ETH zürich

ETH AI CENTER

Near-Optimal Multi-Agent Learning for Safe Coverage Control


Manish Prajapat

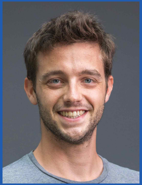
Matteo Turchetta


Melanie N. Zeilinger[†]


Andreas Krause[†]

What is Safe Multi-agent Coverage Control?

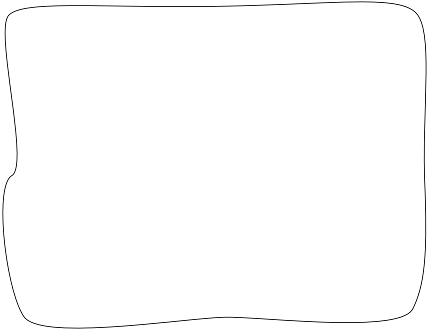

What is Safe Multi-agent Coverage Control?

- Density
  - Spatially distributed events

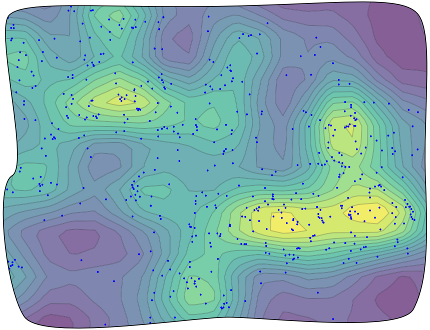

What is Safe Multi-agent Coverage Control?

- Density
  - Spatially distributed events
- Multiple agents
  - A set of agents coordinate in the process

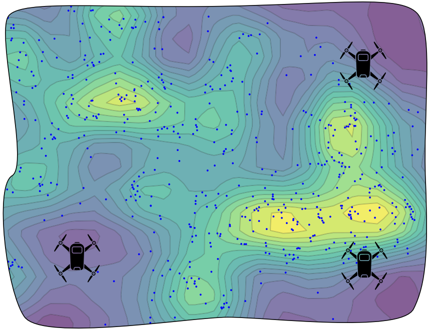

What is Safe Multi-agent Coverage Control?

- Density
  - Spatially distributed events
- Multiple agents
  - A set of agents coordinate in the process

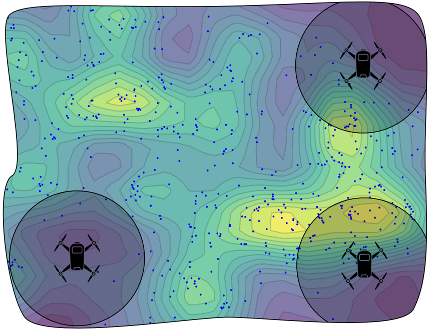

What is Safe Multi-agent Coverage Control?

- Density
  - Spatially distributed events
- Multiple agents
  - A set of agents coordinate in the process
- Coverage Control
  - Agents navigate to cover the space "as best as they can"

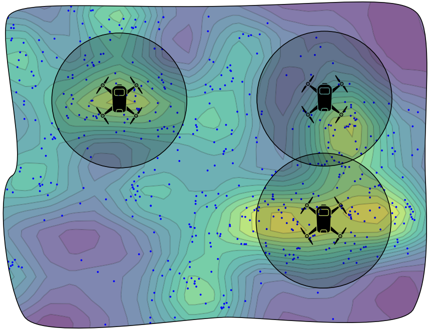

What is Safe Multi-agent Coverage Control?

- Density
  - Spatially distributed events
- Multiple agents
  - A set of agents coordinate in the process
- Coverage Control
  - Agents navigate to cover the space "as best as they can"
- Safe
  - Constraints by the environment

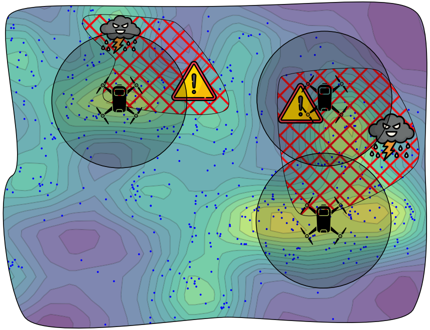

What is Safe Multi-agent Coverage Control?

- Density
  - Spatially distributed events
- Multiple agents
  - A set of agents coordinate in the process
- Coverage Control
  - Agents navigate to cover the space "as best as they can"
- Safe
  - Constraints by the environment
  - Safe execution

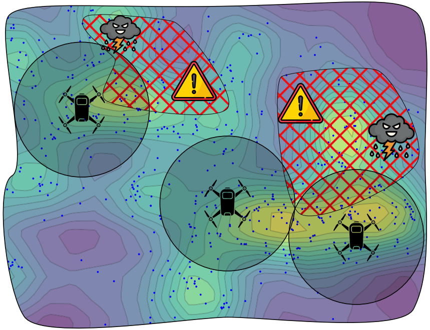

What is Safe Multi-agent Coverage Control?

