CAREER: Breaking the Barriers in Wireless Network Information Theory: A Deterministic Approach

1 Proposal Summary

Wireless communication is one of the most vibrant areas in the communication field today. Over the past decade, we have witnessed quite a few successful solutions in the wireless industry, for example second-generation (2G) and third-generation (3G) digital wireless standards with more than half a billion subscribers worldwide. Quite amazingly, Shannon's theory for point-to-point communication systems is the scientific foundation upon which the designs of all of these technologies are based.

Extending Shannon's theory to network setting has been one of the greatest challenges in information theory over the past few decades. It is expected that progress in this area will produce a significant breakthrough in the design of distributed wireless networks of the future, such as ad-hoc and sensor networks. So far, most research efforts have approached this problem using the same generality and accuracy used by Shannon for point-to-point systems. However, meeting such a standard has proven to be extremely difficult, to the extent that the capacity of most basic networks is still unknown.

In this proposal, we present a *new approach* to wireless network information theory. Our methodology is to seek *approximate solutions* accompanied by guarantees on the gap to optimality. We believe that much *broader progress* can be made if we change the focus. Not only we can make progress in approximating the capacity, we can also study the information-theoretic impact of other fundamental practical limitations that have often been ignored in information theory. At the heart of our approach is our development of simple, *deterministic* channel models that capture the main features of the wireless medium, and are utilized to approximate more complex models.

The proposed research will build on the PI's recent work, which demonstrates the power of this approach to approximate the capacity of general Gaussian relay networks. While the capacity of relay networks with even one relay has remained unsolved for 40 years, using the deterministic approach the PI *has been able to approximate the capacity of relay networks* (with an arbitrary number of relays) within a constant number of bits, independent of the channel parameters.

Intellectual Merit. Information theory is well poised to tremendously impact the design of future communication systems. But, the mathematical difficulty of solving network information theory problems has been a barrier to progress. In this research, we propose a new approach, based on using simpler, deterministic models for wireless channels, to find approximate solutions to these problems. The proposed research is divided into *four areas* to utilize this powerful approach to **1)** approximate the capacity of wireless networks; **2)** study the impact of incomplete network knowledge on the capacity; **3)** design protocols that are robust against adversaries; and **4)** utilize recent physical layer coding techniques for cross-layer protocol design.

Broader Impacts. The proposed research will result in simpler abstractions of the wireless medium that can be used to *bridge* the gap between research in the communication society and the networking society. The proposed educational activities are designed to integrate these ideas in *teaching communications and information theory courses* to make them more accessible to students in the networking area as well as students from the Computer Science Department. The proposed educational activities also include a significant outreach program to women and under-represented minorities in the form of presenting *tutorials* and participating in *workshops*. The proposed outreach plan also aims to raise high school students' interest in engineering by *speaking in classrooms* and designing an *inspiring weblog* to disseminate the discoveries in the field to the broader communities.

2 Career Development Plan: Research Plan

2.1 Motivation

Shannon's masterpiece, A Mathematical Theory of Communication (1948) [1], is one of the most influential works in the history of communications. In this seminal work, Shannon characterized the fundamental limit for reliable communication over a point-to-point communication channel, the so-called *channel capacity*, and provided the architectural system design for achieving it.

Extending this theory to the network setting has been one of the greatest challenges in information theory in the last four decades. This problem became significant as the result of the emergence of new wireless technologies, such as wireless ad-hoc and sensor networks, in which no centralized infrastructure is deployed to manage the communications. Inspired by the power of Shannon's result for point-to-point systems, most prior work has naturally focused on characterizing network capacity, with the hope that this characterization would reveal the optimal architectural system design for communication over these networks. However, after nearly 40 years, we are still facing the following issues associated with this challenging situation:

- *Characterizing network capacity is still an open problem, even for very basic networks.*

  In the last four decades, there have been many significant efforts to characterize the capacity region of some canonical examples in network information theory. Two of the most influential results are the multiple access channel capacity region [2, 3] and the degraded broadcast channel capacity region [4, 5]. However, most other problems related to characterizing the network capacity are notoriously difficult and remain unsolved. Even when we focus only on an important class of linear Gaussian channels, the problem remains quite challenging.

- *The difficulty of characterizing network capacity is a barrier that has prevented our addressing the fundamental limitations that arise due to the distributed nature of wireless networks.*

  One of the main features of the distributed wireless networks is that the nodes must operate in a distributed fashion and implement transmission strategies with local information about the network. These are important practical constraints that need to be considered carefully when extending Shannon's theory to networks. However, since the problem of characterizing network capacity remains unsolved even if we ignore these limitations, these fundamental constraints have often been ignored in network information theory and their impacts on network capacity have not been addressed properly.

- *Most recent discoveries in wireless network information theory are too abstract and not accessible to other communities.*

  Among those who study wireless network information theory, most have devoted their efforts to the study of the fundamental limits of communication over the wireless medium and to the development of physical layer communication techniques to more efficiently exploit this medium. Despite much progress in information theory, most research in other communities is still performed on the basis of overly-simplified physical layer models; nodes can communicate at a prescribed rate if the signal to interference plus noise ratio (SINR) is above a certain threshold. This is mainly due to the fact that, other than the "bit-pipe" abstraction of the physical layer, we have not been able to propose new abstractions that can directly export our physical layer advances for use at higher layers.

