ICQ: A Quantization Scheme for Best-Arm Identification Over Bit-Constrained Channels

Fathima Z. Faizal, A.G., Manjesh K. Hanawal, Nikhil Karamchandani

Presented by: Adway Girish
Information Theory Laboratory

EPFL

EPFL Information Processing Group

August 25, 2023

Team

Fathima Zarin Faizal,
EECS, MIT

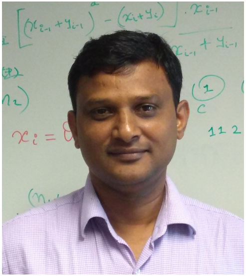

Prof. Manjesh Hanawal,
IEOR, IIT Bombay

Prof. Nikhil Karamchandani,
EE, IIT Bombay

Qualcomm Innovation Fellowship India 2022 Winners

Outline

1. Classical Best-Arm Identification

Outline

Outline

1. Classical Best-Arm Identification
2. A Distributed Variant
3. Proposed Solution ICQ

Outline

Outline

Best-arm identification

Best-arm identification

- K arms, each with a reward distribution

Best-arm identification

- K arms, each with a reward distribution (bounded*)

Best-arm identification

- K arms, each with a reward distribution (bounded*)

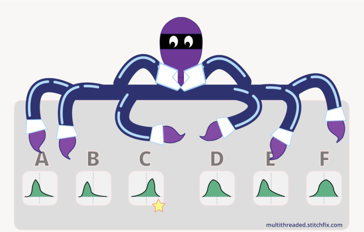

[https://multithreaded.stitchfix.com/blog/2020/08/05/bandits/]

Best-arm identification

- K arms, each with a reward distribution (bounded*)

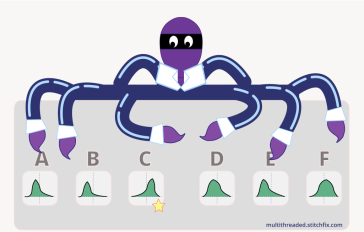

[https://multithreaded.stitchfix.com/blog/2020/08/05/bandits/]

- Objective: find best arm

Best-arm identification

- K arms, each with a reward distribution (bounded*)

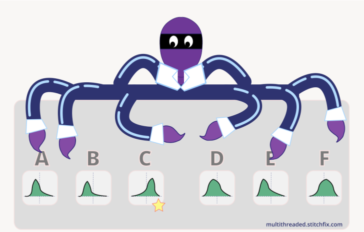

[https://multithreaded.stitchfix.com/blog/2020/08/05/bandits/]

- Objective: find best arm — highest mean reward

Best-arm identification: Fixed confidence

- K arms, each with a reward distribution (bounded*)

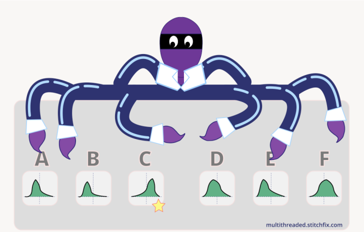

[https://multithreaded.stitchfix.com/blog/2020/08/05/bandits/]

- Objective: find best arm — highest mean reward
- Fix confidence level $\delta \in (0, 1)$

Best-arm identification: Fixed confidence

- K arms, each with a reward distribution (bounded*)

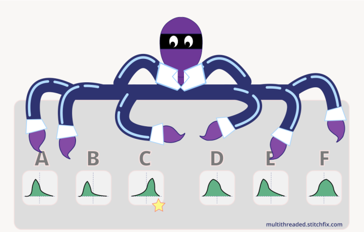

[https://multithreaded.stitchfix.com/blog/2020/08/05/bandits/]

- Objective: find best arm — highest mean reward
- Fix confidence level $\delta \in (0, 1)$; want

$$\mathbb{P}(\tau_\delta < \infty, J_{\tau_\delta} \neq J^*) < \delta$$

Best-arm identification: Fixed confidence

- K arms, each with a reward distribution (bounded*)

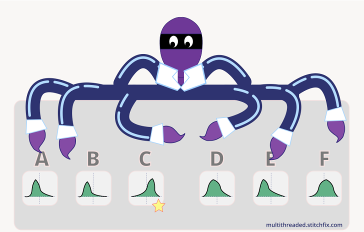

[https://multithreaded.stitchfix.com/blog/2020/08/05/bandits/]

- Objective: find best arm — highest mean reward
- Fix confidence level $\delta \in (0, 1)$; want

$$\mathbb{P}(\tau_\delta < \infty, J_{\tau_\delta} \neq J^*) < \delta, \quad \text{“success w.h.p.”}$$

