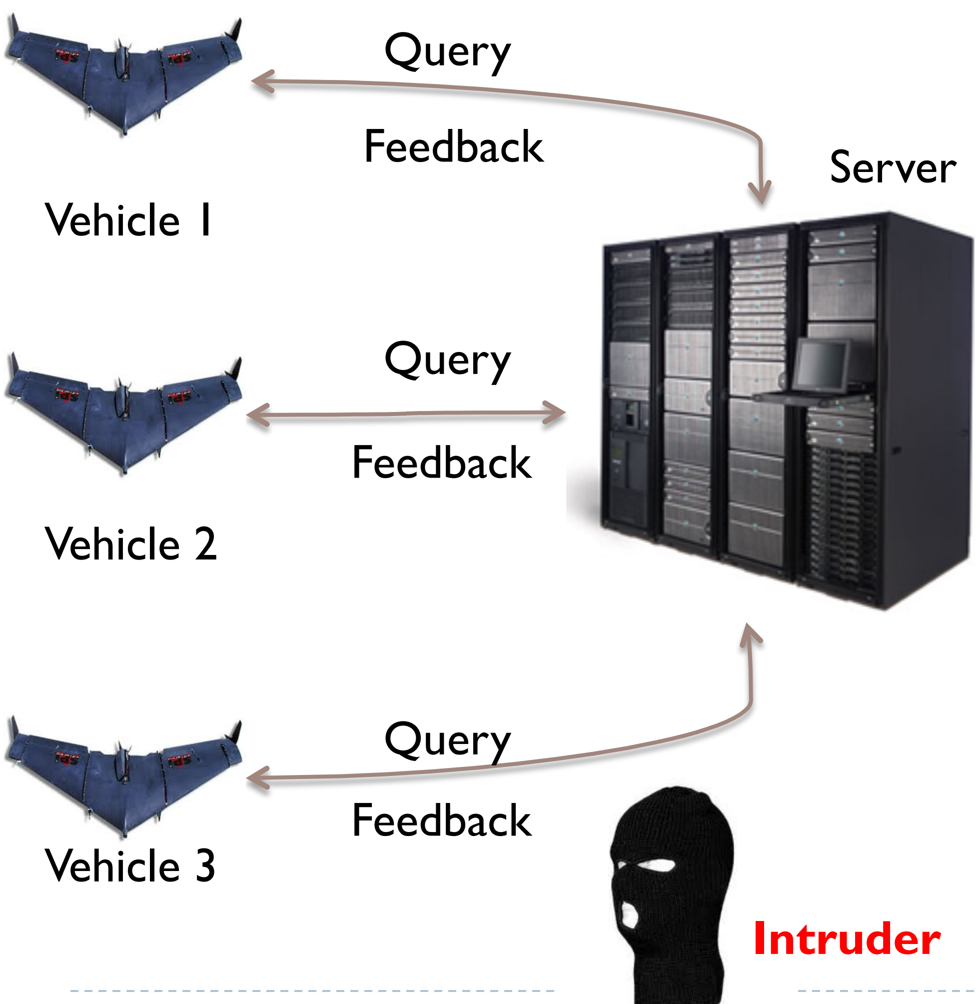

- N vehicles require move in a group while protecting their locations
- **Private consensus on velocities**
- Private data: initial velocity
- Vehicles send query, server computes feedback, vehicles' local controller updates state based on feedback
- **Eventually all vehicles move with same velocity (consensus)**
- Intruder can **see all messages** as well as **server's states** (honest but curious)
- **Privacy: Intruder cannot guess vehicles initial velocity, and therefore its current position,** with any high degree of confidence

Private Iterative Consensus Problem

- Requirements: Achieve **consensus** while **protecting privacy** of $\mathrm{v}(0)$
- Randomization for privacy ➔ probabilistic convergence
  - A mechanism is **(1-b) probability** and **r-radius accuracy** if for any initial state $\mathrm{v}(0)$ all executions starting from $\mathrm{v}(0)$, converges to within r of $\bar{v}(0)$ with probability **(1-b)**
- Differential privacy **[Dwork et al. 2006]**
  - Privacy of individuals participating in statistical databases
- We adapt differential privacy to continuous computations:
  - Two initial states $\mathbf{v}(\mathbf{0})$ **and** $\mathbf{v}'(\mathbf{0})$ **are $\boldsymbol{\delta}$-adjacent** if there exists j such that $v_j(0) - v'_j(0) \le \delta$ and for all other agents $v_i(0) = v'$

  - If for **any sequence of observations $\boldsymbol{\beta}$** & any **$\boldsymbol{\delta}$-adjacent ini tes** $\mathrm{v}(0)$, $\mathrm{v}'(0)$

$$Pr[Execution\ from\ v(0)\ produces\ \beta]\ /Pr[Execution\ from\ v'(0)\ produces\ \beta] \le e^{\varepsilon\delta}$$

  - Then, the randomized mechanism has ε-differential privacy.