- Density
  - Spatially distributed events
- Multiple agents
  - A set of agents coordinate in the process
- Coverage Control
  - Agents navigate to cover the space "as best as they can"
- Safe
  - Constraints by the environment
  - Safe execution

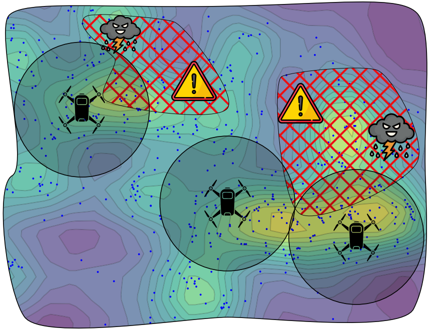

Real-world applications: Bio-diversity monitoring, Swarm robots, 3D scene reconstruction, etc.

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \underbrace{\sum_{x^i \in X}}_{\textit{agents' location}} \underbrace{\sum_{v \in D^{i-}}}_{\textit{disc of agent } i} \underbrace{\rho(v)}_{\textit{density}}$$

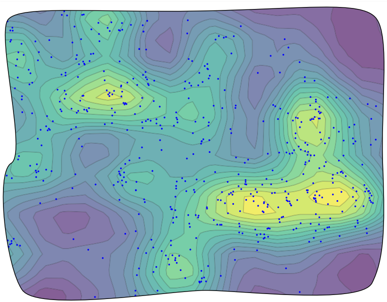

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \sum_{\underbrace{x^i \in X}_{\textit{agents' location}}} \; \sum_{\underbrace{v \in D^{i-}}_{\textit{disc of agent } i}} \underbrace{\rho(v)}_{\textit{density}}$$

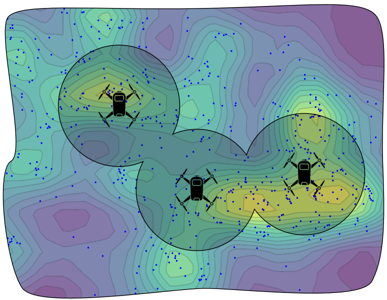

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \sum_{\underbrace{x^i \in X}_{\textit{agents' location}}} \sum_{\underbrace{v \in D^{i-}}_{\textit{disc of agent } i}} \underbrace{\rho(v)}_{\textit{density}}$$

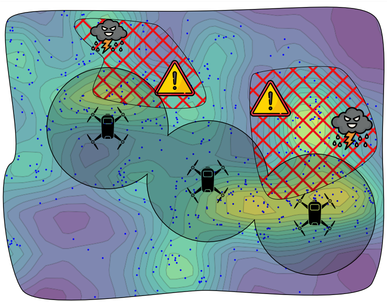

Goal: $\max_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t) \geq 0}_{\textit{constraint}} \; \forall t$

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \underbrace{\sum_{x^i \in X}}_{\textit{agents' location}} \underbrace{\sum_{v \in D^{i-}}}_{\textit{disc of agent } i} \underbrace{\rho(v)}_{\textit{density}}$$

Challenges

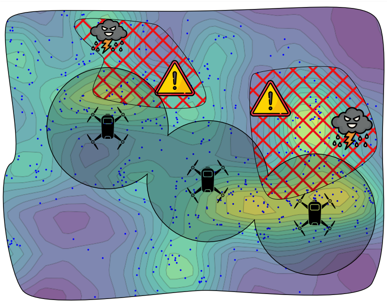

Goal: $\max_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t) \geq 0}_{\textit{constraint}}\ \forall t$

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\text{coverage}} = \underbrace{\sum_{x^i \in X}}_{\text{agents' location}} \; \underbrace{\sum_{v \in D^{i-}}}_{\text{disc of agent } i} \underbrace{\rho(v)}_{\text{density}}$$

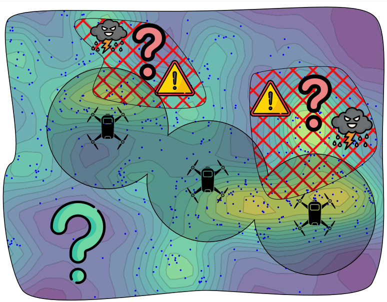

Goal: $\max_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t) \geq 0}_{\text{constraint}} \; \forall t$

Challenges

- A-priori unknown density
- A-priori unknown constraints

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \underbrace{\sum_{x^i \in X}}_{\textit{agents' location}} \underbrace{\sum_{v \in D^{i-}}}_{\textit{disc of agent } i} \underbrace{\rho(v)}_{\textit{density}}$$

