Performance of Network-Coding in Multi-Rate Wireless Environments for Multicast Applications

Luiz Filipe M. Vieira*, Archan Misra†, Mario Gerla*

* Computer Science Department, University of California, Los Angeles
† IBM T J Watson Research Center, Hawthorne, NY
{luiz, gerla}@cs.ucla.edu, archan@us.ibm.com

***Abstract*— This paper investigates the interaction between network coding and link-layer transmission rate diversity in multi-hop wireless networks. By appropriately mixing data packets at intermediate nodes, network coding allows a single multicast flow to achieve higher throughput to a set of receivers. Broadcast applications can also exploit link-layer rate diversity, whereby individual nodes can transmit at faster rates at the expense of corresponding smaller coverage area. We first demonstrate how combining rate-diversity with network coding can provide a larger capacity for data dissemination of a single multicast flow, and how consideration of rate diversity is critical for maximizing system throughput. We also study the impact of both network coding and rate diversity on the dissemination latency for a class of quasi real-time applications, where the freshness of disseminated data is important. Our results provide evidence that network coding may lead to a latency-vs-throughput tradeoff in wireless environments, and that it is thus necessary to adapt the degree of network coding to ensure conformance to both throughput and latency objectives.**

[1]

I. INTRODUCTION

There is an increasing interest in understanding the potential performance gains accruing from the use of network coding in multi-hop wireless environments. In particular, many military battlefield scenarios exhibit two characteristics that appear to motivate the use of network coding: a) the reliance on bandwidth-constrained, ad-hoc wireless links (e.g. using MANETs formed by vehicle-mounted radios in urban insurgencies) and b) the need to disseminate information (e.g., maps, mission commands) to multiple recipients. The initial results on the power of network coding **NC**, such as the original demonstration in [1] of how in-network mixing of packets by intermediate nodes helps to achieve a communication capacity that is not achievable solely through routing, were obtained for the case of a lossless, wireline network. More recently, several groups have investigated the potential performance gains realized by network coding for both unicast (e.g., [5]) and multicast (e.g., [2]) traffic in wireless environments, for a variety of application scenarios. All of these approaches fundamentally aim to exploit the wireless broadcast advantage (**WBA**) by using, whenever possible, a single link-layer broadcast transmission (of a packet formed by a linear combination of individual packets) to reach multiple neighboring nodes. By saving on the number of independent transmissions needed, network-coding approaches effectively reduce the fraction of time the wireless channel is held by a single transmitting node and thereby help to increase the overall network throughput.

We believe that there is another degree of freedom in wireless environments, namely *link-layer rate diversity*, that network coding approaches have so far failed to exploit. Most commodity wireless cards are now capable of performing adaptive modulation to vary the link rate in response to the signal-to-interference levels at the receiver. Link rate diversity typically exhibits a *rate-range tradeoff*: if the same transmission power is used for all link transmission rates, then, in general, the faster the transmission rate, the smaller is the transmission range (although, the rate-distance variation in real life is somewhat irregular (e.g., see [3])). While this rate diversity has been extensively exploited for unicast traffic and is often standardized, its use in link-layer broadcasting is relatively limited. For example, while the current IEEE 802.11a/b/g standards mandate the transmission of the control frames (e.g. RTS/CTS/ACK) at the lowest rate (e.g., 6 Mbps for IEEE 802.11a), transmission rates for broadcast data are typically implementation-specific. Recently, however, there has been some work (e.g., [4]) that demonstrates that effective exploitation of such rate diversity by *routing* algorithms for link-layer broadcasts can result in significant (often 6-fold) reduction in the broadcast latency and increase in the achievable throughput.

In this paper, we investigate the impact that the use of such rate-diversity for link layer broadcasts may have on the performance of network coding. In addition, we shall also study the *relative importance* of network coding and link-layer transmission rate diversity. It is easy to conceptualize how the rate-range tradeoff inherent to all link-layer broadcasts might impact the performance of various network coding strategies. Without consideration of rate diversity, network coding algorithms operate on an implicit "more-is-better" assumption: since each broadcast transmission takes the same time, encod-