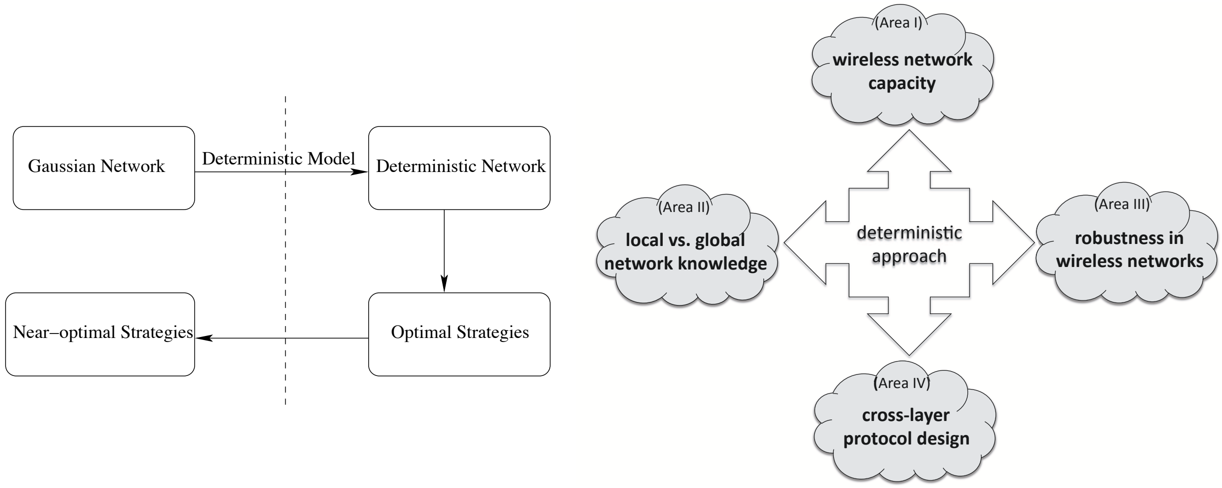

(a) (b)

Figure 1: Illustration of (a) the deterministic approach and (b) the proposed research areas.

In summary, although information theory is well poised to tremendously impact the design of future communication systems, the difficulty of solving network capacity problems has been the main barrier to making further progress, addressing practical challenges, and impacting other communities. In this proposal we present an innovative approach to wireless network information theory that we believe will enable significant progress in addressing all three of the aforementioned challenges.

2.2 Overview of the proposal

To date, there have been many efforts to solve network information theory problems with the same generality and accuracy used by Shannon for point-to-point systems. However, meeting such a standard has proven to be very difficult. We advocate that the exact characterization of capacity in the most general form is not essential for impacting real-world problems, and we believe that much broader progress can be made if we seek *approximate solutions* with a guarantee on the gap between these solutions and optimality.

Approximation methods are commonly used in all fields of science, mathematics, and engineering to tackle difficult problems. Approximation techniques are also used to reveal insights that are maybe opaque in general and precise solutions. For example, analyzing the motion of several planets orbiting a star (i.e. n-body problem [6]) is one of the classical problems in astronomy, which is still unsolved. However, in 1680s, Newton was able to use perturbation methods to approximate the orbit of the Moon [7] and prove that it moves noticeably differently from a simple Keplerian ellipse [8] because of the competing gravitation of the Earth and the Sun.

We also aim to use approximation methods to break the barriers in wireless network information theory. One of the main reasons behind the difficulty of research in wireless network information theory is the complexity of wireless medium. In communications, the linear additive Gaussian channel model is commonly used to capture the fundamental features of a wireless channel. However, solving network information theory problems is quite difficult, even within the realm of Gaussian models.

We present a new deterministic channel model that is analytically simpler than the Gaussian model but still captures the main features of the wireless medium, such as broadcast and superposition. The motivation to study such a model is that, in contrast to fixed, point-to-point channels where noise is the only source of uncertainty, the signal interactions are also a critical source of uncertainty in multiuser networks. Therefore, for a first level of understanding, our focus will be on such signal interactions rather than the received noise.

Based on our deterministic model, *we propose a unified, systematic approximation framework in wireless network information theory.* As illustrated in Figure 1 (a), our approach has three steps; first, we replace the Gaussian model with the deterministic model which is more tractable. Next, we analyze the resulting deterministic network to find optimal strategies. Finally, we use the insights obtained from this analysis to obtain near-optimal strategies for the original Gaussian problem.

The proposed approach will have multiple advantages. By changing the focus, we hope to make progress in approximating the capacity, study the impact of commonly ignored practical limitations in wireless networks, and bridge the gap between research in the communication society and the networking society. As illustrated in Figure 1 (b), *the proposed research-plan develops a unified framework to effectively utilize the deterministic approach to make significant progress in the following four research areas:*

- **Area I:** The development of a systematic framework for approximating the capacity of wireless networks
- **Area II:** The determination of the effect of incomplete network knowledge on the capacity
- **Area III:** The design of communication protocols that are robust against malicious adversaries
- **Area IV:** The utilization of physical-layer coding techniques for designing protocols at higher network layers.

We believe that research in these areas will enable us to make major contributions in addressing the three main challenges in wireless network information theory that were discussed in Section 2.1.

The proposed research will also be integrated with educational activities that complement it and further its goals. The principal investigator (PI) plans to revise two courses in communications and a course in network information theory at Cornell. The first, more basic course will be on digital communications with emphasis on using simple concrete examples to teach the key principles. The second, more advanced course will be focused on wireless communication systems, with the goal of presenting a comprehensive understanding of recent advances in wireless communications and discussing the key concepts behind each one of them. In addition to these two courses in communications, a network information theory course will also be designed. Because of our emphasis on the use of deterministic models, we also hope to make this course accessible to students in the networking area as well as students from the Computer Science Department. The proposed educational activities also include a significant outreach program to women and under-represented minorities in the form of presenting *tutorials* and participating in *workshops*. The proposed outreach program also aims to raise high school students' interest in engineering and math by *speaking in classrooms* and designing an *inspiring weblog* to illustrate the discoveries in the field.

