

On the Complexity of Cooperative Peer-to-Peer Repair for Wireless Broadcasting

Gene Cheung, *Member, IEEE*, Danjue Li, *Student Member, IEEE*, and Chen-Nee Chuah, *Senior Member, IEEE*

Abstract—The well-known NAK implosion problem for wireless broadcast can be addressed by leveraging cooperative peer-to-peer connectivity to repair corrupted data. This paper studies the Cooperative Peer-to-Peer Repair (CPR) framework for multimedia broadcast. We show that CPR can be formulated as an optimization problem that minimizes the number of iterations it takes to wirelessly disseminate a desired message from peers with the content to peers without it. Complicating the problem are transmission conflicts, where pre-specified sets of links cannot simultaneously transmit due to interference. In this paper, we formalize the CPR minimum delay problem and prove that it is NP-hard.

Index Terms—Wireless broadcast, peer-to-peer, complexity.

I. INTRODUCTION

A NEW and promising distribution model for 3rd Generation Partnership Project (3GPP) networks is Multimedia Broadcast Multicast Service (MBMS) [1], where a piece of widely interested multimedia content (message) is broadcasted to large groups of 3G clients listening collectively in a pre-assigned broadcast channel. While it is clear that efficient usage of network resources is a benefit, avoiding the NAK implosion problem (a scenario where server is overwhelmed by floods of individual retransmission requests from clients) is a major reason why broadcasting servers typically do not perform retransmissions on request in the event of packet losses due to wireless transmission failures. Even with the use of Forward Error Correction (FEC) to correct predictable channel noise, temporary wireless link failures are unavoidable, leaving groups of clients without the desired message at a given time.

Fortunately, many modern wireless devices are multi-homed and each contains multiple wireless interfaces, so that one can connect to a wireless wide area network (WWAN), like a 3G network, and to a wireless local area network (WLAN), like a wireless ad-hoc peer-to-peer network, simultaneously [2]. In such setting, a “have not” wireless peer can request retransmission of a message from a neighboring “have” peer listening to the same broadcast. Given a group of cooperative wireless peers willing to repair neighbors’ dropped message, the problem is: how to schedule rounds of retransmissions within a group, so that the time required to complete repair




Fig. 1. Example of connectivity graph for cooperative peer-to-peer repair.

to all peers is minimized? Care must be taken so that pre-specified sets of interfering transmission links are not activated simultaneously. We call this problem the *Cooperative Peer-to-Peer Repair* problem (*CPR*).

In this paper, we present the following theoretical result: *CPR* is NP-hard. We first formalize *CPR* as a discrete optimization problem in Section II. We then present the NP-hardness proof for *CPR* in Section III. We provide concluding remarks in Section IV.

II. COOPERATIVE PEER-TO-PEER REPAIR (CPR)

We formulate *CPR* as follows. A connected graph Θ , modeling the connectivity of wireless peers in WLAN, has a set of nodes \mathcal{N} and a set of *undirected* links \mathcal{L} . Links are labeled from $0, \dots, |\mathcal{L}| - 1$, where link i connecting nodes m and n is represented by $i \leftrightarrow (m, n)$. At start time $t = 0$, each node $n \in \mathcal{N}$ has color $C_{0,n} \in \{0, 1\}$, where 0 (*blue*) means node n is in need of the desired message, and 1 (*white*) means the node has the message. As done in [3], a *conflict matrix* \mathbf{I} of dimension $|\mathcal{L}| * |\mathcal{L}|$ dictates which links cannot be activated at the same time due to interference; in particular, $I_{i,j} = 1$ if link i and j cannot be activated simultaneously, and $I_{i,j} = 0$ otherwise. Matrix \mathbf{I} is by definition symmetric. We assume \mathbf{I} has the *unicast conflict property*: assignments of 1’s and 0’s so that two links stemming from the same node are in conflict. This is in compliance with standard 802.11 MAC behavior for unicast mode, where a node can be in communication with at most one other node at the same time.

At each iteration t , we select links, each connecting a white node to a blue node, such that no two selected links are in conflict according to \mathbf{I} . By next iteration $t + 1$, blue nodes

Manuscript received June 2, 2006. The associate editor coordinating the review of this letter and approving it for publication was Prof. Carla-Fabiana Chiasserini.

