
AC 2011-1800: ADMINISTERING A DIGITAL LOGIC CONCEPT INVENTORY AT MULTIPLE INSTITUTIONS

Geoffrey L. Herman, University of Illinois at Urbana-Champaign

Geoffrey L. Herman is a PhD Candidate in Electrical and Computer Engineering and a Mavis Future Faculty Fellow at the University of Illinois at Urbana-Champaign. His research interests include cognitive science, identifying and assessing common student misconceptions and difficulties in electrical and computer engineering topics, blended learning (integrating online teaching tools into the classroom), intelligent tutoring systems, and music signal processing. He is a winner of the 2011 Educational Research and Methods Division Apprentice Faculty Grant. He has been recognized with the Olesen Award for Excellence in Undergraduate Teaching for the Department of Electrical and Computer Engineering and the Ernest A. Reid Fellowship for engineering education. He has served as a graduate affiliate for the Center for Teaching Excellence, as the seminar chair for the Dean's Graduate Student Advisory Committee, and as the information chair for the American Society for Engineering Education's Student Constituent Committee.

Michael C. Loui, University of Illinois, Urbana-Champaign

Michael C. Loui is Professor of Electrical and Computer Engineering and University Distinguished Teacher-Scholar at the University of Illinois at Urbana-Champaign. His interests include computational complexity theory, professional ethics, and the scholarship of teaching and learning. He serves as Executive Editor of *College Teaching*, and as a member of the editorial board of *Accountability in Research*. He is a Carnegie Scholar and an IEEE Fellow. Professor Loui was Associate Dean of the Graduate College at Illinois from 1996 to 2000. He directed the theory of computing program at the National Science Foundation from 1990 to 1991. He earned the Ph.D. at M.I.T. in 1980.

Administering a Digital Logic Concept Inventory at Multiple Institutions

Abstract

A concept inventory is a standard assessment tool that evaluates how well a student's conceptual framework matches the accepted conceptual framework of a discipline. In this paper, we describe our development cycle to create a digital logic concept inventory. We show that the concept inventory is a reliable and valid instrument even when administered at multiple institutions across the United States of America and can be used to evaluate the effectiveness of different pedagogies.

1. Introduction

Instructors in electrical and computer engineering and in computer science have developed innovative methods to teach digital logic circuits. These methods attempt to increase student learning, satisfaction, and retention. Although there are readily accessible and accepted means for measuring satisfaction and retention, there are no widely accepted means for assessing student learning. Rigorous assessment of learning is elusive because differences in topic coverage, curriculum and course goals, and examination content prevent direct comparison of teaching methods. Because of these difficulties, computing educators have issued a general call for the adoption of standard assessment tools to critically evaluate and compare the various teaching methods^{4,17}.

Measuring students' conceptual knowledge upon completion of a course is a common accepted benchmark for comparing student learning^{14,19,21}. Conceptual knowledge is often favored because all courses should teach a fundamental set of concepts even if they emphasize design or analysis to different degrees. Increasing conceptual learning is also important, because students who can organize facts and ideas within a consistent conceptual framework are able to learn new information quickly and can apply what they know in new situations⁵. If instructors can accurately assess their students' conceptual knowledge, they can target instructional interventions to remedy common problems. To properly assess conceptual learning, many engineering disciplines have developed concept inventories (CIs)^{1,2,7,15,16,19-22}. CIs are multiple-choice assessment tools that evaluate how well a student's conceptual framework matches the accepted conceptual framework of a discipline or common faulty conceptual frameworks.

As has been seen in the physics education community and the success of the first CI – the Force Concept Inventory (FCI) – CIs have power to motivate the adoption of new teaching techniques¹⁴. The FCI provided evidence that interactive engagement pedagogies increased learning compared with traditional lectures at all levels of instruction of introductory physics⁸. Since those pedagogical studies took place, introductory physics education has rapidly adapted new interactive pedagogies¹⁸.

In response to the call for standard assessment in computing education and to enable the comparison of conceptual learning, we have developed the Digital Logic Concept Inventory (DLCI). In this paper, we report on the creation of the DLCI, particularly the public dissemination of version $\beta 1.0$ at several institutions. The construction and dissemination of the

DLCI were supported through funding under the Course, Curriculum, and Laboratory Improvement Program of the National Science Foundation.

To maintain the integrity of the instrument, we do not provide the whole DLCI. The DLCI is available only by request to the authors.

