

Journal of Medical Imaging

MedicalImaging.SPIEDigitalLibrary.org

Toward real-time remote processing of laparoscopic video

Zahra Ronaghi
Edward B. Duffy
David M. Kwartowitz

Toward real-time remote processing of laparoscopic video

Zahra Ronaghi,^a Edward B. Duffy,^b and David M. Kwartowitz^{a,*}

^aClemson University, Department of Bioengineering, 301 Rhodes Research Center, Clemson, South Carolina, 29634-0905, United States

^bClemson University, Clemson Computing and Information Technology, Barre Hall, 120 McGinty Court, Clemson, South Carolina 29634, United States

Abstract. Laparoscopic surgery is a minimally invasive surgical technique where surgeons insert a small video camera into the patient's body to visualize internal organs and use small tools to perform surgical procedures. However, the benefit of small incisions has a drawback of limited visualization of subsurface tissues, which can lead to navigational challenges in the delivering of therapy. Image-guided surgery uses the images to map subsurface structures and can reduce the limitations of laparoscopic surgery. One particular laparoscopic camera system of interest is the vision system of the daVinci-Si robotic surgical system (Intuitive Surgical, Sunnyvale, California). The video streams generate approximately 360 MB of data per second, demonstrating a trend toward increased data sizes in medicine, primarily due to higher-resolution video cameras and imaging equipment. Processing this data on a bedside PC has become challenging and a high-performance computing (HPC) environment may not always be available at the point of care. To process this data on remote HPC clusters at the typical 30 frames per second (fps) rate, it is required that each 11.9 MB video frame be processed by a server and returned within 1/30th of a second. The ability to acquire, process, and visualize data in real time is essential for the performance of complex tasks as well as minimizing risk to the patient. As a result, utilizing high-speed networks to access computing clusters will lead to real-time medical image processing and improve surgical experiences by providing real-time augmented laparoscopic data. We have performed image processing algorithms on a high-definition head phantom video (1920 × 1080 pixels) and transferred the video using a message passing interface. The total transfer time is around 53 ms or 19 fps. We will optimize and parallelize these algorithms to reduce the total time to 30 ms. © 2015 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JMI.2.4.045002](https://doi.org/10.1117/1.JMI.2.4.045002)]

Keywords: laparoscopy; image-guided surgery; medical imaging; high-performance computing; high-throughput networking.

Paper 15060PRR received Mar. 23, 2015; accepted for publication Nov. 3, 2015; published online Dec. 14, 2015.

1 Introduction

Minimally invasive surgery reduces patient trauma, hospitalization stay, and recovery time through reducing incisions to small keyholes. In laparoscopic procedures, a camera is used for real-time visualization of the surgical field and guidance of the procedures. However, the small field-of-view of the laparoscope as well as small incision size will result in a small overall visual field of the underlying tissues of interest. The reduced visual field along with organ motion due to natural and surgical effects can limit the surgeon's ability to target specific subsurface locations with high accuracy and precision, therefore, requiring external information such as medical imaging to complement laparoscopic video.^{1,2}

Image-guided surgery (IGS) uses preoperative and intraoperative images to map a surgical region of interest, thus providing surgeons with visualization of subsurface structures and accurate positional information during a surgical procedure.^{3,4} The accurate coregistration of an IGS system with laparoscopic video can reduce the impact of limited surgical access and allow for a resection with higher specificity and tighter margins, thus sparing more healthy tissues. As a result, augmentation and mapping imaging data to the operative field-of-view can improve the accuracy and efficiency of surgical procedures.^{5–8}

After the introduction of the Computer Motion AESOP and the Intuitive Surgical daVinci, both in 1998, the use of robotics in medicine has continued to grow. Some of the clinical benefits of medical robotics are improvement in dexterity and accuracy of manipulation compared with traditional laparoscopy.⁸ Currently, the primary robotic surgery devices on the market are the daVinci-si and daVinci-xi systems. For this work, we will look at the daVinci-si robot, with an high-definition (HD) vision cart, which uses two parallel 1080p (1920 × 1080 × 3) high-definition video cameras to generate stereoscopic vision. These two views are then displayed using a pair of high-definition monitors set up in a stereopticon-like device, thus allowing the clinician the ability to perceive depth.^{9,10} If captured raw, these video streams generate approximately 360 MB of data per second.

