The ACO Program at Georgia Tech

A Brief History and a Few Personal Reflections

by

R. Gary Parker[1]

Doctoral-degree-granting *inter*disciplinary academic programs have become common in the contemporary research university, though even as few as 15 years ago that wasn't typically the case, at least not at Georgia Tech. Such programs tend to be popular since they often prove to be fairly low-overhead creations and can, once approved, serve to enhance the visibility of the participating home academic units that support them, albeit sometimes in only cosmetic ways. Some interdisciplinary programs are rigorous and enforce highly restrictive admissions requirements, attributes that appeal to particularly talented students and that are demanded by a first-rate faculty; alternately, others are known to exist under operational models that settle for standards aligned in the opposite direction. The latter typically occurs as the number of academic units sponsoring a program increases and/or when the academic and research quality of the participating units may be problematic. On the other hand, creating and sustaining a competitive, high-caliber interdisciplinary program—let alone an elite one—requires the devotion and efforts of a confident and highly regarded faculty.

The University System of Georgia—the governing body that approves academic degrees for public institutions in the State—presently lists on its website[2] nine interdisciplinary doctorates sanctioned to be awarded by the Georgia Institute of Technology[3]; of a total of forty doctorates currently authorized to be conferred by the Institute, these nine comprise nearly a

[1]R. Gary Parker was a founding member of the ACO program, serving on its coordinating committee from the program's inception until his retirement from Georgia Tech in December of 2011. He is presently Professor Emeritus in the Stewart School of Industrial and Systems Engineering.

[2]www.info.usg.edu

[3]www.grad.gatech.edu/inter

quarter. Most of the listed interdisciplinary PhDs were established fairly recently (largely in the last ten years); the oldest was formed in 1991 and is the subject of this narrative. Identified as the interdisciplinary PhD in Algorithms, Combinatorics, and Optimization (ACO), existing records appear to establish that it was the first such doctorate at Tech. What follows is a story of how that degree program came about, how it has managed to not only endure but to thrive, and why an arguable case can be made that its success corroborates its candidacy as a model of academic interdisciplinarity at its very best.

In the Beginning

By the mid-1980s a rapidly growing area of research was emerging in the domain of discrete mathematical structures residing at the intersection of the fields of combinatorics, discrete and combinatorial optimization, and the analysis of algorithms. Clear was that serious research in each of the three fields was becoming increasingly dependent on knowledge of all of them, a claim that was easy to support because many of the most significant advances were being produced by researchers working in more than one, if not all three, of the areas. The activity was also becoming apparent at Georgia Tech.

During the period alluded to above—through the 70s and into the early-80s—a fair number of faculty at Tech were directly engaged, though not in a particularly unified way, in teaching and research activities clearly related to discrete and combinatorial topics, largely concentrated in three distinct academic schools. In the School of Mathematics (SoM) Richard Duke, Ted Hill, Tom Morley, Kevin Phelps, and Jon Spingarn were active—Duke and Phelps were most prominent—through their primary interests, which included familiar areas such as graph theory, combinatorics, coding theory, combinatorial design, and classical probability theory. In the School of Industrial and Systems Engineering (ISyE) a sizeable number of faculty, including, at one time or another, John Bartholdi, John Jarvis, Donna Llewellyn, R. Gary Parker, Ron Rardin, Craig Tovey, and John VandeVate, taught courses and conducted research in integer programming, network flows, combinatorial scheduling theory, and discrete optimization. Topics in computability, complexity theory, and the analysis of algorithms were, then, covered in the School of Information and Computer Science (ICS)—the College of Computing had not yet been invented—primarily by young faculty, probably most

prominent among whom were Jeremy Spinrad, and, somewhat later, Jim Burns and H. Venkataswaran.

In truth, however, while there was substantial, and growing, activity at the Institute in those "old days," much that went on regarding the three subject areas of primary interest in this narrative was confined to the respective home academic units, i.e., ICS, ISyE, and SoM. It was an era when crossing academic boundaries at Georgia Tech—particularly ones existing at the college level—was not, putting it charitably, routine. Nontrivial interactions between like-minded faculty as well as their students—or prospective students—pursuing research naturally related to discrete structures was typically limited to attending lectures and seminars covering topics of shared interest, or participating on dissertation committees in the obligatory role of outside reading members. Such a state of affairs, though far from ideal, is usually endurable for faculty members; however, for doctoral students it can prove to be limiting. Motivated to pursue exciting work in discrete and combinatorial topics of interest, PhD students, at the time, often found important course work related to such interests either inaccessible or inconsistently scheduled, following from a lack of even reasonable levels of coordination between academic units where the particular courses were taught. More critically, students who were becoming increasingly attracted to research topics championed by faculty not residing in their home units found that seeking research advisement, accordingly, was often problematic because traditional Institute-level policies frequently got in the way. And, those difficulties became even more acute as the Institute began to attract strong faculty and doctoral students who expressly sought to work in the relevant areas.