2. Background

To better explain the purpose of the DLCI, we first define what a CI is and what it is not.

- A CI is a short, multiple-choice test that can classify a student as someone who thinks in accordance with accepted conceptions in a discipline or in accordance with common misconceptions²⁰.
- A CI is a standard test. It must meet the demands of statistical analysis and be broadly applicable to many programs. A CI covers each concept multiple times to strengthen the validity and reliability of its measurement. This requirement contrasts with a typical classroom exam, which may cover each concept only once during the exam. To be considered a successful instrument, a CI must also be approved by content experts and be widely adopted.
- A CI is not a comprehensive test of everything a student should know about a topic after instruction. A CI selectively tests only critical concepts of a subject¹⁹. If students demonstrate understanding of these critical concepts, then it is reasonable to believe they satisfactorily understand all other concepts of the subject. For example, the FCI tests only a student's knowledge of force after a course in mechanics, which also covers topics such as inertia, momentum, and energy¹⁴.
- A CI may complement but not replace a final examination, because a CI is not comprehensive.
- A CI is not a teacher evaluation. CIs are intended to measure the effectiveness of teaching methods independent of teacher qualifications^{1,8}. As such, a CI can stimulate the adoption of new pedagogies as it provides an objective measure to compare pedagogies.
- A CI evaluates students' conceptual knowledge. For example, the DLCI measures how much a student's conceptual framework matches the accepted conceptual framework of the discipline. It is not intended to evaluate students' problem solving skills, design skills, analytical skills, or interpersonal skills.

3. Creation of the DLCI

The process for developing a CI is often described as a three to five step process^{19,20}. We present our CI development model (see Figure 1) which synthesizes these other development models.

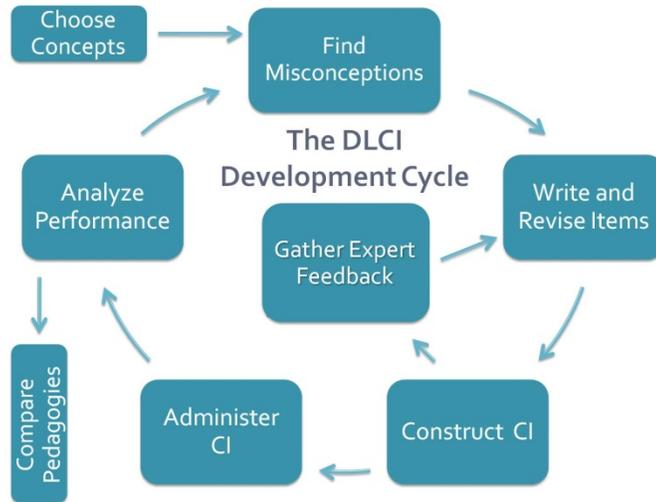


Figure 1: Flowchart for the development of the DLCI

Step 1. Choosing concepts: CI developers must carefully choose which concepts will be assessed in a CI to ensure appropriate content validity. To create a standardized, validated assessment tool, domain experts must widely acknowledge that the tool assesses what it claims to assess. By soliciting the opinions of experts from the beginning of our development process, we can show that our CI assesses core concepts and establish that our CI has appropriate content validity³. By selecting a set of concepts that is included in most courses, we can also encourage the adoption of the CI at many institutions.

A CI is typically administered as both a *pre-test* at the beginning of a course and a *post-test* at the end, to measure the conceptual “gain” created by instruction. A CI measures this gain without comprehensively testing all significant course topics, but it provides a quick snapshot of students' beliefs about core concepts of a discipline. For example, many CIs contain between 20 and 30 multiple-choice questions. As a result, the scope of the test must be determined carefully to include an indicative subset of core concepts that are important and that distinguish students who have a strong conceptual understanding from those who do not. This subset of concepts must be viewed as important, difficult, and central by the *instructors*, so they will adopt the CI.

Step 2. Identifying misconceptions: While instructors can often identify the topics that students struggle to understand, their knowledge can sometimes be incomplete or inaccurate. For example, the FCI revealed that students struggle to learn the *force* concept much more than their instructors initially believed¹⁴. Instructors also may not fully know which misconceptions are prevalent among their students or even how students misunderstand core concepts. Students must be interviewed to determine which topics are truly difficult and why students fail to understand these core concepts correctly. If possible, these interviews should catalogue specific, identifiable misconceptions about standard problems in the discipline.