Processing of the volume of data generated by the stereo video feeds would be demanding for a local system, thus high-performance computing (HPC) hardware is needed,^{11–13} which can be both expensive and large in size. Even though the price of computing hardware continues to fall, the price of HPC hardware can limit it to only hospitals with adequate financial resources, while size considerations can make hardware difficult to place in any operating room. As the size of the data generated by surgery increases from enhanced video

*Address all correspondence to: David M. Kwartowitz, E-mail: robodoc@clemson.edu

resolutions and new instrumentation in the operating room, the need to process and store this data is increasingly becoming a problem that requires efficient connectivity to HPC systems at remote locations.

We will develop image processing algorithms for augmentation of laparoscopic video with preoperative and intraoperative medical images on HPC clusters. Computing clusters are often not available in hospital environments, thus data will need to be transported from the point of care to be processed and the results will need to be returned in real time. To realize the goal of rapid processing of laparoscopic data, we propose the development of a protocol where data are captured, transmitted, processed, returned, and displayed within 1/30th of a second for the left and right eye synchronizations of the daVinci.

Clemson University has recently deployed the OpenFlow software-defined network (SDN) communication protocol on its network and is in the process of leveraging the new protocol and its connection to the Internet2 Innovation Platform. SDN enables programmable network control for high-bandwidth and dynamic applications.¹⁴ The new network will assist researchers in overcoming the network limitation to flexibly establish 10 GB or higher end-to-end connectivity, within a scientific demilitarized zone (DMZ), thus creating a frictionless connection between machines.¹⁵ As a result, using high-speed networks to access computing clusters will lead to real-time or near real-time medical image processing with great flexibility.

2 Methods

Prior to the design of the overall system, network latency, bandwidth, and speed were tested in a series of time-of-flight experiments. For our study, we developed three separate applications that run simultaneously, which serve the purposes of video acquisition (`video_client`), video display (`video_server`), and image processing (`video_server_palmetto`). Currently, a single PC is used for the acquisition and display of images (`video_client` and `video_server`). This machine is outfitted with a Piccolo Alert (Euresys Inc., San Juan Capistrano, California) SD frame grabber board, and connectivity is provided over a 10 Gbps Ethernet connection. Our video processing module (`video_server_palmetto`) is running on the Palmetto compute cluster, which is rated number 4 among academic research clusters. Palmetto contains 1978 compute nodes (20,728 cores) and 598 NVIDIA Tesla GPU accelerators.¹⁶ The distance between the client (our local computer) and server (Palmetto Cluster) is 9 miles. The connectivity from the outside to the Palmetto cluster is 10 Gbps. The communication within the Palmetto cluster is provided over InfiniBand or Myrinet. There is a 100 Gbps connection to the Internet2 innovation platform, which will provide rapid connectivity between campuses. We have tested our transfers using our local area network and plan on connecting to multiple institutions through Internet2. Preliminary network tests were conducted using the workflow shown in Fig. 1.

In our workflow, we have currently tested different protocols and methods such as the secure file transfer protocol (SFTP), user datagram protocol (UDP), and message passing interface (MPI) for the transfer of captured frames from the client machine (our local computer) to the HPC server (Palmetto cluster). We have currently measured the round-trip time for captured images of 57 KB and have compared our results using the 1 Gbps and updated 10 Gbps network connections. We are trying to minimize network overhead and latency, so

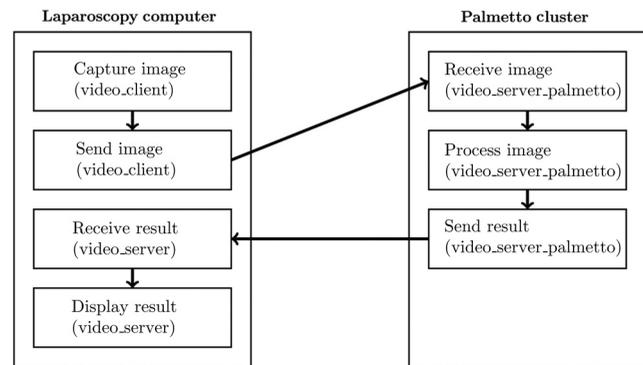


Fig. 1 Workflow for network tests, demonstrating the connectivity between the local laparoscopy computer (left) and the HPC Palmetto cluster (right).

we can use the HPC resources to improve the visualization and accuracy of laparoscopic video and provide surgeons with real-time augmented laparoscopic data during surgical procedures.