Successive Elimination

Round 0, $S = \{1, 2, 3\}$

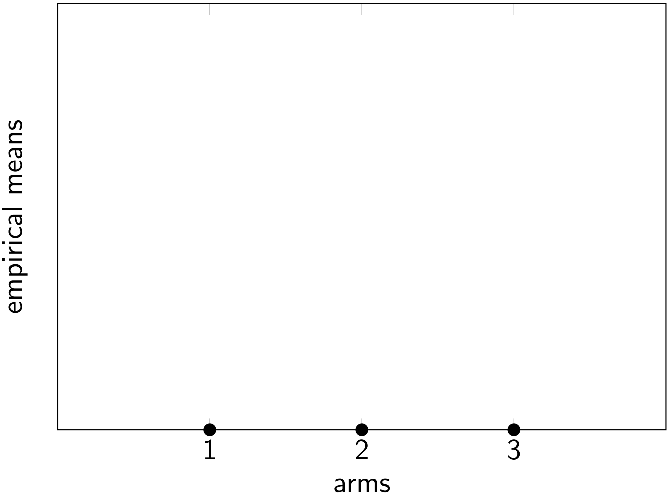

Successive Elimination

Round 1, $S = \{1, 2, 3\}$

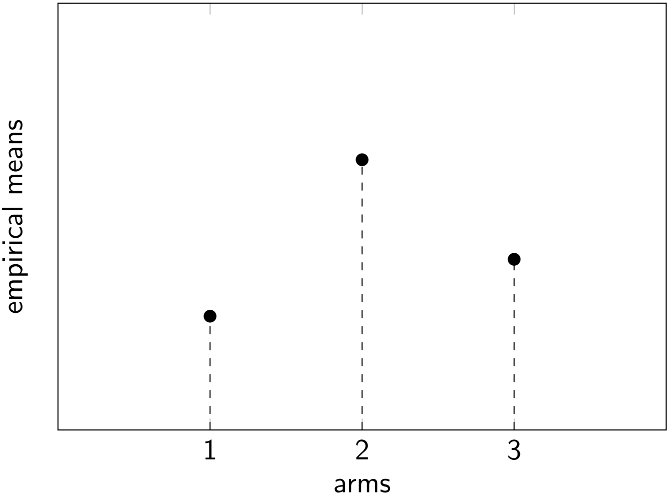

Successive Elimination

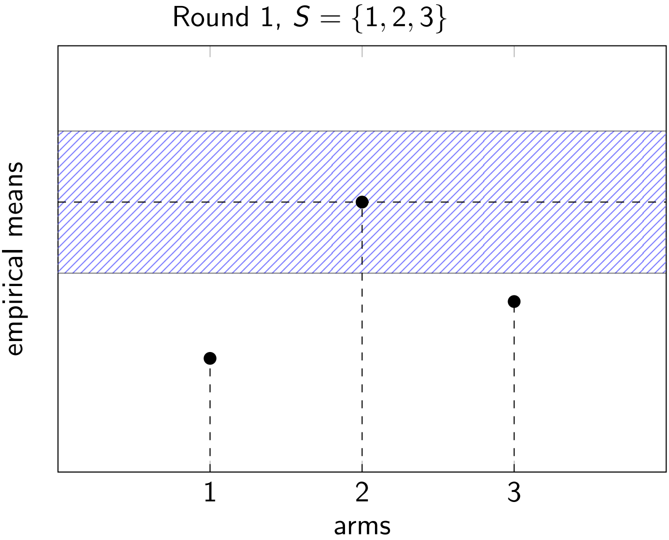

Successive Elimination

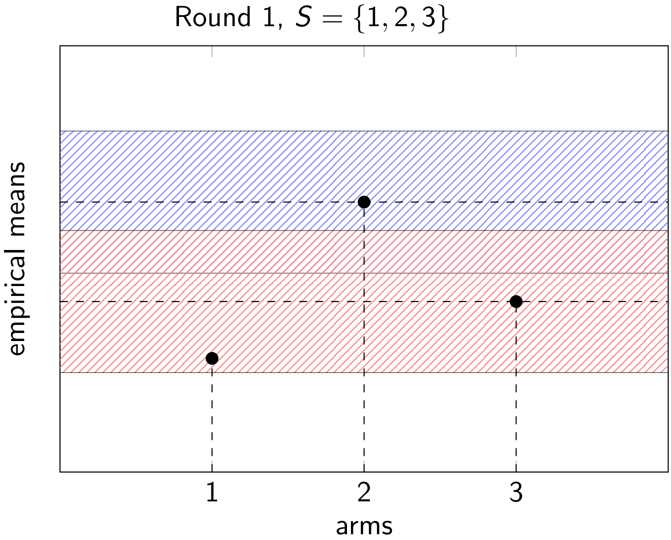

Successive Elimination

Round 1, $S = \{1, 2, 3\}$

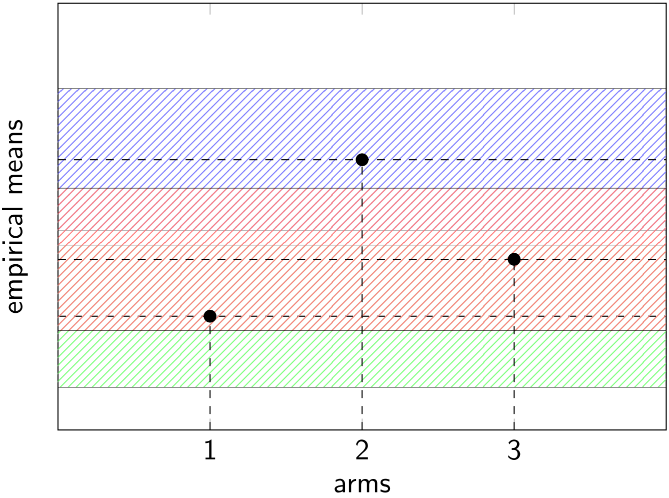

Successive Elimination

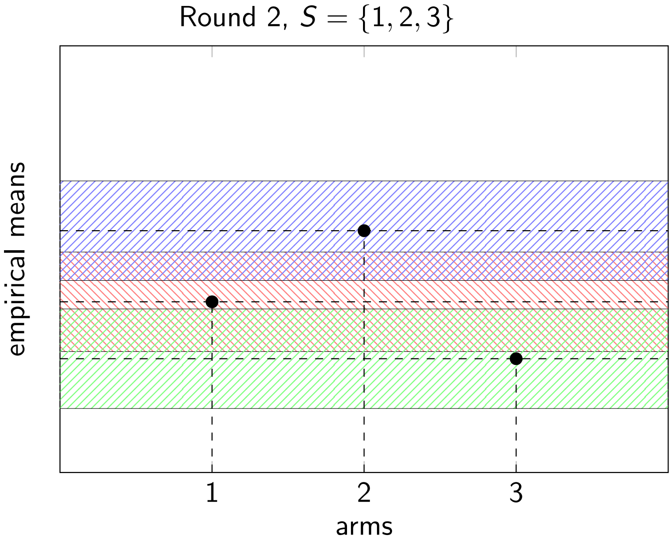

Successive Elimination

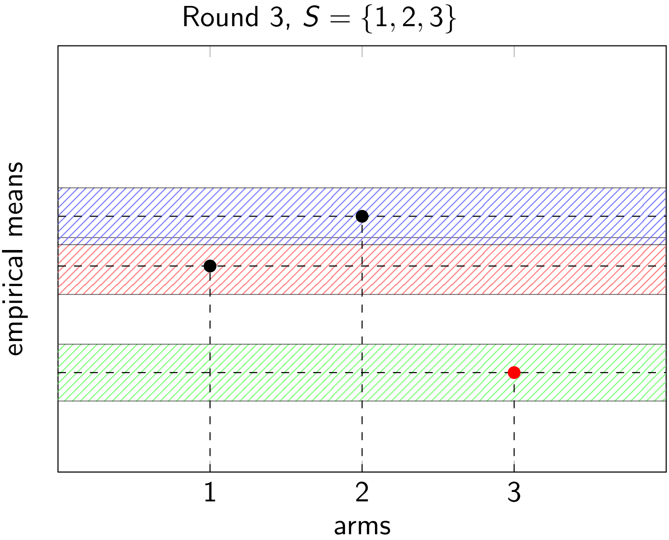

Successive Elimination

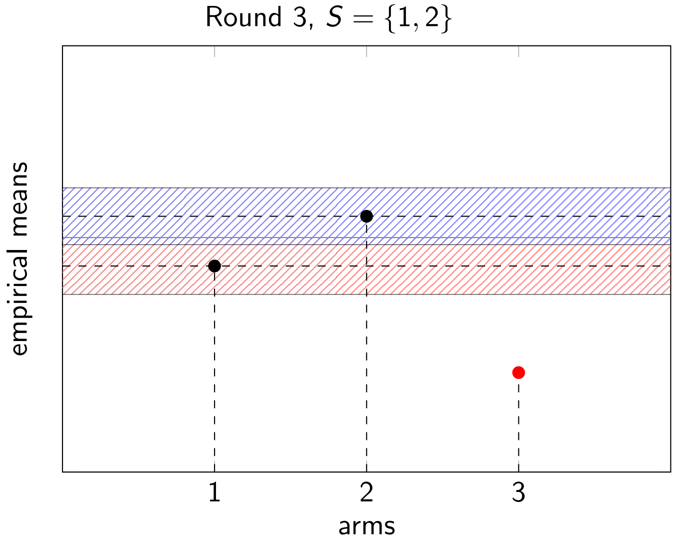

Successive Elimination

Round 4, $S = \{1, 2\}$

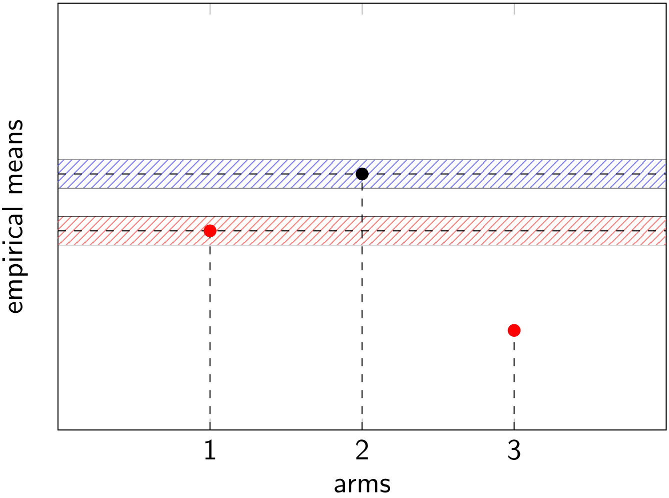