In 1978 the School of ISyE named a new chair, Michael Thomas. Aiming to build on existing strength in the discipline of operations research (OR) in the School, Thomas began to recruit faculty from what were, at the time, the best programs in the country; a couple of those hires are noteworthy to this story. One was Craig Tovey (already identified), a newly minted PhD from Stanford's famed Department of Operations Research.[4] Hired in 1981, Tovey's formal background included degrees in mathematics, computer science, and optimization, which coupled with his demonstrated

[4]At the time, Stanford had a stand-alone OR department; in 1996, however, it would be merged with other programs at the university to form the Department of Management Science and Engineering (MSE). Notably, 2004 ACO graduate Amin Saberi (advised by Vijay Vazirani and Milena Mihail) would join the MSE faculty.

ability/inclination to work with faculty and students whose interests also spanned those disciplines, would come to represent something of a *model* of precisely the sort of teaching and research needs that were beginning to motivate the creation of an ACO-like academic program—whether he knew it or not.[5] The other key Mike Thomas hire was George Nemhauser, a world-class figure in integer programming and combinatorial optimization. Lured away from Cornell's prestigious OR program in 1985—Mike Thomas had actually been George's PhD student at Johns Hopkins in the mid-60s—Nemhauser would fill ISyE's first endowed position (if not one of the first of its kind at Georgia Tech) as the A. Russell Chandler III Chair. And, whether he knew it or not, he, too, would play a particularly crucial role in the establishment of the ACO program at Tech.

Shortly after joining the ISyE faculty, and possibly motivated by his experience with the Center for Applied Mathematics at Cornell, matched with what he observed as an opportunity given existing, as well as emerging, faculty and doctoral student interdisciplinary interests at the Institute, Nemhauser engaged with Richard Duke on the matter of creating a standalone academic framework that would facilitate a more structured pursuit of those interests. Mike Thomas, who was making plans at the time to move into the administration of the newly appointed president of Georgia Tech, John Patrick Crecine, was also supportive of the notion, at least as a concept. That support by Thomas is relevant to this story since, indeed, he would become Crecine's acting executive vice president in 1989—a position that would later morph into that of provost—and any formal programmatic arrangement at the level anticipated, vis-a-vis an ACO-based format, would have required his approval. Accordingly, a small working committee was formed, charged with producing a draft plan for the creation of such a program; representing the three primary academic units that would be involved, that original planning committee consisted of Jim Burns (ICS), Richard Duke (Math), and R. Gary Parker (ISyE).

Upon approval of the draft plan produced by the committee, work on a formal proposal (to the Institute) commenced in 1989, or thereabout. The intent of that effort was to create an independent interdisciplinary doctoral program that would fix its own curriculum, admit its own students, and

[5]Craig Tovey would, indeed, ultimately take on adjunct and/or joint appointments in Tech's College of Computing.

grant its own degree, all free of influence by the respective home academic units sponsoring the program, other than as it would pertain to common institutionally based requirements that are typically applied to all PhDs, e.g., administration of a comprehensive exam, satisfactory completion of an academic minor, adherence to basic dissertation requirements, etc. There had been no such program at Georgia Tech. However, one did exist elsewhere. Created a year or two earlier, Carnegie Mellon University supported the only formal degree-granting activity of the sort being envisioned at Tech. Tapping into their experience, input was sought from professional acquaintances at CMU—theirs was a joint endeavor with mathematics, computer science, and the business school, which housed CMU's operations research strength—that proved beneficial. Referred to as ACO as well, the CMU model and the one that would ultimately be created at Tech remain, as this is written, the only academic programs of their kind in existence in the US.

But while work on a full proposal for an ACO academic program proceeded, the Georgia Institute of Technology would undergo—and not without a fair amount of pain—a major restructuring under a grand plan proposed by President Crecine. Relevant for ACO was that as part of that restructuring effort, the College of Computing (CoC) would be established, replacing the School of ICS. In addition, at about that time (1989/90), Ellis Johnson, a leading figure in integer programming, would leave IBM's Thomas J. Watson Research Center and come to Tech to assume the newly established Coca Cola Chair position in ISyE, and Robin Thomas would join the faculty in the School of Mathematics.

Under the oversight of an expanded committee—the initial draft-forming committee of Duke, Parker, and Burns, would be augmented by the addition of Ellis Johnson, George Nemhauser, Robin Thomas, Craig Tovey, and H. Venkateswaran—a finished proposal was brought forth and formally approved by the faculties of each of the three sponsoring academic units (CoC, ISyE, and SoM), followed by the Institute Graduate Committee and Tech's Academic Senate. Finally, during its 1990-91 academic year (AY) meeting the proposal was approved by the Board of Regents of the University System of Georgia, which meant that the ACO program would be open for business at the beginning of Fall Quarter 1991,[6] fully authorized to admit students to

[6]Georgia Tech operated on the quarter system until August, 1999, at which point it converted to semesters.

pursue the degree: *Doctor of Philosophy with a Major in Algorithms, Combinatorics, and Optimization.* Two students would be admitted to the program that fall: Dan Sanders (home unit SoM) and Ismael de Farias (home unit ISyE). In the fall of 1992, Gruia Calinescu would be the first ACO student admitted with home unit CoC.

The Early Years

The initial ACO program faculty was formed, with membership essentially self-determined, easily following from obvious pedagogical and research interests that existed across the three participating academic units. The program's administration, as called for in the original proposal, was conducted by a coordinating committee consisting of a small subset of its faculty; membership on the committee was approved by Mike Thomas, who, as intimated, was by then Tech's executive vice president. Richard Duke was named as the program's director, also serving as the chair of the coordinating committee. The latter, with more or less uniform representation from the three sponsoring units, was responsible for program admissions decisions, the administration of the program's doctoral examinations, full oversight of its curricula, and the task of monitoring the progress of all of the program's students. The coordinating committee also approved requests for membership to the ACO program faculty. The program director reported to the office of the vice president.