3 Results

Data from the time-of-flight tests for video frames using the SFTP file transmission protocol showed an increase in transmission time as the frame number increased for both 1 Gbps and 10 Gbps networks, except when the protocol had completed sending previous data which appear as valleys in the graphs. However, the round-trip time decreased by a factor of 5 to 7 when using the 10 Gbps network instead of the 1 Gbps network. Figure 2(a) shows the file transfer using the 1 Gbps connection. The round-trip time is around 5 s for the first file and 35 s for the last. Figure 2(b) indicates the transmission using the 10 Gbps network. The round-trip time starts at 1 s for the first frame and 5 s for the 100th frame.

After connecting directly to the compute node, the round-trip time decreased to 4 ms and remained around the same value for all the frames. Figure 3 indicates the round-trip time in seconds between our local computer (client) and the Palmetto cluster (server).

SFTP runs over transmission control protocol (TCP), which offers end-to-end connection, ordered data transmission, and is reliable. However, the connection-oriented communication causes delays. On the other hand, UDP creates a connectionless communication that will decrease the transfer time.¹⁷ This could result in packet loss, depending on the network bandwidth and connectivity. We transferred the files using UDP, and the transmission time decreased to 0.7 ms with no packet loss using the 10 Gbps network. Figure 4 indicates the round-trip time using UDP.

UDP has a packet size limit of 64 KB, and larger files have to get reassembled before processing. TCP has the same size limit. UDP and TCP are both in the transport layer of computer networks.¹⁷ As a result, we focused on the application layer and protocols that use UDP or a TCP connection, but handle the reassembly of large files. We tested our transfers using MPI. MPI is a message passing system and communications protocol that can be used for parallel programming.¹⁸ We transferred the files using this method, and the results indicate that the transmission time is around 1.2 ms. Figure 5 shows the round-trip time using MPI.

We also tested our transfers using MPI for 6-MB files to measure the transmission time for HD videos. The results indicate

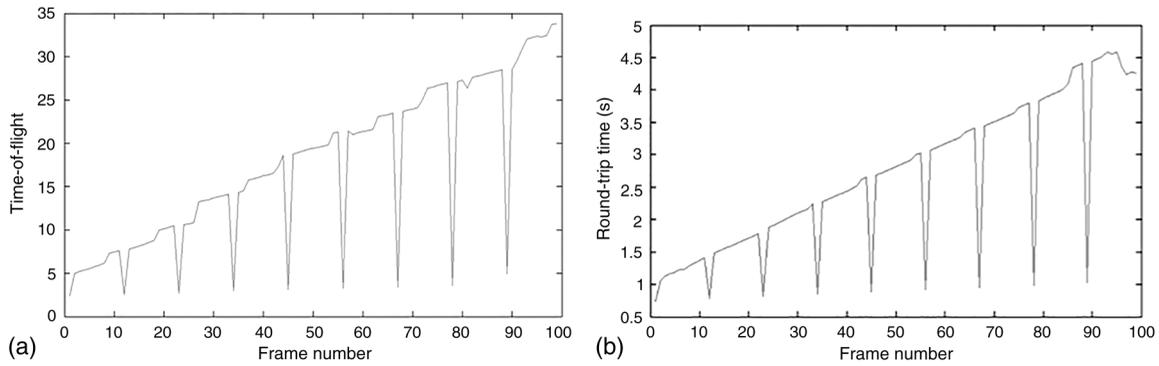


Fig. 2 Round-trip time for SFTP using (a) 1 Gbps, and (b) 10 Gbps. Note the difference in y-axis scales between the graphs.

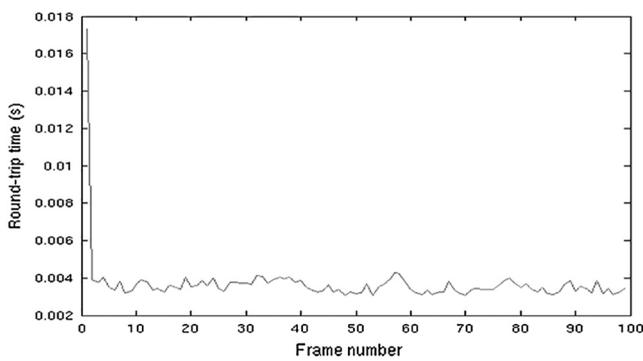


Fig. 3 Round-trip time for SFTP through compute node using 10 Gbps.

