




TEL +1 778 782 3925
 FAX +1 778 782 5876
 sfu.ca/vpacademic

Simon Fraser University
 Strand Hall 3100
 8888 University Drive
 Burnaby BC
 Canada V5A 1S6

MEMORANDUM

ATTENTION: Senate	TEL
FROM: Peter Keller, Vice-President, Academic and Provost, and Chair, SCUP	
RE: Full Program Proposal for a Professional Master of Science in Computing Science (SCUP 16-50)	
DATE: December 12, 2016	TIME

At its December 7, 2016 meeting, SCUP reviewed and approved the full program proposal for a Professional Master of Science in Computing Science in the School of Computing Science within the Faculty of Applied Sciences, effective Fall 2017.

Motion:

That Senate approve and recommend to the Board of Governors the full program proposal for a Professional Master of Science in Computing Science in the School of Computing Science within the Faculty of Applied Sciences, effective Fall 2017.

c: F. Popowich

Simon Fraser University
Maggie Benston Centre 1100
8888 University Drive
Burnaby, BC V5A 1S6

TEL 778.782.3042
FAX 778.782.3080

gradstudies@sfu.ca
www.sfu.ca/grad

MEMORANDUM

ATTENTION Senate Committee on University Priorities (SCUP) DATE November 23, 2016
FROM Wade Parkhouse, Chair of Senate Graduate Studies Committee (SGSC) No. GS2016.36
RE: Proposal for Professional Master of Science in Computer Science



At its meeting of November 7, 2016, SGSC approved the full program proposal for a Professional Master of Science in Computer Science and is recommending it to SCUP for approval, **effective Fall 2017**.

Motion:

That SCUP approve and recommend to Senate the full program proposal for a Professional Master of Science in Computer Science in the School of Computing Science within the Faculty of Applied Sciences.

New courses to be offered approved under delegated authority:

CMPT 706 – Design and Analysis of Algorithms for Big Data
CMPT 756 – Big Data Systems



MEMORANDUM

Attention Dr. Wade Parkhouse Date October 18, 2016
Dean, Graduate Studies

From Dr. Mirza Faisal Beg mfbeg@sfu.ca
Faculty of Applied Science, Graduate Studies Committee

Re: 1) Professional Master of Science Program (conversion from special cohort to regular program)
2) MSE PhD Qualifying examination course MSE 890

The faculty of Applied Sciences Graduate Studies Committee would like to send the following two items for consideration by SGSC. These have been approved by FGSC by electronic vote.

- 1) Professional Master of Science in Computer Science in the area of 'Big Data'
- 2) ~~Creation of a MSE PhD Qualifying Examination via a course MSE 890~~

Documents for the above items are attached with this memo. I request you to please place these on the agenda for the next SGSC meeting.

Cc: Dr. Greg Mori, Director, School of Computing Science
Dr. Glenn Chapman, Director, School of Engineering Science
Dr. Farid Golnaraghi, Director, School of Mechatronic Systems Engineering



SIMON FRASER UNIVERSITY
ENGAGING THE WORLD

Professional Master of Science in Computer Science

Full Program Proposal

December 12, 2016
School of Computing Science

PART A:

EXECUTIVE SUMMARY

The School of Computing Science proposes the creation of a *Professional Master of Science in Computer Science* degree program. This program was initially offered as a cohort special arrangements master's program in Big Data through the Office of Graduate Studies at SFU. The special cohort program is now in its third year and has been highly successful, with a dramatic increase in demand combined with the successful placement of students in workplace positions. Given this, it is time to convert it to a regular program, building on the information gained from the two special arrangement offerings. While the proposed *Professional Master of Science in Computer Science* (which will be referred to as the Program in the rest of this document) will continue to have Big Data as a specialization area, it will allow for future specializations in addition to Big Data, as identified in discussions with the current Big Data advisory committee, industry employers, and academics.

Our goal is to train computational specialists who can construct models, develop algorithms, and write software using state-of-the-art graduate-level techniques in computer science. This includes the ability to extract actionable knowledge from Big Data as well as other applications of computer science. Program graduates will know how to handle, analyze, and visualize massive data sets; they will be able to ask feasible questions and answer these questions efficiently using computational methods. The Program trains students to design, develop, and use complex systems to solve real world problems.