The proposal is organized as follows. In Section 2.3, we describe the deterministic model and illustrate how it captures the main features of the wireless medium. This information provides the foundation for the material that will be covered in the later sections. In Sections 2.4-2.7 we discuss the four research areas that we plan to study by using our deterministic approach. Finally, in Section 3, we propose the educational activities that complement the research component of this proposal.

2.3 Description of the Deterministic Model

There are three main distinguishing features of wireless medium: channel strength, broadcast, and superposition. Because of these effects, links in a wireless network are never isolated; instead they

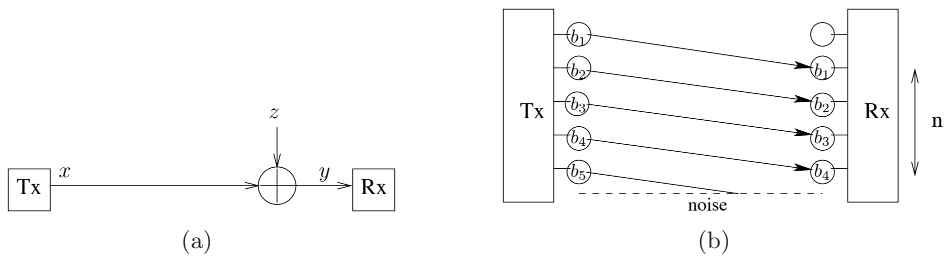

Figure 2: (a) Gaussian model and (b) deterministic model for point-to-point channel.

interact in seemingly complex ways. In this section, we illustrate how to capture all of these effects deterministically, and we present our linear deterministic channel model.

2.3.1 Modeling signal strength

Consider a point-to-point Gaussian channel with a signal-to-noise ratio denoted by SNR, shown in Figure 2 (a). Consider the binary-expansion of the real-valued input of the Gaussian channel:

$$x = 0.b_1 b_2 b_3 b_4 b_5 \dots .$$

Intuitively speaking, this sequence of bits goes through the channel and some of them are received above the receiver's noise level and some are received below. One way to capture the effect of noise deterministically is by truncating all bits that are received below noise level. Therefore, we can think of a transmitted signal 'x' as a sequence of bits at different "signal levels," with the highest signal level in 'x' being the most significant bit (MSB) and the lowest level being the least significant bit (LSB). In the deterministic channel model, shown in Figure 2 (b), the receiver only gets the 'n' most significant bits of 'x'. The correspondence between the number of bits received in the deterministic channel, 'n', and the SNR in the complex Gaussian channel is $n \leftrightarrow \lceil \log \mathsf{SNR} \rceil^+$.

2.3.2 Modeling superposition

Consider the Gaussian multiple access channel (MAC), shown in Figure 3 (a). The deterministic model for this channel is constructed similarly to the point-to-point channel (Figure 3 (b)), with n_1 and n_2 bits received above the noise level from users 1 and 2, respectively. To model the superposition of signals at the receiver, the bits received on each level are added modulo two. Addition modulo two, rather than normal integer addition, is chosen to make the model more tractable. As a result, the levels do not interact with one another. This way of modeling interaction is in some sense similar to the collision model. In the collision model, if two packets arrive simultaneously at a receiver, both are dropped; similarly, here, if two bits arrive simultaneously at the same signal level, the receiver gets only their modulo two sum, which means it cannot figure out any of them. However, in this case, the most significant bits of the stronger user remain intact, contrary to the simplistic collision model in which the entire packet is lost when there is collision. This is reminiscent of the familiar capture phenomenon in CDMA systems, i.e., the strongest user can be heard even during simultaneous transmissions by multiple users.

To demonstrate how well our deterministic MAC approximates the Gaussian MAC, we can compare their capacity region. The capacity region of the Gaussian MAC (Figure 3 (a)) is,

$$\begin{aligned} R_i &\leq \log(1+\mathsf{SNR}_i), \quad i=1,2 \\ R_1 + R_2 &\leq \log(1+\mathsf{SNR}_1+\mathsf{SNR}_2) \end{aligned} \tag{1}$$

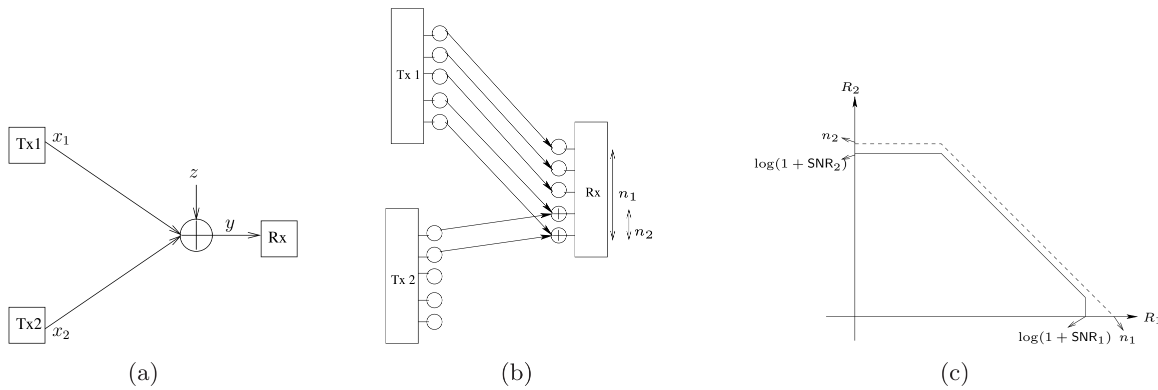

Figure 3: (a) Gaussian model for MAC, (b) deterministic model for MAC, and (c) capacity region of Gaussian MAC (solid) and deterministic MAC (dashed).

