Data exchange over arbitrary wireless networks

Ivo Kubjas

June 1, 2015

Motivation

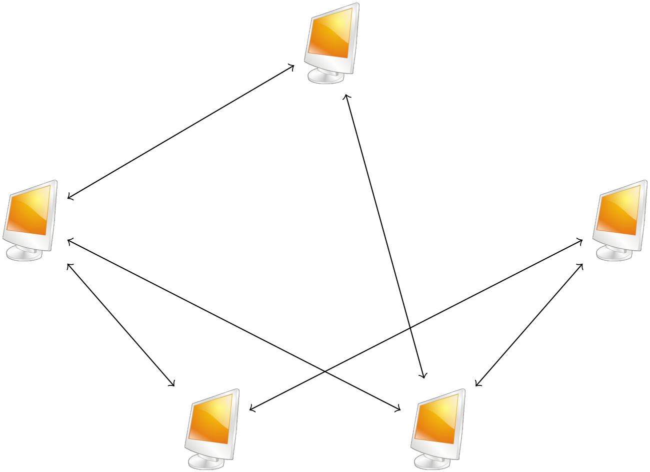

Figure: The network of 5 nodes

Motivation

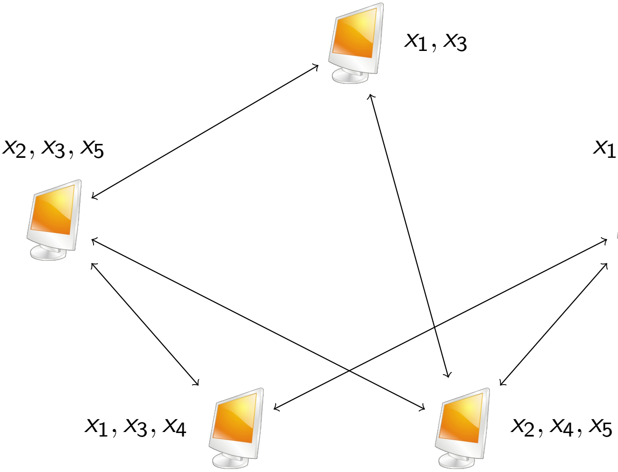

Figure: The sets associated with nodes

Motivation

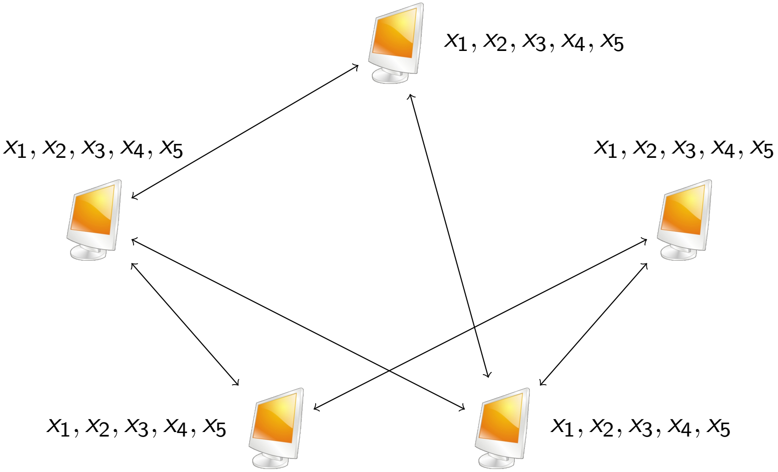

Figure: Reconciled sets

Multicast channel

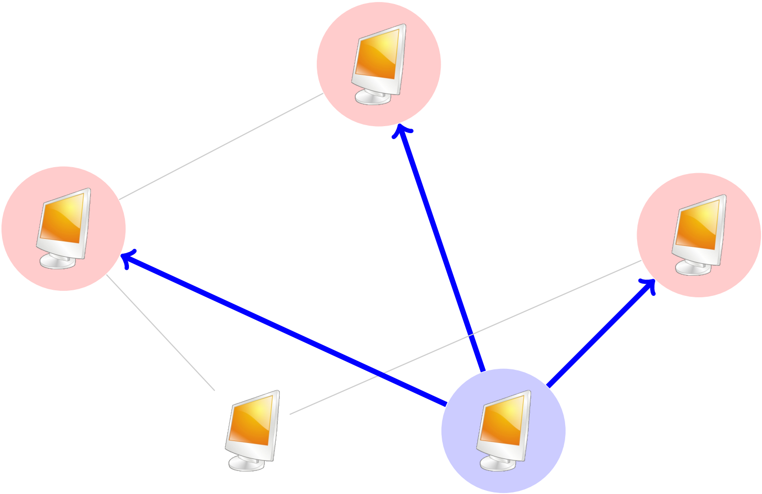

Figure: The node in blue circle transmits and the nodes in red circles receive the message

Index coding

Transmitter wants to transmit bits $x_1, \ldots, x_n$ such that each of the nodes recovers bit x_i while nodes are having bits $\{x_j, j \neq i\}$ as side information.

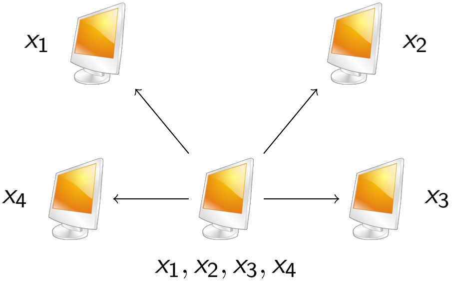

Index coding

Transmitter wants to transmit bits $x_1, \ldots, x_n$ such that each of the nodes recovers bit x_i while nodes are having bits $\{x_j, j \neq i\}$ as side information.

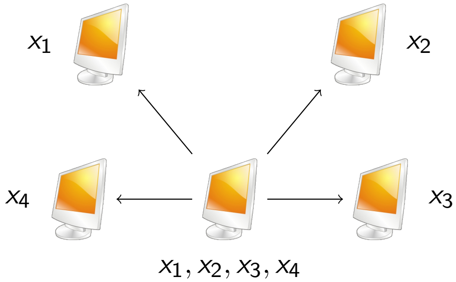

Solution (Bar-Yossef, Birk, Jayram, Kol '06)

$\mathcal{G}$ is a side information graph if there is edge between i-th and j-th node if node X_j has bit x_i. Then the transmitter needs to send $\text{minrank}_2(G)$ *bits.*

minrank_2

Definition