---

[1] Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. Research was sponsored by CAPES.
1-4244-1513-06/07/$25.00 2007 IEEE

ing a larger number of packets (for a correspondingly larger set of neighbors) into a single packet *always* results in a more efficient use of the wireless channel. In reality, the existence of the rate-range tradeoff often invalidates this assumption. For example, assume that a node n has a set of packets $\{P_1, P_2, \ldots, P_N\}$ targeted for its neighbors $\{n_1, n_2, \ldots, n_N\}$, where the neighbor indices are arranged in non-increasing order of the link transmission rates. Moreover, let R_i be the link rate between the node-pair (n, n_i). In this case, it is possible that combining the first i packets (transmitted at the rate R_i) proves to be more effective than combining the first $i{+}1$ packets, because the additional multiplexing gain achieved is negated by the need to use a disproportional smaller rate R_{i+1} for the packet broadcast. Our goals in this paper are thus to answer the following questions:

1) How does the consideration of transmission-rate diversity affect the maximum throughput that may be achieved by linear network coding in wireless environments, i.e., how sensitive are the achievable throughput curves to the impact of link-rate heterogeneity?
2) How does the throughput achieved by a combination of rate-diverse transmissions and network coding differ in practice from that achieved by pure routing-based strategies that are rate-diversity aware?
3) How may network coding and link-layer transmission rates be tuned to achieve satisfactory performance for scenarios where we are concerned about both throughput and latency?

The last item mentioned above introduces a new dimension to the practical utility of network coding in wireless environments. In particular, we shall see how the existence of network coding seems to suggest a *tradeoff between throughput and latency*, a phenomenon that is not typically true for pure-routing based disemmination schemes.

Given the closely-coupled interactions between the degree of encoding, the resultant transmission rate and the contentions on the wireless channel, we focus in this paper on the case of single-source multicast[2] problem. Note that the current paper is not *constructive*, i.e., it does not address the design of specific network-coding algorithms that are better at taking advantage of the rate diversity available in a specific network. Instead, our goal is to understand the fundamental interactions between transmission rate diversity and network coding.

A. Contributions of This Paper

This paper makes the following contributions towards understanding the basic performance of network-coding for broadcast/multicast applications in wireless environments:

- It uses extensive simulation-based studies, to evaluation the relative performance of network coding algorithms vs. pure routing-based broadcasting strategies in rate-diverse wireless networks.
- It studies the impact of NC on not just the network throughput, but also the *dissemination latency*. Such a study suggests the possible existence of a throughput-vs-latency tradeoff that depends on the 'degree of network coding employed; we believe that such a tradeoff will have important implications in the design of suitable network coding strategies for latency-sensitive information dissemination, such as the timely broadcast of sensor values.
- It suggests how a joint (throughput,latency) utility function may be used to express the conflicting objectives and provides some initial results on how the degree of NC and the choice of link-layer transmission rates may be tuned to optimize such a utility function.

The rest of the paper is organized as follows. Section 2 describes a motivating example to establish the interplay between NC and rate diversity. Section 3 describes related work. In Section 4, we used simulation-based studies to study how variations in the level of NC and transmission rates affect the throughput. In Section 5, we analyze the corresponding impact of network coding on the latency observed by multiple receivers. In Section 6, we then consider a utility function that combines the sensitivity to both latency and throughput metrics, and observe the level of NC needed to optimize the resulting utility. Finally, Section 7 concludes the paper with our main conclusions and an enumeration of our current research directions.

II. MOTIVATION

We first use the classic 'butterfly' network example to understand the relative merits of rate diversity and network coding, and the potential gains that may accrue from a judicious combination of both. Consider the 5-node wireless topology in Figure 1. The links between the nodes are all 11Mbps, except for the link between node 1 and 2 which is 1 Mbps. Node 1 wants to broadcast packet A and node 2 wants to broadcast packet B. Note that, in a wireless environment, the links are not independent; for example, node 1 uses the same interface to simultaneously reach both neighbors 3 and 4. For simplicity, let us assume that each packet is of size 11Mbits.

