REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY) 23-02-2017	2. REPORT TYPE Final Report	3. DATES COVERED (From - To) 16-Sep-2013 - 15-Sep-2016

4. TITLE AND SUBTITLE
Final Report: Communicating under Adversarial Attacks: Models, Codes, and Fundamental Limits (10.1.2 Mobile Ad Hoc Networks)

5a. CONTRACT NUMBER
W911NF-13-1-0455

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611102

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHORS
Aaron Wagner

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
Cornell University
Office of Sponsored Programs
373 Pine Tree Road
Ithaca, NY 14850 -2820

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
64059-NS.6

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT
This project considered how coding can be used to achieve resilient communication in the face of adversarial attack. Prior work on this topic makes one of two rather limiting assumptions: first, that the adversary operates at the physical layer and therefore injects errors that are akin in a certain sense to noise, and second, that source-channel separation can be assumed, meaning that data compression can be performed separately from coding against adversarial errors. In reality, the adversary may be able to subjugate a node in the network and thereby inject errors at a much higher layer of abstraction. Also, prior work of the PI has shown that the performance

15. SUBJECT TERMS
resilient communication, adversarial communication, error-correcting codes, data compression, rate-distortion

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Aaron Wagner
UU	UU	UU	UU		19b. TELEPHONE NUMBER 607-255-1017

Report Title

Final Report: Communicating under Adversarial Attacks: Models, Codes, and Fundamental Limits (10.1.2 Mobile Ad Hoc Networks)

ABSTRACT

This project considered how coding can be used to achieve resilient communication in the face of adversarial attack. Prior work on this topic makes one of two rather limiting assumptions: first, that the adversary operates at the physical layer and therefore injects errors that are akin in a certain sense to noise, and second, that source-channel separation can be assumed, meaning that data compression can be performed separately from coding against adversarial errors. In reality, the adversary may be able to subjugate a node in the network and thereby inject errors at a much higher layer of abstraction. Also, prior work of the PI has shown that the performance penalty associated with enforcing source-channel separation can be arbitrarily large in this application.

This project provides code constructions that jointly perform data compression and error correction to provide superior resiliency against adversarial attack, which are novel in several respects. In addition to jointly performing data compression and error correction and targeting errors at a higher layer in the network stack, the codes gracefully degrade with increasing strength of the adversary and rely on regular, instead of modulo, integer arithmetic.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received		Paper
02/21/2017	4	Xiaoqing Fan, Oliver Kosut, Aaron B. Wagner. Variable Packet-Error Coding, IEEE Transactions on Information Theory, ( ): . doi:
07/28/2016	3	Ebad Ahmed, Aaron B. Wagner. Coding for the Large-Alphabet Adversarial Channel, IEEE Transactions on Information Theory, ( ): . doi:
TOTAL:		**2**

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received		Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received		Paper
08/01/2014	1.00	Xiaoqing Fan, Aaron B. Wagner, Ebad Ahmed. Polytope codes for large-alphabet channels, 2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton). 02-OCT-13, Monticello, IL. : ,
09/16/2015	2.00	Yuguang Gao, Aaron B. Wagner. Must one learn the channel to communicate at capacity?, 2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton). 30-SEP-14, Monticello, IL, USA. : ,

TOTAL: **2**

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received		Paper

TOTAL:

Number of Manuscripts:

Books

Received Book

TOTAL:

Received Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

NAME	PERCENT_SUPPORTED	Discipline
Xiaoqing Fan	0.72	
Yuguang Gao	0.33	
Ibrahim Issa	0.13	
FTE Equivalent:	**1.18**	
Total Number:	**3**	

Names of Post Doctorates

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

NAME	PERCENT_SUPPORTED	National Academy Member
Aaron Wagner	0.12	
FTE Equivalent:	**0.12**	
Total Number:	**1**	

Names of Under Graduate students supported

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:...... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:...... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):...... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:...... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:...... 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

(4) Statement of the Problem Studied

This project focused on the development of coding schemes for resilient communication. We considered a communication scenario in which a transmitter and receiver are connected by multiple, independent paths. We assume that an adversary may take over some, but not all, of these paths, and arbitrarily alter the data transmitted along those paths. The goal is then to design codes for this problem that allow the decoder to satisfactorily reproduce the data source even in the face of path errors introduced by the adversary.