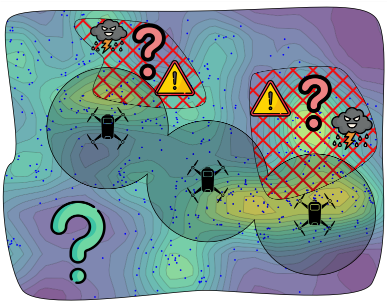

Goal: $\max_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t) \geq 0}_{\textit{constraint}}$ $\forall t$

Challenges

- A-priori unknown density
- A-priori unknown constraints
- Even with known functions $\rightarrow$ NP Hard problem ☹

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \sum_{\underbrace{x^i\in X}_{\textit{agents' location}}} \sum_{\underbrace{v\in D^{i-}}_{\textit{disc of agent } i}} \underbrace{\rho(v)}_{\textit{density}}$$

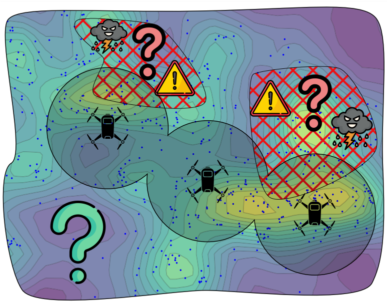

Goal: $\max_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t)\geq 0}_{\textit{constraint}}\ \forall t$

Challenges

- A-priori unknown density
- A-priori unknown constraints
- Even with known functions $\rightarrow$ NP Hard problem ☹

Submodular function

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \underbrace{\sum_{x^i \in X}}_{\textit{agents' location}} \underbrace{\sum_{v \in D^{i-}}}_{\textit{disc of agent } i} \underbrace{\rho(v)}_{\textit{density}}$$

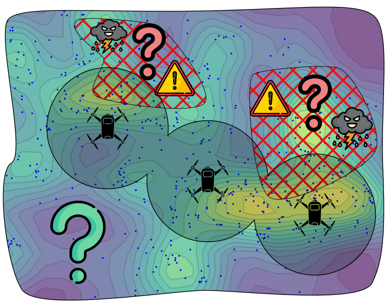

Goal: $\max_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t) \geq 0}_{\textit{constraint}} \ \forall t$

Challenges

- A-priori unknown density
- A-priori unknown constraints
- Even with known functions $\rightarrow$ NP Hard problem ☹

Submodular function $\rightarrow$ Greedy is Near-Optimal,

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \underbrace{\sum_{x^i \in X}}_{\textit{agents' location}} \underbrace{\sum_{v \in D^{i-}}}_{\textit{disc of agent } i} \underbrace{\rho(v)}_{\textit{density}}$$

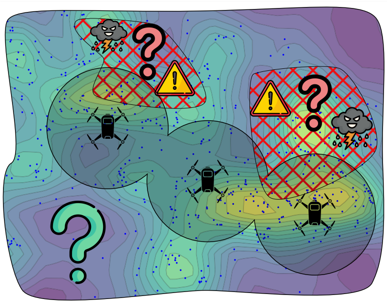

Goal: $\max_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t) \geq 0}_{\textit{constraint}} \ \forall t$

Challenges

- A-priori unknown density
- A-priori unknown constraints
- Even with known functions $\rightarrow$ NP Hard problem ☹

Submodular function $\rightarrow$ Greedy is Near-Optimal,

$$\underbrace{F(X_t;\rho)}_{\textit{Our goal}} \geq (1 - \frac{1}{e}) \underbrace{F(X_\star;\rho)}_{\textit{optimal clairvoyant}} - \epsilon_\rho$$

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\text{coverage}} = \underbrace{\sum_{x^i \in X}}_{\text{agents' location}} \underbrace{\sum_{v \in D^{i-}}}_{\text{disc of agent } i} \underbrace{\rho(v)}_{\text{density}}$$

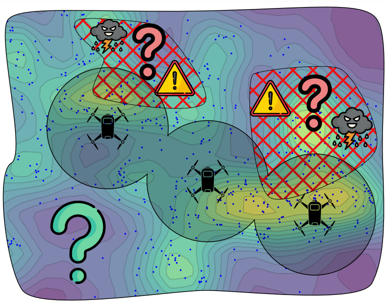

Goal: $\max_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t) \geq 0}_{\text{constraint}}$ $\forall t$

Challenges

- A-priori unknown density
- A-priori unknown constraints
- Even with known functions $\rightarrow$ NP Hard problem ☹

Submodular function $\rightarrow$ Greedy is Near-Optimal,

$$\underbrace{F(X_t;\rho)}_{\text{Our goal}} \geq \left(1-\frac{1}{e}\right) \underbrace{F(X_\star;\rho)}_{\text{optimal clairvoyant}} - \epsilon_\rho$$