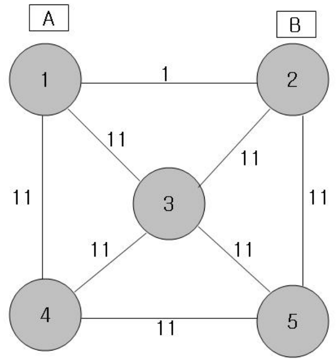

Fig. 1. Network Example

A pure routing-based and rate diversity-*unaware* strategy tries to schedule the dissemination of the broadcast packets so as to avoid collisions among contending links. It is easy to verify that an optimal transmission schedule consists of first having node 1 broadcast packet A to nodes 2, 3 and 4. To ensure that all these neighboring nodes are able to receive this packet, node 1 has to transmit with the lowest rate among

---

[2] The multi-source, network coding problem is part of our future work.

the links (1-2,1-3 and 1-4); in this case, 1 Mbps (link 1-2). Therefore, it spends 11 time units. Following that, node 2 transmits B to 1, 3 and 5. Node 3 subsequently transmits B to 4, following this with a transmission of packet A to 5. The scheduling solution is shown below. The total time unit is 24.

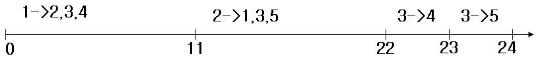

Fig. 2. Broadcast Scheduling

We can improve it by using Network Coding. In the usual Network Coding solution, Node 3 sends A XOR B to node 4 and 5, in a single transmission. Node 4 has A and can recover B from A XOR B. Node 5 has B and can recover A from A XOR B. The scheduling solution is shown below.The total time unit is 23 instead of 24.

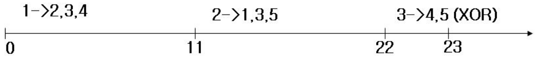

Fig. 3. Broadcast Scheduling

Let us now consider the rate diversity aware case, where different nodes may employ different rates for their broadcast transmissions. First, consider a *pure routing* based approach, where, as before, the network must schedule the transmissions to avoid contentions among interfering links. It is then easy to verify that the entire broadcast dissemination can be completed in 4 time units. Namely, 1 first transmits A to 3 and 4 (taking a total of 1 time unit). Node 2 then similarly transmits B to 3, 5 (taking an additional 1 time unit). Following this, node 3 broadcasts A to 2 and 5 at 11 Mbps, and follows up with a broadcast of packet B to 1 and 4, again at 11 Mbps.

Interesting enough, combining network coding with rate diversity can reduce the overall transmission latency even further. To illustrate this, consider the following network coding-based transmission strategy. Node 1 first sends the packet A to 3 and 4 using the 11 Mbps transmission rate (node 2 cannot receive at this high rate..). Next, Node 2 sends packet B to 3 and 5 using the faster rate (11Mbps). Then node 3 sends (at 11 Mbps) the XOR message to 4 and 5, and also to nodes 1 and 2. Node 1 will retrieve B by applying XOR(A, A XOR B), as it is already aware of its own packet A. An identical reasoning applies for node 2. Figure 4 illustrates the transmission schedule in this case. Note that the total time consumed by this combination is 3 time units.

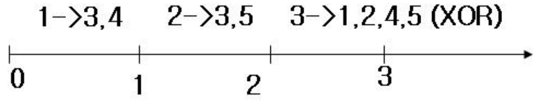

Fig. 4. Broadcast Scheduling

This canonical example serves to illustrate two important points. First, we have established that a combination of network coding and transmission rate diversity may prove to be mutually beneficial, resulting in an overall network throughput that is higher than that achievable by either strategy alone. Second, the example suggests that the gains from exploiting rate diversity (a reduction from 24 time units to 4) *may* be more spectacular than the gains accruing purely from network coding (a reduction from 24 time units to 23). Of course, we need to obtain the quantitative nature of the improvements in more practical, generalized topologies.

III. RELATED WORK

The research around network coding was motivated by the seminal paper [1], which demonstrated that, in general, the use of in-network encoding of packets could attain an optimal capacity that cannot be realized via any feasible routing-only scheme. For multicast traffic, the ‘capacity’ is defined as the maximum data rate that a sender can send to *all members* of a set of receivers. It is given by the minimum of maximum flow (s, t) between sender s and each receiver t. It was shown that network coding can achieve multicast capacity. Li, Yeung, and Cai [8] showed that it is sufficient for the encoding function to be linear. In addition to throughput, network coding offers additional benefits, such as robustness (by allowing nodes to receive potentially multiple copies of a single packet).