Highly qualified personnel trained in cutting-edge computer science techniques are in great demand. As an example, Big Data remains a fast-growing area of computer science that has transformed the information technology industry. We have seen evidence of this demand not only in the co-op placements of our students over the first two years, but also in the placement of our first cohort students who have already completed the program. We will continue to train students to fill this need, and the changing needs of the employers seeking computer science students with professional master degrees.

Our intended audience is recent graduates, international students, and professionals who already have degrees in computer science or a related technological field, but who wish to advance their knowledge in a rapidly growing area of information technology. It is not intended to prepare students for PhD studies; that is the focus of existing MSc degrees offered at SFU and elsewhere.

1. Proposed credential to be awarded

Professional Master of Science in Computer Science

2. Location of program

The new Program will be offered in the School of Computing Science at the Burnaby campus.

3. Academic unit(s) offering proposed program

School of Computing Science in the Faculty of Applied Sciences

4. Anticipated program start date

Fall 2017

5. Anticipated completion time

Four or five terms (16 or 20 months) to complete the Program. The duration depends on whether the student takes a 4 month or 8 month co-op placement during the degree.

6. Summary of proposed program

a) Aims, goals and/or objectives of the proposed program

The goal of the Program is to address the talent gap in the area of Big Data and related information technology sectors. There continues to be a rapidly growing demand for data scientists in industry, science, medicine, and government. According to a McKinsey Global report, *"By 2018, the United States alone could face a shortage of 140,000 to 190,000 people with deep analytical skills as well as 1.5 million managers and analysts with the know-how to use the analysis of Big Data to make effective decisions."* IDC estimates that the worldwide spending on business analytics alone was \$90 billion in 2011. EMC's Chuck Hollis says of big data analytics: *"The race is now on to acquire - and maximize the productivity of the key talent behind this wave: data scientists and their supporting data science teams"*. Recruiting, retaining and training big data analytics professionals is becoming increasingly difficult among employers across all regions, sectors, and industries, according to 2015 paper published by [Canada's Big Data Consortium, Closing Canada's Big Data Talent Gap](#).

- <http://www.ryerson.ca/content/ryerson/provost/partnerships/talentgap/WhitePaper.html>

While there is currently a dramatic demand for Big Data specialists, there are other application areas of computer science showing impressive growth. For example, computer security (also known as cyber security) has been identified as a priority area by many organizations. Similarly, there is a great deal of interest in health informatics. The proposed program will allow for other specializations, in addition to Big Data, to deal with the rapidly changing technology landscape. As noted in the 2016 BC Technology Strategy (<https://bctechstrategy.gov.bc.ca>), the "tech sector is a key driver of growth for the provincial economy, expanding faster than the economy at large." The degree will support the "talent" pillar of

the *#BCTECHStrategy* in conjunction with work-integrated learning. Note that work-integrated learning is one of the “strategy actions” of the 2016 document.

b) Anticipated contribution of the proposed program to the mandate and strategic plan of the institution

Earlier this year, SFU embarked on its university-wide big data initiative. As noted at <https://www.sfu.ca/vpresearch/bigdatainitiative.html>, one of the university priorities is to build “on our strategic investments in advanced computing expertise and infrastructure [to engage] our students in data-intensive research to expand their analytical capacities and meet a significant talent gap.” The structure and curriculum of the proposed Program will allow students to take courses that address research issues and combine them with practical computing skills to solve real world problems.

c) Linkages between the educational goals and the curriculum.