Shortly after the ACO program officially launched, an event was held to celebrate its founding. Identified as the Workshop on Algorithms, Combinatorics, and Optimization—dubbed WACO—a number of highly regarded figures in the US and abroad, representing the core ACO-relevant strengths of the three units, were invited to campus for a series of talks. Designed to showcase the program, and Tech's strength in supporting such an endeavor, it was evident to everybody—notwithstanding CMU's slightly older program—that the notion of an ACO program at Tech was, at its base, something of an experiment—a well thought out one perhaps, but an untested concept all the same. A host of obvious and crucial questions loomed: Could the program attract capable students? Could its graduates get positions? Could they get positions on upper-tier faculties and at cutting-edge research labs? Could its faculty stay interested?

The ACO program curriculum was, and remains, a stern one. The original program core consisted of six courses: an algorithms course, two courses in optimization, one in graph theory, one in algebra, and a course in classical probability theory. None of the six were watered down versions of the stated material—service-like courses often created to cater to *multi*disciplinary interests and activities, but lightened in order to facilitate accessibility for participating students possessing differing backgrounds. Indeed, all of the courses, with the exception of the one covering classical probability,[7] were precisely the same ones likely to be taken by the respective non-ACO doctoral students residing in each of the three participating units. The comprehensive examination for ACO students covered every course in the core; it was not an easy exam.

Importantly, the demands of the ACO program were not by-products of an artificially contrived curriculum, a plan gratuitously designed in order to satisfy a suitably high threshold on rigor for its own sake. Rather, they were the result of a conscious and focused effort to create an academic framework that prepared students seeking to work on hard problems—ones arising in domains where legitimate success plainly relied on skills grounded in the core subject matter that was at the heart of what defined ACO; indeed, what motivated the program's creation. Still, it was an arduous program to navigate, particularly in terms of preparing for its comprehensive examination. A working rule-of-thumb held that in order to safely earn a passing score on the exam, a student needed—absent cheating, or dumb luck—to be proficient in two of the three representing areas and at least competent in the third. A funny-but-not-funny joke among some of the ACO faculty held that the level of preparation/ability required to get over the ACO comprehensive exam hurdle, if not to successfully complete the full program, was on the order of three-halves—which may have been only a lower bound—that demanded by a "normal" PhD.

Some of the early headwinds that confronted the ACO program as it worked to gain traction, stemmed from, for lack of a better description, struggles with self-identification. The program had a natural appeal, at least at the surface level, but one that tended to mask its rigorous scholastic

[7]Math 6221 (Advanced Classical Probability Theory) was created especially for ACO students; ultimately that requirement was switched out in favor of Math 7018: Probabilistic Methods in Combinatorics.

demands. Would-be applicants who thrived on mathematical puzzles and games, or who had toyed with the traveling salesman problem, or who could skillfully churn out computer code, and such, sometimes assumed the program was a good match, but quickly sobered up. Usually. Others, even if admissible—and in possession of an appropriate level of awareness regarding the program's level of difficulty—wondered where an ACO degree would be employable. Some, of course, looked over the requirements—perhaps upon having heard the 3/2 joke referenced above—and decided that it might be better for their health to simply stick with a conventional PhD offered in their home academic unit. By whatever measure, enrollment in the ACO program was meager in the early years of its existence.

During the first five years following its approval, a half-dozen matriculants entered the ACO program at Tech; in 1993 Dan Sanders, advised by Robin Thomas, became its first graduate. By fall of 1996, just as the Atlanta Olympics were leaving town, three ACO degrees would be awarded.

Maturation

Around the mid-90s, the ACO program would begin to experience a major alteration to its faculty ranks, particularly through the College of Computing and the School of Mathematics. Saugata Basu,* Neil Calkin, Dana Randall,* Prasad Tetali,* and Xingxing Yu would join the faculty in SoM; Howard Karloff, Leonard Shulman, and Vijay Vazirani would come to CoC; some of these individuals (identified with asterisks) would ultimately take joint appointments in SoM and CoC, if not permanently switching to one of the two. Renato Monteiro would join ISyE, becoming the ACO program's first Hispanic faculty member.[8] Enrollment would increase in a notable way as well: recruiting classes from fall 1996 to fall 2000 would total 16 students, nearly triple the number from the program's first five classes (fall 1991-fall 1995).

Also noteworthy was that the program began to have success in attracting the attention of very strong applicants, some among the most highly sought. That competition wasn't confined to CMU, vis-a-vis its role as the only other formal ACO program, but involved, at least as much, if not more,

[8]The first African-American on the ACO program faculty was Earl Barnes, who, like Ellis Johnson, had left IBM-Watson to join ISyE in 1989.

powerhouse academic programs in discrete mathematics, OR, and CS theory that existed at institutions such as Princeton, Stanford, MIT, Berkeley, and Cornell. Well known in the high-stakes recruiting game is that attracting the very best doctoral students, certainly those in the elite class, is not overly influenced, within any reasonable understanding of the process, by the complexion of stipend and/or fellowship offers, but, rather, by a program's academic quality, the reputation of its faculty, and the track record of its placement potential. Not surprising is that the ACO program at Tech lost more often than it won in those early recruiting competitions, but what was new, by the mid- to late-90s, was that it began to appear on the short list of top applicants. That wasn't a small thing. And, as faculty strength steadily grew, and with it, reputational credibility, the program did begin to enroll some exceptionally talented students.

