Experimental Analysis of Denial-of-Service Attacks on Teleoperated Robotic Systems

Tamara Bonaci
University of Washington
Department of Electrical Engineering
185 Stevens Way
Seattle, WA, 98195
tbonaci@uw.edu

Junjie Yan
University of Washington
Department of Electrical Engineering
185 Stevens Way
Seattle, WA, 98195
junjiey@uw.edu

Jeffrey Herron
University of Washington
Department of Electrical Engineering
185 Stevens Way
Seattle, WA, 98195
jeffherr@uw.edu

Tadayoshi Kohno
University of Washington
Department of Computer Science and Engineering
185 Stevens Way
Seattle, WA, 98195
yoshi@cs.washington.edu

Howard Jay Chizeck
University of Washington
Department of Electrical Engineering
185 Stevens Way
Seattle, WA, 98195
chizeck@uw.edu

ABSTRACT

Applications of robotic systems have had an explosive growth in recent years. In 2008, more than eight million robots were deployed worldwide in factories, battlefields, and medical services. The number and the applications of robotic systems are expected to continue growing, and many future robots will be controlled by distant operators through wired and wireless communication networks.

The open and uncontrollable nature of communication media between robots and operators renders these cyber-physical systems vulnerable to a variety of cyber-security threats, many of which cannot be prevented using traditional cryptographic methods. A question thus arises: what if teleoperated robots are attacked, compromised or taken over?

In this paper, we systematically analyze cyber-security attacks against Raven II®, an advanced teleoperated robotic surgery system. We classify possible threats, and focus on denial-of-service (DoS) attacks, which cannot be prevented using available cryptographic solutions. Through a series of experiments involving human subjects, we analyze the impact of these attacks on teleoperated procedures. We use the Fitts' law as a way of quantifying the impact, and measure the increase in tasks' difficulty when under DoS attacks.

We then consider possible steps to mitigate the identified DoS attacks, and evaluate the applicability of these solutions for teleoperated robotics. The broader goal of our paper is to raise awareness, and increase understanding of emerging cyber-security threats against teleoperated robotic systems.

ICCPS '15 April 14 - 16, 2015, Seattle, WA, USA

ACM 978-1-4503-3455-6/15/04 ...$15.00.
http://dx.doi.org/10.1145/2735960.2735980

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General - Security and protection (e.g., firewalls); C.3. [**Special-purpose and Application-based Systems**]: Process control systems; C.4. [**Performance of Systems**]: Reliability, availability and serviceability

Keywords

Cyber-physical systems; Teleoperated robotic systems; Cyber-security thereats; Denial-of-service attacks; Fitts' law

1. INTRODUCTION

Interest in robotics and robotic applications is rapidly growing. In 2008, the total number of robots worldwide was estimated to 8.6 million [2], and the last year marks an all-time record in the robot sales[4]. Many predict that robots today are in a stage similar to that of personal computers in the 1970s, and that the application of robotics will only continue to grow, reaching the number of a couple of billions in a few decades [16].

In the future, teleoperated robots will likely be expected to combine the existing publicly available networks with temporary *ad-hoc* and satellite networks to send video, audio and other sensory data to remote operators [27]. Such teleoperated systems will be used to provide immediate medical relief in under-developed rural areas, areas of natural and human-caused disasters, and in battlefield scenarios [20]. The question arises, however: what if teleoperated robotic systems are compromised? Recent examples, such as Stuxnet worm, specifically designed to target programmable logic controllers, which was blamed for ruining a significant part of Iran's nuclear centrifuges [14], exemplify possible issues when a cyber-physical system is being explicitly targeted.

To date, security has not been a primary concern for teleoperated robotic systems. Yet the problem has recently been

recognized as important [12, 45, 46]. The main efforts thus far have focused on ensuring private communication between an operator and a robot [45, 46] and on the ability to verify the robot-side code [12].

At the moment, however, there is little understanding of what the actual cyber-security threats against teleoperated robotics are, and what the impact and implications of these attacks might be. This lack of understanding of the actual threats is a function of two factors. It is not known: (1) how easy would it be for an attacker to compromise a teleoperated robotic system, and (2) what the applications of such an attack might be. Moreover, not being able to answer these questions makes it hard to understand what the challenges to improving security of teleoperated systems are, much less to address them.

In this paper, we seek to evaluate one specific class of cyber-security attacks, namely the *denial-of-service (DoS) attacks*. We focus on this class of attacks because they cannot be mitigated using available cryptographic solutions, and various proposed telerobotic-specific solutions [12, 45, 46] are also unable to prevent these attacks. Moreover, teleoperated robotic systems operating in either natural or man-made catastrophes may be required to operate in DoS-like conditions, due to the fact that remaining communication resources (after the catastrophe) will likely be clogged by benign, but concerned survivors and well-wishers.

Our work is experimental, along the lines of much past work that explored the security and privacy properties of emerging technologies, including modern automobiles [24, 11] and medical devices [17, 18]. Through an empirical analysis of a commercially-available robotic surgery platform, the *Raven II*®, we provide an informed understanding of the impacts, consequences and risks of denial-of-service attacks. We make the following specific contributions:

Experimental vulnerability analysis: We focus on *denial-of-service* (DoS) attacks, which cannot be mitigated using available off-the-shelf solutions. Through a series of experiments involving human subjects, we analyze practical examples of DoS attacks, and assess the level of the actual impact on a teleoperated surgical procedure.

Risk assessment and new security metrics: In assessing the impact of DoS attacks, we consider several attack levels (benign, intermediate, and severe) over several tasks, and quantify the impact using the following metrics: (i) the overall procedure (trial) time, (ii) the subjective assessment of difficulty, and (iii) the Fitts' indices of difficulty and of performance. We further make a distinction between different components of telerobotic tasks, such as moving a robot end effector tool, and grasping or dropping a rubber block. Extending Fitts' law, we propose a new metric to quantify the impact of attacks on these components of telerobotic tasks.

Defense directions: We consider several well-established methods to prevent and mitigate DoS attacks, and analyze their feasibility to teleoperated robotic systems.

Opportunities and challenges specific to teleoperated procedures: During our experimental analysis, we observe several challenges and opportunities specific to teleoperated robotics. The most interesting is that many human operators exhibit a significant learning effect when performing tasks under DoS attacks. This learning effect is evidence that human operators are capable of adapting to unfavorable network conditions.