In wireless environments, network coding has been demonstrated to offer several benefits, such as improved energy efficiency [6](by reducing the number of distinct transmissions) and higher throughput. For unicast applications, [5] has recently demonstrated that the judicious use of network coding can improve the overall wireless network throughput. Random linear coding for multi-hop wireless multicast applications has been studied in [10], which showed how such randomized coding could improve the overall download latency for file-sharing applications. A linear programming formulation was used in [7] to compute the theoretically maximum throughput that may be achieved for a wireless multicast flow under network coding. However, all of these analysis do not consider the impact of transmission rate diversity at the link layer.

The use of such link-layer transmission rate diversity for broadcast and multicast routing was first explored in [4]. This paper introduced a rate-diversity aware broadcast tree construction heuristic, called WCDS, that was shown to reduce the broadcast latency (defined as the worst case dissemination delay of a packet to a group of receivers) by 3-5 times, compared to conventional diversity-unaware routing strategies.

To study the impact of multicast throughput, both with and without NC, we have implemented random linear coding over ODMRP [9], an existing multicast wireless routing protocol that uses a forwarding mesh (as opposed to a tree) for robust dissemination of multicast packets.

IV. INTERACTION OF NETWORK CODING AND RATE DIVERSITY AND IMPACT ON THROUGHPUT

In this section we show how the combination of network coding and rate diversity affects the performance of practical, distributed broadcasting protocols. Our simulations are performed using the Qualnet simulator. At the MAC layer, we used 802.11b. The rates ranged from 1Mbps (nodes at most 483.741m apart) to 11Mbps (283.554m). The network terrain had size 1000m x 1000m. The nodes were static.

In the experiments illustrated here, we have one source sending multicast packets for 120secs. The network size was 50 nodes. (We also performed the simulations with a network of 20 nodes and observed qualitatively similar results– accordingly, we only present the results for the larger-sized 50 node wireless network.)

To vary the "degree of network coding" used, we vary the number of consecutive packets that a node attempts to linearly combine before a single transmission. We refer to this "level of coding" with the symbol NC. Thus, the case $NC = 0$, corresponds to the case of pure-routing, where each received packet is forwarded immediately by an intermediate forwarding node. As the value of NC increases, the degree of network coding is higher, as the network nodes implicitly try to achieve greater bandwidth savings by encoding a proportionately higher number of packets. A large value of NC could, however, result in significantly high forwarding delay (as a node could wait indefinitely until it received an appropriate number of consecutive packets). To keep this delay bounded, a practical implementation uses a transmission timer, which is reset at every encoded transmission: in case of timer expiry, a node will immediately transmit by combining the number of packets currently available, even if this is smaller than NC. (To avoid complicating our studies unnecessarily, we do not vary the value of the transmission timer.)

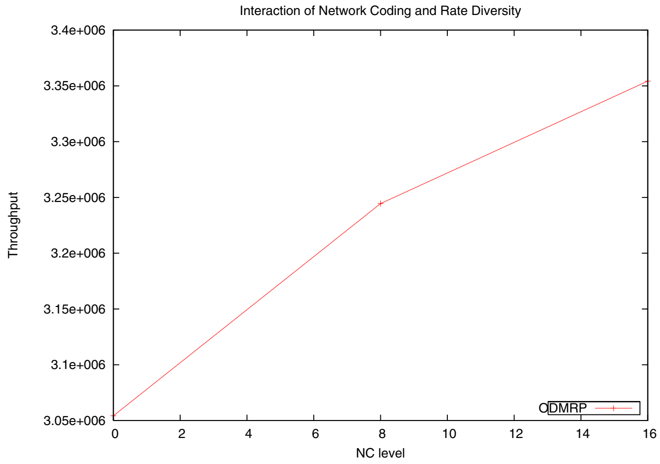

Fig. 5. Throughput vs. level of network coding.

Figure 5 shows the variation in throughput vs. level of network coding for ODMRP, without any transmission-rate diversity (the link-layer broadcasts all occur at a fixed rate of 11 Mbps). The throughput is plotted for $NC = \{0, 8, 16\}$. We can observe, as expected, that, the level of network coding increases, the throughput increases. The greater gain is observed when we move from a non-network coding situation (NC=0) to the NC level equals 8. Further increases in the 'degree of NC' do not provide as spectacular improvements in capacity, indicating that the benefits of network coding saturates at some modest value of NC.