Successive Elimination

Round 4, $S = \{2\}$

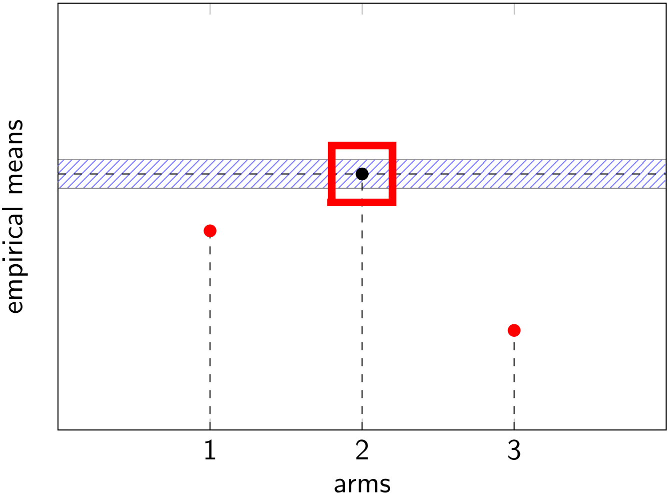

Successive Elimination

- Confidence width $U'(i, \delta) = \sigma \sqrt{\frac{2 \log(4Ki^2/\delta)}{i}}$

Successive Elimination

- Confidence width $U'(i, \delta) = \sigma \sqrt{\frac{2 \log(4Ki^2/\delta)}{i}}$
- Gives $\bigcup_{i=1}^{\infty} \bigcup_{j=1}^{K} \mathbb{P}\Big(|\hat{\mu}_{j,i} - \mu_j| \geq U'(i, \delta) \Big) < \delta$

Successive Elimination

- Confidence width $U'(i, \delta) = \sigma \sqrt{\frac{2 \log(4Ki^2/\delta)}{i}}$
- Gives $\bigcup_{i=1}^{\infty} \bigcup_{j=1}^{K} \mathbb{P}\left(|\hat{\mu}_{j,i} - \mu_j| \geq U'(i, \delta) \right) < \delta$
- Success w.h.p.!