Step 3. Write concept inventory items and draft the concept inventory: Using data from Step 2, CI developers should construct multiple-choice questions (items) whose incorrect answers (distracters) correspond to students' common misconceptions. Developers should construct a CI from these items. For the sake of the reliability of the CI, they should ensure that the CI tests

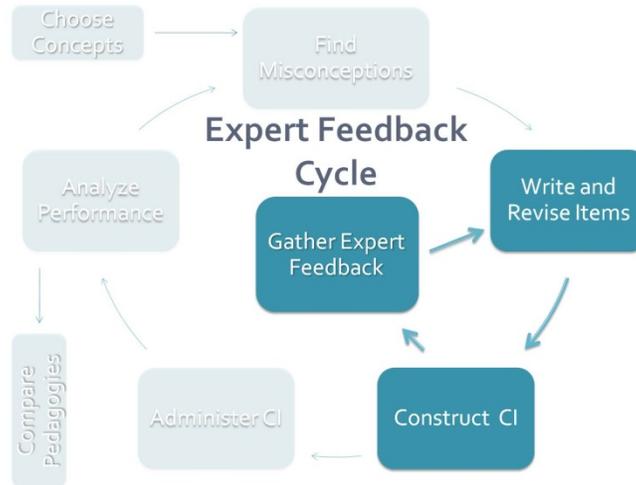


Figure 3: Expert feedback cycle for the development of the DLCI

coast) and three small private colleges (one in the Midwest and two in the South). We have also administered the DLCI in both Electrical and Computer Engineering courses and Computer Science courses. These institutions provide a stratified, representative sampling of students across the country. A total of 688 students have taken version $\beta 1.0$ of the DLCI.

5. Reliability and Validity

According to Classical Test Theory, an assessment tool called a test is intended to estimate an examinee's ability or true score T along a single attribute³. An examinee's score after a single administration of an instrument is called their observed score X . The observed score is assumed to be comprised of the examinee's true score plus some level of error E .

$$X = T + E$$

Instrument developers must try to find an examinee's true score by limiting the error of the measurement. The instrument's error can be reduced by ensuring that the instrument provides consistent measurements of each examinee and ensuring that the instrument measures what it claims to measure³. This first requirement is called the *reliability* of the instrument. The second requirement is called the *validity* of the instrument³. In order for a CI to be broadly applicable and useful, it must be both reliable and valid³.

Reliability is often estimated by three methods: test-retest reliability, split-half reliability, and the Cronbach α . Test-retest reliability estimates the reliability of the CI measurement by requiring students to take the CI multiple times in close succession³. Let a_i be student i 's observed score on the first administration of the CI, and let b_i be student i 's observed score on the second administration of the CI. Furthermore, let A ($A = \langle a_1, a_2, \dots, a_{I-1}, a_I \rangle$) and B ($B = \langle b_1, b_2, \dots, b_{I-1}, b_I \rangle$) be the sample population's scores on the first and second administrations respectively, where I is the number of students who took both administrations of the CI. The reliability of the CI is estimated by the correlation coefficient between A and B . Although test-retest reliability often provides a good estimate for the reliability of instruments, it

is often not performed. Test-retest is rare, because it is time consuming, and its results are confounded because students typically learn a little bit while taking an instrument the first time³.

Alternatively, reliability can be estimated through split-half reliability³. Split-half reliability is similar to test-retest reliability, except the instrument is split randomly into two sub-tests: SubT1 and SubT2. Split-half reliability treats each sub-test as a separate administration of the instrument and then estimates the reliability of the instrument by correlating examinees' observed scores on the two sub-tests just like in test-retest reliability. We define our correlation coefficient to be r where the relative strength of different correlations are defined in Table 1. We should expect that a good CI will demonstrate strong to very strong, positive correlations with statistical significance set at $p = 0.01$.