This paper is organized as follows: in Section 2, we give an overview of recent results in robotic surgery and in the security of relevant cyber-physical systems. In Section 3, we present the *Raven II* teleoperated surgical system. In Section 4, we give a brief overview of Fitts' law. In Section 5, we present an attacker model, and in Section 6, we describe the conducted experiments. In Section 7, we analyze the experimental results, and in Section 8, we discuss possible steps to mitigate the attacks. Finally, in Section 9, we discuss possible next steps.

2. BACKGROUND AND RELATED WORK

2.1 Robotic Surgery Systems

The use of robots in surgery dates to 1985, when an industrial robot, *Puma 560*, was used for needle placement in brain biopsy [21]. In 1988, the *Probot*, developed by the Imperial College, was used to perform prostate surgery [21]. The first indirectly-controlled surgical system, where the surgeon controlled a robot using a computer, was developed in the 1990s by the Stanford Research Institute (SRI) [49]. For the next several years, the development of robotic surgery was enabled by the advent of three commercial systems: the *Aesop* and the *Zeus* (Computer Motion) and the *da Vinci* (Intuitive Surgical) [21]. The *da Vinci* is currently the only FDA approved system for use on humans. It uses a dedicated (and private) communication network, since a surgeon and a robot are placed within the same operating room.

In September 2001, the *Zeus* system was used to perform the first transatlantic telesurgery, operating from New York City on a patient in Strasbourg, France [31]. A few years later, in 2005, the *da Vinci* system was used to perform the first transcontinental telesurgery, between a surgeon in Sunnyvale, CA and patients in Cincinnati and Denver [49].

Next generation teleoperated surgical robots are envisioned to be used in a variety of scenarios, including battlefields and natural disasters [41]. In these circumstances, a robot's portability becomes important, as well as its ability to operate with limited power resources and in challenging climates and environments, often without a basic infrastructure. Despite these extreme operating conditions, the robots will be expected to carry out missions while maintaining specific performance requirements, in particular guaranteeing patient safety.

The *Raven II*, a portable surgical robot, has been evaluated in several extreme environments scenarios, including the desert HAPs/MRT experiment [27] and underwater habitation module experiment NEEMO 12 (mission16) [13]. In these experiments, the following network states were recognized as critical for reliable performance [28]: (i) communication latency, (ii) jitters, (iii) packet delays and out-of-order arrivals, (iv) packet losses, and (v) devices failures.

In addition to these stochastic but benign network patterns, human operator-robot communication over publicly available networks expose telerobotic procedures to a slew of new problems that are not present in dedicated and private networks. Due to the open and uncontrollable nature of the communication medium, it becomes easy for malicious entities to jam, disrupt, or take over the communication between a robot and an operator. Thus, in order to develop *safe and secure* telerobotic systems, it is necessary to ensure that these systems are *information secure*, in addition to maintaining all of the systems' performance, reliability, and privacy requirements.

2.2 Security of Most Relevant Cyber-Physical Systems

Security of cyber-physical systems has been a rapidly growing research area, with researchers focusing on the topics of monitoring and estimation (e.g., [35, 9]), networked control systems verification (e.g., [44, 15]), as well as robust communication, consensus and distributed computation (e.g., [43, 36]). We give a brief overview of recent security results for the CPS most relevant to teleoperated robotic systems: networked control, automotive, medical systems.

2.2.1 Security of Networked Control Systems

It has been shown that some of the attack classes against networked control systems, wireless sensor networks, and multi-agent systems can be mitigated by relying on the system's dynamics (see, e.g., [6, 8, 34]). In [10, 35, 34], the authors assumed that the system's dynamics are linear, and showed that a simple optimal controller and a Kalman filter can be used to guarantee the desired probability of detecting attacks, such as replay, false data injection and integrity attacks, given a certain model. In [6], the authors considered a networked control system with linear dynamics under a denial-of-service (DoS) attack. They proposed that a semi-definite programming approach can be used to find a causal feedback controller that ensures that the given networked system operates properly (i.e., that ensures the system's objective function is minimized) while maintaining the system's security and power constraints. The approach proposed in this paper can be applied to telerobotic systems where the linear system dynamics assumption is not satisfied.

2.2.2 Security of Automotive Systems

Automobiles are becoming highly computerized and increasingly 'connected', as well as semi-autonomous and autonomous. Recent research has shown that although automotive computer standards indeed describe mechanisms to improve security, these mechanisms are not universally implemented on all the computers in modern cars [24, 11]. In [24], through an analysis of the security properties of all the critical computerized components of a car, the authors found that an attacker connected to the vehicle's internal computer networks can affect the state of all the analyzed components. In [11], the authors provided a experimental study of an external attack surface on a modern automobile, and they discussed structural characteristics of the automotive ecosystem that make addressing the identified vulnerabilities challenging. Some of the observed challenges can be easily avoided in teleoperated robotic systems, given a relatively early design phase telerobotics is currently at.

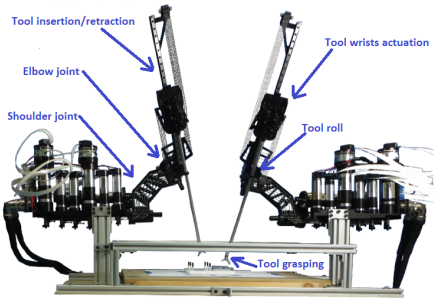

Figure 1: The Raven II system consists of two 7-degrees-of-freedom surgical manipulators. The motion axes of the robot are: shoulder joint, elbow joint, tool insertion/retraction, tool roll, tool grasping, tool wrist 1 actuation and tool 2 wrist actuation.

2.2.3 Security of Medical Systems

Security and privacy issues related to tele-medical applications were first recognized in the mid-1990s [30, 47]. After the establishment of the Health Insurance Portability and Accountability Act (HIPPA) [3], patients privacy became a primary concern, with researchers focusing on the confidentiality of transmitted and stored patient data. More recently, it has been observed that many modern implantable medical devices, including pacemakers and implantable cardioverter-defibrillators, are vulnerable to a variety of attacks, which allow attackers to wirelessly obtain private patient information and change device settings in ways that can directly impact patient health [17, 18].

2.2.4 Security of Teleoperated Robotic Systems

Very recently, security concerns regarding telerobotic surgery systems have emerged, with a focus on system verification [12], communication reliability [45, 46], as well as private and authenticated communication [25]. In our recent work [7], we identified and experimentally evaluated the scope and impact of several classes of cyber-security attacks against teleoperated robotic surgery. We further demonstrated experimentally that some of the existing cryptographic methods may be readily applicable to teleoprated surgery, without negative impacts on system's performance and real-time operation requirements.