The educational goals of the Program are to provide students with capabilities in the following areas of applied computer science:

- Analysis of scalability of algorithms to big data
- Data warehouses and online analytical processing.
- Efficient storage of big data including data streams.
- Scalable querying and reporting on massive data sets.
- Scalable and distributed hardware and software architectures.
- Software as a service.
- Cloud Computing (e.g. Amazon EC2, Google Compute Engine)
- Big data programming models: map-reduce, distributed databases, software for implementing streaming and sketching algorithms.
- Dealing with unstructured data such as images, text or biological sequences.
- Scalable machine learning methods such as online learning.
- Data mining: methods for learning descriptive and predictive models from data.
- Distributed algorithms over very large graphs and matrices.
- Social media analysis.
- Visualization methods and interactive data exploration

Additional capabilities will be identified in the future in conjunction with discussions with students and with employers. The required and recommended graduate courses listed in part D appendix 1 provide basic coverage of all of the areas listed above, allowing flexibility for more in-depth investigations into the topics based on student interest. Students are required to take an algorithms

course designed specifically for the specialization area (Big Data), since several of the topics focus on different aspects of algorithms. The graduate course in machine learning is required and is taken in the first semester, as it forms an essential basis for many of the topics listed above. The two programming labs provide hands on experience with all of the topics mentioned above that involve implementation. The Big Data systems course then deals with the hardware and associated software issues.

The mandatory co-op component of the Program will provide students with the real-world experience of taking many of the research issues that are introduced in their classes, and seeing how they can be applied to a practical problem.

d) Delivery methods

The Program will be delivered face-to-face. Students will complete at least 30 units of graduate work: a minimum of 15 credits of courses relevant to the field, 12 credits of specialized lab work, and at least 3 credits of co-op. There will not be a capstone. The Program will be delivered using a cohort model, but there will be flexibility for optional courses (primarily in the final term) to deal with individual student interests.

e) Program strengths

This Program fills a large employer need for professional graduates that can participate in the design, development, maintenance and use of information technology systems in general, and with the current and future needs of big data. It provides students with a compulsory 4 or 8 month paid internship (co-op placement), putting students in the workplace within 8 months of starting the Program. It helps fill a widely discussed talent gap in Canada and around the world. Students are provided with extensive support beyond traditional programs in the classroom, in the lab, and in the transition to the workplace. This bridge between the university and the workplace will also strengthen the research profile of the university and increase the visibility of the School of Computing Science in cases where the real world data and problems are brought back into the classroom and research labs.

f) An overview of the level of support and recognition from other post-secondary institutions, and relevant regulatory or professional bodies, where applicable and plans for admissions and transfer within the British Columbia post-secondary education system.

Before we launched the special cohort professional masters program in Big Data in the fall of 2014, we had had discussions with several industry partners including IBM, SAP, Simba, Demonware, BC Genome Centre, and Hootsuite among others, and we observed that all of them are keen to hire highly qualified personnel proficient in Big Data analytics. Companies like Software AG, Oracle

Corporation, IBM, Microsoft, SAP, EMC, HP and Dell have spent more than \$15 billion on software firms specializing in data management and analytics. This interest has been supported by the 100% placement rate for co-op for the 39 students in the current cohort. In the January meeting with our Big Data advisory committee, it was also suggested that the curriculum be expanded to allow for additional streams or specializations. Note, however, that the Program is not intended to be the basis for a PhD.

We have been in close communication with the organizers of UBC's new Master of Data Science program, with the goal of ensuring that our Program and their new program are complementary. Similarly, we are in regular contact with Ryerson University concerning the development of their new program. SFU has been a key participant in Canada's Big Data Talent Gap study, and is a member institution of Canada's Big Data Consortium.

Although the Program is not designed as one for which there would be transferring within the British Columbia post-secondary education system, students could request to have graduate courses taken elsewhere to be considered towards the credit requirements of the proposed degree (as is the case for other graduate programs).

The Program will not require the approval of any external regulatory or professional bodies in the area of computer science.

To show the level of support and recognition for the program, Appendix 3 also contains links to a selection of promotion videos featuring our students and companies in which they worked during their co-op term(s).

g) Related programs in the institution or other British Columbia post-secondary institutions.

The School of Computing Science already offers a research-based Master of Science in Computing Science. However, the existing MSc does not incorporate the 12-credits of specialized lab programming, and the high level of instructional support from domain specialists. The Program has two additional lecture courses (3 credits each) that have been designed for the big data specialization area. The existing MSc does not have a compulsory co-op component. Associated with this co-op placement is a large amount of support for resume preparation, interview skills, and communication provided by the FAS co-op office.

With respect to other graduate computing degrees in BC, the new UBC Master of Data Science Program will complement nicely the SFU Program – the UBC program targets the large pool of applicants without degrees in computing science, while the SFU Program is designed for students who already have a degree in computing science or a related field. So the programs do not overlap in the target populations of applicants.

