A New Paradigm for Shortest Link Scheduling in Wireless Networks: Theory And Applications

F. Al-dhelaan, **P.-J. Wan**, and H.Q. Yuan

wan@cs.iit.edu

Outline

- Overview
- An Adaptive Zero-Sum Game with Retirement
- The General Paradigm
- Discussion

Model of Links

- m node-level communication links

Model of Links

- m node-level communication links
- E_l: communication primitives of the l-th link for $l \in [m]$

Model of Links

- m node-level communication links
- E_l: communication primitives of the l-th link for $l \in [m]$
  - radio-level links in MCMR wireless networks

Model of Links

- m node-level communication links
- E_l: communication primitives of the l-th link for $l \in [m]$
  - radio-level links in MCMR wireless networks
  - streams in MIMO wireless networks

Model of Links

- m node-level communication links
- E_l: communication primitives of the l-th link for $l \in [m]$
  - radio-level links in MCMR wireless networks
  - streams in MIMO wireless networks
- $E := E_1 \cup E_2 \cup \cdots \cup E_m$

Model of Links

- m node-level communication links
- E_l: communication primitives of the l-th link for $l \in [m]$
    - radio-level links in MCMR wireless networks
    - streams in MIMO wireless networks
- $E := E_1 \cup E_2 \cup \cdots \cup E_m$
- Data rates $b : E \to \mathbb{R}^+$

Model of Links

- m node-level communication links
- E_l: communication primitives of the l-th link for $l \in [m]$
  - radio-level links in MCMR wireless networks
  - streams in MIMO wireless networks
- $E := E_1 \cup E_2 \cup \cdots \cup E_m$
- Data rates $b : E \to \mathbb{R}^+$
- $\mathcal{I}$: collection of all independent subsets of E.

Shortest Link Scheduling

- Link demands $d : [m] \to \mathbb{R}^+$

Shortest Link Scheduling

- Link demands $d : [m] \to \mathbb{R}^+$
- A *link schedule* of d:

$$\Gamma = \{(I_j, x_j) \in \mathcal{I} \times \mathbb{R}^+ : j \in [k]\}$$

s.t.

$$d(l) \leq \sum_{j \in [k]} x_j \sum_{e \in E_l \cap I_j} b(e), \forall l \in [m];$$

$\|\Gamma\| := \sum_{j=1}^{k} x_j$: the *length* of $\mathcal{S}$

Shortest Link Scheduling

- Link demands $d : [m] \rightarrow \mathbb{R}^+$
- A *link schedule* of d:

$$\Gamma = \left\{ \left(I_j, x_j \right) \in \mathcal{I} \times \mathbb{R}^+ : j \in [k] \right\}$$

s.t.

$$d\left(l\right) \leq \sum_{j \in [k]} x_j \sum_{e \in E_l \cap I_j} b\left(e\right), \forall l \in [m];$$

$\|\Gamma\| := \sum_{j=1}^{k} x_j$: the *length* of $\mathcal{S}$

- $\chi^*\left(d\right) =$ minimum length of all link schedules of d

Shortest Link Scheduling

- Link demands $d : [m] \to \mathbb{R}^+$
- A *link schedule* of d:

$$\Gamma = \left\{ \left(I_j, x_j \right) \in \mathcal{I} \times \mathbb{R}^+ : j \in [k] \right\}$$

s.t.

$$d\left(l\right) \leq \sum_{j \in [k]} x_j \sum_{e \in E_l \cap I_j} b\left(e\right), \forall l \in [m];$$

$\|\Gamma\| := \sum_{j=1}^{k} x_j$: the *length* of $\mathcal{S}$

- $\chi^*\left(d\right) =$ minimum length of all link schedules of d
- **Shortest Link Scheduling (SLS)**: find a shortest link schedule of d

Prior Art

- Greedy method

Prior Art

- Greedy method
  - limited to protocol Interference model with uniform data rates

Prior Art

- Greedy method
  - limited to protocol Interference model with uniform data rates
- Ellipsoid method

Prior Art

- Greedy method
  - limited to protocol Interference model with uniform data rates
- Ellipsoid method
- Arbitrary wireless network