Existing work either assumes source-channel separation, meaning that data compression is performed separately from the error-control coding used to counteact the adversarial errors, or it assumes that the adversary can attack all of the paths (or, equivalently, that there is only one path to begin with), but that the adversary is somehow unable to arbitrarily alter the data along the paths. That is, the adversary is constrained in how it alters the transmission along one path, say by only altering it during a prespecified fraction of the time. Work completed as part of this project (but initiated prior to the start of the project) has shown that schemes that enforce source-channel separation perform worse than schemes the perform joint source-channel coding by an arbirarily large factor is some cases. The work performed as part of this project has also shown that the problem of coding to counteract adversaries that can completely alter the traffic along a subset of paths is very different from that of coding to counteract an adversary that can partially alter the data along all of the paths.

Another novel aspect of the project is that it sought code designs
that degrade gracefully as the power of the adversary increases.
Existing work on this problem sets a single criterion for acceptable performance and ensures that this criterion is met for all possible adversary actions. This approach can be quite pessimistic, in that it does not provide any improved performance guarantees if there happens to be no jamming, which could be the typical scenario for the system. This project sought is to construct systems that provide the highest level of performance when there are no errors along any of the paths, a slightly lower level of performance when there is an error along one path, etc.

(5) Summary of the Most Important Results

All of our results assume that the encoder wishes to transmit a source to a decoder over several paths. The adversary can take over a small number paths and alter the traffic along them in an arbitrary way. Neither the encoder nor the decoder know which of the paths the adversary will attack or how it will alter the traffic along those paths. Performance is measured by the data rates that the scheme requires along each of the paths and also by the distortion of the reconstruction at the receiver as a function of the number of paths attacked by the adversary.

(5.1) Source-Channel Separation

As noted above, we showed that source-channel separation is not optimal for this problem, even if one is only concerned with the distortion when the adversary takes over the maximum number of paths (which is assumed to be less half the overall number of paths, to avoid degeneracy). More precisely, separation is optimal if the source is i.i.d. uniform and binary and the distortion measure is Hamming distance or if the source is i.i.d. Gaussian with mean-squred error (MSE) distortion, but not if the source is i.i.d. uniform and binary and the distortion is measured using the erasure distortion measure. In the erasure case, the ratio of the distortion achievable via separation to the optimal distortion can be arbitrarily large. This work was initiated prior to the start of this project and completed during the course of it.

This establishes that source-channel separation is not optimal for this problem that the problem formulation should model the source as well as the adversarial errors. Of course, source-channel separation affords great simplification in the design process, and as such, it is an attractive approach to design even if it is suboptimal. We expect that the code designs developed by this project would be only applied to high-priority communications, however, for which added resiliency against attack outweighs the added design complexity and cost.

(5.2) Lossless Reconstruction and the Role of a Secret Key

Consider the special case of the problem in which the adversary erases, but does not otherwise alter, traffic along some of the paths. This problem has been studied extensively and is called the "multiple descriptions" (MD) problem in the information theory literature. The MD problem is easier than the problem in which the adversary injects errors in the sense that any performance that is achievable under the latter is achievable the former.

We showed that in two special cases the two problems actually coincide; that is, our problem can likewise be reduced to an MD problem. The first is when the encoder and decoder share a secret key (that is unknown to the adversary). In this case, the encoder can "sign" the packets in a way that makes alterations detectable by the decoder. The second is when the decoder's goal is to losslessly reproduce functions of the source. In this case we operationally relate the two problems: that is, a code for the MD problem can be translated into a code for our problem with the same performance and vice versa.

(5.3) The Utility of Polytope Codes

In the general case, however, the problem does not seem to be reducible to the MD problem and it seems to require new code constructions. Our work during the early part of this project focused on binary sources with the erasure distortion measure. The erasure distortion measure is quite generic in the sense that it is meaningful for many different types of data, and it is also a useful starting point when considering new problems. We assume that the rates of the different paths are the same, and the decoder wishes to reproduce the binary source completely and without distortion if the adversary does not alter any of the traffic, and the decoder wishes to reconstruct as much of the source as it can, again without error, when the adversary does alter traffic.

For this instance we showed that a variation of the "polytope" code constructions of Kosut, Tong, and Tse outperform designs based on traditional (maximum-distance separable (MDS)) constructions. Kosut et al. originally developed their codes for networking coding subject to adversarial attacks on the nodes in the network. Their original constructions were tailored to specific network topologies and not easily generalizable, however. They also involved very long blocklengths and complicated encoding. We simplified the code construction, making the operation of the code more transparent, simplifying the encoding, reducing the required blocklength, and making the code more generalizable. The application of the polytope coding idea to joint source-channel coding is also novel.