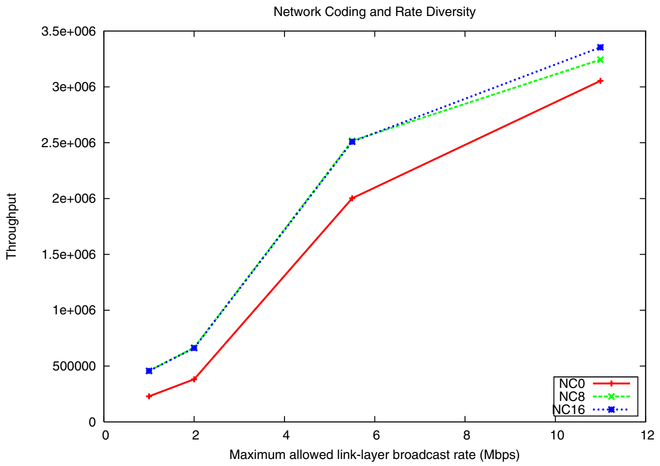

Fig. 6. Throughput vs. rate diversity

Figure 6 shows the throughput vs. rate diversity (maximum transmission rate allowed in Mbps) for different levels of NC. A rate value of 5.5 Mbps on the x axis implies that the nodes were allowed to transmit at all link rates, up to a *maximum* of 5.5 Mbps. The figure leads to several important observations:

- In all cases, as expected, the use of network coding (a higher value of NC) results in an increase in the achieved throughput.
- It appears to be clear that *incorporation of link-layer rate diversity* has a far greater impact than the use of network coding. For example, for ODMRP, increasing the transmission rate from 1 to 11 Mbps results in 10-fold increase in throughput (for $NC = 0$); however, increase NC to 16 only results in a two-fold increase in the throughput (at 1 Mbps).

Taken together, the observations suggest that the design of network coding strategies and link diversity-aware routing strategies *must be done jointly*. In particular, it appears that the gains from network coding can be somewhat modest for wireless multicast flows, unless the power of link layer rate diversity is adequately harnessed.

V. LATENCY VARIATION WITH NETWORK CODING

It is well known that network coding can result in higher *per-packet* delay, as the essence of the algorithm is to delay the transmission of already-arrived packets until additional packets have been collected and can be combined into a single broadcast transmission. However, it has also been shown that relatively file-insensitive applications, such as file transfer, can benefit significantly from network coding: the overall transfer latency of files (the time till the reception of the last packet) can be significantly reduced. It is also well known that per-packet latency and jitter are critical metrics for many real-time or interactive multimedia applications, such as VoIP or video conferencing. In such cases, it is apparent that network coding may not be a good choice.

There is, however, an interesting *an emerging class* of *quasi-real time* applications that lies between the extremes of file-sharing and VoIP/video conferencing, and that are *relevant to military battlefield environments.* Consider, for example, an image of a city street being disseminated by a camera sensor to multiple armored carriers and soldiers in an urban battlefield. Minimizing per-packet latency is not critical, as the sensor's data is useful only when the entire image can be reconstructed. However, there is an intrinsic value associated with the *freshness* of the data: for effective decision making, the sensor data must be disseminated rapidly, within a specific bound (most likely, in the order of tens of seconds, in this case).

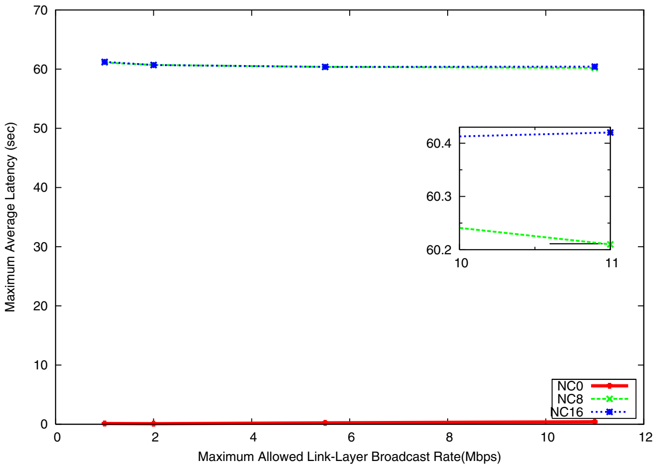

Fig. 7. Maximum Average Latency vs. rate vs. level of NC

In this section we thus study the impact of network coding on the packet-level latency observed for multicast flows, and the variation in this latency as the link-layer transmission rate is varied. Figure 7 shows the Maximum of the average latency (i.e., the maximum, among all 50 receivers, of their individual per-packet latencies) as a function of rate diversity (the maximum permitted link layer broadcast rate) for various NC level, when the dissemination is performed using ODMRP. Figure 8 shows the corresponding "Average (among nodes) of the average(among packets) latency" as a function of rate diversity and 'degree of network coding'.