Let $\mathcal{G}$ be a directed graph of n vertices without self-loops. We say that a $0-1$ matrix $A = (a_{ij})$ fits $\mathcal{G}$ if for all i and j:

- $a_{ii} = 1$,
- $a_{ij} = 0$ whenever (i, j) is not an edge of $\mathcal{G}$.

minrank_2

Definition

Let $\mathcal{G}$ be a directed graph of n vertices without self-loops. We say that a $0-1$ matrix $A = (a_{ij})$ fits $\mathcal{G}$ if for all i and j:

- $a_{ii} = 1$,
- $a_{ij} = 0$ whenever (i, j) is not an edge of $\mathcal{G}$.

Definition

$\text{minrank}_2(G) := \min \{\text{rank}_2(A) : A \text{ fits } \mathcal{G}\}$.

Side information graph

Example

The side information graph for graph $\mathcal{G}$ from previous example is

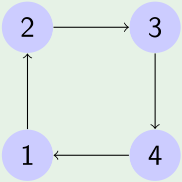

Matrix

$$A = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 1 \end{bmatrix}$$

fits $\mathcal{G}$ and $\text{rank}_2(A) = 3$. Thus, the transmitter needs at least 3 transmissions.

Data Exchange Protocol

Each node holds some items x_i and they wish to recover all items for every node.

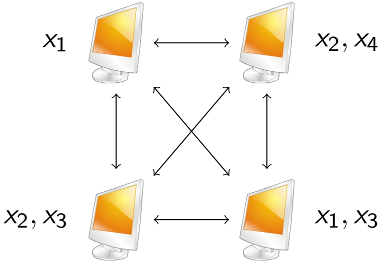

Data Exchange Protocol

Each node holds some items x_i and they wish to recover all items for every node.