3. RAVEN II® SURGICAL ROBOT

The *Raven II*, teleoperated surgical robot is shown in Figure 1. It is a research platform used for investigation of advanced robotic-assisted surgery techniques [39, 19]. It is the first experimental system to support both software development and experimental testing for surgical robotics. It was developed at the University of Washington with NSF support, and is now manufactured and distributed by Applied Dexterity [1]. It is currently used as a research tool by 15 institution in the U.S., Canada, France, United Kingdom, Denmark, Israel and South Korea.

The *Raven II* system consists of two 7-degrees-of-freedom (DOF) surgical manipulators, divided into three main sub-

systems: the static base that holds all seven actuators, the spherical mechanism that positions the tool, and the tool interface. The motion axes of the robot are: *shoulder joint* (rotational), *elbow joint* (rotational), *tool insertion/retraction* (linear), *tool roll* (rotational), *tool grasping* (rotational), *tool wrist 1 actuation* (rotational), and *tool wrist 2 actuation* (rotational). DC motors mounted to the base actuate all motion axes. The motors of the first three axes have power-off brakes to prevent tool motion in the event of a power failure. Each manipulator has a total (moving plus non-moving) mass of approximately 10 kg, which includes the motors, gear heads and brakes. A tool interface allows quick changing of tools, and transmits motion to the tool rotation, grasp and wrist axes. The links and control system support a 3-axis wrist.

The *Raven II* software is based on open standards, including Linux and the Robot Operating System (ROS) [38]. The low-level control system includes real-time Linux processes (modified by the RT-Prempt Config kernel patch), running at a deterministic rate of 1000 Hz. Key functions running inside the 1000 Hz servo-loop are: (i) coordinate transformations, (ii) forward and inverse kinematics, (iii) gravity compensation, and (iv) joint-level closed-loop feedback control. The link between the control software and the motor controllers is a USB 2.0 interface board, designed with eight channels of high-resolution 16-bit digital-to-analog conversion for control signal output to each joint controller, and eight 24-bit quadrature encoder readers. The board can perform a read/write cycle for all 8 channels in 125 microseconds. The two *Raven II* arms are controlled by a single PC with two USB 2.0 boards.

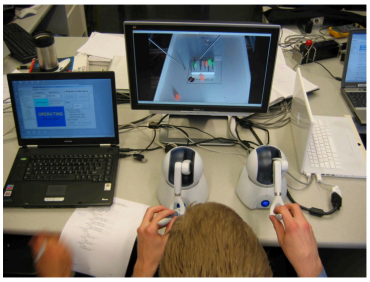

Figure 2: An example surgical control console, used with the Raven II system. The console consists of three main parts: surgical GUI, shown on the laptop screen, surgical video transmission on the LCD monitor, and two Omni haptic devices [Picture credit: [40]].

An example of a surgeon control console for the Raven II is shown in Figure 2. Control inputs and robot feedback, which includes video and haptic information, are transmitted using a communication standard for surgical teleoperation, the Interoperable Telesurgery Protocol (ITP) [23]. The ITP allows communication between heterogeneous surgical consoles (masters) and manipulators (slaves), regardless of their individual hardware and software. In this protocol, messages between a surgeon and a manipulator are exchanged using the User Datagram Protocol (UDP). Possible negative effects of UDP's unreliability, which may include out-of-order arrivals, packet duplications and losses, are reduced by transmitting surgeon inputs in small increments, based on the assumption that surgical tool motions are continuous. Surgical messages consist of the following information: (i) position and orientation increments, (ii) an indicator variable "button state", defining actuation of end effectors, (iii) a variable "surgeon mode", used to coordinate indexing between master and slave robot, (iv) message sequence number, and (v) checksum [23].

4. FITTS' LAW

Fitts' law is an empirical model, developed in 1954 as a way to model the performance of human physical pointing tasks. It is often used in human-computer interaction (HCI) research to characterize subjects' performance during simple movement tasks under different speed-accuracy conditions.

Fitts' law characterizes subjects performance in terms of the duration of point-to-point reaching movements, and defines the *movement time*, T, as a function of movement distance to the target, D, and the width of the target, W [22, 32]:

$$T = a + b \log_2 \left(\frac{D}{W} \right) = a + b(ID) \tag{1}$$

where intercept parameter a represents the non-movement time needed to start and stop (finish) the trial, and slope parameter b is taken to represent an inherent inverse of the device's speed, i.e.:

$$b \propto \frac{1}{speed}$$

Parameter ID represents the *Fitts' index of difficulty*, defined as a function of target distance, D and target width, W.

In [29], the original Fitts' index of difficulty from equation (1), ID, was redefined as:

$$ID_{Shannon} := \log_2 \left(\frac{D+W}{W} \right) = \log_2 \left(1 + \frac{D}{W} \right) \tag{2}$$

Formulation (2) is typically referred to as the *Shannon formulation*, due to its similarity with the Shannon's channel capacity theorem, characterizing the channel capacity as a function of signal and noise powers [32]. This analogy between Fitts' law and the channel capacity theorem is further extended by typically expressing the Fitts' index of difficulty in *bits*.

For a given experiment, the indices of difficulties, IDs, are determined in the experiment design phase (computed from the chosen target distances and widths), and the movement time is measured for every experimental trial. The metric of interest is the *index of performance*, IP, which quantifieshow movement times change with task difficulty. It has two competing definitions; *the direct division component*:

$$IP_1 := \frac{ID}{T} \tag{3}$$

and *the version derived from linear regression*, using equation (1):

$$IP_2 := \frac{1}{b} \tag{4}$$

5. THREAT MODEL

Telerobotic surgery systems are expected to be used in natural disasters, as well as in man-made catastrophes. In these extreme situations, robots may have to operate in low-power and harsh conditions, with some, potentially lossy connection to the internet. The last communication link will likely be a wireless link to a drone or a satellite, where a drone or a satellite will further provide a connection to a trusted facility, as depicted in Figure 3.

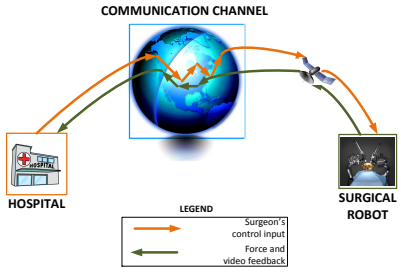

Figure 3: Visualization of a typical telerobotic surgery setup. Dashed lines indicate wireless links, solid lines pre-established network connections (either wired or wireless). Orange color indicates surgeon's control messages, and green color robot's feedback messages.