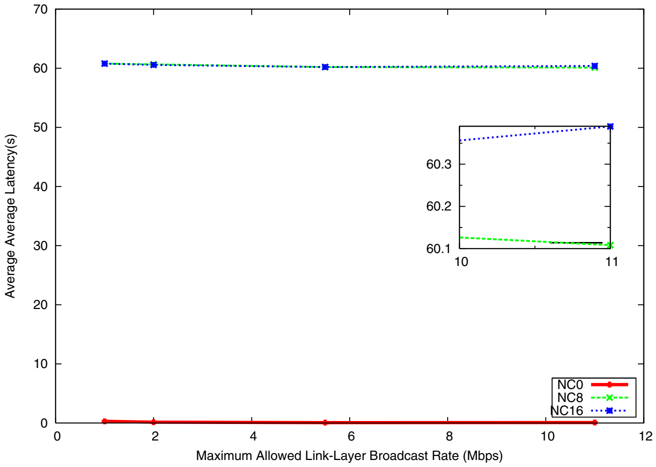

Fig. 8. Average of Average Latency vs. rate vs. level of NC

Several interesting observations emerge from the figure and table.

- It is clear that both the maximal and average latency values are *significantly higher* in the presence of network coding, compared to the pure-routing case (NC=0, which appears as a flat line close to the x axis). Moreover, when the level of NC is quite high, further increases in NC do not appreciably affect either latency metric (the lines $NC = 8$ and $NC = 16$ almost overlap with one another). For example, when NC increases from 8 to 16, we observe a very negligible increase (from 60.21s to 60.42s).
- While the latency observed decreases with an increase in the maximum permitted link-layer broadcast rate, the rate of decrease is minuscule. For example, for $NC = 16$, the maximum latency is 61.22sec. at 1 Mbps, and only decreases to 60.42sec. when the transmission rate is increased to 11 Mbps.

The above results suggest that a minimal level of network coding might indeed be beneficial for this new category of 'quasi-real time' applications, as small values of NC appear to provide observable gains in throughput, without resulting in unacceptable high dissemination latency. However, network coding must be combined with proper use of link-layer rate diversity, to achieve significantly higher throughput; the observed latency is, however, fairly insensitive to the use of such rate diversity.

VI. UTILITY FUNCTION-BASED BALANCING OF THROUGHPUT AND LATENCY OBJECTIVES

We observed that Network Coding (as applied to the ODMRP multicast dissemination protocol) maximizes throughput but suffers from significantly higher latency. For the 'quasi real-time' applications under consideration, we want to ensure that we can appropriately bound both these metrics. To express these twin, and conflicting, objectives can be addressed by an appropriate level of network coding, we first define a very simple utility function $U(w)$, where

$$U(w) = w.(1/L) + (1 - w).T \tag{1}$$

In Equation1, L represents the average latency across all nodes, T means the throughput, and w is a weight value that varies according to the application requirements.

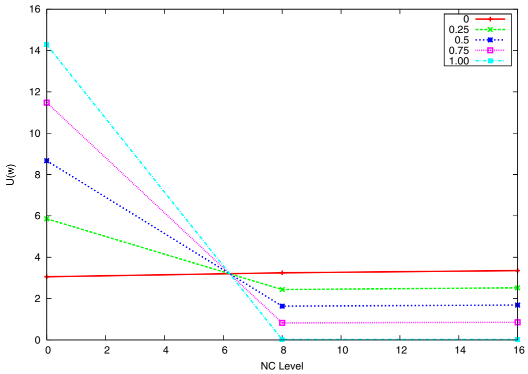

Fig. 9. Utility function vs. rate vs. level of NC

We can then study how the choice of NC affects this utility function $U(w)$ and the appropriate level of NC that proves to be optimal for a given $U(w)$. Figure shows the utility function

$U(w)$ for different values of w with different levels of NC. The values for latency (in seconds) and throughput(in Mbps) were taken from the prior experiments with ODMRP using an 11Mbps maximum rate. As expected, we observe that there appears to be an 'optimal level' of NC for a given $U(w)$: while a larger NC results in a significant rise in the packet latency without corresponding benefits in throughput, a smaller NC reduces the latency but also suffers from a more dramatic reduction in the throughput. More importantly, for our chosen scenarios, it appears that the choice of the optimal NC was reasonably insensitive to the choice of w: $NC = 8$ seemed to perform well across all parametric variations.