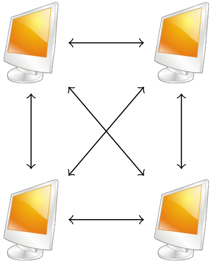

Solution (El Rouayheb, Sprintson, Sadeghi)

Let $\mathbb{A}$ *be a family of matrices corresponding to the items the nodes have and* B_i *be a matrix denoting the items i-th node has. Then the number of transmissions is* $\tau = \min_{A \in \mathbb{A}} rank(A)$ *such that* $\operatorname{rank}\left(\begin{bmatrix} A \\ B_i \end{bmatrix}\right) = n$*,* $\forall i = 1, \ldots, k$*.*

Set reconciliation protocol using possession matrices

Let $\mathcal{P}_\ell$ be the indices of possesseed bits by node ℓ.

Definition

For each node ℓ, the possession matrix $\mathbb{A}_\ell$ is a $n \times n$ matrix over $\mathbb{F} \cup \{\star\}$, where '$\star$' is a symbol which can take any value in $\mathbb{F}$. It is defined as

$$(\mathbb{A}_\ell)_{i,j} = \begin{cases} \star & \text{if } j \in \mathcal{P}_\ell \\ 0 & \text{otherwise} \end{cases} .$$

The possession matrix of the graph is the $(kn \times n)$-dimensional matrix

$$\mathbb{A} = \begin{bmatrix} \mathbb{A}_1 \\ \mathbb{A}_2 \\ \vdots \\ \mathbb{A}_k \end{bmatrix},$$

where $\mathbb{A}_i$ is the possession matrix family corresponding to the node i, $i \in [k]$.

The members of the matrix family $\mathbb{A}_i$ denote the transmission matrices.

Example

Let the set of all items be $\mathbf{X} = (x_1, x_2, x_3, x_4)$. If the node v_1 has items $X_1 = \{x_1, x_3\}$, then the possession matrix of this node is

$$\mathbb{A}_1 = \begin{bmatrix} * & 0 & * & 0 \\ * & 0 & * & 0 \\ * & 0 & * & 0 \\ * & 0 & * & 0 \end{bmatrix}.$$

For a matrix $A_1 \in \mathbb{A}_1$ such that

$$A_1 = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

the transmitted messages are non-zero elements of the vector $A_1\mathbf{X}^T = (x_1 + x_3, x_3, 0, 0)^T$.

Definition

Given $\boldsymbol{A} \in \mathbb{A}$, the j-th $n \times n$ sub-matrix of $\boldsymbol{A}$ will be denoted as $\boldsymbol{A}_j$.

Definition

Given the matrix family $\mathbb{A}_i$, the operator $\Gamma(\cdot)$ replaces the symbols '$\star$' in the maximal number of the first rows with linearly independent canonical vectors, and replaces the symbols '$\star$' in the remaining rows with zeros.

max-rank

Let $A_{\max} \in \mathbb{A}_i$ have the maximal rank within $\mathbb{A}_i$. Then $\mathrm{rank}(A_{\max})$ is the largest possible number of transmissions of independent items. The transmission vector $A_{\max}\mathbf{X}^T$ is enough to recover X_i.

max-rank

Let $A_{\max} \in \mathbb{A}_i$ have the maximal rank within $\mathbb{A}_i$. Then $\operatorname{rank}(A_{\max})$ is the largest possible number of transmissions of independent items. The transmission vector $A_{\max}\mathbf{X}^T$ is enough to recover X_i.

Definition

The max-rank of the matrix family $\mathbb{A}$ is defined as

$$\text{max-rank}(\mathbb{A}) = \max_{A \in \mathbb{A}} \operatorname{rank}(A).$$

- During a single round, the messages are transmitted to nearest neighbours.
- Several rounds are required for full reconciliation.
- For an adjacency matrix D, the i-th power D^i denotes the number of paths of length i between the nodes.