Prior Art

- Greedy method
    - limited to protocol Interference model with uniform data rates
- Ellipsoid method
- Arbitrary wireless network
    - Arbitrary wireless network

Prior Art

- Greedy method
  - limited to protocol Interference model with uniform data rates
- Ellipsoid method
- Arbitrary wireless network
  - Arbitrary wireless network
  - Approx.-preserving reduction to **MWIS**: Given $w : E \to \mathbb{R}_+$, find $I \in \mathcal{I}$ maximizing $w(I) := \sum_{e \in I} w(e)$

Prior Art

- Greedy method
    - limited to protocol Interference model with uniform data rates
- Ellipsoid method
- Arbitrary wireless network
    - Arbitrary wireless network
    - Approx.-preserving reduction to **MWIS**: Given $w : E \to \mathbb{R}_+$, find $I \in \mathcal{I}$ maximizing $w(I) := \sum_{e \in I} w(e)$
    - Impractical

The New Paradigm

- $\mathcal{A}$: μ-approximation algorithm for **MWIS**

The New Paradigm

- $\mathcal{A}$: μ-approximation algorithm for **MWIS**
- $\varepsilon \in (0, 1/2]$: trade-off between accuracy and efficiency

The New Paradigm

- $\mathcal{A}$: μ-approximation algorithm for **MWIS**
- $\varepsilon \in (0, 1/2]$: trade-off between accuracy and efficiency
- **LS**(ε): iterative $(1+\varepsilon)\,\mu$-approx. alg. for **SLS** by making $O\left(\varepsilon^{-2} m \ln m\right)$ calls to $\mathcal{A}$

The New Paradigm

- $\mathcal{A}$: μ-approximation algorithm for **MWIS**
- $\varepsilon \in (0, 1/2]$: trade-off between accuracy and efficiency
- **LS**(ε): iterative $(1+\varepsilon)\,\mu$-approx. alg. for **SLS** by making $O\left(\varepsilon^{-2} m \ln m\right)$ calls to $\mathcal{A}$
  - **Augmentation Stage**: Compute a LS Γ of ϕd for some $\phi > 0$ by successively selecting an approximate MWIS and its duration

The New Paradigm

- $\mathcal{A}$: μ-approximation algorithm for **MWIS**
- $\varepsilon \in (0, 1/2]$: trade-off between accuracy and efficiency
- **LS**(ε): iterative $(1+\varepsilon)\,\mu$-approx. alg. for **SLS** by making $O\left(\varepsilon^{-2} m \ln m\right)$ calls to $\mathcal{A}$
  - **Augmentation Stage**: Compute a LS Γ of ϕd for some $\phi > 0$ by successively selecting an approximate MWIS and its duration
  - **Scaling Stage**: return $\frac{1}{\phi}\Gamma$.

Technical Issues Yet to Be addressed

- How to choose the scaling factor ϕ?

Technical Issues Yet to Be addressed

- How to choose the scaling factor ϕ?
- What is an appropriate weight of E?

Technical Issues Yet to Be addressed

- How to choose the scaling factor ϕ?
- What is an appropriate weight of E?
- How much is the duration of an IS?

Augmentations as A Game

m links (agents) vs. scheduler (adversary)

Augmentations as A Game

m links (agents) vs. scheduler (adversary)

- Each iteration $\leftrightarrow$ a game round

Augmentations as A Game

m links (agents) vs. scheduler (adversary)

- Each iteration $\leftrightarrow$ a game round
- In each round,

Augmentations as A Game

m links (agents) vs. scheduler (adversary)

- Each iteration $\leftrightarrow$ a game round
- In each round,
  - a link earns a profit in service,

Augmentations as A Game

m links (agents) vs. scheduler (adversary)

- Each iteration $\leftrightarrow$ a game round
- In each round,
  - a link earns a profit in service,
  - the scheduler incurs a loss in schedule length.