Algorithmic questions

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \underbrace{\sum_{x^i \in X}}_{\textit{agents' location}} \underbrace{\sum_{v \in D^{i-}}}_{\textit{disc of agent i}} \underbrace{\rho(v)}_{\textit{density}}$$

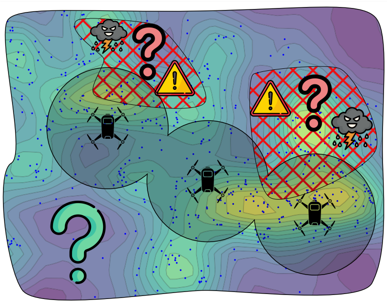

Goal: $\max_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t) \geq 0}_{\textit{constraint}} \ \forall t$

Challenges

- A-priori unknown density
- A-priori unknown constraints
- Even with known functions → NP Hard problem ☹

Submodular function → Greedy is Near-Optimal,

$$\underbrace{F(X_t;\rho)}_{\textit{Our goal}} \geq (1-\frac{1}{e}) \underbrace{F(X_\star;\rho)}_{\textit{optimal clairvoyant}} - \epsilon_\rho$$

Algorithmic questions

- Can we always satisfy safety constraints?

ETH zürich ETH AI CENTER

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \underbrace{\sum_{x^i \in X}}_{\textit{agents' location}} \underbrace{\sum_{v \in D^{i-}}}_{\textit{disc of agent i}} \underbrace{\rho(v)}_{\textit{density}}$$

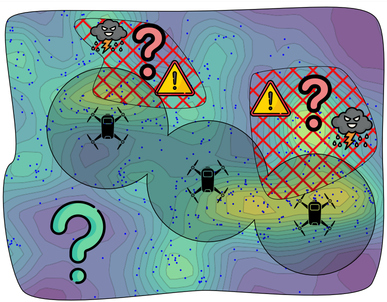

Goal: $\max\limits_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t) \geq 0}_{\textit{constraint}} \; \forall t$

Challenges

- A-priori unknown density
- A-priori unknown constraints
- Even with known functions $\rightarrow$ NP Hard problem ☹

Submodular function $\rightarrow$ Greedy is Near-Optimal,

$$\underbrace{F(X_t;\rho)}_{\textit{Our goal}} \geq (1 - \frac{1}{e}) \underbrace{F(X_\star;\rho)}_{\textit{optimal clairvoyant}} - \epsilon_\rho$$

Algorithmic questions

- Can we always satisfy safety constraints?
- Do we converge? How quickly?

Problem definition

Coverage function:

$$\underbrace{F(X;\rho)}_{\textit{coverage}} = \underbrace{\sum_{x^i \in X}}_{\textit{agents' location}} \underbrace{\sum_{v \in D^{i-}}}_{\textit{disc of agent i}} \underbrace{\rho(v)}_{\textit{density}}$$

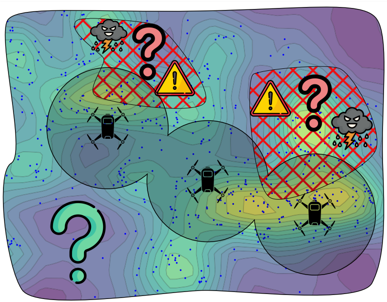

Goal: $\max_{X_t} F(X_t;\rho)$ *s.t.* $\underbrace{q(X_t) \geq 0}_{\textit{constraint}}$ $\forall t$

Challenges

- A-priori unknown density
- A-priori unknown constraints
- Even with known functions $\rightarrow$ NP Hard problem ☹

Submodular function $\rightarrow$ Greedy is Near-Optimal,

$$\underbrace{F(X_t;\rho)}_{\textit{Our goal}} \geq (1-\frac{1}{e}) \underbrace{F(X_\star;\rho)}_{\textit{optimal clairvoyant}} - \epsilon_\rho$$

Algorithmic questions

- Can we always satisfy safety constraints?
- Do we converge? How quickly?
- How far are we from the optimal solution?