In such operating conditions, two attack vectors may be feasible [7] (1) *endpoint compromise*, where either a surgeon's or the deployment endpoints (robotic system) can be compromised, and (2) *network and communication-based attacks*, where an attacker may intercept the existing network traffic, inject new malicious traffic, or both.

Since physical access to a machine would likely be strictly monitored, endpoint compromises are less likely, and therefore less interesting. In this paper, we thus focus on network and communication-based attacks, where we identify the most likely point of attack to be between the network uplink and the *Raven II*.

In [7], based on the impact attacks may have on surgeons, we classified possible attacks into three categories: (a) intention modification, (b) intention manipulation, and (c) hijacking attacks. Intention modification attacks occur when an attacker directly impacts a surgeon's intended actions by modifying his/her messages while packets are in-flight, and a surgeon has no control over them. Similarly, intention manipulation attacks occur when an attacker only modifies feedback messages (e.g., video feed, haptic feedback), originating from a robot. In hijacking attacks, a malicious entity causes the robot to completely ignore the intentions of a surgeon, and to instead perform some other, potentially harmful actions.

In this paper, we focus only on intent manipulation attacks, and more specifically only on denial-of-service attacks, where an attacker renders a robot temporarily or permanently unavailable to a surgeon, since these attacks cannot be prevented using the existing cryptographic methods.

6. EXPERIMENTAL SETUP

6.1 System Setup

To experimentally investigate possible impact of DoS attacks on teleoperated surgical procedures, we establish communication between the surgical control console and the *Raven II* robot through a network hub, as depicted in Figure 4.

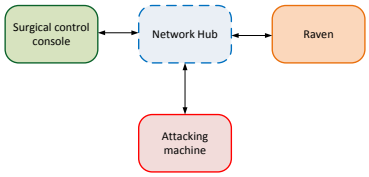

Figure 4: Experimental setup: the attacking machine is running Windows 7 SP3, with attack implementations written in C#.

This allows us to connect an external computer to the same subnetwork, and use it to attack the communication between a surgeon and a robot. Our attacking computer is running Windows 7 SP3, and all of the analyzed attacks are implemented in C#.

6.2 Subjects Demographics

Our analysis is based on the data collected from experiments involving six human participants. All subjects had the same preliminary knowledge about the experiment.

This study was approved by the University of Washington Institutional Review Board approval (#46946 - EB), and all of our subjects were undergraduate and graduate students from the Electrical Engineering and Philosophy departments, ranging in age from 20 to 28 years. We acknowledge that a student's behavior may differ from a surgeon's behavior, but that is acceptable (and an established experimentation method in surgical robotics) since it has been shown that both surgical and non-surgical subjects, upon gaining proficiency, achieve similar results in simple surgical robotic tasks [28].

6.3 Telerobotic Fitts' Law Task

The *telerobotic Fitts' law task* was proposed and developed in [22] as a way of measuring motor control performance of human operators using a teleoperated robot. In this task, a subject uses a robot to move a plastic cylindrical block from a well on the right hand side of the board to a peg on the left hand side of the board, as depicted in Figure 5. One experimental trial consists of moving five blocks.

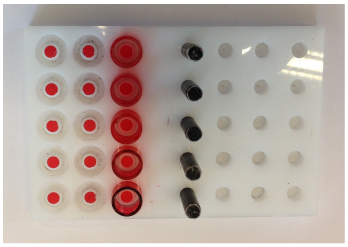

Figure 5: A board used in telerobotic Fitts' law tasks. A subject uses a robot to move the given cylindrical blocks from wells to the pegs. The depicted board represents the "thick" board, "close" configuration scenario.

In our implementation of the telerobotic Fitts' task, we consider two types of pegs, "*thick*" and "*thin*", with widths respectively equal to 8.00 and 4.60 millimeters. For both types of pegs, we consider two board configurations, defined as functions of the center-to-center distance between a block pick-up location and a target peg (also referred to as the *movement amplitude*): the "*close*" and the "*middle*" configurations, with center-to-center distances, respectively, equal to 31.20 and 69.80 millimeters (corresponding to the distance between the closest row of pegs and the closest row of wells (31.20mm) and the second closest row of pegs and of wells (69.80mm)).

For both types of pegs and both board configurations, we consider three attack scenarios: (i) *benign case*, when no attack is mounted, (ii) *intermediate DoS*, and (iii) *severe DoS*. In both intermediate and severe DoS cases, the attacking machine is injecting fake packets, pretending to be either a surgeon or a robot (IP packets spoofing), in accordance with transmission of legitimate surgeon and robot messages, with the goal of creating network layer congestion (*network layer attacks*). The DoS severity is controlled by the number of malicious threads instantiated on the attacking machine. Each malicious thread generates fake packets of the length 256B with frequency of 1000Hz, and sends them both to the *Raven* and to the surgical console.

For the intermediate DoS attack, we instantiate 80 malicious threads and for the severe DoS attack 150. These numbers of threads were obtained through an empirical analysis, where the goal was to find the minimal number of threads such that the impact of the attack is noticeable (*intermediate DoS*), and the maximum number of threads such that the robot is still usable (i.e., the robot is not E-stopping due to too many dropped packets, or too high current levels) (*severe DoS*).

Combining the considered peg types, board configurations and DoS attack severities, each subject was asked to execute 12 different trials, and trials were organized in the following order, where the order was not know to the participants:

1. "thin" board, "close" configuration, no DoS effect,
2. "thin" board, "close" configuration, intermediate DoS,
3. "thin" board, close configuration, severe DoS,
4. "thin" board, "middle" configuration, intermediate DoS,
5. "thin" board, "middle" configuration, severe DoS,
6. "thin" board, "middle" configuration, no DoS,
7. "thick" board, "close" configuration, intermediate DoS,
8. "thick" board, "close" configuration, no DoS,
9. "thick" board, "close" configuration, severe DoS,
10. "thick" board, "middle" configuration, severe DoS,
11. "thick" board, "middle" configuration, intermediate DoS,
12. "thick" board, "middle" configuration, no DoS.

This order of trials was specifically chosen in order to cancel out any inadvertent learning effect. Moreover, before starting the defined experimental sequence, the subjects were given ample time to learn how to use the system and to gain proficiency with it.