A distributed MAB setup

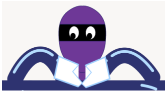

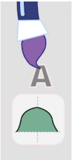

. . .

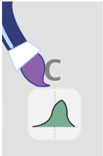

. . .

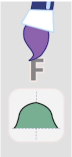

A distributed MAB setup

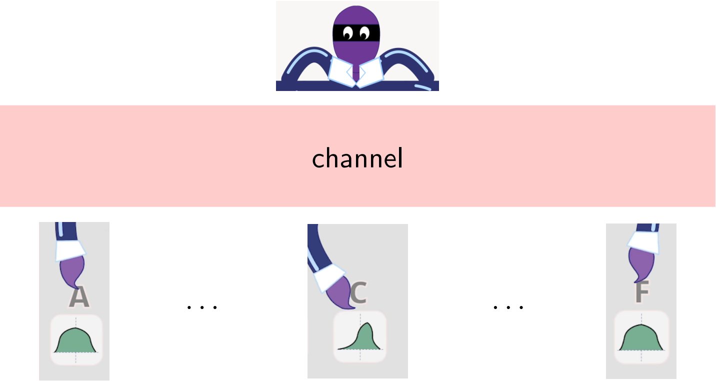

A distributed MAB setup

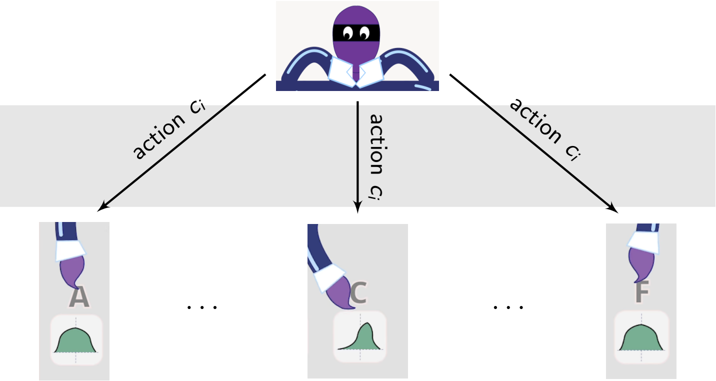

A distributed MAB setup

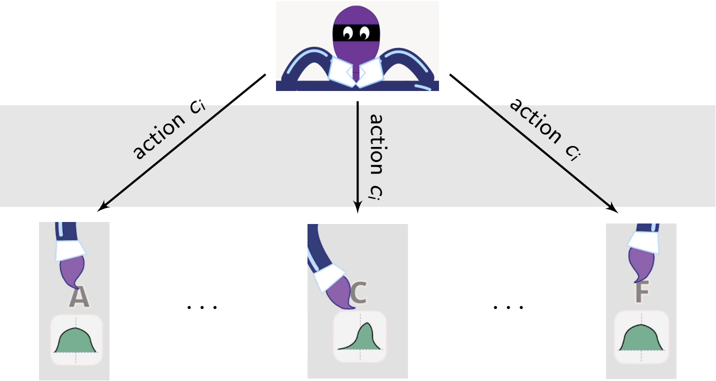

pull arm, encode into $\mathbf{s}_{A,i}$

inactive in comm. round i

pull arm, encode into $\mathbf{s}_{F,i}$

A distributed MAB setup

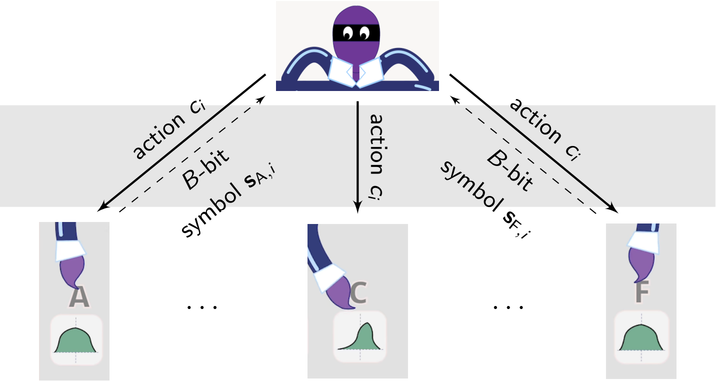

pull arm, encode into $\mathbf{s}_{A,i}$

inactive in comm. round i

pull arm, encode into $\mathbf{s}_{F,i}$

A distributed MAB setup

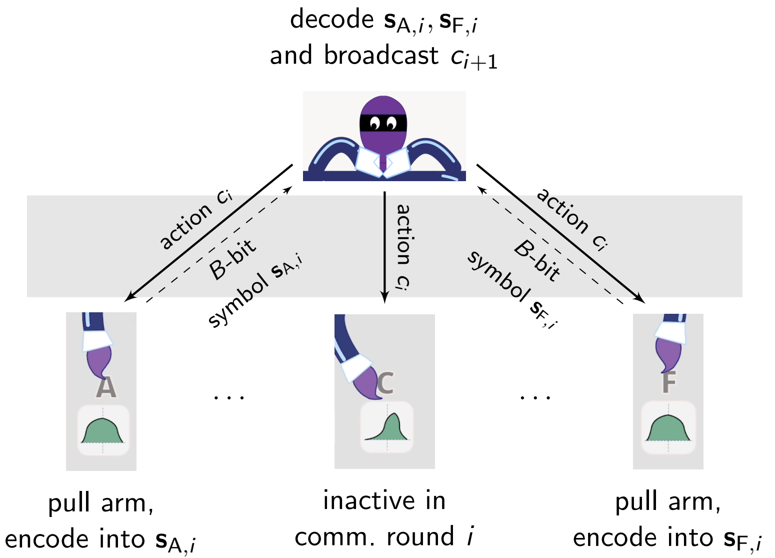

Prior Art

Prior Art

- Conventional MAB algos assume full-precision rewards*

*Lattimore and Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020

Prior Art

- Conventional MAB algos assume full-precision rewards*
- Adaptive quantization schemes with order-optimal regret†‡

*Lattimore and Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020

†Hanna, Yang, and Fragouli. "Solving multi-arm bandit using a few bits of communication". In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022

‡Mitra, Hassani, and Pappas. *Linear Stochastic Bandits over a Bit-Constrained Channel*. 2022. arXiv: 2203.01198 [cs.LG]

Prior Art

- Conventional MAB algos assume full-precision rewards*
- Adaptive quantization schemes with order-optimal regret†‡
- Pure exploration over bit-constrained channels§

*Lattimore and Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020

†Hanna, Yang, and Fragouli. "Solving multi-arm bandit using a few bits of communication". In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022

‡Mitra, Hassani, and Pappas. *Linear Stochastic Bandits over a Bit-Constrained Channel*. 2022. arXiv: 2203.01198 [cs.LG]

§Mitra, Hassani, and Pappas. *Exploiting Heterogeneity in Robust Federated Best-Arm Identification*. 2021. arXiv: 2109.05700 [cs.LG]

Prior Art

- Conventional MAB algos assume full-precision rewards[*]
- Adaptive quantization schemes with order-optimal regret[†‡]
- Pure exploration over bit-constrained channels[§] but no control over B

[*] Lattimore and Szepesvári. *Bandit algorithms.* Cambridge University Press, 2020

[†] Hanna, Yang, and Fragouli. "Solving multi-arm bandit using a few bits of communication". In: *International Conference on Artificial Intelligence and Statistics.* PMLR, 2022

[‡] Mitra, Hassani, and Pappas. *Linear Stochastic Bandits over a Bit-Constrained Channel.* 2022. arXiv: 2203.01198 [cs.LG]