For the most part, new ACO-connected hires during the program's first decade had been confined to beginning- or mid-level positions. While that trend would continue in a natural way, by the year 2000 the program began to add to its ranks high-visibility senior scholars who, by their established reputations, brought a different level of stature to Georgia Tech and, by their natural association, to the ACO program. In 2000 Richard Lipton would leave Princeton to take the Frederick G. Storey Chair in CoC; in 2002 ISyE would attract Bill Cook, who had spent a year visiting at Princeton (after having left Rice), to fill the newly established position of Chandler Family Chair; and Tom Trotter would be named school chair in SoM. Notable entry-level additions included 2004 hires Subash Khot and Matt Baker, who would join CoC and SoM respectively; Eric Vigoda would join the CoC faculty as an associate professor, also in 2004. In 2005 Arkadi Nemirovski would leave the Technion to assume the John Hunter Chair in ISyE.

From 2000 to 2005, the program would enroll 18 students, nearly as many as it had during its first ten years. In the same five-year period, it would graduate 12 students. Becoming clear was not only that the ACO program had staying power, but that it was emerging as an exceptional academic option for top-flight doctoral applicants, many of whom reported that were it not for the program, they would not have considered study at Georgia Tech.

In the fall of 2005, nearly a decade-and-a-half after its creation, the ACO program underwent the Institute's mandated Academic Program Review

(APR) process.[9] The self-study[10] prepared for the 2005 APR—the program's first such review—revealed an academic program that had, at the time, 26 enrolled students and that had awarded a total of 23 PhDs. The ACO program faculty, in the fall of 2005, consisted of 33 members, including:

Five endowed chair holders
An Institute Professor (at the time, one of only six at Georgia Tech)
A Regents' Professor
Three members of the National Academy of Engineering (NAE)
Two winners of the John von Neumann Theory Prize
Two winners of the Dantzig Prize
Two winners of the Fulkerson Prize
An IBM Fellow
A Guggenheim Fellow
An ACM Fellow
Two INFORMS Fellows
Six recipients of the NSF CAREER Award (or its predecessor[11])
Two Sloan Fellowship recipients
Two winners of the Lanchester Prize (three prizes)

The profile above included the first sitting Georgia Tech faculty member to be elected to the NAE (Nemhauser); the first person to have been awarded the von Neumann, Dantzig, *and* Fulkerson Prizes (Nemirovski); and the first two-time winner of the Lanchester Prize (Nemhauser). Importantly, the ACO faculty members producing the distinctions and accolades listed, were also among the most esteemed in their respective home academic units, if not at Georgia Tech, independent of their ACO affiliation.

And, by 2005, ACO graduates from Georgia Tech had begun to amass records that provided credible evidence of the program's ability to produce candidates who could compete at high levels, both in terms of placement

[9] Operating on a staggered timetable, all degree-granting entities at the Institute are required to undergo an outside review of their academic programs every six years; APR visits are conducted by teams of outside visitors, usually four field-relevant individuals, and are relatively formal affairs involving substantial scrutiny. APR processes, which began at the Institute in 2002, are typically conducted out of the provost's office.

[10] *Program Review Self-Study: Algorithms, Combinatorics, and Optimization*, fall 2005.

[11] Prior to the CAREER designation, the analogous award was considered to be the NSF Presidential Young Investigator Award, often referenced as the PYI.

as well as in the context of academic and scholarly recognition through important honors and awards. Including permanent, visiting, and post-doctoral positions, graduates from the ACO program, as of fall 2005, had earned positions at the likes of Berkeley, Stanford, Princeton, Charles University, Cornell, CWI (Centrum Wiskunde and Informatica), the University of Waterloo, CMU, Ohio State, IMA (Institute for Mathematics and Its Applications), Purdue, Microsoft Research (two), the University of Illinois at Urbana-Champaign, Renaissance Technologies (two), Bloomberg LP, the University of Toronto, and the University of Sao Paulo. Awards won by ACO graduates during the period included:

Optimization Prize for Young Researchers from INFORMS (two)
Outstanding Faculty Award from the University of Iowa
Transportation Science Section Dissertation Award from INFORMS
Best Student Paper Award from the ACM Section on E-Commerce
CAREER Award from NSF

Achieving Elite Status

Richard Duke retired in 2006.[12] More than simply one of a small number of its founding faculty members, he was, as one prominent ACO faculty member put it, the heart and soul of the ACO program. With the unassuming style of the Midwesterner that he was, he never drew attention to himself, he never sought the spotlight, and he was loath to accept exclusive personal credit for whatever distinctions the ACO program garnered during the decade-and-a-half span under his direction. Naturally, the task of finding his replacement was viewed, certainly by anybody paying attention, as a particularly critical one in order that the momentum built up during the Duke era not be diminished; that, in fact, it be nurtured in ways that would continue to elevate the program's stature as a growing presence among the elites. It was also the sense, though not mandated in any formal way, that the program's leadership, at least then, should continue to reside in the School of Mathematics. Accordingly, a consensus choice to lead the program

[12]Richard A. Duke would die February 19, 2015, following an 18-month battle with lung cancer. Underscoring his dedicated commitment to both the School of Mathematics and the ACO program, he established, in 2014, the Richard A. Duke Endowment and prior to that, the ACO Endowment Fund.

was conveyed to Gary Schuster, then Georgia Tech's provost, who granted approval: In 2006, Robin Thomas became the second director of the ACO program.