It is easy to show that the capacity region of the deterministic MAC (Figure 3 (b)) is

$$\begin{aligned} R_2 &\leq n_2 \\ R_1 + R_2 &\leq n_1 \end{aligned} \tag{2}$$

where $n_i = \lceil \log \mathsf{SNR}_i \rceil^+$ for $i = 1, 2$. The capacity region of the Gaussian MAC and its corresponding deterministic MAC are plotted in Figure 3 (c). As we notice the capacity regions of these two channels are very close to each other. In fact, it is possible to show that, for all values of SNR_1 and SNR_2 they are within one bit per user of each other [9]. Hence, the deterministic MAC uniformly approximates the Gaussian MAC.

2.3.3 Modeling broadcast

Based on the intuition obtained so far, it is straightforward to think of a deterministic model for the broadcast scenario. Consider a single source broadcasting to two receivers shown in Figure 4 (a). Figure 4 (b) shows the deterministic model for the Gaussian broadcast channel (BC). The user with the stronger channel observes n_1 significant bits from the input. The user with the weaker channel observes only n_2 of the most signicant bits ($n_2 < n_1$).

The capacity region of the deterministic broadcast channel is identical to that of the MAC (2), and it can also be shown to approximate the Gaussian broadcast capacity region to within 1 bit per user, for all channel gains.

2.4 Area I: Developing a systematic framework for approximating the capacity of wireless networks

This area focuses on exploiting the deterministic model as a general capacity approximating tool in wireless network information theory. This work will be built on the results obtained by the PI in the constant-gap approximation of the capacity of Gaussian relay networks [9, 10].

As illustrated in Figure 1 (a), the main steps in our deterministic approach are: 1) replacing the Gaussian channel model with the linear finite-field deterministic model, 2) analyzing the capacity of the resulting deterministic network, and 3) using the insights obtained from this analysis to approximate the capacity of the original Gaussian network. Our approximation of interest sandwiches the capacity in such a way that the approximation error does not depend on network channel gains

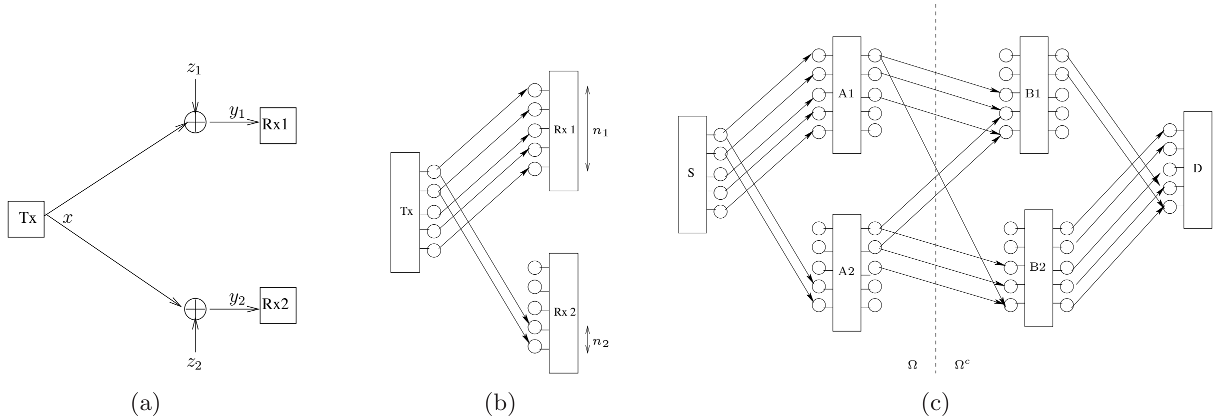

Figure 4: (a) Gaussian model for BC, (b) deterministic model for BC, and (c) an example of a deterministic relay network.

and the signal-to-noise ratios of operation. In this sense, we seek a "uniform" approximation of the capacity.

Using this approach, we recently made considerable progress on single-source single-destination networks (relay networks) [9–13]. Here, one source wants to communicate with a single destination with the help of other nodes in the network, which are called relays. This is a fundamental model to study the role of *cooperation* in networks. The basic paper on this topic is by Cover and El Gamal [14], who focused on the case that there is only one relay in the network and invented a rich set of strategies such as decode-and-forward and compress-forward, which were later generalized to networks with multiple relays, see e.g. [15–24]. Though there have been many interesting and important ideas developed in the above-mentioned papers and many other works in the literature, the capacity characterization of Gaussian relay networks is still unresolved, even in the case in which there is only one relay. In fact, even a performance guarantee, such as establishing how far these schemes are from an upper bound, is unknown, and, hence, the approximation guarantees for these schemes are unclear.

Following our methodology, we first created a deterministic model for an arbitrary relay network [12] (Figure 4 (c)). We exactly characterized the capacity of this deterministic network with an arbitrary number of relays and interconnections [13]. The general capacity result for this class of deterministic networks is:

$$C = \min_{\Omega} rank(G_{\Omega}) \tag{3}$$

where Ω is an arbitrary cut separating the source from the destination (see Figure 4 (c)), and G_{Ω} is the transfer matrix between the vector of all the inputs at the nodes in Ω and the vector of all the outputs at the nodes in Ω^c. In the wireline model, the rank of the transfer matrix is just the number of unit capacity links crossing the cut, so the general capacity result is a natural generalization of the classic Ford-Fulkerson maxow-mincut theorem. Borrowing ideas from network coding [25–30], we found an achievable scheme, which uses random linear codes at the relays. The challenge here is that, due to the superposition property of our model, the signals coming along different routes with different delays mix with each other. We also generalize the argument to the multicast case, in which a single source sends the same information to multiple destinations.