(5.4) Impossibility Results for Polytope Codes

While our polytope code construction beats conventional codes for our problem, it exhibits an undesirable scaling: the the number of paths that can be subject to errors must grow no faster than the square root of the total number of independent paths. Analogies with other coding problems suggest that a linear scaling might be possible. Since the encoding, decoding, and analysis of our code are all rather unorthodox, it is unclear which of these would need to be improved in order to obtain better overall performance.

We showed that this square-root scaling is not attributable to a weakness in either the decoding rule that we use or in our analysis of the code. That is, we showed that the performance guarantees that we provide are the best possible given the encoding that we use. This impossibility argument is quite involved and makes use of techniques from number theory, Euclidean geometry, and linear algebra.

(5.5) How much learning is necessary to communicate at capacity?

All of our prior work on this problem assumes that the channel errors are fully adversarial: an omniscient adversary can take the signal (or packet) transmitted over a link and replace it with an arbitrary signal of its choosing. This is a reasonable model for an adversary that has infiltrated nodes in the network and can alter signals at the highest-levels of the network stack. This model is less appropriate, however, for an adversary that jams at lower layer of abstraction. For such an adversary, the channel seen by the encoder would be better modeled as being stochastic with a distribution that can be controlled by the adversary. We have thus bgun examining such models.

We looked in particular at the single-link scenario in which a discrete memoryless channel is selected at the beginning of the transmission by the adversary, and the choice is unknown to the encoder and the decoder. The encoder and decoder can attempt to learn the adversary's choice via training, however. The question then is how much training is required in order to communicate at capacity, and in particular, whether it is necessary for the encoder and decoder to correctly identify the adversary's choice in order to communicate at capacity.

A priori, there are reasons to believe that it might not be necessary to completely identify the channel in order to communicate at capacity. In order to communicate at capacity, the encoder does not need to know the channel distribution in its entirety; it only needs to know the input distribution that maximizes the mutual information over it, which is evidently a much lower-dimensional object than the channel itself. Likwise, there are universal decoders, such as the Maximum Mutual Information (MMI) decoder, that can provably achieve capacity without any channel knowledge at all.

We showed that, somewhat surprisingly, the encoder and decoder do essentially need to learn the channel in order to communicate at capacity. This shows that one of the ways that an adversary can disrupt communication via jamming is to increase the amount of training that the code must perform, assuming that the code seeks to operate at capacity.

(5.6) Polytope Codes for Gaussian Sources

Our early work focused on discrete sources with the erasure distortion measure. During the later part of the project, we considered more sophisticated source distributions and distortion measures that might model multimedia data such as voice, images, and video.

As a first step in this direction, we studied scalar, i.i.d., Gaussian sources with a quadratic distortion measure. While this is not a satisfactory model of realistic sources, it is an important waypoint on the path toward realistic source models. Assuming that there are two possibilities for the number of paths taken over by the adversary, we characterized the region of achievable distortion pairs (corresponding to the two levels of adversarial attack) and rate vectors (corresponding to the rate along each of the paths) to within a constant gap. The coding scheme that achieves this near-optimality encodes the sourcs into two layers: a "base" layer and a "refinement" layer. The base layer encoding is transmitted with sufficient redundancy so that it can be recovered without error irregardless of the action taken by the adversary. The refinement layer is transmitted with only enough redundancy so that it can be recovered without error under the less-severe attack by an adversary. Under the more-severe attack, the refinement layer can be decoded incorrectly, resulting in a worse reconstruction than if the decoder used the base layer alone. This turns out to perform significantly better than encoding the refinement layer in a way that errors to it can be corrected under the less-severe attack and detected (resulting in the layer being ignored at the receiver) under the more-severe attack.

Technology Transfer

The PI has an ongoing colloboration with MITRE on the subject of this project. Each year, the PI and engineers at MITRE jointly advise Masters-level students who work on an applied component of this research. The students' work culiminates in a debrief presentation on-site at MITRE, where the PI also presents a research seminar. Funding for supplies and travel is provided by MITRE. For 2015-2016, three M.Eng. students (and one undergraduate) considered how voice-recognition technology could be used to embed low-rate text versions of speech signals with very high redundancy, so that the text of the speech can be recovered in the face of jamming, even if the original speech signal cannot. For 2016-2017, one M.Eng. student is working on implementing the codes developed as part of this project and testing them on actual speech signals with and without adversarial errors.

The PI also visited ARL in Spring 2016 to deliver a seminar on the progress to date and receive feedback from ARL researchers and ARO program managers.

Recently the PI has begun a collaboration with Kencast, a small company in Connecticut that develops media streaming technologies focusing on military and public safety applications.