DESIGN AND DEVELOPMENT OF THE DATA SYNCHRONIZATION CASE STUDY

Michael Fuller
Graduate Student,
Dept. of Comp. Science & Software Engineering
Auburn University

Chetan S. Sankar
Thomas Walter Professor
Department of Management
Auburn University

P.K. Raju
Thomas Walter Professor
Dept. of Mechanical Engineering
Auburn University

Abstract

The Laboratory for Innovative Technology and Engineering Education (LITEE) at Auburn University has been developing case studies that engage students in cross-disciplinary learning and require students of multiple disciplines to work together in order to solve a common problem. The Data Synchronization case study is one of the latest case studies from LITEE, which focuses on how a construction company solved the problem of synchronizing data forms among remote sites and the headquarters. Unlike other case studies that handle hypothetical or past events, the Data Synchronization case study focuses on a current real-world problem that existed in this company. Michael Fuller, a graduate student from Auburn University worked directly with the company while living in India during summer 2007 in order to develop this case study. The Data Synchronization case study was made using steps for creating case studies, which involve identifying a company's problem, having students work with the company, writing the case study, and many stages of refinement in order to create the finished product for use in the classroom. This paper further describes this case study and its creation process.

Introduction

Preeminence in technological innovation depends on a wide array of factors, one of which is leadership in engineering research, education, and practice. As other nations increase their investments in engineering research and education, the U.S. risks falling behind in critical research capabilities and, ultimately, the innovations that flow from research[7]. The nation's ability to capitalize on new knowledge resulting from large investments in sciences will depend on contributions from engineering. Engineering research is founded on a disciplined approach to problem solving and the application of sophisticated modeling, design, and testing tools to solve problems. The *Educating the Engineer of 2020* report (2005) calls for system-wide efforts to align the engineering curriculum and engineering profession with the needs of today's global, knowledge-driven economy, with the goal of increasing student interest in engineering careers. It has also been recommended that research be combined with education, thereby training students in critical thinking and research methodologies, as well as providing them with solid engineering skills[7].

As more industries utilize the economic advantages of a global R&D, U.S. engineering teams need to prepare for collaboration across countries and the blurring of national boundaries. Future engineers need to be trained not only in basic engineering skills, but also in managing global research teams[8]. Thus, engineering education needs to be drastically altered to give students opportunities to work in international research teams[1,3,4, 6,7].

Engineering educators should introduce interdisciplinary learning in the curriculum wherever possible and explore the use of case studies of engineering successes and failures as a learning tool[2]. These recommendations lead to the premise upon which this paper is based: New challenges and opportunities are emerging due to the emergence of global R&D teams and future engineers must be given opportunities to learn how to perform effectively in this market.

The LITEE lab has been producing case studies for many years, such as Lorn Textiles, Della Steam Plant, STS-51L and others (www.litee.org), but they have always focused on events that occurred in or involved companies in the United States (Figure 1). In 2007, LITEE was awarded an NSF grant to send students from American universities to India to live, work, and produce a case study. Before the group of five students ventured to India, the project leaders of the LITEE lab were in communication with multiple private companies located in India and IIT Madras, asking them for projects that students may work on. This paper describes the experiences of a graduate student who travelled to India to develop a case study with a large construction company.

Project Selection

In summer 2007, five graduate students from Auburn University conducted a research project in India. Two faculty members assisted them in conducting the project. The private companies and IIT Madras had multiple project options available for students to select, but the students waited until they arrived in India to select a project, because more diversified and interesting options became available when the students spoke directly with each individual project managers at the various organizations.

When the group of students and a faculty member arrived at Larsen & Toubro, L&T, the largest construction firm in India, they were given a brief tour of the campus, and then they met with different project managers to pick individual problems to work on and eventually wrote a case study (Figure 2). The students were taken to the individual departments and spoke directly with the manager of each section. The company managers described the problems they faced, students described their expertise, and through these discussions, students and managers were matched. A few of the proposed projects were extremely technical, such as trying to solve a mini power crisis in India that led to computers shorting out frequently. This project was interesting, but the expertise of the students was not sufficient to work on this problem.