MaCOpt: Multi-agent Coverage Control

Unconstrained case

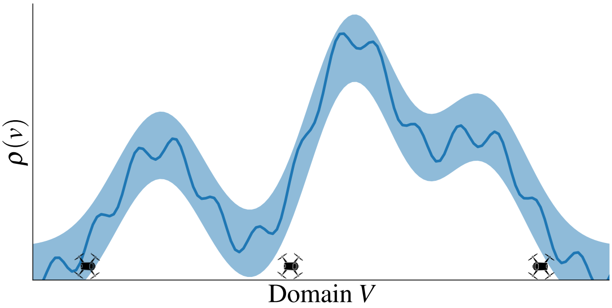

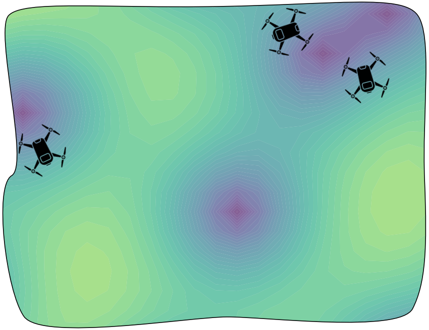

MaCOpt: Multi-agent Coverage Control
Unconstrained case

MaCOpt steps: 1) Greedy UCB

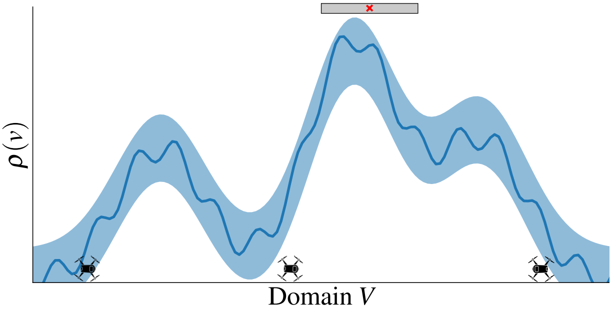

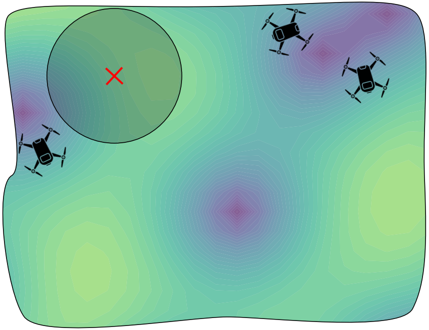

MaCOpt: Multi-agent Coverage Control
Unconstrained case

MaCOpt steps: 1) Greedy UCB

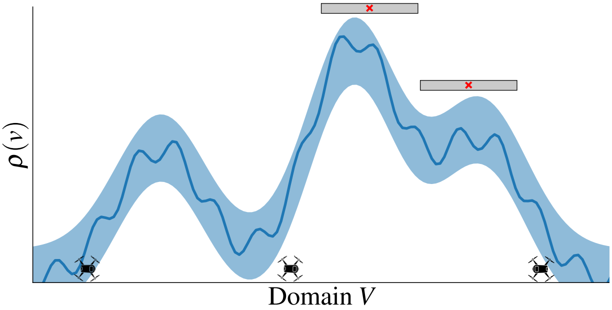

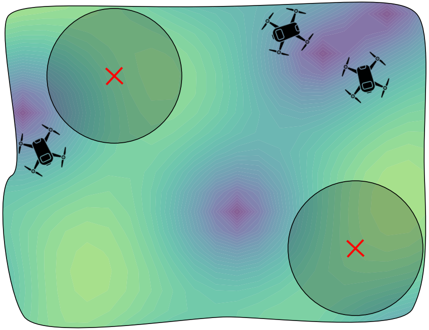

MaCOpt: Multi-agent Coverage Control
Unconstrained case

MaCOpt steps: 1) Greedy UCB

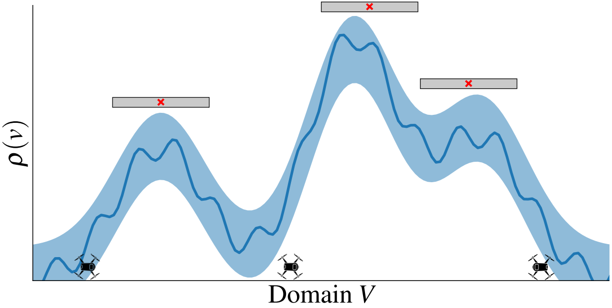

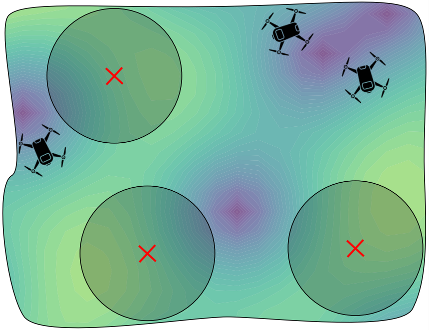

MaCOpt: Multi-agent Coverage Control

Unconstrained case

MaCOpt steps: 1) Greedy UCB 2) Uncertainty sampling

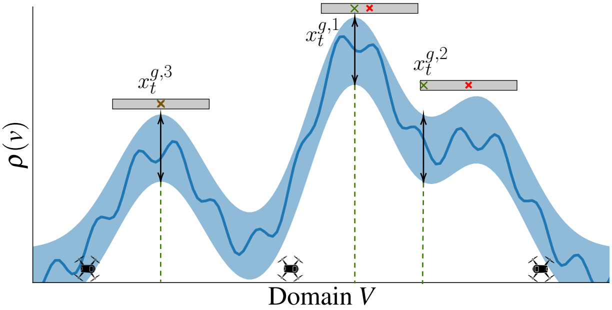