Our Program develops data architects who apply a deep knowledge of computer science to create and support new systems and processes that find value in the vast amounts of information generated today. Our graduates are not just data analysts, they are data architects.

7. Contact information

Fred Popowich

Professor, School of Computing Science

(778) 782-4193

popowich@sfu.ca

PART B: PROGRAM DETAILS

1. Graduation requirements, target audience

Students will complete at least 30 units of graduate work: a minimum of 15 credits of courses relevant to the field, 12 credits of specialized lab work, and at least 3 credits of integrated learning via co-op. The 15 course credits will be graduate courses focussed in the area of specialization. The 12 credits of advanced industry relevant programming is structured into one 6 credit course in the first semester, and a second 6 credit course for the second semester that together will provide hands on experience within a specialization area. All courses contained in the Program have been offered previously, but the plan is to be able to create new lab courses in the future based on changing market demands. The two Big Data specific lecture courses were offered as special topics courses, and will now be converted into permanent courses. The complete Program length is 16 or 20 months, the length varying depending on whether the co-op placement is for four or eight months.

Our intended audience is individuals, both within Canada and internationally, who already have degrees in computing science or a related field, but who wish to advance their knowledge in a rapidly growing area of information technology. It is not intended to prepare students for PhD studies; that is the focus of existing MSc degrees offered at SFU and elsewhere.

2. Admission requirements

To qualify for admission to the Program, a student must satisfy the university admission requirements for a master's program as stated in Section 1.3.3 of the Graduate Admission section of the SFU calendar, and the student must hold a bachelor's degree or equivalent in Computer Science or a related field with a minimum cumulative grade point average (GPA) of 3.00 (on a scale of 0.0 - 4.33) or the equivalent. Students who have a minimum GPA of 3.33/4.33 on the last 60 credits of undergraduate courses also meet the GPA requirements for admission.

The School's Graduate Admissions Committee may, at its discretion, offer M.Sc. admissions to exceptional students who do not have an undergraduate degree in computer science or a related field (qualifying admission). Minimally, we require demonstrated competence in computer science at the third year level equivalent to CMPT 300 (Operating Systems 1), CMPT 307 (Data Structures and Algorithms) and CMPT 354 (Database Systems and Structures).

With a qualifying admission offer, students who do not have the proper background in computer science may take the three courses listed above in the Summer semester before the Fall cohort begins and then join the Program.

We wish to be inclusive of a variety of incoming students, including SFU students from related disciplines who wish to obtain a master's degree; international students; and professionals who wish to advance their knowledge in the area of applied computing.

3. Evidence of student interest and labour market demand

The expected student interest in this program, both domestic and international students is reflected in the rapid growth in application numbers shown on pages 10 and 13. The industry/employment demand for students who obtain this credential is reflected in the current 100% co-op placement rate in the positions shown in section 3a below. The jobs/careers for graduating students are reflected in the positions provided in section 3b below.

Recruiting, retaining, and training big data analytics professionals is becoming increasingly difficult among employers across all regions, sectors, and industries, according to a recent paper published by [Canada's Big Data Consortium, *Closing Canada's Big Data Talent Gap*](#). The paper outlines six strategies for closing the big data gap. SFU is part of the consortium of senior information technology leaders, created by Ryerson University in 2014, collaborating on key issues relate to big data and analytics. Managing Canada's big data capabilities and addressing the growing need for education and training in the field is critical to the country's competitive advantage.