Michael Fuller, the graduate student, was specializing in computer science and engineering, so when the opportunity to work in a software development environment arose, he chose to work on that project.

The Problem

The Senior Manager-System, ECC Division (Engineering Construction & Contracts) at L&T described wanting to solve a problem related to establishing communication to remote worksites

Figure 1: Sample LITEE case study listing.

Figure 2: 2007 U.S. team at IIT Madras.

spread throughout India and synchronize the databases on both ends. The student was given an overview and a computer to use and then told to work on the problem.

Larsen & Toubro is one of India's largest and most respected engineering and construction conglomerates seeking a strong customer-relationship approach and constant quest for top-class quality. With over 350 job sites in India, keeping track of each construction project's progress is essential. That said, Larsen & Toubro have encountered several delays in a construction site's projects that could potentially cost the multi-million dollar business. These delays have lasted anywhere from eight hours to a couple of days for processing work order approval forms.

The current communication system allows each job site to manually package and send information over the Internet. This process requires human intervention and has a long lag time between updates. So if a work order has been submitted, there is a high probability that approval for this form will not take place until either the next login or even the following day. Therefore, Larsen & Toubro wanted to develop a central database that would drastically reduce the processing time for approval for construction projects to commence. Before a work order is approved, all the resources at the remote site sit idle, waiting for approval. So reducing the approval delay will also reduce the amount of time that equipment sits idle at the job site. Larsen & Toubro requested that all solutions must:

- enhance security
- remove manual transmission
- reduce turnaround time
- reduce lag time
- reduce bandwidth, and
- verify information

Researching the Problem

Michael Fuller stayed at the guest house at the Indian Institute of Technology, Madras, and commuted to the L&T campus, about 5 miles apart on a daily basis. Initially, the student researched Web Services and how they operate. Creating the Web Service was easy, because all the required tools were already provided with Microsoft Visual Studio 2005. At L&T, the local network was almost at its capacity, so they wanted to ensure that data transfers were small in size and that they could be verified for accuracy, integrity, size and secure from outside viewers.

Integrity was the first challenge and needed to ensure that the entire data transfer was correct. The issue was solved by introducing a keyed-Hash Message Authentication Code (HMAC) hashing function. Providing a way to check the integrity of information transmitted over or stored in an unreliable medium is a prime necessity in the world of open computing and communications. Mechanisms that provide such an integrity check based on a secret key are usually called "message authentication codes" (MAC). Typically, message authentication codes are used between two parties that share a secret key in order to validate information transmitted between these parties. A variation of the MAC mechanism based on cryptographic hash functions called HMAC, is based on work by Krawczyk, et al[5].

The algorithm devised by the student hashes information with the secret value and then includes the hash with the transfer. On the receiving end, software will again hash the information with the known value and then compare it to the included hash value. If both values match, then the information is free of bit errors and comes from a trusted source.

The **accuracy** problem was not handled by the software developed by the student, but instead was completed by a separate team at L&T. There had been problems in the past with work order forms with faulty information. When a faulty work order form was submitted it would cause a cascade failure, and all other forms that accompanied it would be rejected, and the process to find out if an error occurred might take a day to realize. This problem was solved by forcing the custom ERP software that L&T used to verify that the information was entered correctly before saving it. If a request was made for a tool or part that did not exist then it would not allow the form to be saved until the faulty information was fixed.

In India, the **size** of information matters. High speed Internet connections exist, but in the rural areas, access to the Internet may be limited. The student originally thought he would also have to solve the connection problem, but after a discussion with his supervisor, they said to expect that all sites would at least have a dial-up connection. When working on a dial-up connection or on L&T's congested network, sending a large, plain-text file over the network is not feasible. The solution to this problem was to shrink the information using a two-part strategy.