[§] Mitra, Hassani, and Pappas. *Exploiting Heterogeneity in Robust Federated Best-Arm Identification.* 2021. arXiv: 2109.05700 [cs.LG]

Prior Art

- Conventional MAB algos assume full-precision rewards*
- Adaptive quantization schemes with order-optimal regret†‡
- Pure exploration over bit-constrained channels§ but no control over B
- Our work:

*Lattimore and Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020

†Hanna, Yang, and Fragouli. "Solving multi-arm bandit using a few bits of communication". In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022

‡Mitra, Hassani, and Pappas. *Linear Stochastic Bandits over a Bit-Constrained Channel*. 2022. arXiv: 2203.01198 [cs.LG]

§Mitra, Hassani, and Pappas. *Exploiting Heterogeneity in Robust Federated Best-Arm Identification*. 2021. arXiv: 2109.05700 [cs.LG]

Prior Art

- Conventional MAB algos assume full-precision rewards*
- Adaptive quantization schemes with order-optimal regret†‡
- Pure exploration over bit-constrained channels§ but no control over B
- Our work: order-optimal sample complexity, clear dependence on B

*Lattimore and Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020

†Hanna, Yang, and Fragouli. "Solving multi-arm bandit using a few bits of communication". In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022

‡Mitra, Hassani, and Pappas. *Linear Stochastic Bandits over a Bit-Constrained Channel*. 2022. arXiv: 2203.01198 [cs.LG]

§Mitra, Hassani, and Pappas. *Exploiting Heterogeneity in Robust Federated Best-Arm Identification*. 2021. arXiv: 2109.05700 [cs.LG]

New challenges

New challenges

- Communication complexity

New challenges

- Communication complexity $\rightarrow$ few bits + exponentially sparse

New challenges

- Communication complexity $\rightarrow$ few bits + exponentially sparse
- Quantization error

New challenges

- Communication complexity $\rightarrow$ few bits + exponentially sparse
- Quantization error

$$|\tilde{\mu}_{j,i} - \mu_j|$$

New challenges

- Communication complexity $\rightarrow$ few bits + exponentially sparse
- Quantization error

$$\begin{aligned} |\tilde{\mu}_{j,i} - \mu_j| &\leq |\tilde{\mu}_{j,i} - \hat{\mu}_{j,i}| + |\hat{\mu}_{j,i} - \mu_j| \\ &= \text{quantization error} + \text{error in empirical mean} \end{aligned}$$

New challenges

- Communication complexity $\rightarrow$ few bits + exponentially sparse
- Quantization error

$$
\begin{aligned}
|\tilde{\mu}_{j,i} - \mu_j| &\leq |\tilde{\mu}_{j,i} - \hat{\mu}_{j,i}| + |\hat{\mu}_{j,i} - \mu_j| \\
&= \text{ quantization error } + \text{ error in empirical mean}
\end{aligned}
$$

- We have $\bigcup_{i=1}^{\infty} \bigcup_{j=1}^{K} \mathbb{P}\left(|\hat{\mu}_{j,i} - \mu_j| \geq U'(i, \delta) \right) < \delta$