By a suitable modification of the achievable scheme for the deterministic channel, we were able to construct a scheme that can achieve within a constant gap from the cut-set bound for the Gaussian relay network [9, 11]. For example, in the canonical relay channel with only one relay, the gap is at most 1-bit. To our knowledge, *this is the first result that provides a hard performance guarantee for general relay networks*.

While this result demonstrates the power of our approach in relay networks, we believe that this approach can also be applied to many other important capacity problems in wireless network information theory. It can be effectively used to find outer bound limitations and find new achievable strategies in such a way that there is a uniform, constant gap between them. In the last year, the PI has applied this idea to several canonical networks with multiple sources and multiple destinations, and has successfully characterized a constant-gap approximation of their capacity [31–34]. Our initial results show promise that this approach can potentially turn into a general approximation method in wireless network information theory. To accomplish this goal, we foresee several challenges ahead, and the goal of the proposed research in this area is to address these challenges. More specifically, our objectives are:

Main Objectives

- *Make operational connections between the Gaussian model and the deterministic model*

One of the key steps in the deterministic approach, illustrated in Figure 1 (a), is the last step in which we translate the exact capacity results for the deterministic network into approximation results for the corresponding Gaussian network. So far, we have done this by using the insights and techniques that were obtained while solving the deterministic problem, to develop a separate proof for the corresponding Gaussian network. Hence, it is more of a trial-and-error procedure than a formal procedure. In order to make this more systematic, we need to establish operational connections between the models.

Our goal is *to find a systematic translation of the coding schemes between the models*. In particular, we seek to use an appropriate internal code in the Gaussian network to create the signal levels and also enable the nodes to actually decode the superimposed signals at each signal level, as in the deterministic model. In Gaussian multiple-access channels, lattice codes have already shown to be a powerful tool for decoding the superposition of the received signals [30]. Therefore, it is expected that structured codes, such as lattice codes, will play a significant role in accomplishing our goal.

- *Capture more features of the wireless channel in the deterministic model*

So far, we have focused on channel strength, superposition, and broadcast aspects of wireless channel, and developed a deterministic channel model that captures them. However, to obtain more accurate approximation results, there are several other important features of the wireless medium that must be considered. For example, power gain is one of these aspects that, if modeled properly, can significantly affect the approximation error. To understand this, let's consider a K-user multiple access channel. The sum-rate capacity of this channel under the Gaussian model is known to be

$$R^G_{\text{sum}} = \log\left(1 + \sum_{i=1}^{K} \mathsf{SNR}_i\right) \tag{4}$$

Similar to the two-user case, the sum-rate capacity of the corresponding deterministic channel is:

$$\begin{aligned} R^D_{\text{sum}} &= \max\{n_1, \ldots, n_K\} \\ &= \max\{\lceil \log \mathsf{SNR}_1 \rceil^+, \ldots, \lceil \log \mathsf{SNR}_K \rceil^+\} \end{aligned} \tag{5}$$

By comparing (4) and (5), we note that the worst-case gap between the sum-rates is $\log K$ bits. which happens when all channel gains are the same. This is due to the effect of power gain in wireless channels, which is not captured with finite field operations.

Similarly, the finite field operations of our current mode do not allow proper modeling of the ill-conditioned, multiple-input, multiple-output (MIMO) channels (*i.e.* with determinant close to zero). The development of more accurate, yet simple, models that properly capture these features will allow us to improve our approximation results and apply our approach to other important problems in wireless network information theory.

- *Study the applicability of deterministic approach to other unsolved capacity problems*

As we discussed earlier, our approach yields a good approximation of the capacity of single-sender single-destination networks. The PI has also applied this idea to several canonical multi-source multi-destination examples and has successfully characterized a constant-gap approximation of their capacity [31–34]. While using this approach to make progress in other unsolved problems in wireless network information theory is a promising direction that we plan to work on, an equally important direction is to understand the extent to which we can approximate the capacity of wireless networks by replacing the Gaussian model with a deterministic model.

A potential approach to this problem is by viewing a capacity problem for an acyclic wireless networks as an optimization problem that maximizes the end-to-end mutual information over the (unknown) entropic region [35], subject to complex network constraints. The effect of replacing the Gaussian model with a deterministic model can now be viewed as relaxing the original optimization problem by simplifying the network constraints. So, alternatively, we seek to understand the gap between these two optimization problems. It is anticipated that duality and optimization theory will play significant roles in answering this question.

2.5 Area II: Determining the effect of incomplete network knowledge on the capacity

One of the main practical considerations for decentralized wireless networks is that the nodes have to operate in a distributed fashion with local information about the network. Since characterizing the network capacity is quite difficult, even in the case for which global information is known everywhere, these important issues have often been ignored in information theory. In fact, most works in network information theory implicitly assume that all information about the network (such as network topology and channel gains) is known everywhere, and they analyze capacity only when the communication strategy can be centrally designed. In the PI's view, *this is one of the most overlooked practical considerations in network information theory*, and addressing these issues can substantially widen the applicability of this field in practice.

Our goal in this research area is to develop a framework to study the impact of these practical limitations on network capacity and design distributed algorithms that get close to the optimal performance. We hope that the simplicity of the deterministic model will allow us to avoid the difficulty of analyzing Gaussian channels directly and therefore facilitate substantial progress.

Recently, we took some preliminary steps in this direction [36]. We considered a deterministic network in which the network connectivity (or network topology) and the channel gains at each link are initially not known to any node in the network. We used a simple message passing algorithm that was proposed in [37], so that the nodes can learn some local information about the network. The learning algorithm is as follows: we first use training sequences so that each receiver learns the incoming channel gains. Then, the algorithm proceeds in rounds, where one round consists of a round of messages sent from all receivers, followed by a round of messages sent by all transmitters. In each round, the nodes potentially learn a little bit more about the rest of the network and pass the new information to other nodes. In particular, after the ℓ^{th} round, each node knows the gain of all channels that are at most $(2\ell + 1)$-hop away.