First, the information that would normally be in plain text was converted to a binary format. For example, when storing numbers it would be assumed to save the number 255 as a text version, but converting it to binary converts it to a single byte instead of three. The format also assumes that all known data types are known up front, like integers, floating points, bytes and strings. By knowing the exact format of information, the system can reduce space by using a special format. Before each value, a data description byte was introduced. This byte indentifies the data's type and also the size if necessary. For byte-sized characters, only two bytes would be necessary for identification and storage. More complex objects, such as strings, are stored by using an identification block with size information and then the string. To help protect information from unauthorized viewers, random information is inserted into the data stream that are automatically ignored when read back but will foil people that don’t understand the format or are trying to break the encryption. This binary file format came from the student’s experience with working and reading the Adobe Flash 8 file format, which follows closely with the above rule set.

The second half of the process of shrinking information involved using a compression technique to further compact the data. For this part, the binary information from the previous step was fed into a zip compression function that took the binary data and reduced the size by approximately half. This resulting binary data was then converted to base 64. The base 64 format converted bytes into a string representation. This conversion increased the file size by one third, but was needed to put binary data in a Web Service.

The **security** aspect of the Web Service could have come in two different kinds, either encrypting the Web Service or the information. Trying to encrypt the Web Service would have necessitated writing connecting software that was deemed to be very difficult, therefore the graduate student instead chose to encrypt the data. Before the data was compressed, an encryption technique was used on the data. This function increased the overall size of the data, but since it was later compressed, the size difference was minimal.

Figure 3 shows how information transferred with the new system was secured and verified. The organization of information is like an onion and resembles TCP/IP stack format for packets.

Developing the code for the Web Service, documenting , and testing took about a month. The student spent the second month of the assignment drafting the first versions of a case study describing this problem.

Framing the case study

A few years back, the Della Steam Plant case study was created and it became the standard for how LITEE case studies should be formatted and presented. The format has been modified and adjusted over the years, but it forms the basis for new case studies that are developed by the student team. The section that follows describes the steps taken to develop the Data Synchronization case study. This case study is available for perusal at www.liteecases.com. Students can access the pages of the case study online and can browse between the pages to understand the case study.

Content of the case study

The Data Synchronization case study begins with a quick introduction of the problem at hand. After the introduction, a problem statement is provided where more information and details are presented for the perspective student. The next sections focus more on the individual problems (such as integrity, accuracy, size, and security), with each problem having a page dedicated to it and links to more information. The idea is to fully immerse the student in the subject so that they can explain the problem clearly to others and make informed decisions. After all the problems have been thoroughly explained, the alternatives to solve the problems are provided. The alternatives are setup in a way so that one of them is a wrong solution (no automation), another is the better solution that solves the problems adequately (developing an automated solution similar to the one described in the earlier section) and a few stop-gap solutions (buying a 3rd party software and using itinstead of open software) that do not solve the problems adequately.

Student Assignments

The assignment portion provides scenarios which instructors can use to ensure that students have comprehended the problem. For the Data Synchronization case study, two assignments were developed. The first assignment divided the class into four teams. Three teams were required to defend specific alternatives and convince the fourth team, representing a management group, that a specific alternative was better. Such an assignment makes it