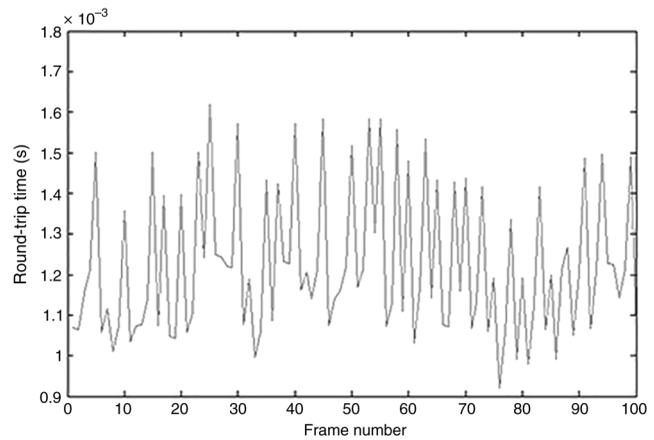


Fig. 5 Round-trip time for MPI using 10 Gbps.

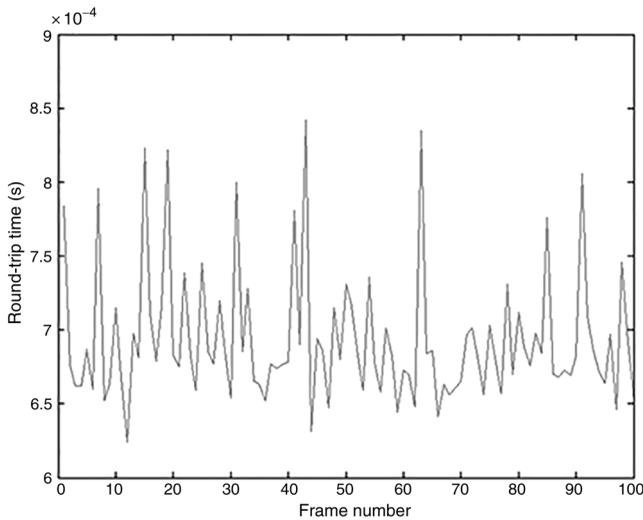


Fig. 4 Round-trip time for UDP using 10 Gbps.

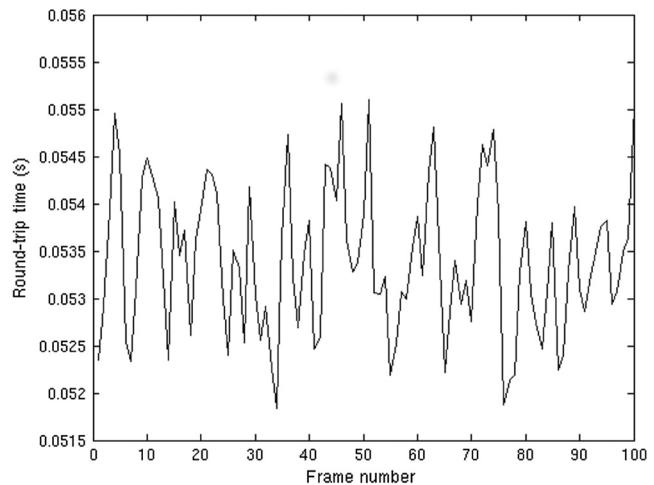


Fig. 6 Round-trip time for MPI using 10 Gbps for 6 MB files.

that the transmission is around 53 ms. Figure 6 shows the round-trip time using MPI for 6-MB files.

Figure 7 shows an original image of an HD head phantom video and the segmented image. The canny edge segmentation was processed on the cluster, and the video was transferred using MPI.