Roadmap

- Overview
- **An Adaptive Zero-Sum Game with Retirement**
- The General Paradigm
- Discussion

Adaptive Zero-Sum Game with Retirement: Initialization

a set A of m agents vs. an adversary

Adaptive Zero-Sum Game with Retirement: Initialization

a set A of m agents vs. an adversary

- S: the set of active agents, initially A

Adaptive Zero-Sum Game with Retirement: Initialization

a set A of m agents vs. an adversary

- S: the set of active agents, initially A
- $P(a)$: the cumulative profit of $a \in A$, initially 0

Adaptive Zero-Sum Game with Retirement: Initialization

a set A of m agents vs. an adversary

- S: the set of active agents, initially A
- $P(a)$: the cumulative profit of $a \in A$, initially 0
- $w(a)$: the weight of $a \in A$, initially 1.

Adaptive Zero-Sum Game with Retirement: Initialization

a set A of m agents vs. an adversary

- S: the set of active agents, initially A
- $P(a)$: the cumulative profit of $a \in A$, initially 0
- $w(a)$: the weight of $a \in A$, initially 1.
- $\phi = \frac{\ln m + \varepsilon}{\varepsilon(1+\varepsilon) + \ln(1-\varepsilon)}$: retirement threshold

Adaptive Zero-Sum Game with Retirement : Rules

Repeat following round while $S \neq \emptyset$:

1. **Generation of profits and loss** by the adversary:

Adaptive Zero-Sum Game with Retirement : Rules

Repeat following round while $S \neq \emptyset$:

1. **Generation of profits and loss** by the adversary:
   - Each $a \in S$ earns a profit $p(a) \in \mathbb{R}_+$ (i.e., $P(a) \leftarrow P(a) + p(a)$) subject to the **Normalization Rule**:

$$\max_{a \in S} p(a) = 1$$

Adaptive Zero-Sum Game with Retirement : Rules

Repeat following round while $S \neq \emptyset$:

1. **Generation of profits and loss** by the adversary:
   - Each $a \in S$ earns a profit $p(a) \in \mathbb{R}_+$ (i.e., $P(a) \leftarrow P(a) + p(a)$) subject to the **Normalization Rule**:

$$\max_{a \in S} p(a) = 1$$

   - The adversary incurs a loss by the **Zero-Sum Rule:**

$$\frac{1}{w(S)} \sum_{a \in S} w(a) p(a)$$

Adaptive Zero-Sum Game with Retirement : Rules

Repeat following round while $S \neq \emptyset$:

1. **Generation of profits and loss** by the adversary:
   - Each $a \in S$ earns a profit $p(a) \in \mathbb{R}_+$ (i.e., $P(a) \leftarrow P(a) + p(a)$) subject to the **Normalization Rule**:

$$\max_{a \in S} p(a) = 1$$

   - The adversary incurs a loss by the **Zero-Sum Rule:**

$$\frac{1}{w(S)} \sum_{a \in S} w(a)\, p(a)$$

2. **Multiplicative Weights Update (MWU)** by the agents:

$$w(a) \leftarrow w(a)\,(1 - \varepsilon p(a)), \forall a \in S.$$

Adaptive Zero-Sum Game with Retirement : Rules

Repeat following round while $S \neq \emptyset$:

1. **Generation of profits and loss** by the adversary:
   - Each $a \in S$ earns a profit $p(a) \in \mathbb{R}_+$ (i.e., $P(a) \leftarrow P(a) + p(a)$) subject to the **Normalization Rule**:

$$\max_{a \in S} p(a) = 1$$

   - The adversary incurs a loss by the **Zero-Sum Rule:**

$$\frac{1}{w(S)} \sum_{a \in S} w(a) p(a)$$

2. **Multiplicative Weights Update (MWU)** by the agents:

$$w(a) \leftarrow w(a)(1 - \varepsilon p(a)), \forall a \in S.$$

3. **Retirement of agents**: $\forall a \in S$, if $P(a) \geq \phi$ then $S \leftarrow S \setminus \{a\}$.

Analysis of The Game

Theorem

The number of rounds $\leq m \lceil \phi \rceil = O\left(\varepsilon^{-2} m \ln m\right)$. *At the end of the game, the cumulative profit of each agent is at least* ϕ, *and the cumulative loss of the adversary is at most* $(1+\varepsilon)\,\phi$.