Table 1: Definitions for the strength of correlations in split-half reliability measurements

Strength of correlation	Range of $ r $
Weak or no correlation	$0 \leq r < 0.4$
Moderate correlation	$0.4 \leq r < 0.6$
Strong correlation	$0.6 \leq r < 0.8$
Very strong correlation	$0.8 \leq r \leq 1$

Perhaps the most commonly used estimate of reliability is the Cronbach α . The Cronbach α essentially finds the average split-half reliability of every possible set of sub-tests³. Let K be the number of items on the instrument. Let σ_X^2 be the variance of the observed total test scores, and let $\sigma_{Y_j}^2$ be the variance of item j for the sample of examinees. Cronbach α is defined as

$$\alpha = \frac{K}{K-1} \left(1 - \frac{\sum_{j=1}^K \sigma_{Y_j}^2}{\sigma_X^2} \right)$$

The Cronbach α varies from -1 to 1 like a correlation co-efficient. A Cronbach α greater than 0.70 is an acceptable level of reliability for a CI. CIs need a high level of reliability to be used as research instruments, but a degree of inconsistency is acceptable because students often apply their conceptual knowledge inconsistently.

The *validity* of the CI, broadly defined, is the degree to which the instrument tests what it claims to test². We claim that the DLCI tests how much a student's conceptual framework in digital logic matches the accepted disciplinary framework. By testing a core set of concepts, we hope to create an estimate of a students' overall conceptual understanding of digital logic concepts. Since there is no established metric for measuring students' conceptual knowledge in digital logic, we must use non-statistical means to establish the validity of the instrument.

Validity can be established through face validity and content validity. Face validity refers to what the test appears to measure. In other words, in order for the DLCI to have face validity, any person who is familiar with digital logic should believe that the DLCI tests digital logic conceptual knowledge at first glance. Content validity is more rigorous and can be established by systematically polling the opinions of experts to see if they believe that the instrument

measures what we claim. The face and content validity of the DLCI has been established in previous studies¹⁰. Here we describe how we created the DLCI to demonstrate how we created a valid instrument and to situate our current work.

6. Results of the DLCI administrations

Table 2 presents the list of DLCI items, the mean score on each item, and what concept each item assessed. Because students were reluctant to take the DLCI multiple times in close succession, we did not measure the test-retest reliability of the DLCI. Instead, we estimated the reliability of the DLCI by using split-half reliability and Cronbach α . First, we paired items together based on what concepts they covered. We then randomly assigned one item in the pair to the first sub-test (SubT1) and the other item in the pair to the second sub-test (SubT2). Each sub-test was composed of 11 items. Item 24 was not included in either sub-test because over 100

Table 2: Mean score, concept coverage, and example split-half reliability assignments

Item	Mean	Concept	SubT1	SubT2
1	0.65	Boolean operators (negated variables)		x
2	0.57	Relationship between states and flip-flops		x
3	0.56	Number bases	x	
4	0.45	Decoders and multiplexers		x
5	0.71	Boolean operators (3-variable XOR)	x	
6	0.56	Relationship between states and flip-flops	x	
7	0.37	Time dependence of state	x	
8	0.85	Boolean operators (negated variables)	x	
9	0.36	Boolean operators (NAND and NOR)		x
10	0.65	Two's complement representation		x
11	0.31	Decoders and multiplexers	x	
12	0.69	RAM inputs and outputs		x
13	0.56	RAM inputs and outputs	x	
14	0.46	Two's complement representation	x	
15	0.59	Time dependence of state		x
16	0.51	Number bases		x
17	0.75	Relationship between states and flip-flops	x	
18	0.82	Relationship between states and flip-flops		x
19	0.69	Boolean operators (NAND and NOR)	x	
20	0.39	Boolean operators (3-variable XOR)		x
21	0.48	Boolean operators (Don't Cares)		
22	0.35	Two's complement representation	x	
23	0.47	Two's complement representation		x
24	0.50	Decoders and multiplexers		
Total	13.26		6.15	6.13

students did not receive Item 24 due to a printing error, and Item 21 was also not included in either sub-test because it lacked a suitable pair. Each item's assignment to each sub-test can be seen in Table 2. Students' scores on SubT1 and SubT2 revealed a statistically-significant, strong-correlation ($r = 0.68$, $p < 0.01$). This split-half reliability check reveals that two DLCI sub-tests that cover similar concepts provide similar estimates of a students' conceptual understanding.

Second, we estimated the reliability of the test with the Cronbach α . The Cronbach α essentially finds the average split-half reliability of every possible set of sub-tests. These sub-tests are created without regard for the concepts covered by each item. We found a strong level of reliability with $\alpha = 0.75$ (the Cronbach α was the same with and without inclusion of item 24). Based on the Cronbach α and the conceptual split-half tests, we believe that the DLCI provides a reliable measurement of students' conceptual understanding of digital logic concepts.