When the baton passed from Richard Duke to Robin Thomas, the ACO program was well-established; concerns from its earlier years, especially those regarding its viability and appeal to applicants of the highest quality, had essentially evaporated. Echoing sentiments held by others in the program, one ACO faculty member—an individual who was highly regarded professionally and who had much experience in interdisciplinary university activities at various institutions—claimed the ACO program at Tech was the only one he knew of that successfully operated at a bar *higher* than that enforced by any of its constituent component programs. By the time the next APR would be undertaken, the program would attract to its ranks a number of impressive new faculty, as well as a host of talented students, many of whom would conduct work that would bring additional—in some cases, particularly noteworthy—recognition to the program.

In 2006, Santosh Vempala would leave MIT and join the faculty in CoC. He would be the first ACO faculty member whose degree was in the discipline, having taken the PhD in ACO from CMU's program; in time though, he would not be alone in possessing that distinction.[13] More importantly, he would be active in the creation of the Algorithms and Randomness Center (ARC), which would prove to be an attractive and innovative research counterpart to the academic entity represented by ACO. He would serve as the first director of the Center.

From 2008 through 2011 the ACO program would add several faculty, largely at the entry level: Nina Balcan (CoC), Santanu Dey (ISyE), Ton Dieker (ISyE), David Goldberg (ISyE), Anton Leykin (SoM), Chris Peikert (CoC), Prasad Raghavendra (CoC), Asaf Shapira (SoM/CoC), and Josephine Yu (SoM). As testimony to the strength of the additions, seven of the individuals would ultimately receive NSF CAREER Awards. Notable at the senior level, in 2010 Zvi Galil would move from Tel Aviv University (where he served as President, and prior to that, Dean of Engineering at Columbia) to join the College of Computing as the John P. Imlay Jr. Dean of Computing, and have his name added to the ACO program faculty.

[13] In 2017, ISyE would hire Mohit Singh, also a graduate of the CMU ACO program.

As documented in the program's second APR,[14] conducted in the fall of 2011, the vitals of the ACO program portrayed an academic model that could reliably be described as thriving. Student interest was healthy, both in terms of the number of applicants seeking admission as well as the quality of each year's applicant pool. The numbers of applicants (per annual recruiting cycle) during the period between ACO's first two APR visits ranged between 60 and 80, and of those admitted, the program began to attract its fair share; in fact, rarely was a top candidate lost to a lesser program. Enrollment figures for fall terms spanning 2006 through 2011 were 28, 38, 35, 32, 37, and 39 respectively. By fall of 2011 47 ACO PhDs had been awarded.

Program faculty distinctions continued to be notable. Updating the list provided previously (re: APR from 2005), by 2011/2012 ACO faculty counted among its ranks five members of the National Academy of Engineering (Cook, Johnson, Nemhauser, Galil, Lipton); a member of the American Academy of Arts and Sciences (Galil); three Fulkerson Prize winners, for a total of four prizes (Nemirovski, Thomas,[15] Vigoda); the Khachiyan Prize (Nemhauser[16]; an Honorary Doctorate in Mathematics from the University of Waterloo (Nemirovski); three Guggenheim Fellows (Lipton, Vazirani, Vempala); three ACM Fellows (Galil, Lipton, Vazirani); three INFORMS Fellows (Cook, Johnson, Nemhauser); four SIAM Fellows (Cook, Johnson, Nemhauser, Tetali); four AMS Fellows (Baker, Randall, Tetali, Thomas); eight CAREER/PYI Award recipients (Ahmed, Balcan, Ergun, Randall, Tovey, Vazirani, Vempala, Vigoda); and three Sloan Fellowship recipients (Peikert, Randall, Vempala).

And, important honors and awards continued to earned by ACO students/graduates, including ones that were particularly prestigious:

European Research Council (ERC) Starting Grant[17]
CAREER Award (3)
Herman Goldstine Postdoctoral Fellowship from IBM
NSF Mathematical Sciences Postdoctoral Fellowship
Denes Konig Prize from SIAM[18]

[14] *Program Review Self-Study: Algorithms, Combinatorics, and Optimization*, fall 2011.
[15] Robin Thomas won a second Fulkerson Prize in 2009
[16] First recipient of the prize
[17] ERC grants are considered Europe's version of the NSF CAREER Award, and at least as prestigious.
[18] Awarded to Adam Marcus in 2008 while still a graduate student. The prize is awarded

Humboldt Fellowship
Best Paper Award from ACM-SIAM, SODA 2010
Alfred P. Sloan Fellowship
Best Paper Award from ACM Symposium on PODS 2008
DOE Graduate Fellowship

By 2011—twenty years after its founding—the "steady-state" condition of the ACO program at Georgia Tech was not only healthy, by any of the usual metrics, but was in a position to lay claim to being an elite academic program. Its record of attracting the affiliation of high quality faculty was clear and, importantly, it had reached a point where drawing serious attention of the very best doctoral applicants was not just possible, but an expectation. Indeed, the program's ability to win the recruiting competition for its share of the elite applicants had become much more promising than it had been, only a half-dozen years earlier. The final report of the visiting committee, submitted shortly after the 2011 APR on-campus visit, which was conducted in February of 2012, contained the following assessment regarding the state of the ACO program:

"...the Georgia Tech ACO Program achieved preeminence more broadly than just within cross-disciplinary programs of its own genre; it is a preeminent program in each of the three subareas that it draws from."[19]

Lessons Learned and Looking Ahead

The ACO program would undergo its third APR assessment during AY 2017/18; the respective self-study report[20] would be completed in the fall of 2017; the on-campus visit by outside reviewers would take place the following spring. Readers of the 2017 *Self-Study* will see described an interdisciplinary program that, while closing in on nearly three decades of existence, remains healthy and vibrant; one that continues to boast a faculty that remains

bi-annually to a young researcher who is no more than four years past his or her PhD. Note, also, that Adam was the inaugural recipient of the Konig Prize.