As nodes gather more information, it will potentially be possible to implement better strategies and achieve higher rates. Hence, as shown conceptually in Figure 5, the network capacity region monotonically increases with ℓ until a point ℓ^*, at which the nodes can perform as if they knew everything about the network.

Based on this setup, we can now formulate and study several key problems to help us understand the effect of incomplete, distributed network information on network capacity.

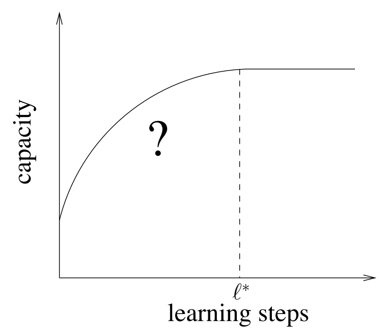

Figure 5: The effect of learning on network capacity.

One basic problem is to *determine the minimum number of learning rounds, ℓ^* , required to guarantee that there is no loss in the capacity.*

Solving this problem would be a major step to understand whether global network information is always necessary to achieve the maximum throughput, and if not, how much information would be sufficient. In a recent work [36], we were able to answer this question for the three-user interference channel shown in Figure 6 (a). We determined ℓ^* for each network connectivity and showed that, quite surprisingly, for a substantial fraction of all possible network connectivities (14 out of 64), just one round of learning is sufficient to be sum-rate optimal. An interesting open problem is to understand this phenomenon in more general networks and determine the minimum learning effort required to guarantee that there is no loss in the capacity.

In practice, completing each round of a learning algorithm requires a considerable portion of the total available resources and can not be ignored [38, 39]. Therefore, it is critical to account for learning overhead in defining the capacity. To do so we aim to study a more general problem that is to *characterize the network capacity as a function of the number of learning rounds, i.e.* $C(\ell)$ (see Figure 5). By solving this problem we can characterize the effect of incomplete network knowledge on network capacity and develop a unifying analytical framework to account for learning overhead.

Finally, we hope to use the tools that will be developed from research in Area I to map the outcomes of studying these problems into good approximation results under the Gaussian model.

To summarize, Area II focuses on addressing the fundamental limitations that arise due to the lack of global network knowledge in distributed wireless networks. The proposed research seeks to shed light on several important problems in this area, more specifically the objectives are:

- determining the minimum network knowledge required to guarantee the optimal performance;
- quantifying the effect of learning on network capacity;
- developing cooperative strategies in the form of exchanging network information among the users and quantifying their benefits; and
- developing distributed communication protocols that operate with incomplete network information at the users.

2.6 Area III: Designing communication protocols that are robust against malicious adversaries

An important concern with network coding in distributed networks is that a single corrupted packet injected by a malicious user might end up contaminating all the information reaching a destination, thereby preventing decoding. Therefore, distributed networks are more susceptible to Byzantine (i.e., arbitrary) attacks from compromised nodes. In this area, we focus on studying the fundamental limits of correcting adversarial errors in wireless networks and developing distributed communication strategies that are robust to such errors.

To be more specific, we can formulate the problem as following: there is a source, Alice, who communicates over a network to a receiver Bob. There is also an attacker Calvin, hidden somewhere in the network. Calvin aims to minimize the flow of information from Alice to Bob. Calvin knows the communication strategy of all nodes, observes some or all of the transmissions, and can jam z links to inject his own transmission. The questions is: what is the maximum rate of robust and reliable communication from Alice to Bob?

For wireline networks, there has been significant progress towards developing a solution for this problem over the last 4 years. The information-theoretic bounds on the throughput have been studied in [40], and several error correction strategies have been proposed in [41–44]. In particular, it has been shown that when all links have unit capacity and the adversary can covertly jam z links, the information flow from Alice to Bob is at most $C - 2z$, where C is the network capacity without any adversaries. Also, distributed *linear network coding* strategies have been proposed to achieve this optimal performance [45, 46].

Wireless networks, on the other hand, are much more vulnerable to Byzantine attacks and provide a less robust communication medium. In fact, the broadcast, superposition, and channel-strength variations in wireless channels give the adversary much more power to attack and even inject errors. All of these factors make the problem much more challenging for wireless networks.

Recently, we have made some preliminary progress on this problem [47]. As the first step, we focused on the effect of non-equal link capacities in the network and considered a basic scenario in which all communication links are noiseless, point-to-point channels with non-equal capacities. Although this problem is just a slight generalization of that all equal capacity links, it turns out that it is much more challenging. This is intuitively due to the fact that we should now design the protocols much more carefully to minimize the dissemination of the burst of errors injected by the adversary over the high-capacity links.

We have shown that, quite surprisingly, *linear network coding strategies are sub-optimal* in this case, and much more gain can be obtained if, instead, the internal nodes attempt to detect the erroneous packets (which is a non-linear operation) before encoding and forwarding them [47]. We have also been able to find a class of such strategies that achieves the optimal performance in some canonical networks. An interesting open problem is to extend these results and find the optimal strategies and characterize the capacity in more general networks.