Roadmap

- Overview
- An Adaptive Zero-Sum Game with Retirement
- **The General Paradigm**
- Discussion

Augmentations As An Adaptive Zero-Sum Game with Retirement

- Each link $l \in [m] \leftrightarrow$ an agent

Augmentations As An Adaptive Zero-Sum Game with Retirement

- Each link $l \in [m] \leftrightarrow$ an agent
- Each iteration $\leftrightarrow$ a game round

Augmentations As An Adaptive Zero-Sum Game with Retirement

- Each link $l \in [m] \leftrightarrow$ an agent
- Each iteration $\leftrightarrow$ a game round
- Cumulative profit of $l \leftrightarrow$ proportion of its demand served by Γ

Augmentations As An Adaptive Zero-Sum Game with Retirement

- Each link $l \in [m] \leftrightarrow$ an agent
- Each iteration $\leftrightarrow$ a game round
- Cumulative profit of $l \leftrightarrow$ proportion of its demand served by Γ
- Active set $S \leftrightarrow$ the set of links not fully served by $\frac{1}{\phi}-$

Algorithm Design

Algorithm **LS**(ε):

// **initialization**

$\Gamma \leftarrow \emptyset$, $P \leftarrow \mathbf{0}$, $w \leftarrow \mathbf{1}$, $S \leftarrow [m]$; $\phi \leftarrow \frac{\ln m+\varepsilon}{\varepsilon(1+\varepsilon)+\ln(1-\varepsilon)}$;

// **link schedule augmentations**

while $S \neq \emptyset$ do

    // **augmentation**

    $\vdots$

    $\Gamma \leftarrow \Gamma \cup \{(I, x)\}$;

    // **updates**

    $\vdots$

// **scaling**

return $\frac{1}{\phi}\Gamma$.

Augmentation: Selection of IS

$I \leftarrow$ the IS of $\bigcup_{l \in S} E_l$ output by $\mathcal{A}$ w.r.t. $\overline{w}$;

Augmentation: Selection of IS

$I \leftarrow$ the IS of $\bigcup_{l \in S} E_l$ output by $\mathcal{A}$ w.r.t. $\overline{w}$;

- $\overline{w}$: the weight function on E defined by

$$\overline{w}(e) = \frac{w(l)}{d(l)} b(e); \forall e \in E_l, l \in [m].$$

Augmentation: Selection of IS

$$I \leftarrow \text{the IS of } \bigcup_{l \in S} E_l \text{ output by } \mathcal{A} \text{ w.r.t. } \overline{w};$$

- $\overline{w}$: the weight function on E defined by

$$\overline{w}(e) = \frac{w(l)}{d(l)} b(e); \forall e \in E_l, l \in [m].$$

- Weak duality:

$$\chi^*(d) \geq \frac{w(S)}{\mu \overline{w}(I)}.$$

Augmentation: Duration of IS

$I \leftarrow$ the IS of $\bigcup_{l \in S} E_l$ output by $\mathcal{A}$ w.r.t. $\overline{w}$;
$x \leftarrow \frac{1}{\max_{l \in S} \frac{b(E_l \cap I)}{d(l)}}$;
$\Gamma \leftarrow \Gamma \cup \{(I, x)\}$;

Augmentation: Duration of IS

$$\boxed{\begin{array}{l} I \leftarrow \text{the IS of } \bigcup_{l \in S} E_l \text{ output by } \mathcal{A} \text{ w.r.t. } \overline{w}; \\ x \leftarrow \dfrac{1}{\max_{l \in S} \frac{b(E_l \cap I)}{d(l)}}; \\ \Gamma \leftarrow \Gamma \cup \{(I, x)\}; \end{array}}$$

- each $l \in S$ earns a profit $x \frac{b(E_l \cap I)}{d(l)}$

Augmentation: Duration of IS

$$
\boxed{\begin{array}{l}
I \leftarrow \text{the IS of } \bigcup_{l \in S} E_l \text{ output by } \mathcal{A} \text{ w.r.t. } \overline{w}; \\
x \leftarrow \dfrac{1}{\max_{l \in S} \frac{b(E_l \cap I)}{d(l)}}; \\
\Gamma \leftarrow \Gamma \cup \{(I, x)\};
\end{array}}
$$