[19] Visiting committee report following the ACO program review; submitted spring 2012.

[20] *Program Review Self-Study: Algorithms, Combinatorics, and Optimization*, fall 2017.

world-class in various concentrations. They will see a program that attracts exceptional doctoral students, who upon graduation compete for positions at first-tier universities and cutting-edge research facilities.

Interestingly, the size of the ACO program faculty has varied little over the course of the program's life, particularly after its initial start-up years. In both the 2005 and 2011 APR years, the program's faculty list consisted of 33 appointed members; in the fall of 2017 the number stood at 36. Of course the complexion of the faculty roster has changed substantially over the years, all the while preserving, if not enhancing, its overall strength. From the names appearing in the 2005 APR *Self-Study*, only a third would remain on the roster by fall of 2018. Still, through the normal comings and goings that are natural on university faculties, the ACO program has continued to be staffed by talented individuals who understand and appreciate what it has become, and who respect the intent of its fundamental mission, which has not changed since it was formally articulated in the first program review in 2005. The visiting review committee for the 2017 APR summarized its assessment in the following way:

"As with the previous review committee, we were all very well aware of the highest academic reputation of this academic program before our visit. Georgia Tech can be very proud of having had the vision and enthusiasm a quarter century ago to start this program, and its impressive list of alumni validates its overwhelming success and preeminence."[21]

Still, if the APR assessments, mandated to be conducted every six years, had become non-threatening exercises for the ACO program by 2017, such a view never morphed into complacency, nor invited an inclination to discount the need to remain vigilant in terms of staying relevant. The ACO program's anointed level of preeminence was genuine, but so too was its faculty's appreciation of how it had been built. Indeed, the ACO program faculty—certainly through the program's superior leadership—has always assumed a healthy and mature sense of professional responsibility, grounded in the understanding that the climb to elite status for a serious academic program can be arduous, whereas the fall from a position of prominence is often precipitous.

[21] Visiting committee report following the ACO program review; submitted spring 2018.

Ultimately, the ACO program at Georgia Tech became something bigger and better than merely the sum of the strengths of its three academic tribes. In the specific disciplines of theoretical computer science, operations research (particularly optimization), and discrete mathematics, the competitor programs reside at the likes of Stanford, MIT, Berkeley, Princeton, and Cornell. Employing language from the visiting committee report following the 2017 APR, those top universities "...can only compete for the best doctoral students due to their name recognition for their individual programs." That the ACO programs at Tech and CMU compete with those institutions—all famous, all exerting powerful "halo effects"—for top doctoral students accordingly, is more than noteworthy. That, after nearly thirty years, they remain the only two programs of their named genre (in the US) may speak most persuasively to how difficult it can be to create and, far more importantly, to sustain academic programs of their type and stature, particularly for the long haul. How did it happen? For this writer, one answer seems obvious: the existence of a nontrivial level of necessity, a committed faculty of believers who recognized such necessity, and the gift of devoted and principled leadership.

Tracing when pedagogical and research interests from the three disciplines—the A, the C, and the O—began to intersect in substantive ways, in a general sense, is probably an ill-defined, if not uninteresting, exercise. Certainly, influential scholarly papers appeared in the 50s and 60s, and no doubt even before, that portended formal study in ACO-like contexts, but much tended to be produced in isolation. Nonetheless, an academic archaeologist might accurately point to the late-60s, early- to mid-70s as a particularly pivotal period in terms of seeding the notion of formal ACO-related academic programs, following the appearance of a host of major research results as well as several relevant comprehensive books and monographs covering the stated subject matter. Impactful research contributions authored by figures with names such as Stephen Cook, Jack Edmonds, Ray Fulkerson, Richard Karp, and others would prove to be seminal. Important and influential books included the likes of: *The Design and Analysis of Computer Algorithms* by Aho, Hopcroft, and Ullman (1974); *Combinatorial Optimization: Networks and Matroids* by Lawler (1976); *Computers and Intractability* by Garey and Johnson (1979); and slightly later, *Combinatorial Optimization: Algorithms and Complexity* by Papadimitriou and Steiglitz (1982).

But while those references were valuable, and informed a growing popu-

lation of interested doctoral students and faculty research efforts at the time, the academic infrastructure at Georgia Tech lagged. There was no course at the Institute titled combinatorial optimization until 1980, when the graduate course ISyE 7676 was created (courses in network flows and integer programming did exist at the time, also taught in ISyE). In the School of Mathematics, the only relevant graduate course was Math 6012 (Combinatorial Methods), which would be created in 1983. In ICS, a graph theory course (ICS 6530) appeared as early as in the Institute's 1975 course catalog, with ICS 6155 (Analysis of Algorithms) and ICS 6156 (Complexity of Computation) appearing in 1977. Still, despite the availability of a few important courses, effective coordination regarding their scheduling was virtually nonexistent, which tended to render their putative availability somewhat problematic.