When SFU launched its Professional Master's Program in Big Data in 2014 as a three-year special cohort initiative, it provided a way to start to meet the demand for big data specialists. Given the dramatic increase in student applications, the increase in the number of employers working with the co-op program, and the successful placement of

students in industry positions, it is clear that the special cohort program should be continued as a regular program upon the completion of its final year as a special cohort.

a. List of 2016 co-op employers and partners

- Aeroinfo
- Amazon
- Appnovation
- Bcjobs.ca
- BuildDirect
- Communications Research Ctr
- CommunitySift
- Demonware
- EOTPRO
- FuseForward
- Left
- Microsoft
- MITACS Kinematic Soup
- PHEMI
- Royal Bank of Canada
- Salesforce
- Samsung Electronics Canada
- SAP Canada
- Schneider Electric
- SES
- Simba
- Stats Canada
- Vancity
- Vancouver Coastal Health

b. List of positions and employers for graduated students


- Product Support Specialist - SAP
- Software Developer, Big Data - SAP
- Agile Developer - SAP
- Senior Software Developer - D2L
- Software Developer - IBM
- Software Engineer - Splunk
- Data Scientist - Hootsuite Media Inc.
- Developer - Visier
- Data Scientist - BuildDirect
- PHP Web Developer - FuturePay
- PHEMI – Data Scientist
- Big Data Developer – Irdeto

Although students from the second cohort have not yet completed their studies, as of November 2, 2016, there are already 9 full time offers or hirings from companies including PHEMI, RBC, Samsung, Scotiabank, and Schneider Electric.

c. List of NOC codes for positions for graduated students

<u>2171</u>	Information systems analysts and consultants
<u>2172</u>	Database analysts and data administrators
<u>2173</u>	Software engineers and designers
<u>2174</u>	Computer programmers and interactive media developers

d. Application Numbers by Country of Citizenship 2014-2016



- “Other” in 2014 includes (in alphabetic order): Bangladesh, Brazil, Cameroon, Republic of Congo, Germany, Ghana, Hong Kong, Russian Federation, Saudi Arabia, Thailand, Venezuela.
- “Other” in 2015 includes: Brazil, Colombia, Ecuador, Egypt, Germany, Ghana, Greece, Grenada, Hong Kong, Iceland, Indonesia, Republic of Korea, Mexico, Nepal, Oman, Philippines, Russian Federation, Spain, Sri Lanka, Syrian Arab Republic, Thailand, Turkey, UK, US, Vietnam, Zambia.
- “Other” in 2016 includes: Bangladesh, Brazil, Colombia, Cuba, Ethiopia, Germany, Hong Kong, Jamaica, Kazakhstan, Kenya, Republic of Korea, Mexico, Russian Federation, Sierra Leone, Sri Lanka, Taiwan, Tunisia, UK, US, Vietnam, Zambia.

4. Eligibility for scholarships, awards, and financial aid

Students will not be able to apply for awards associated with the university's operating budget. They will be eligible to apply for private awards, and the School will be approaching external entities for funding for student awards.

Students admitted to the Program can make use of the financial aid program. The School of Computing Science program manager has contacted the financial aid office concerning this eligibility.

5. Program evaluation and academic/administrative oversight

The steering committee for the Program consists of the program director (currently Dr. Fred Popowich), the computing science graduate program director (currently Dr. Binay Bhattacharya), the director of the School of Computing Science (currently Dr. Greg Mori), one instructor in the program (currently Dr. Martin Ester), and Academic Program Coordinator Katie Knorr.

The School has already established an external advisory committee for the special cohort program, and this advisory committee will be expanded. The external advisory committee currently has five members on three-year terms. Please see appendix 8 for a current list of members, along with the invitation letter template. Each year, we anticipate needing to replace two people. The target size for the committee is six.

The Program will be reviewed internally using the same mechanisms that are used to review the other graduate programs in the School of Computing Science. Program changes will be administered through the existing School of Computing Science Graduate Program Committee.

Co-op employers and students are given surveys on a regular basis with respect to Program content and structure.

6. Main competitors outside BC

Given that the focus of the Program is on big data, we have provided a summary of all "Big Data" oriented programs, including some analysis of similarities and differences including tuition at the following location.

http://www.sfu.ca/~popowich/private/appendix_bigdataprogams.htm

With respect to master degree programs focusing specifically on big data,

- Carleton University in Ottawa offers a collaborative master's in data science which has a longer duration, and is combined with other disciplines outside of computer science. Due to the interdisciplinary nature of their program, it does not provide for an opportunity to cover as many of the computer science topics to the same depth as the SFU Program