New challenges

- Communication complexity $\rightarrow$ few bits + exponentially sparse
- Quantization error $\rightarrow$ new confidence widths

$$\begin{aligned} |\tilde{\mu}_{j,i} - \mu_j| &\leq |\tilde{\mu}_{j,i} - \hat{\mu}_{j,i}| + |\hat{\mu}_{j,i} - \mu_j| \\ &= \text{ quantization error } + \text{ error in empirical mean} \end{aligned}$$

- We have $\bigcup_{i=1}^{\infty} \bigcup_{j=1}^{K} \mathbb{P}\left(|\hat{\mu}_{j,i} - \mu_j| \geq U'(i,\delta) \right) < \delta$
- We want $\bigcup_{i=1}^{\infty} \bigcup_{j=1}^{K} \mathbb{P}\left(|\tilde{\mu}_{j,i} - \mu_j| \geq U(i,\delta) \right) < \delta$

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using B bits

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

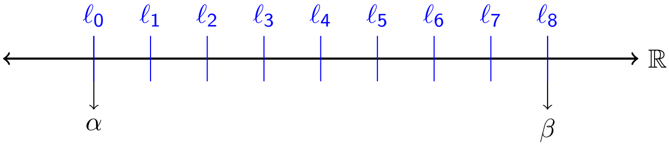

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

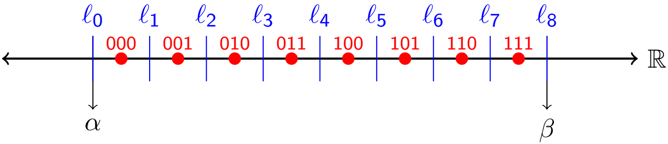

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

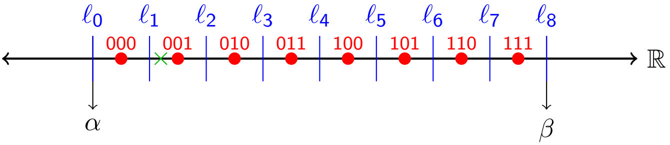

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

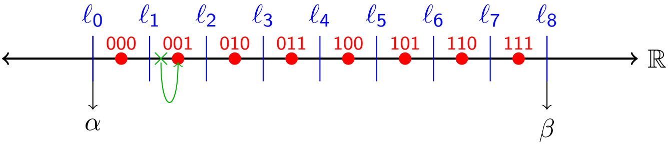

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

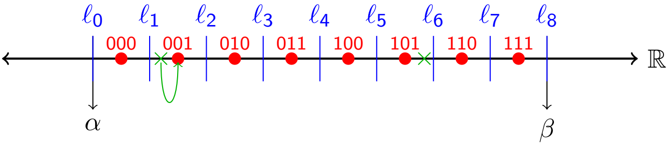

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

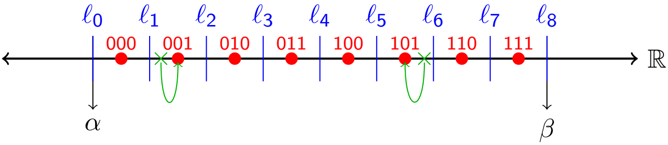

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

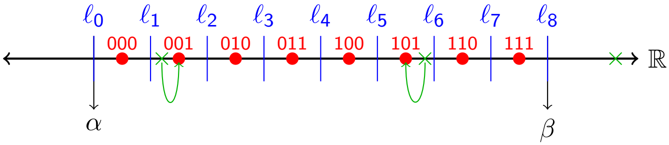

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

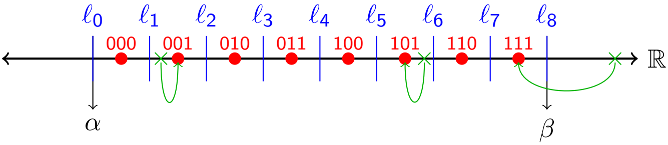

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

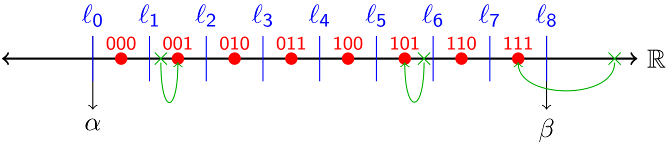

- Quantization error for points in $[\alpha, \beta] \leq$

ICQ: Quantization

- Quantize $[\alpha, \beta]$ using $B = 3$ bits

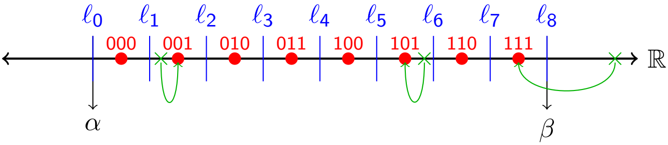

- Quantization error for points in $[\alpha, \beta] \leq \dfrac{\beta - \alpha}{2 \cdot 2^B}$