So far, we have done preliminary research on the effect of channel strength variation in robust network coding. It is expected that the problem will be much more difficult when the broadcast and superposition constraints of wireless medium are also considered. We hope to utilize the deterministic channel model effectively to overcome these challenges and understand the fundamental limits of error correction in wireless networks. In particular, we aim to focus on single-source singe-destination networks and study the following important problems:

- characterizing the limits of error correction in deterministic networks;
- developing error correction schemes that achieve the optimal performance; and

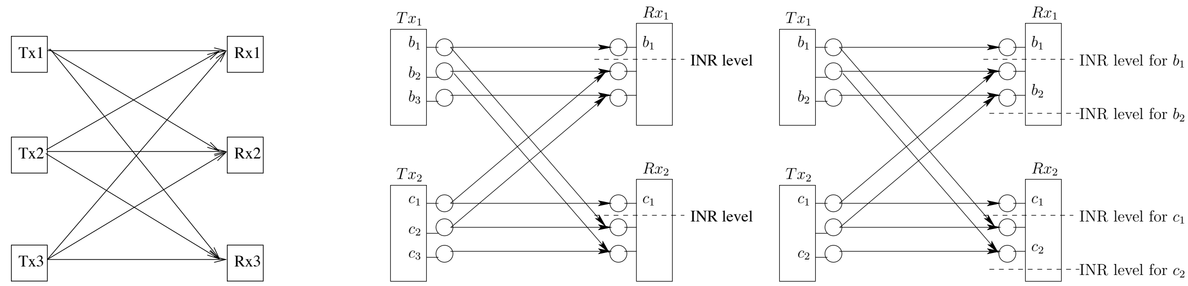

(a) (b) (c)

Figure 6: (a) Three user interference channel, (b) achieving sum-rate of 2 by treating interference as noise, and (c) achieving sum-rate of 4 by a more careful assignment of signal levels.

- translating the results back to Gaussian networks.

2.7 Area IV: Utilizing physical-layer coding techniques for protocol design at higher network layers

Protocol design for wireless networks is an active research area in several fields, including communications and networking. In the communication society, most efforts have been on developing physical layer communication techniques that achieve a throughput near the fundamental limits of the physical channel. On the other hand, other fields typically use the simple "bit-pipe" abstraction of the physical layer and focus more on other aspects of protocol design, such as power allocation, scheduling and delay.

It is well understood that the simple bit-pipe mode is inadequate and does not reflect all degrees of freedom available at the physical layer. Hence, there has been an increasing interest in cross-layer approaches to protocol design (*i.e.*, approaches that jointly consider the physical layer and higher network layers), see e.g. [48–51]. These cross-layer approaches shift the research landscape away from optimizing the performance of individual layers and, instead, treat optimization as a problem for the entire stack. Although much performance gain can be achieved by these approaches, they also add the complexity of the physical layer to the protocol design. Furthermore, once the layering is broken, the luxury of designing a protocol in isolation is also lost [52].

We contend that it is possible to overcome these challenges by using the deterministic model as a finer abstraction of the physical layer. While this model captures more features of the physical layer to be utilized at higher levels, it still keeps the separation between the layers. Hence, we can now integrate recent advanced techniques that have been developed in the communication community, such as successive interference cancellation and interference alignment [53, 54], in protocol design at higher layers, without worrying about the details of implementing the techniques.

We believe our approach will significantly affect both the networking and communication communities. It has a positive effect on the network community in that it proposes a more systematic way for cross-layer protocol design. It also affects the communication community by making the recent discoveries more accessible to other communities, which, in turn, makes them more likely to impact the real world.

As an example, consider the problem of sum-rate maximization in wireless networks. In this problem, we have a wireless network with several transmitter-receiver pairs and the goal is to find the communication protocol that achieves the maximum total throughput. To deal with the interference among users, most works in the networking community use the "treat interference as noise" model which simply captures the effect of interference as an additive noise at the receiver. Therefore, under this simplistic model, the only physical layer degree of freedom left to optimize is the transmit-power at each node. This problem, which is known as "sum-rate maximization with power control", has been extensively studied over the past 15 years (see e.g. [55–61]) and some of the discovered power control algorithms have significantly impacted the digital cellular technology.

However, there are several other physical layer degrees of freedom that are not captured in the "treat interference as noise" model and that can be utilized to significantly improve the spectral efficiency and increase the sum-rate. We will now use a simple example to illustrate how it is possible to use the deterministic model, which captures these other degrees of freedom, to design more efficient protocols.

Consider the network shown in Figure 6 (b). Here Tx_1 wants to communicate with Rx_1 and Tx_2 wants to communicate with Rx_2. In the "treating interference as noise" model, first each transmitter chooses a power level to send the information. Then, each receiver gets the bits that are received above the highest level of the interference or and the noise, called the INR level. In our example it is easy to verify that the maximum achievable sum-rate under this model is 2 bits, and Figure 6 (b) shows how to achieve it.

However, it is possible to achieve higher rates by choosing the transmission signal levels more carefully. As shown in Figure 6 (c), we can achieve a sum-rate equal to 4 bits by transmitting over signal levels that don't interfere with each other at the receivers. It can be show that any such strategy can be implemented over the Gaussian model by using superposition coding at the transmitters and successive interference cancellation at the receivers, and it achieves a rate within a constant gap to the rate in the deterministic network.

Therefore, by using the deterministic model, we can convert the problem of sum-rate maximization by superposition coding and successive interference cancellation, into a discrete optimization problem over signal levels, and obtain approximate solutions. Similarly, we can capture the effect of other physical layer coding techniques and use them to design more efficient protocols.

To summarize, the research in this area aims to utilize the deterministic model as a new abstraction of the physical layer to study

- Cross-layer protocol design for wireless networks, by using advanced physical layer coding techniques, such as:

1) superposition coding	3) successive interference cancellation
2) lattice coding	4) interference alignment

2.8 Summary of objectives

We propose a systematic approach to make progress on several fundamental problems in wireless network information. During the five-year span of the proposal we intend to effectively use our approach to pursue a thorough study in four important research areas in wireless network information theory; 1) to develop a systematic framework to approximate the capacity of wireless networks, 2) to determine the effect of access to local vs. global network information on the capacity, 3) to design communication protocols that are robust against malicious adversaries, and 4) to utilize recent advances in physical layer coding techniques for cross-layer protocol design.