- During a single round, the messages are transmitted to nearest neighbours.
- Several rounds are required for full reconciliation.
- For an adjacency matrix D, the i-th power D^i denotes the number of paths of length i between the nodes.

Lemma

Let D be the adjacency matrix of the graph. Let ℓ be the smallest positive integer such that $\sum_{i=1}^{\ell} D^i$ is a positive matrix. Then the network is ℓ-solvable.

- During a single round, the messages are transmitted to nearest neighbours.
- Several rounds are required for full reconciliation.
- For an adjacency matrix D, the i-th power D^i denotes the number of paths of length i between the nodes.

Lemma

Let D be the adjacency matrix of the graph. Let ℓ be the smallest positive integer such that $\sum_{i=1}^{\ell} D^i$ is a positive matrix. Then the network is ℓ-solvable.

Example

If the network is  , then the adjacency matrix is $D = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix}$ and the network is 2-solvable.

Lemma

Let $\mathbb{A}$ be the possession matrix of the graph, D be the adjacency matrix of the graph and E be a $(n \times n)$-dimensional all-ones matrix. After performing one round of the protocol, the new possession matrix $\mathbb{A}_+$ is related to $\mathbb{A}$ as

$$\mathbb{A}_+ = (D \otimes E)\mathbb{A}.$$

Lemma

Let $\mathbb{A}$ be the possession matrix of the graph, D be the adjacency matrix of the graph and E be a $(n \times n)$-dimensional all-ones matrix. After performing one round of the protocol, the new possession matrix $\mathbb{A}_+$ is related to $\mathbb{A}$ as

$$\mathbb{A}_+ = (D \otimes E)\mathbb{A}.$$

Example

If $\begin{matrix} X_1 = \{x_1\} \\ X_2 = \{x_2\} \\ X_3 = \{x_3\} \end{matrix}$, then $\mathbb{A} = \begin{bmatrix} * & 0 & 0 \\ * & 0 & 0 \\ * & 0 & 0 \\ 0 & * & 0 \\ 0 & * & 0 \\ 0 & * & 0 \\ 0 & 0 & * \\ 0 & 0 & * \\ 0 & 0 & * \end{bmatrix}$ and $\mathbb{A}_+ = \begin{bmatrix} * & * & 0 \\ * & * & 0 \\ * & * & 0 \\ * & * & * \\ * & * & * \\ * & * & * \\ 0 & * & * \\ 0 & * & * \\ 0 & * & * \end{bmatrix}$.

The set reconciliation protocol construction

- Using the adjacency matrix, it is possible to obtain the number of items each node **should** have after a round.
- Similarly, for a specific matrix family member $\boldsymbol{A} \in \mathbb{A}$, with

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A}_1 \\ \boldsymbol{A}_2 \\ \vdots \\ \boldsymbol{A}_k \end{bmatrix},$$

  it is possible to obtain the number of items each node **could** have after a round.
- If the latter two are equal, then the items are reconciled in a single round.
- Iterating over the required number of rounds, a protocol for full set reconciliation is obtained with the minimal number of rounds and transmissions.

Theorem

Let $\mathcal{G}$ be an underlying directed graph of an r_0-solvable network defined by the adjacency matrix $\boldsymbol{D}^T$. Let $\mathbb{A}$ be the corresponding possession matrix of the network. Then there exists an iterated data exchange protocol with r rounds, for any $r \geq r_0$, and τ transmissions, where

$$\tau = \sum_{i=1}^{r} \left(\min_{\boldsymbol{A}^{(i)} \in (\boldsymbol{D}^{i-1} \otimes \boldsymbol{E}) \cdot \mathbb{A}} \left\{ \sum_{j=1}^{k} \text{rank} \left(\boldsymbol{A}_j^{(i)} \right) \right\} \right) \quad (1)$$