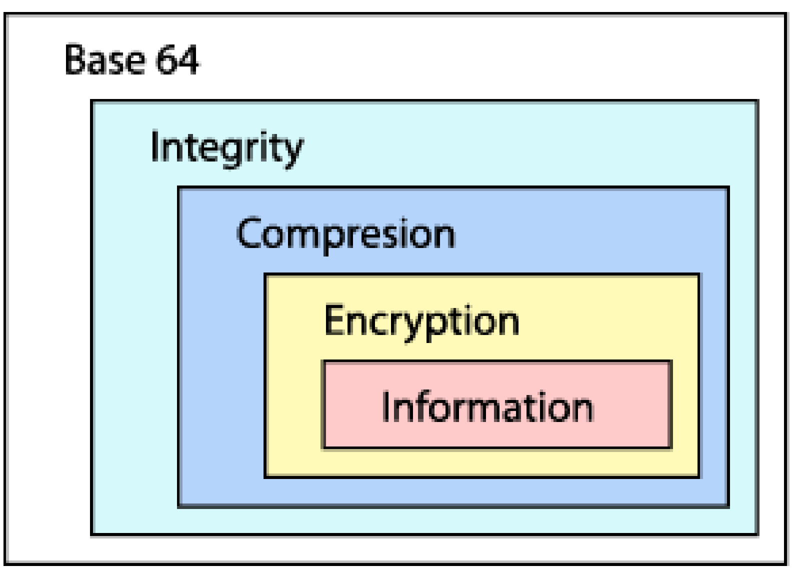

Figure 3: Layers of information for data transfers, based on the TCP/IP stack.

possible for students to fully know their subject and develop the skills to help communicate their ideas to others.

The second assignment is based upon a fictional event where a remote job site near the coastline is waiting for a work order to be released before they can pour concrete (Figure 4). The issue is that a cyclone is expected to land at the job site and if the concrete is poured too late, then it will not have sufficient time to harden. The assignment is crafted to force students to take the side of fixing or ignoring the problem and the results of their actions. Implementation of the appropriate alternative is essential in order to solve the above problem.

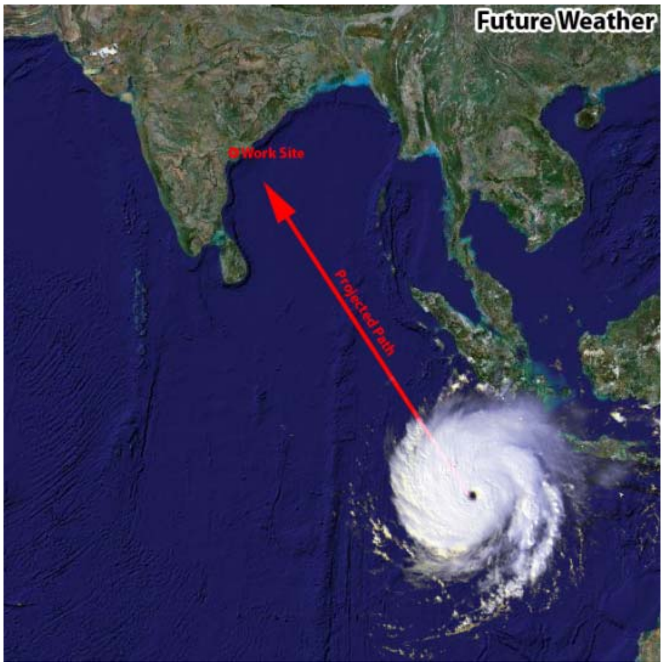

Figure 4: Motivation for assignment 2.

Supplementary materials

The case study also includes supplementary information for students. A reference section was included to explain the difficult technical concepts. For example, an encryption demonstration is attached which allows students to test out a simple encryption scheme and thereby get a better understanding of how it works. Other examples included are hashing functions, such as MD5, CRC-32, and a HMAC-MD5 function, which could all be used to check integrity of the data. With these examples students can dynamically see how hashing works. Other topics covered include:

- Base (Base 2, Base 16, Base 64)
- Compression
- Encryption
- Endians (Data formats)
- Flash Format
- Integrity
- Internet
- Packing
- Synchronization
- Verification
- Web Services

Search function

A new function was introduced in this case study that allows students to find information faster. A search menu was added that shows results that matched the requested search query (Figure 5).

Site map

The last portion of the case study is the site map. Students can use the site map to see all available pages and quickly identify content. The data is organized into categories and pages, and pages may show up multiple times in different categories. For example, all the content related to encryption is located under a single column for quick reference.