Not only did it become obvious that it was necessary to create at least some sort of facilitating framework for interested doctoral students who were being attracted to each of the three relevant academic units, pressure to serve those interests was increasingly being exerted by new faculty whose backgrounds and research interests coincided with those aims as well. As related at the outset of this narrative, a core group of existing Georgia Tech faculty took up the matter in a non-cosmetic way, with the foresight to create not simply an interest area or a concentration track, but a formal degree option that would stand alone.

But it's not so easy to make a go of such a thing in universities, at least not back in the day at Georgia Tech (or probably anywhere for that matter). Historically, academic tribal barriers tend to be difficult to breach—permissive perhaps of courteous interactions as part of general committee work, participation in generic seminar activities, and such, but far less so regarding programmatic endeavors that require heavier lifting and substantive levels of responsibility, including risk. Also frequently needed is a willingness to confront (or ignore) naysayers—particularly those of the administrative class—who are adept at trafficking in the it'll-never-work-here varieties of roadblocks that are ready-made for stifling academic innovation and creativity. When these sorts of headwinds, bureaucratic or otherwise, pick up, excitement and high-minded plans can begin to evaporate, and enthusiasm to actually do something can dull quickly.

Even if, in the face of impediments like those alluded to above, work on serious programmatic matters—which creation of something like the Tech ACO

degree most certainly was—ambles along, one of the most harmful threats is that program ownership can be a casualty. If the parties involved are not more or less uniformly dedicated to the task at hand in an interdisciplinary endeavor, the prospects for success, certainly ones of the long term variety, can get very wobbly; even if programs materialize under a flimsy support system, without dedicated and effective oversight their half-life is often fairly short. But by virtue of some weighted combination of the emerging pedagogical and research climate evolving around the subject matter, an institutional infrastructure and landscape that existed at Georgia Tech at the time, and the initiative taken by a few interested faculty with shared disciplinary trust and respect, a serious academic program got created; importantly, it was a program that was not simply owned by its faculty participants, but became something where direct association was highly valued. The report from the visiting review committee following the 2005 APR described the effect accurately:

"The ACO faculty have in many cases more in common with their ACO comrades across campus than with their intra-college colleagues."[22]

Still, even with an energized and dedicated faculty, a highly competitive program has to have strong and trusted leadership. Here, again, the ACO program at Georgia Tech made the right decisions. Richard Duke was an obvious choice to lead the program at the point of its founding. Not only was he respected in his discipline, he was open to interdisciplinary interactions even before the idea of an ACO-like project emerged. But, he was also attentive to the need to stay atop programmatic and administrative details that, though often amounting to not infrequently occurring tedious and thankless chores, simply have to be undertaken. In short, he believed in the program and shepherded it through its experimental years—a period where, frankly, the verdict regarding its viability was still out—on into stability and, ultimately prominence. Then, upon Richard Duke's retirement, the program could not have selected a better successor.

Robin Thomas is arguably one of the most prominent members of the ACO faculty. His distinguished scholarly reputation was well-established long before he entertained any notion of taking over the directorship of ACO.

[22]Visiting committee report following the ACO program review in 2005.

Indeed, few would have faulted him for staying away from the administrative obligations that he most certainly had observed throughout Richard Duke's impactful tenure at the helm. But that he did assume the director's position in 2006, and that he did so not with some caretaker-like mindset, but with vigor and new ideas, corroborated his commitment to the program and to the task of maintaining its character, and to elevating its position to the status of elite. Remarkable is that an academic program of ACO's class has had only two directors—each of whom were arguably optimal fits at the time—during its nearly three decades of existence; notable is that that leadership trajectory provides strong evidence underscoring stability, and, importantly, can rightly be viewed as a major factor in influencing the program's rise to preeminence.

The 2017 APR *Self-Study* reported that at the beginning of the fall term of 2017, 32 students were enrolled in the ACO program. To that point in the program's life, 81 degrees had been awarded. Noteworthy is that applicants for fall 2017 admission numbered 91, the largest in the program's history. Faculty identified with the program also continue to distinguish themselves; a listing of honors and awards that have been added, or modified, relative to the prior listing that was documented in the 2011 APR appear below:[23]

Seven endowed chair holders
National Academy of Engineering (4)
John von Neumann Theory Prize (2)
Dantzig Prize (2)
Knuth Prize
ACM Fellow (5)
INFORMS Fellow (2)
SIAM Fellow (2)
American Academy of Arts and Sciences Fellow (2)
CAREER Award (or PYI) (12)
Khachiyan Prize (2)
Golden Goose Award
Honorary Doctorate, Catholic University of Louvain

[23]Most information is accurate relative to when the 2017 APR *Self-Study* was prepared (late fall); some departures were either imminent or were pending at that point, an attribute that's not reflected in the referenced listing.