- Queen’s University in Kingston offers a Master of Management Analytics, with a heavy focus on modeling and statistics, with a lighter focus on computational aspects. It does not have a compulsory co-op education component, and is priced at a much higher level than the SFU Program.
- Dalhousie University offers the “Master of Applied Computer Science (MACS) [as] a 16-month graduate degree that prepares students for dynamic careers in the software industry.” Their degree does not have a mandatory co-op requirement. Furthermore, their program does not have the intensive instructor led big-data or specialist lab components associated with the SFU Program.
- St. Mary’s University has a Master of Science in Computing & Data Analytics. It is “16-month professional program designed to meet the complex challenges associated with Big Data.” Their program involves 6 credits of work in “applied projects, paid internship, or a thesis, whereas our program involves 12 credits of applied projects within the programming courses, in addition to 3 credits of paid internship.
- The University of Windsor provides a “Master of Applied Computing” as a “course-based professional graduate program ... [to provide] a solid foundation and knowledge of industry oriented practical aspects of Computer Science, preparing [students] to take up positions in the growing software industry in Canada and around the world.” It is the same duration as the SFU program, but has a slightly lower course requirement. As with Dalhousie, their program does not have the intensive instructor led big-data or specialist lab components associated with the SFU Program.

With respect to professional MSc programs in Computer Science outside of British Columbia, the following programs stand out.

- The Master of Science in Applied Computing from the University of Toronto, like the proposed SFU program, is a 16 month program incorporating a paid internship. Launched in 2010, the U of T program has an entrepreneurial focus, with a required course on Technical Entrepreneurship. Their current tuition is \$17,500 for domestic students (Canadian citizens and permanent residents) and \$39,000 for international students.
- The University of Washington offers a Professional Master’s Program in Computer Science since 1996, but it is designed for fully employed individuals who intend to keep working while engaged in their graduate studies. It consists of 8 courses, rather than the 9 associated with the SFU Program, with students generally completing their studies in two and a half years. Their current tuition is (USD) \$13,875 per year for either domestic or international students.
- The Professional Master of Computer Science (MCS) from the University of Illinois has three possible options, one of which is an MCS in Data Science. The tuition for their 32 credit hour degree is (USD) \$19,200, which incorporates a combination of core courses supplemented with advanced courses and electives. Their program does not have a compulsory co-op component.

PART C: RESOURCES

1. Enrolment Plan

The following table summarizes the enrolment over the past three years. With the launch of the Program in the fall of 2017, we anticipate a further increase in the number of applications, but intend to keep the enrollment number at the 2015 and 2016 level of approximately 40 students, since it matches our delivery capacity. Our past trend in application numbers, in conjunction with our current cohort sizes shows that we should be able to meet the expected enrollment targets.

Year Started	2014	2015	2016
Students enrolled (completed)	(13)	39	42
Applications processed	98	238	365

Our enrolment plan includes a commitment to ensure increased participation of women in Computer Science. We focus on female representation in the images we use on social media and in our advertising. Our newest recruitment video features a strong female student in the main role and has otherwise good female representation as well. Currently, one of the two University Research Associates teaching our courses is female, who is also acting as a great role model in her interaction with students and industry. For our third special cohort, the number of female applicants increased from 22% to 29%. The number of female students enrolled in the current cohort is also 29%, up from 23% in the previous year. Currently, 23% of all computing science graduate students are female.

2. Resources required and/or available to implement the program (financial and personnel) including any new faculty appointments

The lab courses will focus on hands-on learning of various models, algorithms, and software. The labs themselves will be run by limited term faculty members, professors of professional practice, and/or university research associates who are knowledgeable in the lab content. The students will have considerable help from teaching assistants. The lab courses will be run in the Vancouver Institute for Visual Analytics (VIVA) lab space in what is now known as the SFU Big Data Hub. Since the lab courses provide the premium experience in this program, our plan is to continue to use the VIVA lab as the location for the lab courses, as has been done during the special cohort delivery. The VIVA lab has recently acquired sophisticated, state-of-the-art computer hardware and software. The VIVA lab which is located in Big Data Hub also includes breakout discussion areas and useful rooms where the instructor or TAs can conduct impromptu discussions with groups of students. These rooms have built-in laptop projectors and other amenities.