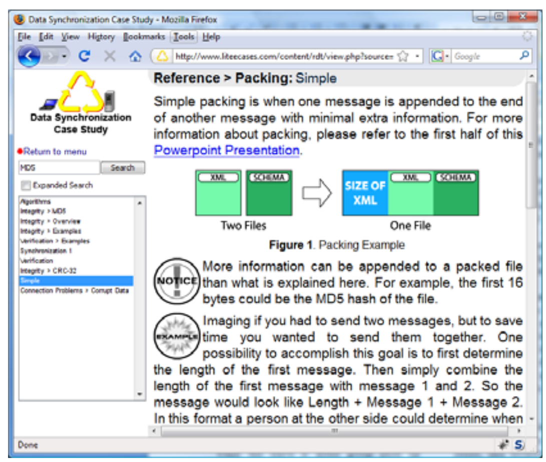

Figure 5: Search pane and results.

Lost sections

The author of the case study had to exclude other information that focused exclusively on other issues, such as security and integrity. These information were removed, but the technical content stayed in the references section since it was still useful in analyzing the overall case study.

Multiple solutions

The Data Synchronization case study presents multiple alternatives so that the students can choose an appropriate one.

Picking alternatives

For the Data Synchronization case study, the alternatives were designed so that the students would analyze the situation with India in mind. Because software engineers are paid less in India compared to the U.S., it made sense for a few alternatives to focus on training employees to enhance the process. The training aspect incorporated some technical aspects to reach the company's goals, such as encryption. Employees could encrypt by hand or use a software tool that would handle the encryption process.

Each alternative costs the company different sums of money. Some of the cost can come from training personnel, but there is another cost to be considered: programming hours. L&T employs many talented software developers. If these employees are requested to work on this project, then it costs L&T nothing more, but it might slow down other projects. L&T could ask a third party to develop the required applications, but then the cost of development and support will cost money, not man hours.

The idea of the alternatives is to show students not only what is possible when the cost of labor is relatively cheap when compared to the U.S., but also the trade-offs involved when a company uses people instead of technology to solve problems. Better training has the potential to improve efficiency, but still doesn't solve a few of the issues related to time delay. In the end, alternatives are ways to guide students to think in a specific direction and help better understand the overall concept in solving the problems faced by this company.

Presentation

Before leaving India, the graduate student who worked on this project presented his findings to the executives and managers at L&T and to the faculty members from Auburn University and IIT Madras. This made it possible for the student to show what he did, to demonstrate the case study, and finally to see if the management approved of his work.

The presentation for the Data Synchronization case study was held at L&T's Chennai headquarters in the presence of the Vice President of Operations and Human Resources. The student presented a mini synopsis of the case study with each page cut down to a few bullet points. The management liked the research project and the presentation and gave the go-ahead to finalize the case study.

Refinement

Every case study, when initially written, is not a finished product. After returning from India, the case study went through major rewrites before it was finalized.

Rewrites

The Data Synchronization case study didn't have as many re-writes as the other cases, but there were substantial changes. The content was altered so that it could be better accessed by students. One issue that the case study format introduced was the over-nesting of menu items. Over-nesting is when a menu item is located in a difficult place to access. To solve menu nesting, some content items were combined and others moved to a more convenient location.

The case study was also simplified so that the salient points were brought out. The Data Synchronization case study as originally written had two extra case studies included that had their own separate menu items. However, these were eventually removed from the case study because students might become confused when looking at the menu and not know which case study to examine. The extra case studies were originally moved into different menu links and later removed from the case study entirely.

Other than removing extra content, the case study's main focus stayed the same during editing, but more information was added to better explain the points. Diagrams and more detailed explanations were added where necessary to better detail the problem and solution.

Multiple editors

The case study, after months of edits and weekly update reports, was almost finished. Then a student that had not worked on the case and had a strong editing capability was requested to examine and find all the flaws that the designer missed. This part was a bit painful to the developer because he got back sheets of paper with marks indicating what alterations needed to be made. This process was repeated a few times until no further changes were required.