7. EXPERIMENTAL ANALYSIS

7.1 Fitts' Law Analysis

7.1.1 Fitts' Indices of Difficulty

For this analysis, we used experimental results from six subjects The Fitts' indices of difficulty were computed using distances D_1 and D_2 respectively equal to 31.20mm for the "close" and 69.80mm for the "middle" board configuration. Widths W_1 and W_2 were computed as the difference between the block diameter and the peg diameters. For "thin" and "thick" pegs, we obtained:

$$\begin{aligned} W_1 &= w_r - w_{p,1} = 12.8 - 4.60 = 8.20\text{mm} \\ W_2 &= w_r - w_{p,2} = 12.8 - 8 = 4.8\text{mm} \end{aligned}$$

Combining all board configurations and all peg widths, from equation (2) we obtained four different indices of difficult, as presented in Table 1.

Config / **Width**	"close", 31.20mm	"middle", 69.80mm
"thin", 8.20mm	2.2645	3.2498
"thick", 4.80mm	2.9069	3.9581

Table 1: Fitts' indices of difficulty, ID, for two different pegborad configurations ("close" and "middle") and two different peg widths ("thin" and "thick").

7.1.2 Data Preprocessing

In order to compute indices of performance, IP, for the 12 experimental trials, we combined data collected from six subjects, and computed the mean, μ and the standard deviation, σ, of the movement times. We then discarded outliers, defined as those where the movement time is greater than $\mu + 0.5\sigma$, or less than $\mu - \sigma$. Most outliers occurred during the first experimental trial, which seem to imply that subjects were still adjusting to the setup.

7.1.3 Fitts' Indices of Performance

After discarding the outliers, we combined the trials corresponding to the same attack scenario, and through linear regression found intercept and slope parameters, a and b for each of the attack scenarios. The obtained results are presented in Table 2.

Attack scenarios	**Slope** b	**Intercept** a	IP_2 [b/s]
No Dos	0.1086	3.7401	9.2116
Intermediate DoS	0.1167	5.56123	8.5692
Severe Dos	1.1054	2.7539	0.9047

Table 2: Parameters of the Fitts' model, a and b, and Fitts' indices of performance, IP, for three considered DoS scenarios (no attack, intermediate attack, severe attack).

We observe that the DoS attack has a significant impact on the overall task, since the index of performance decreases under both attack scenarios, especially under the severe DoS case, where it decrease by more than a factor of ten.

Experimental trial	Subject 1 (s)	Subject 2 (s)	Subject 3 (s)	Subject 4 (s)	Subject 5 (s)	Subject 6 (s)	ID	a	b	IP1	IP2	zeta
Thin, Close, No attack	6.784	3.576	5.054	4.118	8.12	3.964	2.907	3.7401	0.1086	0.569	9.212	0.938
Thin, Close, Intermediate	8.53	4.542	6.456	4.288	8.426	6.298	2.907	5.5613	0.1167	0.389	8.569	0.955
Thin, Close, Severe	10.442	4.648	5.524	4.324	9.924	6.746	2.907	2.7539	1.1053	0.431	0.905	0.524
Thin, Middle, No attack	7.582	2.802	4.636	2.998	7.56	4.332	3.958	3.7401	0.1086	0.795	9.212	0.914
Thin, Middle, Intermediate	11.372	5.712	6.33	4.31	9.438	7.132	3.958	5.5613	0.1167	0.548	8.569	0.936
Thin, Middle, Severe	13.52	5.898	6.932	4.81	8.938	6.224	3.958	2.7539	1.1053	0.512	0.905	0.434
Thick, Close, No attack	6.9	3.84	3.56	3.654	4.944	4.228	2.265	3.7401	0.1086	0.717	9.212	0.922
Thick, Close, Intermediate	5.406	5.732	5.83	4.01	7.32	6.962	2.265	5.5613	0.1167	0.493	8.569	0.943
Thick, Close, Severe	5.82	4.84	6.928	4.37	8.486	6.126	2.265	2.7539	1.1053	0.487	0.905	0.462
Thick, Middle, No attack	4.282	3.932	4.062	3.376	5.776	5.212	3.25	3.7401	0.1086	0.949	9.212	0.897
Thick, Middle, Intermediate	6.39	4.804	5.638	5.644	7.05	7.344	3.25	5.5613	0.1167	0.657	8.569	0.923
Thick, Middle, Severe	9.096	4.98	6.166	6.28	8.952	8.596	3.25	2.7539	1.1053	0.555	0.905	0.386

Table 3: Duration of experimental movement times (in seconds), Fitts' indices of performance IP_1 and IP_2, Fitts' model parameters a and b, and parameter ζ for all subjects over twelve experimental trials. Green color denotes the case when no DoS attack was mounted, yellow color the 'intermediate' DoS attack, and red color the 'severe' DoS attack. 'Thin' and 'thick' denote peg types, and 'close' and 'middle' configurations of the experimental board.

Experimental trial	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6
Thin, Close, No attack	10	9	4	6	13	0
Thin, Close, Intermediate	15	12	11	10	16	4
Thin, Close, Severe	22	16	15	12	20	6
Thin, Middle, No attack	5	4	5	11	11	4
Thin, Middle, Intermediate	19	19	7	13	17	6
Thin, Middle, Severe	28	17	10	14	13	5
Thick, Close, No attack	7	21	15	14	9	8
Thick, Close, Intermediate	9	18	16	13	20	8
Thick, Close, Severe	7	20	12	14	21	10
Thick, Middle, No attack	12	4	14	11	13	4
Thick, Middle, Intermediate	18	16	14	11	17	10
Thick, Middle, Severe	25	19	15	14	17	11

Table 4: Subjective assessments of trial difficulties. Presented is the DoS difficulty score, representing the aggregated difficulty values for four evaluated questions about task difficulties, where the difficulties ranged from 0 (easy) to 7 (hard). Green color denotes the case when no DoS attack was mounted, yellow color the 'intermediate' DoS attack, red color the 'severe' DoS attack, and blue color the case where an anomaly occurred. 'Thin' and 'thick' denote peg types, and 'close' and 'middle' configurations of the board.

7.1.4 The Telerobotic Fitts' Law

Fitts' law was originally developed for *simple movement tasks under different speed-accuracy conditions*. This assumption may not be satisfied for telerobotic Fitts' tasks[1], where in addition to a movement task, a subject needs to perform additional *fine motor tasks*, which involve picking up a cylindrical block or putting that block down on the appropriate peg.

The impact of these additional fine motor tasks on the overall task is represented by the intercept parameter a of the Fitts' model (equation (1)). When the Fitts' index of difficulty, ID, is equal to 0, no movement should be necessary to execute the tasks, since $ID = 0$ implies that the distance to the target is $D = 0$. Yet, in many such cases, measured movement time is typically strictly positive, $T > 0$, and the model (1) implies that time is governed purely by the intercept parameter.