Building on lessons learned from the special cohort offerings, we limit enrollment in each lab section to 20 spots and have the instructors and the teaching assistants all spending the time in the lab with the students. This will ensure that we can offer a truly premium lab experience by ensuring we have a high-quality lab for the students admitted into this program. A letter of support from VIVA is attached in appendix 5. For the labs' software and hardware requirements, we intend to use the following resources:

- Premium workstations and displays available in the VIVA lab.
- Amazon compute cloud EC2 / Google Compute Engine.
- Local Hadoop cluster for programming assignments.
- Other cloud computing experimental testbeds already in use in our School for research.

The main need is for 1) limited term faculty members (or grant tenure-track faculty or university research associates) to take charge of and teach each of the two lab courses and lecture courses, and 2) support related to developing employment opportunities for the students, and preparing them for their co-op placements.

The overall need for faculty members can be split up into two parts:

- Five faculty members are required to teach the classroom courses. No new faculty will be required as there are existing faculty members who can teach these courses on a regular schedule.
- Limited term faculty members (or equivalent as mentioned above) are needed to teach the lab courses. These are innovative lab courses, and the course development necessitates the creation of a series of programming assignments that will impart a hands-on experience in programming with the tools commonly used in their area of specialization, which is currently big data analytics. Each of the 6-credit lab courses will require 12 hours per week for students (using the typical 0.5 credit per weekly lab hour). A limited term faculty member will be assigned to each 6-credit lab course, where each section will have a maximum capacity of 20 students. There will also be substantial teaching assistant support in the labs to assist students with their assignments. Thus, for an ongoing commitment to these lab courses, one faculty member will be required for each of the two lab courses.

To support the mandatory co-op component of the Program, the School of Computing Science will be working with the FAS graduate co-op office. Based on the experience from placing the last two cohorts of students into their mandatory co-op placements, we anticipate requiring a dedicated full time co-op program coordinator, and a full time co-op student advisor. These two staff members will help all the students prepare their applications, assist with the interview process, monitor the placements, deal with Program marketing and opportunity development (in order to successfully place 40 students, the identification of over 100 possible positions is required). In circumstances where we are unable to find a direct industry placement for a student, faculty members

having industrial research projects will be able to hire such students into their labs at SFU to collaborate with industry on funded projects. Depending on cohort size, the Program may require additional support from an office assistant from the FAS graduate co-op office.

This new Program will not result in the reduction of our other graduate programs in the School of Computing Science. It will likely result in the increase in the size of some of our existing graduate courses, but we have budgeted additional teaching assistant support to deal with the increased class size of these courses whether they be offered by Computing Science or by other units.

3. Faculty member's teaching/supervision

The teaching and supervision needs of the Program will be provided by the faculty members in the table below which also shows their research expertise, complemented by two limited term faculty members (currently two University Research Associates shown in italics below) who can be dedicated to this specific program and funded by the revenue received from the Program. Note that the two limited term appointees will be working with a permanent faculty member in the School to ensure continuity. The Program will not affect regular faculty members' involvement in other programs, however, their interest in the special cohort offering has already played a key role in its success, and it can play an important role in helping attract new tenure-track faculty members to new computing science faculty positions at SFU.

Data mining	Martin Ester	Jian Pei	Ke Wang
Machine learning	Greg Mori	Oliver Schulte	
Big data systems	Jiannan Wang	<i>Ryan Shea</i>	
Big data algorithms	Valentine Kabanets	Andrei Bulatov	
Programming	Greg Baker	<i>Apala Guha</i>	

4. Proposed tuition and other program fees including a justification

Tuition for the Program will be charged on a program basis, to be \$7,337 per semester for domestic students and \$9,555 for international students effective September 1, 2017. Appendix 9 contains additional information about tuition levels for different academic years, and relates it to the tuition charged under the previous special arrangement cohorts through the office of the Dean on Graduate Studies. Students will normally complete the

Professional Master of Science in Computer Science

program in 4 semesters, so the total program cost for domestic students will be \$29,350, and for international students \$38,219. There will be a continuing fee per term, one half of the regular fee, for those students who take longer than 4 terms. The tuition fees are summarized in the following table. Not all students will need to take a continuing semester; only those that take a reduced course load or an extended co-op placement.