- **For what class of adversaries can we find mechanisms for solving PIC problem? How accurately? With what cost?**

A Randomized Mechanism

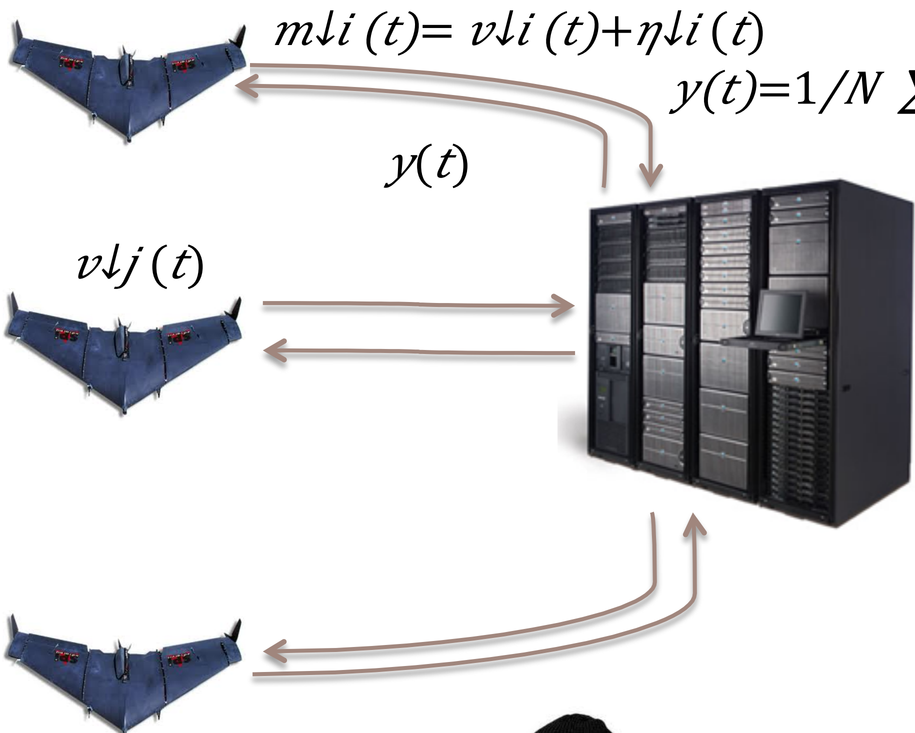

- ε: Security parameter
- $\lambda \in (0,1)$, $q \in (0, \lambda)$
- $\eta_i(t) \sim Lap(q$

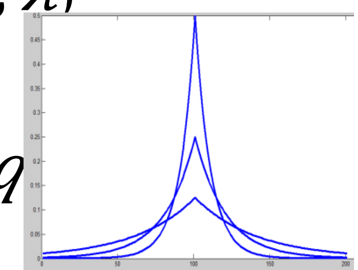

- **Decaying noise**
- y(t) feedback from server
- **Local control law**
- $v_i(t+1) = (1-\lambda) v_i(t) + \lambda y(t)$

Main Result

- **Theorem.** For any ε, N and b, our mechanism achieves (1-b) **probability** accuracy for **radius** $O(1/\varepsilon\sqrt{bN})$ and **ε-differential privacy.**

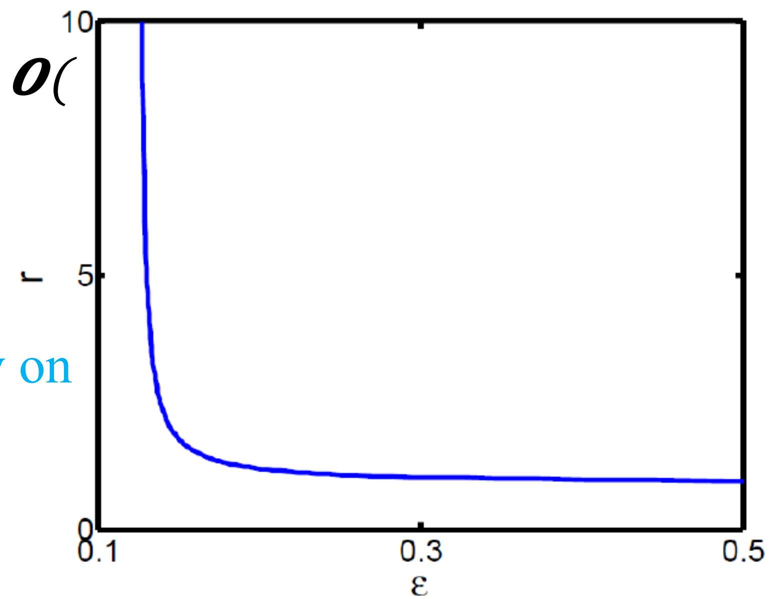

- **Remark.** For a given level of probabilistic guarantee (1-b), the accuracy depends inversely on privacy (ε), $\sqrt{b}$, and **directly on $\sqrt{N}$**
- **Lessons.**
  - Initially distort private data with large noise & progressively decay the noise level; **noise should converge slower than the system's inertia so as to cover the trail of dynamics**
  - Proof technique relies on **constructing bijection between two sets of executions** starting from adjacent start states