4 Discussion

In the field of computer assisted-surgery, particularly in IGS, the ability to acquire, process, render, and visualize data in real time is essential for the performance of complex tasks, thus minimizing risk to the patient.¹⁹ The contextual validity of data degrades

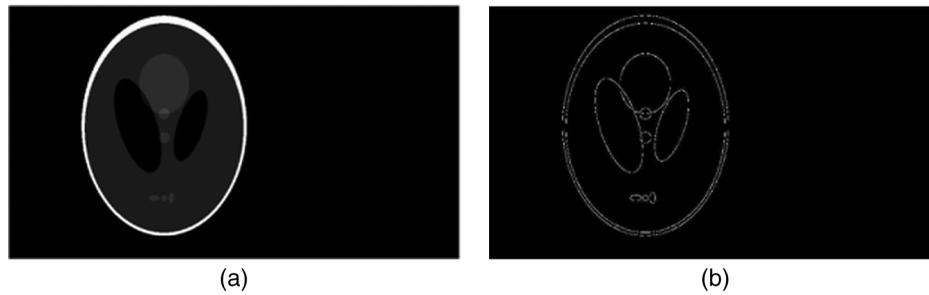


Fig. 7 (a) Original head phantom and (b) segmented images.

quickly as tools within the body move, patient organs move, and resections are performed. It is critical that the data be acquired, processed, and returned in a timely manner.²⁰

Initial results from the use of SFTP ruled out its use due to high latency in time-of-flight measurements as well as a noticeable increase in latency over the course of the test. During our tests, we realized that access to the underlying file system and the fact that we must connect to each computing node over SSH through a head node, which is accessed by a large number of users, is causing significant latency issues. We were able to reduce the latency through direct connection to the compute node. However, SFTP runs over TCP which uses acknowledgments and retransmission to prevent packet loss and causes delays. The encryption step of SFTP is also a limiting factor in reducing the round-trip time.

UDP is faster since there is no error recovery and acknowledgment mechanism, therefore, it is used for time-sensitive and real-time applications. The UDP tests resulted in a lower round-trip time and reliable connection with no packet loss due to availability of high bandwidth. However, UDP and TCP are both in the transport layer and have a packet size limit of 64 KB; hence, larger files have to get reassembled before processing. As a result, we focused on the application layer and protocols that use UDP or TCP connection, but handle the reassembly of large files. We tested our transfers using MPI, which is based on a TCP connection and is used for parallel programming, larger files, and HD video frames. Despite MPI using TCP connection, the round-trip time is 1.2 ms comparable with large files. In addition, the round-trip time for MPI is in the order of milliseconds compared with seconds for SFTP.

The total transfer for HD video frames is around 53 ms using MPI. We will optimize and parallelize our algorithms to reduce the time. Through the development of a direct connection, we will achieve more direct control over of the transmission of data as well as develop a communication system between clients and computing nodes to perform real-time medical image processing.

5 Conclusion

The augmentation of live high-definition stereo laparoscopic data with preoperative or intraoperative medical images requires more computational and storage capability than is possible on a single workstation. More advanced visualizations and processing may require new or improved rendering techniques more suited to an HPC environment. Therefore, flexibility in the transmission of data to multiple sites is important. Furthermore, the ability to optimize the data transmission to take advantage of the specialized processing and rendering techniques at multiple

endpoints is the key to allowing for rendering and processing as a service and also for providing optimal data.

Through the use of high-speed access to an HPC resource, we believe that it will be possible to update models and allow for real-time modification of the previously generated object map, based on current surgical data. We have performed canny edge segmentation on an HD head phantom and transferred the video using MPI. The total transfer time is around 53 ms. We will optimize and parallelize these algorithms to reduce the total time to 30 ms. We also plan to develop registration algorithms and send previously generated augmented data to the cluster for adaption and deformation. Success in this portion would lead to the ability of performing surgery at one site (origin), receiving an object map from another site (segmentation/rendering), forwarding this data along with relevant force data to a third site for deformation (model update), and returning this data to the origin for display.

Acknowledgments

This work was supported by the National Science Foundation under Grant ACI-1245936. The authors would like to thank Dr. Kuang-Ching Wang, Dr. Barr von Oehsen, Mr. James Bottum, Ms. Kathryn Mace, Mr. Christopher A. Konger, Mr. Brian Parker, Mr. Joseph Bernard, and Mr. Randall Martin for their help with the network infrastructure and Mr. Fuad Mefleh from the Cutters lab for his assistance.