IBM Faculty Award (2 faculty; 3 awards)
Class of 1934 Distinguished Professor Award for Georgia Tech (3)
Class of 1934 Distinguished Interdisciplinary Activities Award (2)
Neuron Award for Lifetime Achievement in Mathematics
Denes Konig Prize (2)

As prominent and distinguished as the ACO program faculty remains, in a real sense the more impactful profile is the list of achievements produced by graduates of the program at this point in its history, if for no other reason than because it's precisely that record that provides the best marker for a program's reputation. Any doubters of such a claim need only check the priorities commonly listed/articulated by the best and most sought applicants. Accordingly, in 2017, Georgia Tech ACO graduates sit on faculties at the likes of Stanford, Princeton, Dartmouth, CWI, the University of Waterloo (2), Northwestern (2), IIT (Hyderabad), IISc (Bangalore), SUTD-Singapore, the University of Sao Paulo, the University of Rome (2), McGill University, Universidad Catolica de Chile, Mount Holyoke, the University Illinois at Urbana-Champaign, and Davidson. Others have non-academic positions at a number of highly regarded research settings, including at Google, Microsoft, IBM, SAS, Facebook, Amazon, Apple, and others. In addition, a host of new honors and awards have been earned by ACO students; prominent ones are listed below (note that earlier listings, ones that appeared in the 2011 APR *Self-Study*, are not repeated here):

INFORMS Optimization Society Student Paper Prize
International Congress of Mathematicians (Seoul): sectional talk
21st International Symposium on Math Programming: semi–plenary talk
22nd International Symposium on Math Programming: semi–plenary talk
Polya Prize in Mathematics from SIAM
ERC Starting Grant (3)
CAREER Award (4)
Simons Award in Theoretical Computer Science
Simons Postdoctoral Fellowship
A. W. Tucker Prize from the Mathematical Optimization Society

Still, as this is written, the ACO program faces some challenges. Near the top of the list is the matter, common to any established program, of dealing

with departures. Many of the so-called Old Guard faculty have retired or are nearing retirement, some of whom have been exceptionally influential in building and sustaining the program's substantial reputation. And, though a bit more controllable, there have also been unfortunate departures among the ranks of younger faculty, which often includes some of the very best and most talented individuals in academic programs like ACO. The good news, however, is that a number of new hires have been added to the three ACO units' faculties, many with evident star potential, which clearly bodes well for the future.

Even so, a point may be approaching where the ACO faculty's ability—and its willingness—will be tested in terms of staying united and committed to the aim of maintaining the program's presence as an academic powerhouse. Eminently apparent is that new leadership, when that is called for, will need to emerge and be up to addressing this challenge. The program has been fortunate to have been led by exceptional people; whoever steps in to take over will certainly have a big job, but modeling his or her efforts after the truly excellent work of Richard Duke and Robin Thomas would be a good place to start.

Programmatically, the program needs to stay vigilant in looking for ways to remain relevant; appeals to doctoral students that were effective in the 90s and early 2000s, may be less meaningful now. Staying abreast of developing areas of research is obviously important. Prominent, accordingly, are the domains of machine learning, data science, security, and others, where some of the great strengths of ACO faculty—convex optimization, for example—may have new and major roles to play. Failing to get onboard, for instance, with regard to exercising its influence in the aforementioned contexts of machine learning and data science, might prove to be not only a missed opportunity, but, crucially, a threat to the strength of the program's substantial reputation and presence at Georgia Tech.[24]

Final Thoughts

It's interesting to wonder: If not in 1990, would an ACO-like program have been developed ultimately at Georgia Tech? Perhaps. The times were

[24]See committee report following the 2017 APR visit.

certainly ripe for some sort of formal structure that would serve what was clearly a growing interest in the discipline being formed at the intersection of the three defining fields that comprise ACO. But to create a formal academic degree-granting entity—one that has roots squarely planted in well-established disciplines and cultures, each existing in influential academic colleges, while seeking a substantive level of independence from those units—is no small trick. Any number of well-positioned universities had to have noticed the need and could have taken on the task, but that only a couple did, at that point in time, is somewhat notable. That those two remain, after nearly thirty years, the only two of their explicit type, seems especially notable. At the very least, that fact goes a long way in underscoring how difficult it can be to bring together three powerful disciplines—each populated by powerful egos—to break down historical academic barriers, and to find common ground in pursuit of something that is strong and that has real staying power in a highly competitive research university environment.

And, let it be said that even if the high-minded ideas that were floated about three decades ago at Tech were exciting and had genuine merit, doing the heavy lifting of actually moving to the point of program creation, including dealing with all of the tedious, behind-the-scenes details that few want to engage in—and that most certainly can often be thankless—requires a special mix of personalities and mindsets. Had George Nemhauser and Richard Duke not initially interacted in those early years, had a few focused and interested faculty not put pen to paper drafting a cogent proposal for an innovative program that had no peer at Tech, had there not been a sympathetic provost-to-be in the person of Michael Thomas, and, above all, had there not been a level of collegial respect among the early players in the endeavor, it says here that the product known as Georgia Tech's ACO program might well have not gotten off the ground.

But it did happen, and what was produced was a serious doctoral program that not only survived, but evolved into a gem, not simply among Georgia Tech doctoral programs, but among academic peer programs at some of the nation's most highly regarded universities as well. The ACO program at Tech stands as testament to the power of legitimate academic interdisciplinarity, if employed carefully and with genuine commitment. A lot of talented faculty, several of whom have a sound claim to being called famous, currently function in, or have passed through this ACO program; their fingerprints on its rise and the stature it presently enjoys are abundantly evident. For this author, it

was an honor to have been there at the beginning, and to have been permitted to play even a small role in its creation.