Distributed Generalization

- Similar result for fully distributed algorithm (server-less) in which clients exchange information with their neighbors
  - Adversary can see **all messages** as well as **internal states** of a set (**C**) of compromised clients
  - Differential privacy of good clients

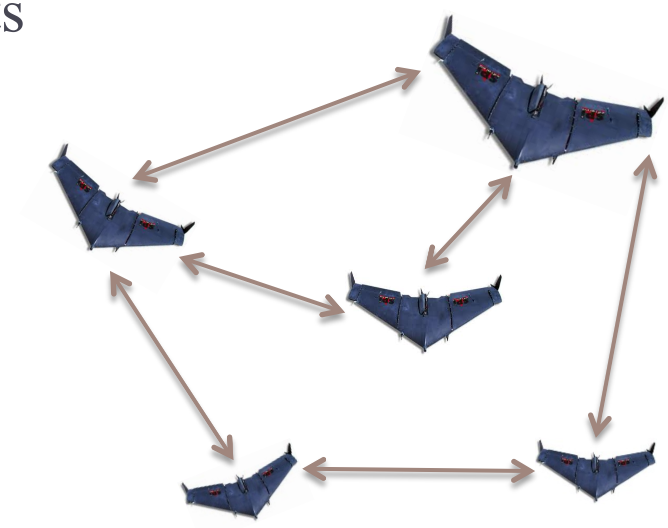

Summary and Next Steps

- A mechanism for private iterative consensus with **eavesdropping/ honest but curious adversaries**
  - Shows trade-off between privacy and accuracy
  - Math framework for rigorous proofs

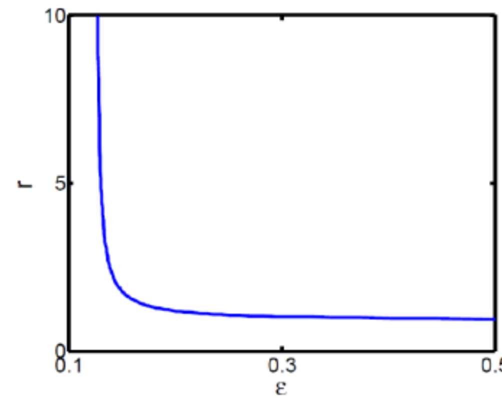

- **Next in PIC**
  - **Lower bound** for same class of adversaries
  - **More powerful adversaries**: State corruption, Time varying C(t),
- **Problems requiring Safety & Convergence**
  - E.g. Maintain safe separation while converging to a formation
  - Monolithic models, distributed systems
- **Synthesizing secure protocols**
- **Proof automation [Barthe et al. 2012, Datta et al. 2012]**

Relation to Science of Security

- Conflicting requirements for Private Consensus problem **Control Systems**
  - Differential Privacy **Security**
  - Accuracy **Control performance**
- Scientific Question: **"Is is possible to have ε-differential privacy and r-accuracy for a given class of adversaries?"**
- Amenable to the standard scientific method? Hypothesis, experimentation, validation?

We believe, yes! **Proof is Data**

- **Hypothesis**: An algorithm or a conjecture for a lower bound
- **Experiments**: Proofs and counter-examples based on specific classes of algorithms. Try to come up with specific algorithms and lower bound proofs

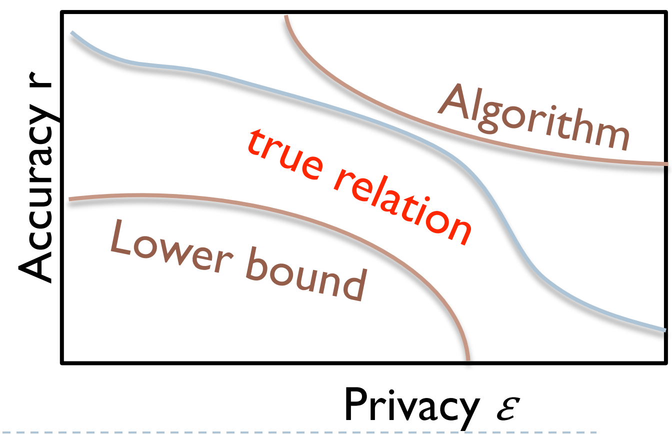

Questions/Comments?