- each $l \in S$ earns a profit $x \frac{b(E_l \cap I)}{d(l)}$
- x is determined by the **Normalization Rule**:

$$x = \frac{1}{\max_{l \in S} \frac{b(E_l \cap I)}{d(l)}}$$

Updates

for each $I \in S$ do

$P(I) \leftarrow P(I) + x\frac{b(E_l \cap \boldsymbol{I})}{d(I)}$; // **update the profit**

$w(I) \leftarrow w(I)\left(1 - \varepsilon x \frac{b(E_l \cap \boldsymbol{I})}{d(I)}\right)$; // **MWU**

if $P(I) \geq \phi$ then $S \leftarrow S \setminus \{I\}$; // **retirement**

Algorithm Analysis

Theorem

The algorithm **LS**(ε) *runs in* $O\left(\varepsilon^{-2} m \ln m\right)$ *iterations and has an approximation bound* $(1+\varepsilon)\,\mu$.

Loss of The Adversary in A Round

By **Zero-Sum Rule**, the loss of the adversary in a round is

$$
\begin{aligned}
& \frac{1}{w(S)} \sum_{l \in S} w(l)\, x \frac{b(E_l \cap I)}{d(l)} \\
= & \frac{x}{w(S)} \sum_{l \in S} \sum_{e \in E_l \cap I} \frac{w(l)\, b(e)}{d(l)} \\
= & \frac{x}{w(S)} \sum_{l \in S} \sum_{e \in E_l \cap I} \overline{w}(e) \\
= & x \frac{\overline{w}(I)}{w(S)} \\
\geq & \frac{x}{\mu \chi^{*}(d)}
\end{aligned}
$$

Cumulative Loss of The Adversary in The End

- On one hand, the cumulative loss of the adversary $\geq \frac{\|\Gamma\|}{\mu\chi^{*}(d)}$.

Cumulative Loss of The Adversary in The End

- On one hand, the cumulative loss of the adversary $\geq \frac{\|\Gamma\|}{\mu\chi^*(d)}$.
- On the hand, the cumulative loss of the adversary $\leq (1+\varepsilon)\,\phi$

Cumulative Loss of The Adversary in The End

- On one hand, the cumulative loss of the adversary $\geq \frac{\|\Gamma\|}{\mu\chi^*(d)}$.
- On the hand, the cumulative loss of the adversary $\leq (1+\varepsilon)\,\phi$
- Thus,

$$\frac{\|\Gamma\|}{\mu\chi^*(d)} \leq (1+\varepsilon)\,\phi.$$

Cumulative Loss of The Adversary in The End

- On one hand, the cumulative loss of the adversary $\geq \frac{\|\Gamma\|}{\mu\chi^*(d)}$.
- On the hand, the cumulative loss of the adversary $\leq (1+\varepsilon)\,\phi$
- Thus,

$$\frac{\|\Gamma\|}{\mu\chi^*(d)} \leq (1+\varepsilon)\,\phi.$$

Cumulative Loss of The Adversary in The End

- On one hand, the cumulative loss of the adversary $\geq \frac{\|\Gamma\|}{\mu\chi^*(d)}$.
- On the hand, the cumulative loss of the adversary $\leq (1+\varepsilon)\,\phi$
- Thus,

$$\frac{\|\Gamma\|}{\mu\chi^*(d)} \leq (1+\varepsilon)\,\phi.$$

$$\frac{\|\Gamma\|}{\phi} \leq (1+\varepsilon)\,\mu\chi^*(d).$$

Roadmap

- Overview
- An Adaptive Zero-Sum Game with Retirement
- **The General Paradigm**
- Discussion

Discussion

- Almost approx.-preserving reduction from **SLS** to **MWIS**

Discussion

- Almost approx.-preserving reduction from **SLS** to **MWIS**
- Arbitrary wireless networks

Discussion

- Almost approx.-preserving reduction from **SLS** to **MWIS**
- Arbitrary wireless networks
- Broader applications of the game to covering problems.