3 Career Development Plan: Education Plan

I believe teaching and research inherently go hand-in-hand with synergistic results. The proposed education plan is designed to complement and further the research goals of this proposal. The education activities are divided into three main categories i.e., 1) course development and teaching; 2) advising students; and 3) the outreach program.

3.1 Course Development and Teaching

During the time frame of this proposal, I plan to teach undergraduate courses on communication theory and probability, as well as graduate courses on digital communications, wireless communications, and network information theory. I was a Graduate Student Instructor for digital communications in the fall of 2006, at UC Berkeley. During that time, I was closely involved with designing the curriculum developing, homework assigning, and preparing a one-hour lecture each week. Also, I was a guest lecturer for an advanced course on the Mathematics of Information at Caltech in the spring of 2009.

In terms of teaching methodology, there are, of course, a variety of models, each with its relative merits and effectiveness, depending on the subject, class size, and the backgrounds of the students [62–66]. However, my emphasis is to have an interactive teaching style that encourages student participation. My goal is to create an intuitive picture of the subject by providing illustrative examples, which I use to ensure that the students acquire a true understanding of the subject and to encourage their participation. I also believe that course projects and homework problems are very important ingredients of teaching, and I plan to spend a considerable amount of time and effort designing appropriate, meaningful homework assignments.

Course on digital communications: In the fall of 2009, I will teach ECE 567, the graduate level course on digital communications. In the fall of 2010, I plan to teach ECE 468, the undergraduate course on digital communications. One of the key challenges in teaching digital communications is to decide what to teach. Unlike many other courses in electrical engineering for which there are clear objectives and outlines, most digital communication courses are taught as a survey of different blocks of a general digital communication system. Instead, I propose to revise the digital communications course to teach principles of communication over four channels with increasing levels of complexity, i.e., 1) erasure channels (for example the Internet); 2) noisy flat Gaussian channels; 3) band-limited channels (e.g., DSL line); and 4) wireless channels.

The first example will illustrate the notion of reliable communication and the effect of feedback for error correction. The second example will be used to teach the Gaussian noise model, signal detection, and different modulation schemes. In the third example, I will teach the concepts of intersymbol interference (ISI), pulse design for band-limited systems, passband and quadrature amplitude modulation, and equalization. Finally, in the fourth example, I will cover the basic modeling of wireless channels and the key parameters involved, fading, and the notion of reliable communication over flat-fading channels. Major advantages of this approach are that the role of each building block in communication systems will be clearly explained to the students, and a concrete, practical example that the students can connect to these building blocks will always be provided.

Course on wireless communications: In addition to the digital communication course, Cornell already has a strong sequence of annual, graduate-level courses in error-correcting codes, digital signal processing, and information theory. Therefore, since most students will be familiar with the coding and signal processing aspects of communications systems from this sequence of courses, most of the emphasis in the wireless course will be placed on the wireless aspects of communication systems. The goal of this course is to present a comprehensive understanding of recent advances in wireless communications and discuss the key ideas behind each one of the advances. After completing this course, the students should have a good understanding of the fundamentals of wireless channels and fading channel models, wireless communication techniques, cellular systems, multi-access and multi-user systems, and diversity techniques, as well as new topics, such as MIMO, opportunistic communication, and cooperative communication.

Course on network information theory: I also expect to teach ECE 662, the advanced network information theory course. This course will be a fresh perspective on the state-of-the-art of the field. The course will begin with the classical wireline networks and progress to more complex models for noisy networks. The approach will be to integrate our discrete deterministic channel model as an effective tool to teach the key information theoretical concepts in more-complex, noisy networks. Because of our emphasis on the use of deterministic models, we hope to also make this course accessible to students in the networking area as well as students from the Computer Science Department. I believe this will be a great opportunity to initiate the study of the algorithmic aspects of wireless network information theory by involving these students in class projects.

3.2 Advising undergraduate students

Cornell attaches great importance to advising undergraduate students to help them achieve their full academic and intellectual potential and reach their career goals. Every Cornell student is assigned a faculty advisor to help her or him plan her or his academic career. I believe real mentoring is more than providing good advice on selecting courses. In fact, as most research demonstrates [67–72], mentoring is critical to a truly successful and life-changing educational experience. I take undergraduate advising very seriously and plan to spend time to work closely with my undergraduate advisees at Cornell, learning about their backgrounds, interests, and personalities, to be able to make them fully aware of all their educational and career opportunities.

I also plan to involve the undergraduates in the research part of this proposal. My one-year appointment at Caltech afforded me an opportunity to define research projects for two junior undergraduates and advise them as they conducted the projects. Several of the research problems discussed in Areas II and IV, such as the problem of utility maximization over deterministic networks, are expected to require basic tools in optimization and combinatorics. These simulation-based problems are well suited for undergraduates to help them learn some fundamental concepts in communication and optimization theory. This is also a great opportunity to motivate them for graduate study and get them interested in the communication and information theory fields.

3.3 Advising graduate students

One of the key roles of a faculty member is to recruit and advise graduate students. The graduate level mentoring relationships are quite different from those at the undergraduate level, i.e., the students have more sophisticated thinking abilities, the relationship is much longer, and its impact on the future of both the advisor and advisee is more significant. Hence, student advising is of critical importance in graduate training. At Cornell, I will be involved with graduate students beginning in the pre-admission stages and continuing until degrees are earned. My approach is to exchange email with prospective students, talk to them about their research interests on the phone, and meet with them on the visit day to make sure I recruit excellent students. This year, two very talented Ph.D. students will enter my program in the fall.

My goal is to train the students to be independent researchers, with high standards of technicality, integrity, and objectivity. I expect to devote a significant portion of my time to technical research