This observation about *pure movement* and *fine motor* parts of the overall task has an impact on the Fitts' index of performance, and index values (3) and (4) may significantly differ, depending on the task. However, the observation that the overall task can be divided into different parts implies that those indices are not competing - they are simply conveying different information.

Based on the assumption that telerobotic Fitts' tasks can be divided into two parts, *pure movement* and *fine motor* tasks, we therefore propose a new metric, ζ, as a way of quantifying the overall task division between these two components, and define it as:

$$\zeta := \frac{IP_2 - IP_1}{IP_2} = 1 - \frac{IP_1}{IP_2} \tag{5}$$

where parameters IP_1 and IP_2 represent indices of performance, defined by equations (3) and (4). Parameter ζ takes on values from the range $[0, 1]$, where value $\zeta = 0$ implies that the considered task consists of pure movement tasks, and value $\zeta = 1$ that the considered task consists only of fine motor tasks, since the index of difficulty, $ID = 0$.

In many telerobotic Fitts' tasks, parameter ζ will be close to 0, implying that the fine motor tasks are only a small

[1] As well as for many other applications of Fitts' law.

component of the overall task, and the larger the value of ζ is (for the same overall task setup), the more prominent the fine motor task is in the overall task.

The defined metric ζ is especially important when evaluating the impact of cyber-security attacks on telerobotic systems, where the same attack may not affect each part equally.

In this experiment, we fit the Fitts' model (equation (1)) for each of the twelve trials. The obtained results are depicted in Table 3. We observe that when no DoS attack is mounted, the major part of the trial corresponds to the fine motor tasks (picking up the cylindrical block and positioning it down on the appropriate peg). This observation is also true for intermediate DoS attacks, where parameter ζ is larger than 0.9 for all considered cases ($\zeta = 0.955, 0.936, 0.943, 0.923$). However, this observation is not true for the cases of severe DoS. Under severe DoS attacks, parameter ζ decreases to the value of 0.5, indicating that, under attack, the movement task becomes more prominent than in benign case. This observation intuitively makes sense, since under attack, overshoot and undershoots are expected to happen quite often.

7.2 Analysis of Subjective Assessments

For each of the 12 telerobotic Fitts' tasks, we asked subjects to evaluate the difficulties of: (i) reaching each of the blocks, (ii) grabbing the blocks, (iii) moving between the pick-up and the put-down locations and (iv) performing the task as a whole, where the allowed difficulties ranged from 0 (easy) to 7 (hard). As before, we analyzed the data from six subjects.

In order to evaluate the collected results, for each DoS attack scenario and each subject, we summed up the four evaluated questions about difficulty, thus obtaining a single number as a representation of the perceived difficulty of an attack. We refer to this number as the *DoS difficulty score* [7]. The obtained results are presented in Table 4.

We observe that subjects always rated the attack scenarios as being more difficult than no attack scenarios. Moreover, the subjects almost always rated the "severe" DoS cases as being more difficult than "no DoS" or "intermediate DoS" cases. The only exceptions seem to be two cases where we start with the "intermediate" attack severity. In those two example, several users both seem to have switched the difficulty of "intermediate" and "severe" DoS case.

An interesting effect can be observed when averaging the DoS difficulty scores and trial times over all users, and comparing them for all twelve trials. The results are summarized in Table 5. For the case where no attack is mounted against the system, the average trial times for different pegboard configurations and different peg widths are decreasing (5.269, 4.985, 4.521, 4.44s), which seems to indicate there is a learning effect present. However, the averaged DoS difficulty scores increases with time (7, 6.667, 12.333, 9.667). This may indicate that the subjects are getting fatigued or annoyed (even though their performance is improving). A similar, but a less prominent trend can be observed with the "intermediate" attack as well, where the average trial times remain approximately constant, yet the DoS difficulty score is increasing (11.33, 13.5, 14, 14.33).

8. DISCUSSION

The evaluated denial-of-service attack was successful in disrupting teleoperated procedures because no mitigation mechanism against it (or many other cyber-security attacks) is currently in place. Recently, approaches have been proposed to ensure private and authenticated communication during a teleoperated surgical procedure [25]. These approaches, however, are not sufficient to prevent DoS attacks.

In the security community, prevention and mitigation of DoS attacks typically relies on a combination of malicious activity detection, network traffic monitoring and malicious traffic blocking [26]. The existing approaches can generally be divided into preventive and reactive methods [33] Some of the existing methods against DoS attacks are: (i) Intrusion Prevention Systems (ISPs) [42], (ii) DoS Defense System (DDS), (iii) blackholing [37], (iv) pipes cleaning [5], and channel surfing [48].

Intrusion Prevention Systems (ISPs)-based approaches are a type of preventive methods. They require a known attack signature in order to be able to stop the attack. These attack signatures typically use packets' content and network behavior as features of interest. However, the problem with DoS attacks is that it is relatively easy for an attacker to flood the communication channel with legitimate packets. For example, an attacker can simply capture one legitimate message between a robot and an operator, copy it multiple times and overflow the network with it. Similarly, in a distributed DoS setting, it may be very hard to distinguish between malicious and legitimate network behavior, since the attack task is spread over a large number of computers.

Blackholing [37] is a reactive DoS mitigation strategy, where all the traffic to the attacked network entity is being rerouted to a non-existent server, typically referred to as the "black hole". The problem with this approach, however, is that in rerouting all the traffic to the attacked network entity, we may end up rerouting the legitimate traffic as well, thus effectively completely preventing communication between the robot and the operator. In order for this approach to be effective, there would have to exist a way to quickly and efficiently distinguish between valid and a malicious traffic, and only reroute the malicious traffic.

Pipe cleaning [5] is another reactive DoS mitigation method. In it, all traffic is passed through a so-called "scrubbing center" where all packets are inspected and only legitimate ones are forwarded. The problem with this approach is the fact that many teleoperated robotic systems require (near) real time operation and communication, and this approach may negatively impact that requirement.