Results from the DLCI administrations reveal that the items span a range of difficulty levels. The difficulty level of an item is estimated by the mean score of the item (μ_i). The items' difficulties are evenly distributed over a range of very difficult items ($\mu_i = 0.31$) to very easy items ($\mu_i = 0.85$). No items are overly difficult and none are trivial.

Table 3 presents the frequency with which each distracter and correct answer (highlighted in gray) was selected by the students (Note: there is a slight discrepancy between the percentages in

Tables 2 and 3, because one institution reported only whether students answered each item correctly or incorrectly. All other institutions provided the specific answers that students selected). The results in Table 3 reveal that all distracters were chosen by the students even on the easier items. Fourteen of 24 items also have compelling and widespread misconceptions where compelling and widespread misconceptions are defined as distracters chosen by more than 20% of students.

7. Conclusions and future work

We have demonstrated that the DLCI is both valid and reliable. The DLCI was constructed to assess concepts deemed important and difficult by a panel of digital logic experts, the distracters were grounded in misconceptions identified through rigorous research, and the DLCI has proven to be reliable even when administered at different types of institutions and different departments across the United States.

Since we have shown that the DLCI is a reliable and valid estimator of students' conceptual understanding in digital logic, there are many new directions for research and development. There are two main categories of future research: pedagogy studies and assessment tool development.

We plan to conduct a large scale pedagogical study where we will use the results from the DLCI to compare the effectiveness of different teaching methods. We will look for statistically significant differences in performance between institutions on the DLCI as a whole, on conceptual subtests, and individual items. We will then collect artifacts of instruction (syllabi, assignments, examinations, etc.) in an attempt to identify how differences in instruction affect conceptual learning.

Table 3: Selection frequency of all answer choices (correct answers are highlighted)

	a	b	c	d	e	f	g
1	17%	5%	12%	66%			
2	57%	26%	14%	3%			
3	5%	25%	56%	14%			
4	28%	13%	45%	13%			
5	72%	20%	6%	1%			
6	4%	9%	12%	12%	56%	7%	
7	37%	12%	16%	35%			
8	13%	85%	1%	2%			
9	36%	12%	39%	12%			
10	13%	14%	11%	62%			
11	28%	2%	2%	3%	37%	16%	11%
12	76%	7%	14%	1%			
13	4%	30%	11%	55%			
14	36%	44%	9%	11%			
15	8%	17%	60%	6%	9%		
16	49%	26%	16%	7%			
17	4%	4%	6%	8%	75%	3%	
18	6%	82%	4%	6%	2%	1%	
19	6%	14%	10%	69%			
20	19%	18%	39%	25%			
21	16%	49%	29%	6%			
22	8%	36%	27%	26%			
23	22%	46%	21%	7%			
24	6%	10%	48%	6%			

We can also investigate ways to improve the DLCI or develop other useful assessment tools. A primary short-coming of CIs is that they are typically given only at the end of an academic term. Consequently, instructors can receive rigorous feedback on the quality of students' conceptual learning only after instruction has ended. This feedback can be useful for reformulating a course for future terms, but the feedback does not benefit the students who took the CI as much. We will investigate whether we can develop shorter, conceptual sub-tests that can be used throughout the semester to provide more timely feedback that can benefit an instructor's current and future students.

With data from multiple institutions we can also begin to check for institution bias in the DLCI. We want to make sure that the DLCI does not favor students from our home institution or any other institution. Our misconceptions research has revealed that different institutions use different terminology or different notation. We will begin to look more rigorously at how these differences affect learning and performance on the DLCI.

Acknowledgements

Thank you to the instructors who gave us access to their students and provided feedback on the DLCI.

This work was supported by the National Science Foundation under grants DUE-0618589. The opinions, findings, and conclusions do not necessarily reflect the views of the National Science Foundation or the University of Illinois at Urbana-Champaign.