ICQ: Confidence intervals

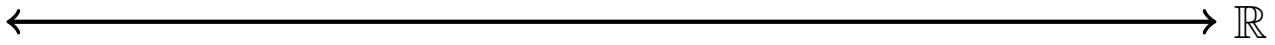

ICQ: Confidence intervals

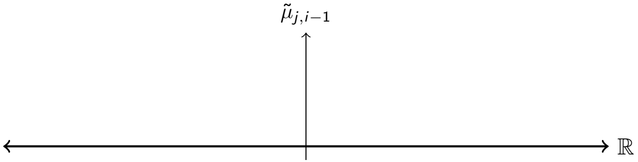

ICQ: Confidence intervals

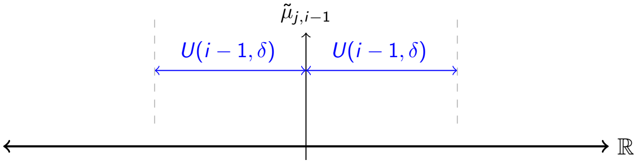

ICQ: Confidence intervals

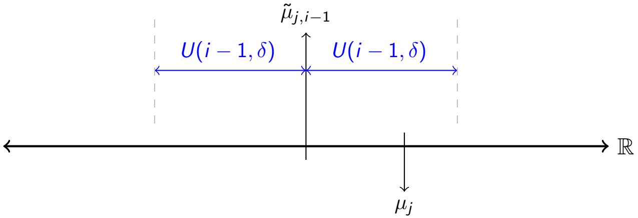

ICQ: Confidence intervals

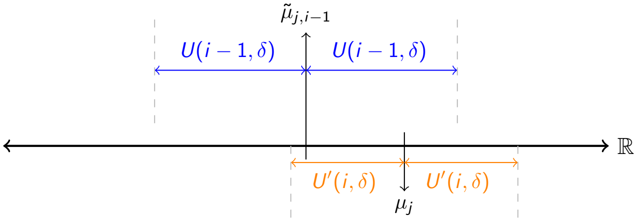

ICQ: Confidence intervals

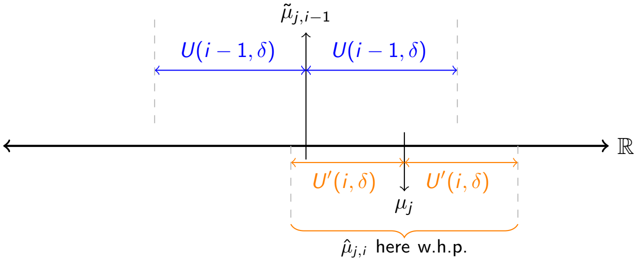

ICQ: Confidence intervals

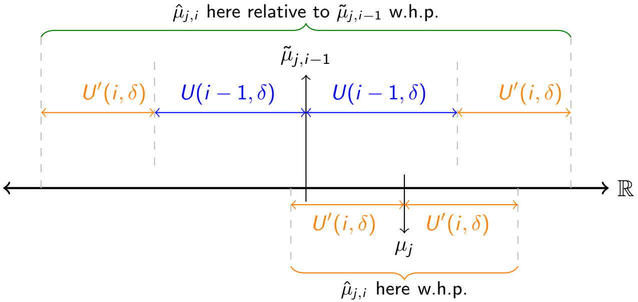

ICQ: Confidence intervals

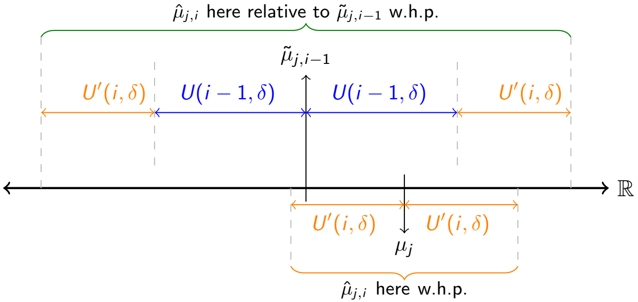

- $$|\tilde{\mu}_{j,i} - \mu_j| \le |\tilde{\mu}_{j,i} - \hat{\mu}_{j,i}| + |\hat{\mu}_{j,i} - \mu_j|$$
  $$=: \text{quantization error} + \text{error in empirical mean}$$

ICQ: Confidence intervals

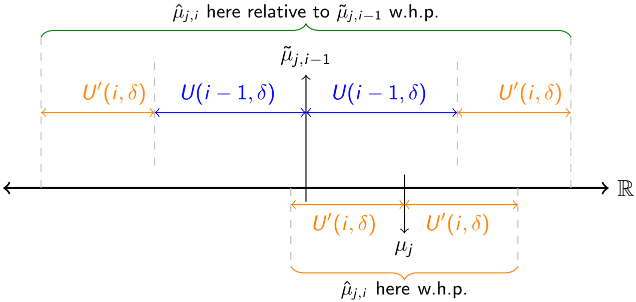

- $$|\tilde{\mu}_{j,i}-\mu_j| \le |\tilde{\mu}_{j,i}-\hat{\mu}_{j,i}| + |\hat{\mu}_{j,i}-\mu_j|$$ $$=: \text{ quantization error } + \text{ error in empirical mean}$$
- Hence, define $U(i,\delta) = \qquad\qquad + $

ICQ: Confidence intervals

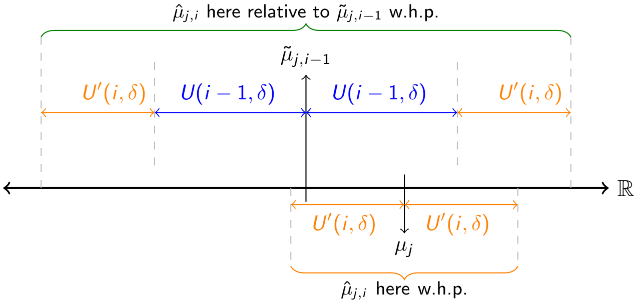

- $$|\tilde{\mu}_{j,i} - \mu_j| \leq |\tilde{\mu}_{j,i} - \hat{\mu}_{j,i}| + |\hat{\mu}_{j,i} - \mu_j|$$
  $$=: \text{quantization error} + \text{error in empirical mean}$$
- Hence, define $U(i,\delta) = \qquad + U'(i,\delta)$