Analyzing the existing DoS mitigation strategies exposes a challenge unique to the security of teleoperated systems, namely a *tension between real-time operation and security* [7]. It therefore may be hard to find one out-of-the box approach to preventing DoS attacks against teleoperated robotic systems, and successfully use it in a variety of telerobotic scenarios. A more feasible approach may be to try to combine some of the existing proactive and reactive approaches. One such *hybrid* approach might be a mechanism to monitor link and network status. Such a mechanism should be able to

Experimental trial	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6	Average time	Average subjective rating
Thin, Close, No attack	6.784	3.576	5.054	4.118	8.12	3.964	5.269	7
Thin, Close, Intermediate	8.53	4.542	6.456	4.288	8.426	6.298	6.423	11.33
Thin, Close, Severe	10.442	4.648	5.524	4.324	9.924	6.746	6.934	15.1667
Thin, Middle, No attack	7.582	2.802	4.636	2.998	7.56	4.332	4.985	6.667
Thin, Middle, Intermediate	11.372	5.712	6.33	4.31	9.438	7.132	7.382	13.5
Thin, Middle, Severe	13.52	5.898	6.932	4.81	8.938	6.224	7.72	14.5
Thick, Close, No attack	6.9	3.84	3.56	3.654	4.944	4.228	4.521	12.333
Thick, Close, Intermediate	5.406	5.732	5.83	4.01	7.32	6.962	5.877	14
Thick, Close, Severe	5.82	4.84	6.928	4.37	8.486	6.126	6.095	14
Thick, Middle, No attack	4.282	3.932	4.062	3.376	5.776	5.212	4.44	9.667
Thick, Middle, Intermediate	6.39	4.804	5.638	5.644	7.05	7.344	6.145	14.333
Thick, Middle, Severe	9.096	4.98	6.166	6.28	8.952	8.596	7.345	16.8333

Table 5: Duration of experimental movement times (in seconds), average DoS difficulty levels and trial times (averaged across subjects). Green color denotes the case when no DoS attack was mounted, yellow color the 'intermediate' DoS attack, and red color the 'severe' DoS attack. 'Thin' and 'thick' denote peg types, and 'close' and 'middle' configurations of the board.

immediately block malicious traffic based on the recognized attack signatures. In addition, it should be able to reactively monitor all network traffic, as well as detect suspicious streams of data, and suspicious increases in the number of out-or-order packet arrivals.

9. CONCLUSION AND FUTURE WORK

In this paper, we focused on denial-of-service (DoS) attacks, based on the observation that these attacks cannot be mitigated using available cryptographic solutions. Using the *Raven II*, we experimentally investigated the impact of DoS attacks of varying severity. Our experimental results indicate there exists a learning effect across the given sequence of experimental trials, implying that human operators are capable of adapting to unfavorable network conditions. This observation, while positive for system defenders, does not imply that DoS attacks are not a problem for teleoperated robotic systems. On the contrary, it urges us to quickly develop efficient DoS mitigation methods, while indicating that in disastrous scenarios, where communication networks may inadvertently be clogged or even DoS-ed, teleoperated robotic systems will remain functional and capable of providing the necessary services.

Next steps in preventing and mitigating DoS attacks on teleoperated robotic systems include investigating the feasibility of the established DoS mitigation methods, including blackholing and sinkholing, as well as pipe cleaning, and assessing their impact on teleoperated procedures. We caution, however, that there exist tensions between cyber security, safety and usability requirements of teleoperated robotic systems which may render some these solutions infeasible. The feasibility of using a monitoring system to prevent DoS attacks against teleoperated robotic systems remains an open question.

10. ACKNOWLEDGMENTS

This work is supported by the National Science Foundation, Grant # CNS-1329751. We gratefully acknowledge Applied-Dexterity and Mr. Andrew Lewis, for their help with the *Raven II*.

11. REFERENCES

[1] Applied Dexterity (last accessed: October 20, 2014).

[2] E. Guizzo: World Robot Population Reaches 8.6 Millions (last accessed: May 14, 2014).

[3] The Health Insurance Portability and Accountability Act of 1996 (HIPPA) Privacy and Security Rules," U.S. Department of Health and Human Services (last accessed: May 14, 2014).

[4] IRF Statistical News - World Robotics News (last accessed: September 27, 2014).

[5] S. Agarwal, T. Dawson, and C. Tryfonas. DDoS Mitigation via Regional Cleaning Centers. Technical report, Sprint ATL Research Report RR04-ATL-013177, 2003.

[6] S. Amin, A. Cárdenas, and S. Sastry. Safe and Secure Networked Control Systems Under Denial-of-Service Attacks. In *Hybrid Systems: Computation and Control*, pages 31–45. 2009.

[7] T. Bonaci, J. Yan, J. Herron, T. Kohno, and H. Chizeck. Experimental Analysis of Cyber Security Attacks on Teleoperated Surgical Robotics, *submitted to the acm transactions on human-computer interaction*.

[8] A. Cárdenas, S. Amin, and S. Sastry. Research Challenges for the Security of Control Systems. In *the Proceedings of the 3rd Conference on Hot Topics in Security*, 2008.

[9] A. A. Cardenas, S. Amin, Z.-S. Lin, Y.-L. Huang, C.-Y. Huang, and S. Sastry. Attacks Against Process Control Systems: Risk Assessment, Detection, and Response. In *the Proceedings of the 6th ACM symposium on Information, Computer and Communications Security*, 2011.

[10] R. Chabukswar, Y. Mo, and B. Sinopoli. Detecting Integrity Attacks on SCADA Systems. In *the Proceedings of the 18th IFAC World Congress*, 2011.

[11] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno. Comprehensive Experimental Analyses of Automotive Attack Surfaces. In *the Proceedings of the USENIX Security Symposium*, 2011.

[12] K. Coble, W. Wang, B. Chu, and Z. Li. Secure Software Attestation for Military Telesurgical Robot Systems. In *the Proceedings of the Military Communications Conference*, 2010.

[13] C. Doarn, M. Anvari, T. Low, and T. Broderick. Evaluation of Teleoperated Surgical Robots in an Enclosed Undersea Environment. *Telemedicine and e-Health*, 15(4):325–335, 2009.

[14] N. Falliere, L. Murchu, and E. Chien. W32. Stuxnet Dossier. *White paper, Symantec Corp., Security Response*, 2011.

[15] H. Fawzi, P. Tabuada, and S. Diggavi. Secure Estimation and Control for Cyber-Physical Systems Under Adversarial Attacks. *IEEE Transactions on Automatic Control*, 59(6):1454–1467, June 2014.

[16] W. Gates. A Robot in Every Home. *Scientific American*, 296(1):58–65, 2007.

[17] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu. They Can Hear Your Heartbeats: Non-Invasive Security for Implantable Medical Devices. *ACM SIGCOMM Computer Communication Review*, 41(4):2–13, 2011.

[18] D. Halperin, T. Heydt-Benjamin, B. Ransford, S. Clark, B. Defend, W. Morgan, K. Fu, T. Kohno, and W. Maisel. Pacemakers and Implantable Cardiac Defibrillators: Software Radio Attacks and Zero-power Defenses. In *the Proceedings of the IEEE Symposium on Security and Privacy*, 2008.