for matrices $\boldsymbol{A}^{(i)}$ which are subject to

$$\forall j \in [k] \,:\, \text{rank} \left(\left[\frac{\left(\text{diag}\left(\boldsymbol{D}^{[j]}\right) \otimes \boldsymbol{I}\right) \cdot \boldsymbol{A}^{(i)}}{\Gamma_j((\boldsymbol{D}^{i-1} \otimes \boldsymbol{E}) \cdot \mathbb{A})} \right] \right) = \text{max-rank} \left((\text{diag}(\boldsymbol{e}_j) \otimes \boldsymbol{I}) \cdot \left(\boldsymbol{D}^{i} \otimes \boldsymbol{E} \right) \cdot \mathbb{A} \right), \quad (2)$$

where the matrices $\boldsymbol{I}$ and $\boldsymbol{E}$ are both $n \times n$.

Application to special case

- General case has greedy approach to every round.

Application to special case

- General case has greedy approach to every round.
- Let $\mathcal{T}_\ell$ be the indices of information bits requested by node ℓ.
  For each $\ell \in \mathcal{V}$, the $n \times n$ *information matrix* $\boldsymbol{P}_\ell = (\boldsymbol{P}_\ell)_{i\in[n],j\in[n]}$ is

$$(\boldsymbol{P}_\ell)_{i,j} = \begin{cases} 1 & \text{if } i = j \text{ and } i \in \mathcal{P}_\ell \\ 0 & \text{otherwise} \end{cases},$$

  and the $n \times n$ *query matrix* $\boldsymbol{T}_\ell = (\boldsymbol{T}_\ell)_{i\in[n],j\in[n]}$,

$$(\boldsymbol{T}_\ell)_{i,j} = \begin{cases} 1 & \text{if } i = j \text{ and } i \in \mathcal{T}_\ell \\ 0 & \text{otherwise} \end{cases}.$$

- The set of in-neighbours for node ℓ is $\mathcal{N}_{in}(\ell)$ and

$$\boldsymbol{A}_{\mathcal{N}_{in}(\ell)} = \begin{bmatrix} \boldsymbol{A}_{i_1} \\ \hline \boldsymbol{A}_{i_2} \\ \hline \vdots \\ \hline \boldsymbol{A}_{i_d} \end{bmatrix}.$$

  where $\mathcal{N}_{in}(\ell) = \{i_1, i_2, \cdots, i_d\}$, and d is an in-degree of ℓ in $\mathcal{G}$.

Application to special case

Theorem

Consider a wireless network defined by the graph $\mathcal{G}$. Let $\mathbb{A}$ be the possession matrix of the network. The minimal number of transmissions needed to satisfy the demands of all nodes in one round of communications is

$$\tau = \min_{\boldsymbol{A} \in \mathbb{A}} \left\{ \sum_{\ell \in \mathcal{V}} \operatorname{rank}(\boldsymbol{A}_\ell) \right\} , \tag{3}$$

where for all $\ell \in \mathcal{V}$

$$\operatorname{rowspace}\left(\left[\frac{\boldsymbol{A}_{\mathcal{N}_{in}(\ell)}}{\boldsymbol{P}_\ell} \right] \right) \supseteq \operatorname{rowspace}(\boldsymbol{T}_\ell) . \tag{4}$$

If the above matrix $\boldsymbol{A} \in \mathbb{A}$ as above does not exist then there is no algorithm that satisfies all requests in one round.

Solving max-rank

Solving max-rank

Theorem (Hartfiel, Loewy '84)

Let B be a partial V_r by V_c matrix with free entries indexed by elements of E. For any cover C of E we have $\operatorname{rank} B^* \leq \operatorname{rank} B \setminus C + |C|$. *Furthermore, there exists a cover C^* of E such that* $\operatorname{rank} B^* = \operatorname{rank} B \setminus C^* + |C^*|$.

Theorem (Geelen '99)

Let B be a partial V_r by V_c matrix with free entries indexed by elements of E, and let $\tilde{B}$ be a completion of B. Then, either $\operatorname{rank} \tilde{B} = \operatorname{rank} B^*$ *or there exists $(i,j) \in E$ and $a \in \{1, \dots, |V_r| + |V_c|\}$ such that $\tilde{B}(i,j;a) > \tilde{B}$.*