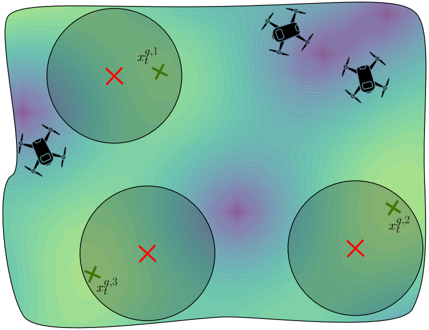

MaCOpt: Multi-agent Coverage Control

Unconstrained case

MaCOpt steps: 1) Greedy UCB 2) Uncertainty sampling

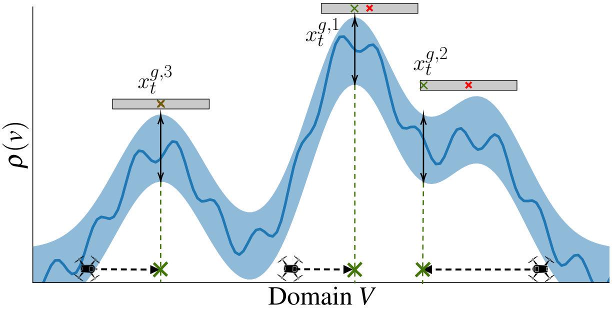

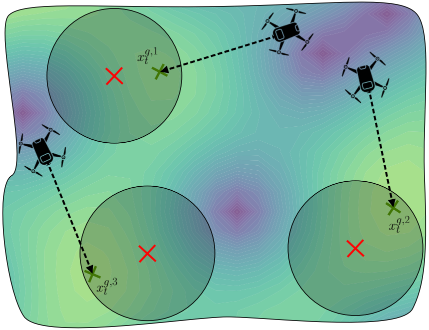

MaCOpt: Multi-agent Coverage Control

Unconstrained case

MaCOpt steps: 1) Greedy UCB 2) Uncertainty sampling

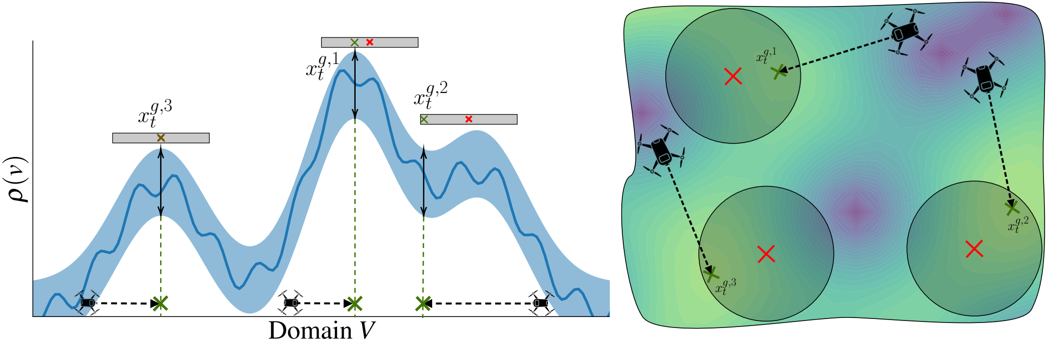

Theoretical result: **Cumulative regret grows sublinear with time**

SafeMaC: Safe Multi-Agent Coverage Control

Multi-Agent extension of Goal-oriented Safe Exploration (GoOSE) (Turchetta et al., 2019)

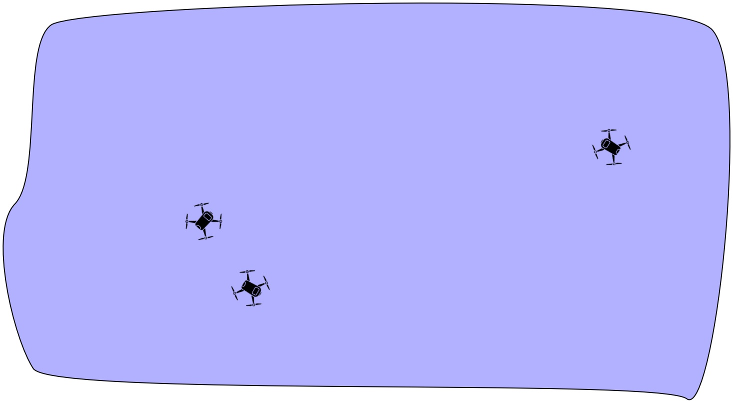

SafeMaC: Safe Multi-Agent Coverage Control

Multi-Agent extension of Goal-oriented Safe Exploration (GoOSE) (Turchetta et al., 2019)