G. Cheung is with Hewlett-Packard Laboratories, Takaido Office Bldg. #3, 3-8-13 Takaido-Higashi, Suginami-ku, Tokyo, 168-0072 Japan (email: gene.cs.cheung@hp.com).

D. Li and C.-N. Chuah are with the Dept. of Electrical & Computer Engineering, University of California, Davis.

Digital Object Identifier 10.1109/LCOMM.2006.060843.




Fig. 2. Example construction for NP-completeness proof for CPR.

of the selected links have received the desired message and changed color to white. The optimization problem is: how to select a set of non-conflicting links in each iteration, so that all nodes are white in the minimum number of iterations? Figure 1 shows an example of connectivity graph $\Theta = \{\mathcal{N}, \mathcal{L}\}$ with initial coloring for CPR.

We write CPR mathematically as follows. Let \mathbf{S} be a $T * |\mathcal{L}|$ binary *link selection matrix*, where $S_{t,l} = 1$ if link l is selected at iteration t and $S_{t,l} = 0$ otherwise, and T is the total number of iterations. Let \mathbf{C} be a $T * |\mathcal{N}|$ binary *color matrix* where $C_{t,n} = 1$ if node n is white at iteration t and $C_{t,n} = 0$ otherwise. Given the first row of \mathbf{C} is initialized to the starting colors of \mathcal{N} , the optimization is:

$$\begin{aligned} \min_{\mathbf{S}, \mathbf{C}} \quad & \text{row}(\mathbf{S}) \quad \text{s.t.} \\ & I_{j,k} = 0 \quad \forall j, k \mid S_{t,j} = S_{t,k} = 1 \\ & C_{t,m} + C_{t,n} = 1 \quad \forall l \mid S_{t,l} = 1, l \leftrightarrow (m, n) \\ & C_{t+1,m} + C_{t+1,n} = 2 \quad \forall l \mid S_{t,l} = 1, l \leftrightarrow (m, n) \\ & C_{t+1,n} = C_{t,n} \quad \exists l \mid S_{t,l} = 1, l \leftrightarrow (m, n) \\ & \sum_n C_{\text{row}(\mathbf{S}),n} = |\mathcal{N}| \end{aligned} \quad (1)$$

where $\text{row}(\mathbf{S})$ is the number of rows in matrix \mathbf{S} . The 1st constraint in (1) states that no two links selected in the same iteration t should be in conflict. The 2nd constraint states that one and only one node of each selected link at iteration t should be white. The 3rd constraint states that both nodes of a selected link at iteration t should be white at iteration $t+1$. The 4th constraint states that color of a node stays the same at iteration $t+1$ if no link connected to it was selected at iteration t . The 5th constraint states that all nodes must be white at iteration $\text{row}(\mathbf{S})$.

We now present the NP-hardness proof for CPR.

III. PROOF OF NP-HARDNESS

We first recast CPR as a decision problem: is there a schedule of non-conflicting links at each iteration, such that all nodes can be whitened in κ iterations? The CPR decision problem is obviously in NP; a solution $(\mathbf{S}^\circ, \mathbf{C}^\circ)$ can be checked against constraints in (1), and $\text{row}(\mathbf{S}^\circ)$ against κ , for feasibility in polynomial time.

We next show that the CPR decision problem is NP-complete via polynomial transformation from a well-known NP-complete problem *Independent Set (IndS)*. The IndS decision problem can be stated¹ as follows (pg.361 of [4]):

¹Notice we adopt the terminology of *nodes* and *links* when referring to CPR, and *vertices* and *edges* when referring to IndS to avoid confusion.

Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and an integer k , is there a set $\mathcal{I} \subset \mathcal{V}$ of k vertices such that no two vertices in \mathcal{I} are connected by an edge?

Figure 2a shows an IndS example with independent set $\{1, 3\}$. We now describe a procedure to construct a CPR instance from an arbitrary IndS instance, so that the output of the CPR decision problem corresponds exactly to the decision in IndS, and hence proving that CPR is at least as hard as IndS.