[19] B. Hannaford, J. Rosen, D. Friedman, H. King, P. Roan, L. Cheng, D. Glozman, J. Ma, S. Kosari, and L. White. Raven-II: An Open Platform for Surgical Robotics Research. *IEEE Transactions on Biomedical Engineering*, 60(4):954–959, 2013.

[20] B. Harnett, C. Doarn, J. Rosen, B. Hannaford, and T. Broderick. Evaluation of Unmanned Airborne Vehicles and Mobile Robotic Telesurgery in an Extreme Environment. *Telemedicine and e-Health*, 14(6):539–544, 2008.

[21] S. Kalan, S. Chauhan, R. Coelho, M. Orvieto, I. Camacho, K. Palmer, and V. Patel. History of Robotic Surgery. *Journal of Robotic Surgery*, 4(3):141–147, 2010.

[22] H. King. *Human-Machine Collaborative Telerobotics: Computer Assistance for Manually Controlled Telesurgery and Teleoperation*. PhD thesis, 2014.

[23] H. H. King, K. Tadano, R. Donlin, D. Friedman, M. Lum, V. Asch, C. Wang, K. Kawashima, and B. Hannaford. Preliminary Protocol for Interoperable Telesurgery. In *the Proceedings of the International Conference on Advanced Robotics*, 2009.

[24] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, and H. Shacham. Experimental Security Analysis of a Modern Automobile. In *the Proceedings of the IEEE Symposium on Security and Privacy*, 2010.

[25] G. Lee and B. Thuraisingham. Cyberphysical Systems Security Applied to Telesurgical Robotics. *Computer Standards & Interfaces*, 34(1):225–229, 2012.

[26] G. Loukas and G. Öke. Protection Against Denial-of-Service Attacks: A Survey. *The Computer Journal*, page bxp078, 2009.

[27] M. Lum, D. Friedman, H. King, T. Broderick, M. Sinanan, J. Rosen, and B. Hannaford. Field Operation of a Surgical Robot via Airborne Wireless Radio Link. In *the Proceedings of the IEEE International Conference on Field and Service Robotics*, 2007.

[28] M. Lum, J. Rosen, T. Lendvay, M. Sinanan, and B. Hannaford. Effect of Time Delay on Telesurgical Performance. In *the Proceeding of the IEEE International Conference on Robotics and Automation*, 2009.

[29] I. MacKenzie. Fitts' Law as a Research and Design Tool in Human-Computer Interaction. *Human-Computer Interaction*, 7(1):91–139, 1992.

[30] L. Makris, N. Argiriou, and M. Strintzis. Network and Data Security Design for Telemedicine Applications. *Informatics for Health and Social Care*, 22(2):133–142, 1997.

[31] J. Marescaux, J. Leroy, M. Gagner, F. Rubino, D. Mutter, M. Vix, S. Butner, and M. Smith. Transatlantic Robot-assisted Telesurgery. *Nature*, 413(6854):379–380, 2001.

[32] C. B. Matlack. *Adaptation for Brain-Computer Interfaces*. PhD thesis.

[33] J. Mirkovic and P. Reiher. A Taxonomy of DDoS Attack and DDoS Defense Mechanisms. *ACM SIGCOMM Computer Communication Review*, 34(2):39–53, 2004.

[34] Y. Mo and B. Sinopoli. Secure Control Against Replay Attacks. In *the Proceedings of the 47th Annual Allerton Conference on Communication, Control, and Computing*, 2009.

[35] Y. Mo and B. Sinopoli. False Data Injection Attacks in Control Systems. In *the Proceedings of the 1st Workshop on Secure Control Systems*, 2010.

[36] F. Pasqualetti, A. Bicchi, and F. Bullo. On the Security of Linear Consensus Networks. In *the Proceedings of the 48th IEEE Conference on Decision and Control*, 2009.

[37] C. Patrikakis, M. Masikos, and O. Zouraraki. Distributed Denial of Service Attacks. *The Internet Protocol Journal*, 7(4):13–35, 2004.

[38] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Ng. ROS: An Open-source Robot Operating System. In *the Proceedings of ICRA Workshop on Open Source Software*, volume 3, 2009.

[39] J. Rosen and B. Hannaford. Doc at a Distance. *IEEE Spectrum*, 43(10):34–39, 2006.

[40] G. Sankaranarayanan, H. King, S.-Y. Ko, M. Lum, D. Friedman, J. Rosen, and B. Hannaford. Portable Surgery Master Station for Mobile Robotic Telesurgery. In *the Proceedings of the 1st International Conference on Robot Communication and Coordination*, 2007.

[41] R. Satava. Future Directions in Robotic Surgery. In *Surgical Robotics: Systems Applications and Visions*, pages 3–11. 2010.

[42] K. Scarfone and P. Mell. Guide to Intrusion Detection and Prevention Systems (IDPS). *NIST special publication*, 800(2007):94, 2007.

[43] S. Sundaram and C. N. Hadjicostis. Distributed Function Calculation via Linear Iterations in the Presence of Malicious Agents - Part I: Attacking the Network. In *the Proceedings of the American Control Conference*, 2008.

[44] A. Tiwari, B. Dutertre, D. Jovanović, T. de Candia, P. D. Lincoln, J. Rushby, D. Sadigh, and S. Seshia. Safety Envelope for Security. In *the Proceedings of the 3rd International Conference on High Confidence Networked Systems*, 2014.

[45] M. Tozal, Y. Wang, E. Al-Shaer, K. Sarac, B. Thuraisingham, and B.-T. Chu. On Secure and Resilient Telesurgery Communications over Unreliable Networks. In *the Proceedings of the IEEE Conference on Computer Communications Workshops*, 2011.

[46] M. Tozal, Y. Wang, E. Al-Shaer, K. Sarac, B. Thuraisingham, and B.-T. Chu. Adaptive Information Coding for Secure and Reliable Wireless Telesurgery Communications. *Mobile Networks and Applications*, 18(5):697–711, 2013.

[47] F. Wozak, T. Schabetsberger, and E. Ammmenwerth. End-to-end Security in Telemedical Networks–A Practical Guideline. *International Journal of Medical Informatics*, 76(5):484–490, 2007.

[48] W. Xu, K. Ma, W. Trappe, and Y. Zhang. Jamming Sensor Networks: Attack and Defense Strategies. *IEEE Network*, 20(3):41–47, 2006.

[49] A. Yoo, G. Gilbert, and T. Broderick. Military Robotic Combat Casualty Extraction and Care. In *Surgical Robotics: Systems Applications and Visions*. 2010.