These observations indicate that the appropriate level of NC may be a function of the network's operating parameters (such as the node density and the number of multicast receivers/group), but may be relatively constant across a broad category of 'quasi real-time' applications. We plan to investigate these dependencies in future work, with a goal of developing a set of design guidelines on the use of appropriate levels of network coding in various military environments.

VII. CONCLUSION

We have demonstrated in this paper that multi-rate link layer broadcasts and network coding can be mutually combined to increase network throughput in multicast applications (at least for the case of a single broadcast or multicast flow). Simulation studies conducted using a practical routing protocol, ODMRP, suggest that the impact of rate diversity on the overall achievable throughput is larger than the independent use of network coding. Accordingly, as part of our ongoing work, we are evaluating how a rate-aware multicast routing protocol, such as WCDS, may be modified and enhanced to take advantage of network coding.

We have also observed through simulations that NC introduces a 'throughput-latency' tradeoff, not just for worst case latency but even in terms of the average delay experienced by packets. We believe that this observation has important implications for the use of network coding in many 'quasi' latency-sensitive broadcast applications, such as the dissemination of sensor data to mission commanders for situational awareness in a battlefield network. Our studies show that the gains achieved by network coding in throughput (the savings in the number of distinct transmissions) often come at the expense of significantly higher latency. This suggests that, for any given application, there exists an 'optimal' level of network coding, such that a greater amount of random linear combination results in an unacceptable degradation of the application latency.

For future work, we plan to further investigate this 'throughput-latency' trade-off. In particular, as we consider multiple simultaneous multicast flows, we shall investigate the notion of *adaptive network coding*, where the level of network coding is increased gradually, in responsive to progressive increase in the wireless network congestion load.

REFERENCES

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow", IEEE Trans. Inform. Theory, 46(4):1204 1216, July 2000.

[2] U. Lee, J-S. Park, J. Yeh, G. Pau and M. Gerla, "CodeTorrent: Content Distribution using Network Coding in VANETs", 1st International Workshop on Decentralized Resource Sharing in Mobile Computing and Networking (MobiShare'06), September 2006.

[3] J. Bicket, D. Aguayo, S. Biswas and R. Morris, "Architecture and Evaluation of an Unplanned 802.11b Mesh Network", Proc. of MOBICOM 2005.

[4] C. T. Chou, A. Misra, J. Qadir, "Low Latency Broadcast in Multi-Rate Wireless Mesh Networks", IEEE JSAC Special Issue on 'Multi-hop Wireless Mesh Networks', Vol. 24(11), November 2006.

[5] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard and J. Crowcroft, "XORs in The Air: Practical Wireless Network Coding", Proc. of ACM SIGCOMM 2006, September 2006.

[6] S. Deb, et al,"Network Coding for Wireless Applications: a Brief Tutorial", International Workshop on Wireless Ad-hoc Networks (IWWAN), May 2005.

[7] J-S. Park, D. Lun, F. Soldo, M. Gerla and M. Medard, "Performance of Network Coding in Ad Hoc Networks", Proc. of IEEE Milcom 2006, October 2006.

[8] S-Y.R. Li, R. W. Yeung, and N. Cai, "Linear network coding", IEEE Trans. Information Theory, Vol. 49(2), Feb 2003.

[9] S.-J. Lee, M. Gerla and C.-C. Chiang, On-demand multicast routing protocol, in: Proceedings of IEEE WCNC'99, New Orleans, LA (September 1999) pp. 1298-1302.

[10] J.S. Park, D. S. Lun, Y. Yi, M. Gerla and M. Mdard, "CodeCast: Network Coding based Ad hoc Multicast Protocol", IEEE Wireless Communications, Vol. 13(5), 2006.