A. Construction of CPR Instance from IndS

We first construct a *IndS conflict binary matrix* \mathbf{J} of size $|\mathcal{V}| * |\mathcal{V}|$, where $J_{i,j} = 1$ if $\exists e_{i,j} \in \mathcal{E}$ and $J_{i,j} = 0$ otherwise. In other words, $J_{i,j} = 1$ iff vertices i and j cannot be selected to the same independent set because they are connected by an edge. A IndS conflict matrix $\{J_{i,j}\}$, $0 \leq i, j \leq 3$, corresponding to the IndS instance in Figure 2a is:

$$\mathbf{J} = \begin{bmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix} \quad (2)$$

We next construct a corresponding CPR instance with graph $\Theta = (\mathcal{N}, \mathcal{L})$ and conflict matrix \mathbf{I} . More specifically, we construct a bipartite graph with $|\mathcal{V}|$ nodes on the left (\mathcal{N}_l) and k nodes on the right (\mathcal{N}_r), so that $\mathcal{N}_l \cup \mathcal{N}_r = \mathcal{N}$. Nodes on the left, each labeled $n \in \{0, \dots, |\mathcal{V}| - 1\}$, are white, and nodes on the right, each labeled $m' \in \{0', \dots, (k-1)'\}$, are blue. We draw a link from left to right for every pair of left-node and right-node. We label a link $l \in \{0, 1, \dots, k|\mathcal{V}| - 1\}$ connecting $n \in \mathcal{N}_l$ and $m' \in \mathcal{N}_r$ as follows:

$$l = n * k + m' \leftrightarrow (n, m') \quad n \in \mathcal{N}_l, m' \in \mathcal{N}_r \quad (3)$$

The constructed CPR instance for our IndS example is shown in Figure 2b. To complete the CPR instance, we construct a CPR conflict matrix \mathbf{I} of size $k|\mathcal{V}| * k|\mathcal{V}|$ from \mathbf{J} as follows. For $0 \leq i, j \leq k|\mathcal{V}| - 1$:

$$I_{i,j} = \begin{cases} 1 & \text{if } J_{\lfloor \frac{i}{k} \rfloor, \lfloor \frac{j}{k} \rfloor} = 1 \\ 1 & \text{else if } \left(\left\lfloor \frac{i}{k} \right\rfloor \neq \left\lfloor \frac{j}{k} \right\rfloor \right) \& (i \bmod k = j \bmod k) \\ 1 & \text{else if } \left(\left\lfloor \frac{i}{k} \right\rfloor = \left\lfloor \frac{j}{k} \right\rfloor \right) \& (i \neq j) \\ 0 & \text{o.w.} \end{cases} \quad (4)$$

where $i \bmod k$ gives the integer remainder of i divided by k .

The CPR conflict matrix \mathbf{I} corresponding to the IndS conflict matrix \mathbf{J} in (2) is as follows.

$$\mathbf{I} = \begin{bmatrix} 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 & 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 & 1 & 1 & 1 & 0 \end{bmatrix} \quad (5)$$

Given the constructed CPR instance, the corresponding decision is: is there a schedule of non-conflicting links, such that all blue nodes can be whitened in 1 iteration?

B. Remarks

First, we show that the constructed *CPR* instance is of polynomial size of the *IndS* instance. The node set \mathcal{N} of bipartite graph for *CPR* is of size $|\mathcal{V}| + k \leq \mathbf{O}(|\mathcal{V}|)$. The number of links is bounded by $\mathbf{O}(|\mathcal{V}|^2)$, and the size of the *CPR* conflict matrix \mathbf{I} is bounded by $k^2|\mathcal{V}|^2 \leq \mathbf{O}(|\mathcal{V}|^4)$. Hence we conclude that the size of the constructed *CPR* instance is $\mathbf{O}(|\mathcal{V}|^4)$, i.e., it is of polynomial size of the *IndS* instance.

Next, we show that the 1's and 0's assigned to \mathbf{I} using (4) satisfy the unicast conflict property. From link labeling (3), we see that links stemming from the same left-node n but arriving at different right-nodes are in conflict with each other due to the 3^{rd} if statement of (4). Similarly, from (3), we see that links arriving at the same right-node m' but stemming from different left-nodes are also in conflict with each other due to the 2^{nd} if statement of (4). Since this covers all links, we conclude that the constructed \mathbf{I} satisfies the unicast conflict property.

Finally, we discuss the intuition behind the construction of the *CPR* instance. Each node $\in \mathcal{N}_l$ of *CPR* corresponds one-to-one to a vertex $\in \mathcal{V}$ of *IndS*. Selecting one of k links stemming from a node $\in \mathcal{N}_l$ of *CPR* means selecting the corresponding vertex $\in \mathcal{V}$ of *IndS* into the independent set. The 1^{st} if statement in (4) prevents selection of links stemming from two nodes representing vertices $\in \mathcal{V}$ that are connected in \mathcal{G} .