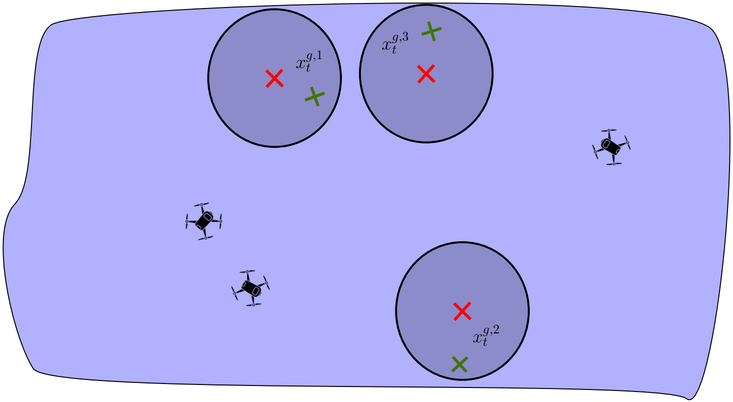

SafeMaC: Safe Multi-Agent Coverage Control

Multi-Agent extension of Goal-oriented Safe Exploration (GoOSE) (Turchetta et al., 2019)

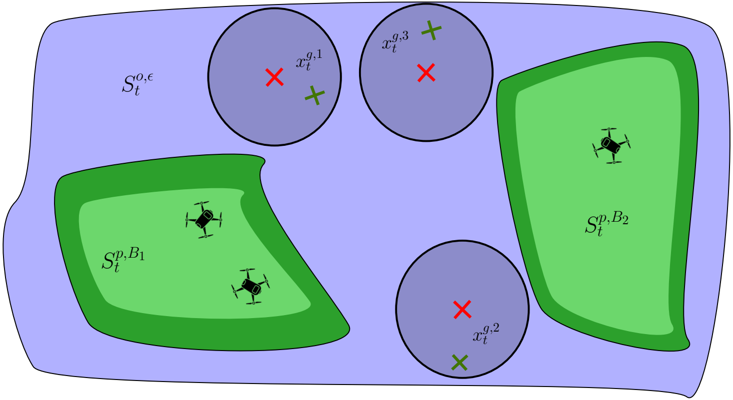

SafeMaC: Safe Multi-Agent Coverage Control

Multi-Agent extension of Goal-oriented Safe Exploration (GoOSE) (Turchetta et al., 2019)

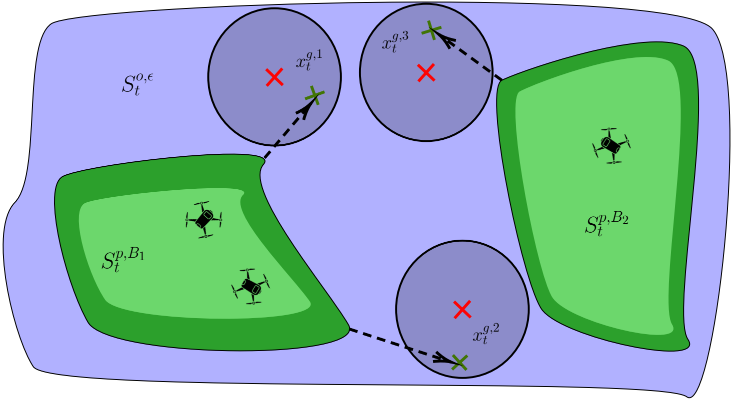

SafeMaC: Safe Multi-Agent Coverage Control

Multi-Agent extension of Goal-oriented Safe Exploration (GoOSE) (Turchetta et al., 2019)

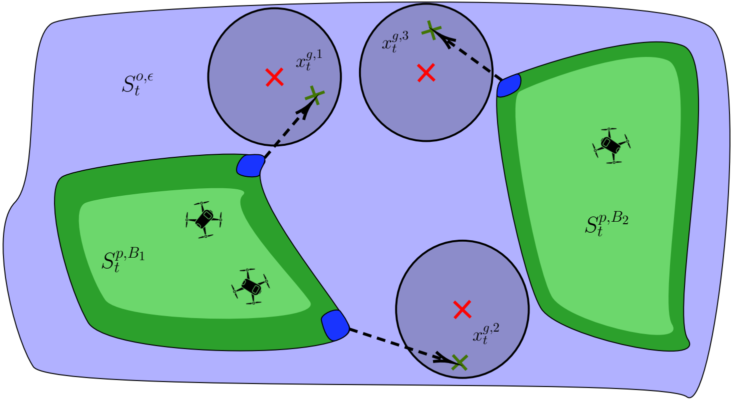

SafeMaC: Safe Multi-Agent Coverage Control

Multi-Agent extension of Goal-oriented Safe Exploration (GoOSE) (Turchetta et al., 2019)