Approval from L&T

Every case study has to have approval from the sponsoring company before release. With the Data Synchronization case study it was no different; CD-ROM versions were created and sent to L&T in India. A faculty member visited the Senior Manager-System, ECC Division and discussed the case study with him. The Senior Manager was very impressed with the case study and agreed to approve it. It was further sent to the Executive Vice President, who finally approved the case study for release.

The Senior Manager, System, ECC Division, was delighted with the quality of the software developed by the student. He stated that L&T had an opportunity to use this code in a project during November-December 2007. A plant in Perungudi, about 20 KM from the headquarters, was mixing cement from various raw materials and loading it into trucks. L&T management requested that the Information Systems team control this process from headquarters. A team led by Manager (Systems), modified the code developed by the student and implemented the system during December 2007.

Implementation of the new system led to several major improvements for L&T. First, because the company's computers are synchronized and they complete the entire process automatically, data from each order is guaranteed to be error-free. Additionally, the computers are able to mix raw materials in the exact quantities needed for each order, thus guaranteeing the quality of each order and eliminating the need for time-consuming and costly quality-assurance testing. The benefits of the new system have greatly increased the company's sales and profits.

During early 2008, the concrete plant was sold to another company, and this company has been discussing with L&T the possibility of implementing a similar data synchronization system across its 57 plants around India.

Case study distribution

One problem with all previous case studies has been getting them to the public. Before 2008, the only way to obtain the CD-ROM copies was to purchase a book or buy them directly from a publisher. In 2008, LITEE started to use www.lulu.com/litee_cases for single-copy distribution, but it still left something to be desired. Another solution was to provide free access to LITEE case studies online. The process to put case studies online isn't simple, but newer case studies are built with this in mind.

Online format

With the shift to online distribution, the case studies of tomorrow have needed major re-adjustments to properly function. The Data Synchronization case study was created with the template to ensure it was online-ready. It did not contain any videos, but used still images instead, so the required alterations were simple and eventually developed into the case study format.

The online distribution system hides the file names through obfuscation, so the case studies search and menu features had to be altered to handle blocked pages. This was achieved by adding in extra JavaScript checks that determined if it's currently being used on the web or from a local folder.

One innovation of the online format has been the modularization of components. Previous case studies that were initially created with a custom website creator had to be later changed to the new format. The modular design enabled the lead case study developer to make critical changes easily and spread the modifications to other developers. Modularization has also led to the separation of the content and presentation layers. The content, when written, is simple and easy to edit, but the resulting pages viewed by students look completely different. This separation allows the developers to focus on the content, and the display is handled by the viewer's computer. It also cuts back on bandwidth requirements by not constantly retransmitting the presentation code on every request.

The purpose of the online format is to have case studies that can be viewed anywhere on machines with different operating systems, such as Linux, Mac, and Windows. Some of the older case studies have required extensive re-writes to get them in the right format, but the benefits are wider distribution and lower costs since CD-ROMs don't have to be produced and shipped.

Online distribution

The LITEE case studies are available for distribution at the website: www.liteecases.com. The first case study placed on the website was the Data Synchronization case study. It wasn't viewable by the general public, because it wasn't released yet, but it was used to test the platform. Once L&T approved the case study, this case study was released and is now available for use by faculty members and students.

Results

The Data Synchronization case study marked a new era for LITEE; it introduced an interesting case study and new technology to help distribute content to students around the world. Because of the technology available today to distribute complex information on-line, the case study format was made online ready. The case study writing experience was enjoyable for the grad student involved for the student because he got to live in India, work on an interesting problem, and document the system.

Acknowledgements

This project was partially funded by NSF DUE #0439706, #0736997, OISE #0623351. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation.

We thank Mr. K.P. Raghavan, Executive Vice President (Corporate Center), ECC Division, Mr. P. Rengarajan, Senior-Manager, System, and Mr. G.D. Sharma, Vice President, Human Resources, L&T for sponsoring and approving this case study. We also thank Dr. Ramachandraiah Professor, Department of Civil Engineering, Indian Institute of Technology, Madras, for coordinating the local arrangements.