References

- [1] Almstrum, V. L., Henderson, P. B., Harvey, V., Heeren, C., Marion, W., Riedesel, C., Soh, L.-K. & Tew, A. E. (2006). Concept inventories in computer science for the topic discrete mathematics. ITiCSE '06: Proceedings of the 11th Annual SIGCSE Conference on Innovation and Technology in Computer Science Education, ACM press, 163-167.
- [2] Allen, K. (2006). *The Statistics Concept Inventory: Development and analysis of a cognitive assessment instrument in statistics*. Doctoral Dissertation, The University of Oklahoma.
- [3] Allen, M. J. & Yen, W. M. (2002). *Introduction to Measurement Theory*. Waveland Press, Inc.
- [4] Ben-Ari, M. (2005). The Concorde doesn't fly anymore. In *the Technical Symposium on Computer Science Education (SIGCSE)*, keynote address.
- [5] Bransford, J. D., Brown, A. L., & Cocking, R. R. (1999). *How People Learn: Brain, Mind, Experience, and School*. National Academy Press, Washington, DC.
- [6] Goldman, K., Gross, P., Heeren, C., Herman, G., Kaczmarczyk, L., Loui, M. C. & Zilles, C. (2010). Setting the scope of concept inventories for introductory computing subjects. *ACM Transactions on Computing Education*, 10, 5:1-29.
- [7] Gray, G. L., Costanzo, F., Evans, D., Cornwell, P., Self, B., & Lane, J. L. (2005). The dynamics concept inventory assessment test: A progress report and some results. In *Proceedings of the American Society for Engineering Education Annual Conference and Exposition*, Portland, Ore., June 12–15, 2005 (Session 3268).
- [8] Hake, R. (1998). Interactive-engagement vs. traditional methods: A six-thousand-student survey of mechanics test data for introductory physics courses, *Am. J. Physics*, 66, 64–74.
- [9] Herman, G. L. & Handzik, J. (2010) A preliminary pedagogical comparison study using the Digital Logic Concept Inventory. In *Proceedings of the 40th ASEE/IEEE Frontiers in Education Conference*, Arlington, VA. October 27-30, 2010, FIG-1-6.
- [10] Herman, G. L., Loui, M. C., & Zilles, C. (2010). Creating the digital logic concept inventory, *Proceedings of the Forty-First ACM Technical Symposium on Computer Science Education*, Milwaukee, Wisc., March 10–13, 2010, 102-106.
- [11] Herman, G. L., Loui, M. C., & Zilles, C. (submitted). Flip-Flops in students' conceptions of state, *IEEE Transactions on Education*.
- [12] Herman, G.L., Loui, M.C., & Zilles, C. (in press). Students' misconceptions about medium scale integration circuits, *IEEE Transactions on Education*.
- [13] Herman, G. L., Kaczmarczyk, L., Loui, M. C., & Zilles, C. (2008). Proof by incomplete enumeration and other logical misconceptions, in *ICER '08: Proceeding of the Fourth international Workshop on Computing Education Research*, ACM, 59-70.
- [14] Hestenes, D., Wells, M. and Swackhamer, G. (1992) Force Concept Inventory, *The Physics Teacher*, 30, 141–158.

- [15] Jacobi, A., Martin, J., Mitchell, J. & Newell, T. (2003), A concept inventory for heat transfer. In *Proceedings of the Thirty-Third ASEE/IEEE Frontiers in Education*, Boulder, CO, November 5-8, T3D-12-16.
- [16] Krause, S., Decker, J.C., & Griffin, R. (2003). Using a materials concept inventory to assess conceptual gain in introductory materials engineering courses. In *Proceedings of the Thirty-Third ASEE/IEEE Frontiers in Education*, Boulder, CO, November 5-8, T3D-7-11.
- [17] Lister, R. (2008). After the gold rush: Toward sustainable scholarship in computing, in *the Proceedings of the Tenth Australasian Computing Education Conference (ACE2008)*, Wollongong, Australia, January 2008, 3-17.
- [18] Mestre, J. P. (2005). Facts and myths about pedagogies of engagement in science learning *Peer Review*, 7, 24-27.
- [19] Olds, B. M.; Streveler, R. A.; Miller, R. L. & Nelson, M. A. (2004). Preliminary results from the development of a concept inventory thermal and transport science. In *Proceedings of the 2004 American Society for Engineering Education Annual Conference and Exposition*, (Session 3230).
- [20] Rowe, G. & Smail, C. (2008). Development of an electromagnetics course concept inventory. In *Proceedings of the 2008 American Society for Engineering Education Annual Conference and Exposition*, AC 2008-242-266.
- [21] Steif, P. S. & Dantzler, J. A. (2005) A Statics Concept Inventory: Development and psychometric analysis, *Journal of Engineering Education*, 94, 363-371.
- [22] Wage, K.E., Buck, J.R., Wright, C. H. G., & Welch, T.B. (2005). The signals and systems concept inventory. *IEEE Transactions on Education*, 48 (3), 448-461.