Each node $\in \mathcal{N}_r$ corresponds to a unique, successful selection of an independent vertex $\in \mathcal{V}$. We track to see if all k nodes $\in \mathcal{N}_r$ can be whitened in 1 iteration. There is no under-counting, since the unicast conflict property prevents joint selection of links stemming from two nodes $\in \mathcal{N}_l$ going to the same node $\in \mathcal{N}_r$. There is also no over-counting, because the unicast conflict property also prevents joint selection of links stemming from the same node $\in \mathcal{N}_l$ going to different nodes $\in \mathcal{N}_r$.

We now state the proof formally as a theorem.

Theorem 1: The *CPR* decision problem is NP-complete.

Proof: We prove the theorem by showing that “yes” to the constructed *CPR* instance implies “yes” to the original *IndS* instance, and vice versa. Suppose the output of the constructed *CPR* decision problem is “yes”. That means the solution composes of *exactly* k non-conflicting links, originating from k *distinct* left-nodes $\in \mathcal{N}_l$ and terminating at all k blue right-nodes $\in \mathcal{N}_r$. We know there are *exactly* k selected links, because all k blue right-nodes $\in \mathcal{N}_r$ are whitened in 1 iteration, and links terminating at the same right-node $\in \mathcal{N}_r$ are conflicting due to the unicast conflict property. We know

these k links are originated from k *distinct* left-nodes because links originated from the same left-node $\in \mathcal{N}_l$ are conflicting, again due to the unicast conflict property. Finally, these *distinct* k left-nodes $\in \mathcal{N}_l$ of the k selected links must correspond to k independent vertices $\in \mathcal{V}$ in *IndS* due to (3) and 1^{st} if statement of (4). Therefore, we conclude that “yes” to the *CPR* decision problem corresponds to “yes” in the original *IndS* decision problem.

Suppose there is an independent set $\{v_0^o < \dots < v_{k-1}^o\} \subset \mathcal{V}$ of size k in the *IndS* instance. We can correspondingly select a set of k non-conflicting links $\in \mathcal{L}$, where for each v_i^o , we pick link $v_i^o * k + i$. First, we know each of these k links connect to a different right-node $\in \mathcal{N}_r$ due to (3). Second, we know each of these k links connects to a different left-node $\in \mathcal{N}_l$, again due to (3). Hence, given each of these k links has a different left-node and a different right-node, they do not violate unicast conflicts (2^{nd} and 3^{rd} if statements in (4)). They also do not conflict due to 1^{st} if statement of (4), since vertices v_i^o 's, where $v_i^o = \lfloor \frac{v_i^o * k + i}{k} \rfloor$, do not conflict in \mathcal{G} by assumption. Therefore, this set of k non-conflicting links $\in \mathcal{L}$ converts all k blue nodes $\in \mathcal{N}_r$ in 1 iteration. We can now conclude that since both directions of the implication have been proven, Theorem 1 is also proven. ■

Corollary 1: The *CPR* optimization problem is NP-hard.

Proof: Given the *CPR* decision is NP-complete as stated in Theorem 1, it follows by definition of NP-hardness [4] that the corresponding optimization is NP-hard. ■

IV. CONCLUSIONS

In this paper, we studied the Cooperative Peer-to-Peer Repair (*CPR*) framework for multimedia broadcast. We formulated *CPR* as an optimization problem that minimizes the number of iterations it takes to wirelessly disseminate a desired message from peers with the content to peers without it. We proved that *CPR* is NP-hard.

REFERENCES

- [1] “Technical Specification Group Services and System Aspects; Multimedia Broadcast/Multicast Service; Protocols and Codecs (3GPP TS.26.346 version 6.1.0), June 2005.
- [2] P. Sharma, S.-J. Lee, J. Brassil, and K. Shin, “Distributed communication paradigm for wireless community networks,” in *Proc. IEEE International Conference on Communications 2005*, vol. 3, pp. 1549–1555.
- [3] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu, “Impact of interference on multi-hop wireless network performance,” in *Proc. ACM International Conference on Mobile Computing and Networking (MobiCom) 2003*.
- [4] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Dover, 1998.