References

1. V. Vitiello et al., "Emerging robotic platforms for minimally invasive surgery," *IEEE Rev. Biomed. Eng.* **6**, 111–126 (2013).
2. D. J. Mirota, M. Ishii, and G. D. Hager, "Vision-based navigation in image-guided interventions," *Annu. Rev. Biomed. Eng.* **13**, 297–319 (2011).
3. W. Grimson et al., "Image-guided surgery," *Sci. Am.* **280**(6), 54–61 (1999).
4. T. M. Peters, "Image-guidance for surgical procedures," *Phys. Med. Biol.* **51**(14), R505 (2006).
5. H. Fuchs et al., "Augmented Reality Visualization for Laparoscopic Surgery," *LNCS* **1496**, 934–943 (1998).
6. M. Baumhauer et al., "Navigation in endoscopic soft tissue surgery: perspectives and limitations," *J. Endourol.* **22**(4), 751–766 (2008).
7. D. Van Krevelen and R. Poelman, "A survey of augmented reality technologies, applications and limitations," *Int. J. Virtual Reality* **9**(2), 1 (2010).
8. S.-L. Lee et al., "From medical images to minimally invasive intervention: computer assistance for robotic surgery," *Comput. Med. Imaging Graphics* **34**(1), 33–45 (2010).
9. D. M. Kwartowitz et al., "Towards image guided robotic surgery: multi-arm tracking through hybrid localization," *Int. J. Comput. Assisted Radiol. Surg.* **4**(3), 281–286 (2009).

10. D. M. Kwartowitz, S. D. Herrell, and R. L. Galloway, "Toward image-guided robotic surgery: determining intrinsic accuracy of the daVinci robot," *Int. J. Comput. Assisted Radiol. Surg.* **1**(3), 157–165 (2006).
11. R. Shams et al., "A survey of medical image registration on multicore and the GPU," *IEEE Signal Process. Mag.* **27**(2), 50–60 (2010).
12. S. S. Samant et al., "High performance computing for deformable image registration: towards a new paradigm in adaptive radiotherapy," *Med. Phys.* **35**(8), 3546–3553 (2008).
13. O. Fluck et al., "A survey of medical image registration on graphics hardware," *Comput. Methods Programs Biomed.* **104**(3), e45–e57 (2011).
14. Open Networking Foundation, "Software defined networking (SDN) definition," 2015, <https://www.opennetworking.org/sdn-resources/sdn-definition> (24 November 2015).
15. Brocade, "100 gigabit ethernet: the way forward for research networks," 2014, <http://info.brocade.com/rs/brocade/images/CDE14%20100GbE%20BRIEF%20Brocade.PDF> (24 November 2015).
16. Cyberinfrastructure Technology Integration (CITI) group at Clemson University, "Overview of the Palmetto cluster," 2015, <http://citi.clemson.edu/palmetto/pages/userguide.html> (24 November 2015).
17. J. F. Kurose and K. W. Ross, *Computer Networking: A Top-Down Approach Featuring the Internet*, 3rd ed., Addison Wesley, Upper Saddle River, New Jersey (2004).
18. W. Gropp, E. Lusk, and A. Skjellum, *Using MPI: Portable Parallel Programming with the Message-Passing Interface*, Vol. **1**, MIT Press, Cambridge, Massachusetts (1999).
19. A. Eklund et al., "Medical image processing on the GPU—past, present and future," *Med. Image Anal.* **17**(8), 1073–1094 (2013).
20. Y. Kawasaki et al., "High-performance computing service over the internet for intraoperative image processing," *IEEE Trans. Inf. Technol. Biomed.* **8**(1), 36–46 (2004).

Zahra Ronaghi is a doctoral student in the Department of Bio-engineering at Clemson University. She received her BS and MS degrees in electrical engineering in 2010 and 2013, respectively. Her research interests include robotic surgery, image-guided surgery, high-performance computing, and computer networks. She is a member of SPIE.

Edward B. Duffy is a computational scientist in Clemson University's Cyberinfrastructure Technology Integration (CITI) group. He has more than 10 years of experience working in high-performance computing (HPC). As director of Advanced Computing Research & Methods for CITI, he is responsible for overseeing and HPC application support, research, and outreach. His primary research interests include optimization for high-performance applications, compiler technologies, and programming languages.

David M. Kwartowitz is an assistant professor of bioengineering at Clemson University in Clemson, South Carolina. He has a PhD in biomedical engineering from Vanderbilt University in 2007. His research interests include robotic surgery, image-guided surgery, image-based soft tissue analysis, instrumentation, computer networks, and wearable computing. He sits on the program committee for SPIE: Medical Imagings Conference on Image-Guided Procedures, Robotic Interventions, and Modeling. He is a member of SPIE.