References

1. Accreditation Board for Engineering and Technology, Inc., www.abet.org/criteria. html, 2009.

2. Educating the Engineer of 2020: Adapting Engineering Education to the New Century, National Academy of Engineering,

http://www.nap.edu/catalog/11338.html, Accessed Jan. 10, 2006.

3. Engardio, P., Bernstein, A., and Kripalani, M. (2003). "The New Global Job Shift," Business Week, 2003, Issue 3818: 50.

4. Felker, G.B. (2003). "Southeast Asian Industrialisation and the Changing Global Production System," Third World Quarterly - Journal of Emerging Areas, 24(2): 255-282.

5. H. Krawczyk, M. Bellare, and R. Canetti (1997). RFC2104 - HMAC: Keyed-Hashing for Message Authentication, Retrieved from http://www.faqs.org/rfcs/rfc2104.html, January 22, 2009.

6. Kripalani, M. and Engardio, P. (2003). "The Rise of India," Businessweek, Dec. 8, 2003, 66-76.

7. National Academy of Engineering, The Engineer of 2020: Visions of Engineering in the New Century, 2005, http://books.nap.edu/catalog/10999.html, Accessed Jan. 5, 2006.

8. Pennoni, C.R. (1998). "Managing Your Career in an Era of Change," Journal of Professional Issues in Engineering Education and Practice, 124(3): 75-78.

9. Wulf, W.A., and Fisher, G.M.C. "A Makeover for Engineering Education," Issues in Science and Technology, 18(3): 35-39, Spring 2002.

Biographical Information

Michael Fuller, a graduate student at Auburn University, has been working with the LITEE Laboratory since 2005, designed the new XML case study format, designed and built liteecases.com and is working on methods to bring LITEE case studies to mobile platforms. He can be reached at fullemg@auburn.edu.

Chetan S. Sankar is the Thomas Walter Professor of Management at Auburn University. He has received more than 2.5 million dollars from ten National Science Foundation and Economic Development Administration grants to develop exceptional instructional materials that bring real-world issues into classrooms and to help develop geospatial systems in coastal areas. He has won awards for research and teaching excellence from the Society for Information Management, NEEDS, Decision Sciences Institute, American Society for Engineering Education, American Society for Mechanical Engineering, International Network for Engineering Education & Research, and the Project Management Institute. He is a co-director of the Laboratory for Innovative Technology and Engineering Education (LITEE). He is the editor-in-chief of the *Decision Sciences Journal of Innovative Education* and the managing editor of the *Journal of STEM Education: Innovations and Research*. He can be contacted at sankacs@auburn.edu.

P.K. Raju is the Thomas Walter Distinguished Professor of Mechanical Engineering at Auburn University. He has made significant research contributions in acoustics, noise control, nondestructive evaluation, and engineering education, resulting in award-winning and significant breakthroughs. He has received a total of $8.4 million in funding, including grants from industries, the United Nations, the National Science Foundation, and other U.S. Federal agencies. He has published 17 books, 8 book chapters and 160 papers in journals and conference proceedings. He received several awards for his work in the area of engineering education from NEEDS & John Wiley, ASME, ASEE and others. He has been an invited or a keynote speaker at several engineering education conferences organized in USA, France, Germany, Chile, Singapore and in India. He is the director of the Laboratory for Innovative Technology and Engineering Education (LITEE) and the Auburn Engineering Technical Assistance Program (AETAP). He is a Fellow of the American Society for Engineering Education, Fellow of the American Society of Mechanical Engineers, Fellow of the Institution of Engineers India and Fellow of the Acoustical Society of India. In order to promote engineering education research, he also edits and publishes the *Journal of STEM Education: Innovations and Research* (www.jstem.org). He can be reached at rajupol@auburn.edu.