ICQ: Confidence intervals

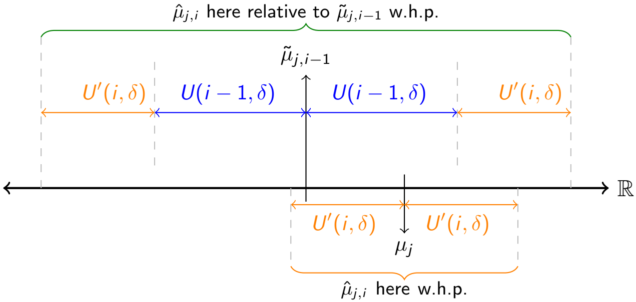

- $$|\tilde{\mu}_{j,i} - \mu_j| \le |\tilde{\mu}_{j,i} - \hat{\mu}_{j,i}| + |\hat{\mu}_{j,i} - \mu_j|$$
  $$=: \text{quantization error} + \text{error in empirical mean}$$
- Hence, define $U(i,\delta) = \frac{1}{2^B}\left[U'(i,\delta) + U(i-1,\delta)\right] + U'(i,\delta)$

Theoretical results

Theoretical results

- ICQ-SE "succeeds w.h.p"

Theoretical results

- ICQ-SE "succeeds w.h.p"
- With probability $\geq 1 - \delta$, the sample complexity

$$\tau_\delta^{\text{ICQ-SE}} \leq \mathcal{O}\left(1 + \frac{1}{2^B}\right) \tau_\delta^{\text{SE}}$$

Theoretical results

- ICQ-SE "succeeds w.h.p"
- With probability $\geq 1 - \delta$, the sample complexity

$$\tau_\delta^{\text{ICQ-SE}} \leq \mathcal{O}\left(1 + \frac{1}{2^B}\right)\tau_\delta^{\text{SE}}$$

- Only constant factor overhead!

Numerical experiments

$K = 5$ arms; means from Beta$(\gamma, 1-\gamma)$ distribution, $\gamma \sim$ Unif$([0,1])$

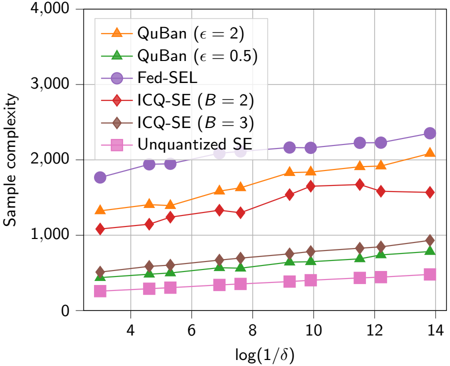

QuBan: Hanna, Yang, and Fragouli. "Solving multi-arm bandit using a few bits of communication". In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022

Fed-SEL: Mitra, Hassani, and Pappas. *Exploiting Heterogeneity in Robust Federated Best-Arm Identification*. 2021. arXiv: 2109.05700 [cs.LG]

Numerical experiments

$K = 5$ arms; means from Beta$(\gamma, 1-\gamma)$ distribution, $\gamma \sim$ Unif$([0,1])$

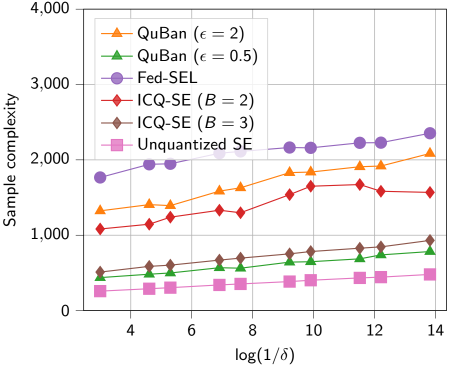

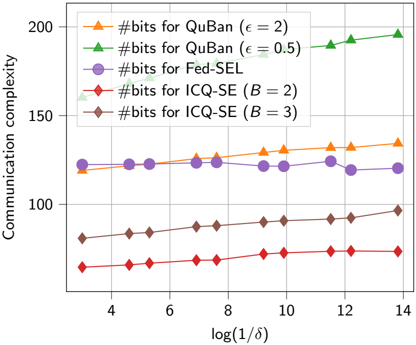

QuBan: Hanna, Yang, and Fragouli. "Solving multi-arm bandit using a few bits of communication". In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022

Fed-SEL: Mitra, Hassani, and Pappas. *Exploiting Heterogeneity in Robust Federated Best-Arm Identification*. 2021. arXiv: 2109.05700 [cs.LG]

Summary and future directions

- Proposed ICQ:

Summary and future directions

- Proposed ICQ:
    - Extends confidence-bound-based algorithms to bit-constrained settings

Summary and future directions

- Proposed ICQ:
  - Extends confidence-bound-based algorithms to bit-constrained settings
  - Updating confidence widths to account for quantization errors

Summary and future directions

- Proposed ICQ:
  - Extends confidence-bound-based algorithms to bit-constrained settings
  - Updating confidence widths to account for quantization errors
  - "Order-optimal" sample complexity

Summary and future directions

- Proposed ICQ:
  - Extends confidence-bound-based algorithms to bit-constrained settings
  - Updating confidence widths to account for quantization errors
  - "Order-optimal" sample complexity
- Future directions:

Summary and future directions

- Proposed ICQ:
  - Extends confidence-bound-based algorithms to bit-constrained settings
  - Updating confidence widths to account for quantization errors
  - "Order-optimal" sample complexity
- Future directions:
  - Lower bound

Summary and future directions

- Proposed ICQ:
  - Extends confidence-bound-based algorithms to bit-constrained settings
  - Updating confidence widths to account for quantization errors
  - “Order-optimal” sample complexity
- Future directions:
  - Lower bound
  - More distributed variants

Thank you!