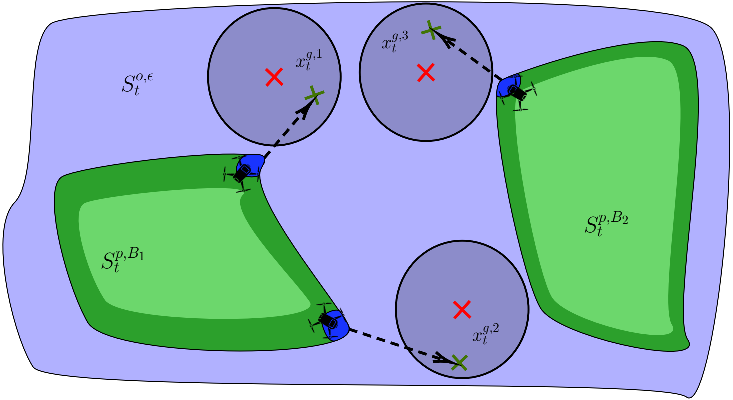

SafeMaC: Safe Multi-Agent Coverage Control

Multi-Agent extension of Goal-oriented Safe Exploration (GoOSE) (Turchetta et al., 2019)

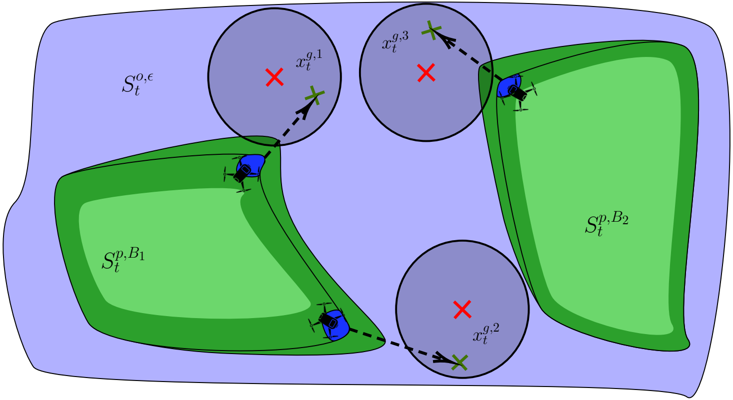

SafeMaC: Safe Multi-Agent Coverage Control

Multi-Agent extension of Goal-oriented Safe Exploration (GoOSE) (Turchetta et al., 2019)

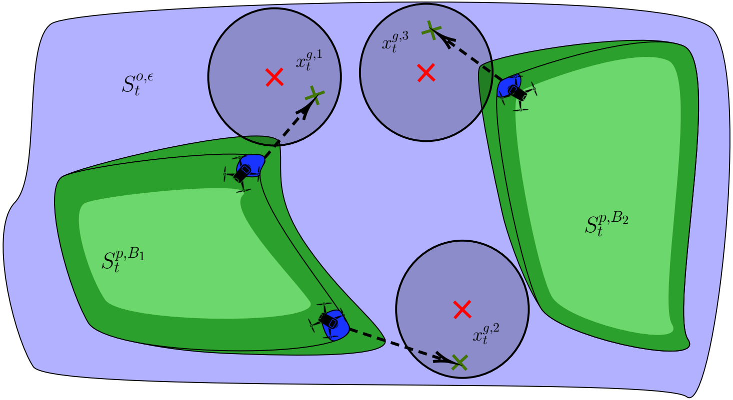

Theoretical Results:

- **Guarantees that SafeMaC is safe with high probability**
- **Achieves near-optimal coverage in finite time**

Conclusion

Experiments on biodiversity monitoring and obstacle avoidance environments

Conclusion

Experiments on biodiversity monitoring and obstacle avoidance environments

- MaCOpt **up to 40%** more coverage as compared to UCB
- SafeMaC **up to 50%** more sample efficient as compared to two-stage algorithm

Conclusion

Experiments on biodiversity monitoring and obstacle avoidance environments

- MaCOpt **up to 40%** more coverage as compared to UCB
- SafeMaC **up to 50%** more sample efficient as compared to two-stage algorithm

If you are interested in:

- Multi-agent learning
- Submodular optimization
- Safety
- Bayesian optimization

See you at our poster @NeurIPS 2022 !!!

Scan for paper !!!

Conclusion

Experiments on biodiversity monitoring and obstacle avoidance environments

- MaCOpt **up to 40%** more coverage as compared to UCB
- SafeMaC **up to 50%** more sample efficient as compared to two-stage algorithm

If you are interested in:

- Multi-agent learning
- Submodular optimization
- Safety
- Bayesian optimization

See you at our poster @NeurIPS 2022 !!!

Thank you for your attention !!!

Scan for paper !!!