	Term 1	Term 2	Term 3	Term 4	Total Tuition	Continuing Term**	Total with continuing term
Domestic	\$7,337	\$7,337	\$7,337	\$7,337	\$29,350	\$3,669	\$33,018
International	\$9,555	\$9,555	\$9,555	\$9,555	\$38,219	\$4,777	\$42,996

The tuition rates will not only offset the following expenses associated with the delivery of the Program, but will also provide graduate fellowships to Computing Science thesis graduate students in existing programs. For instance, last year we were able to offer entrance scholarships of \$6,500 each to Computing Science thesis graduate students to make us more competitive in admitting the best students to enhance our graduate research program, and increase the quality of the existing program streams. This in turn benefits the professional master's students in the Program as they will receive high quality mentoring from some of these students who will be their teaching assistants.

Fixed Program Expenses	Cost per unit	# of Units	Total	
Program Coordinator (Full time)	\$71,867	1.0	\$71,867	
Recruitment and Marketing	\$15,000	1.0	\$15,000	
Course Development	\$10,000	1.0	\$10,000	
Subtotal of Fixed Expenses				\$96,867
Variable Expenses (per lab section)				
University Research Associates	\$70,542	1.5	\$105,813	
Grant Tenure Track Faculty	\$117,570	0.5	\$58,785	
Teaching Assistants (Lab courses)	\$6,760	2.0	\$13,520	
Teaching Assistants (Other courses)	\$1,352	10.0	\$13,520	
EAL Curriculum	\$1,500	0.5	\$750	
Hardware/Software			\$25,000	
VIVA Lab Space and Tech Support			\$40,000	
Computing Science Graduate Fellowships	\$6,500	20.0	\$130,000	
Subtotal of Variable Expenses (per lab section)				\$387,388

The program management and support expenses are broken into two categories. The fixed expenses will generally not vary based on the cohort size, while the variable expenses are dependent on the cohort size, specifically the number of lab sections that are needed for the cohort. We are assuming a lab section size of 20 students.

Fixed program expenses include the salary plus benefits of the program coordinator, technical support for the software and hardware needed by the big data students, recruitment and marketing expenses along with course development. The Program Coordinator will be an APSA staff member who will oversee all aspects of a Program including, budget, recruitment/marketing, admissions, student support, and general administration. The salary is \$58,907 plus 22% benefits = \$71,866. Note that the projected expenses are based on those incurred (and expected) during the three years of the special cohort offering. Course development is budgeted as an ongoing expense as it will take an ongoing effort from the faculty teaching courses relevant to the Program in order to provide updated assignments and projects for the lab courses. For each lab course we have budgeted \$10,000 towards continued development of revised materials assignments for the lab courses (including material such as video lectures) to give students truly effective hands-on experience in unique areas.

University Research Associates (URAs) and grant tenure track faculty are essential to the success of the premium lab courses that form the backbone of our proposed program. We follow Policy R50.01 on URAs. The salary for the limited term faculty will be \$60,000 per URA plus 17.57% benefits, which amounts to \$70,542 each. Based on past experience, we found that the optimal level of support was having one URA leading each lab section of a 6-credit programming course. Given that we have two programming courses, we have a need for two URAs for each group of 20 students. We are estimating that grant tenure track faculty will need a compensation of \$100,000 per year plus 17.57% benefits. Note that these URAs will also play a key role in the admission process, doing an initial review of applications that are submitted by students wanted to be admitted to the Program, and then working with the Program steering committee and the School of Computing Science graduate admissions committee to determine who should be given offers. As such, dealing with the workload of the large number of applicants will not overload the existing committees.

Teaching assistants will be paid at \$1352 per base unit (for PhD students; the rate is a bit lower for Masters students). We have allocated 2 TAs at 5 base units = \$6760 x 2 = \$13,520 for the lab courses. This is more TA support than normal for our graduate courses, but is warranted due to the premium lab coursework for the program. There is also a budget for additional teaching assistance support for the 15 credits of non-lab courses taken by students in the program.

We had original planned to have a cohort size of at 20 students, which we planned to support with 2 part-time (50%) staff. As we move forward, the plan is to have a cohort size of 40 students supported by two full time co-op staff. The first is a co-op coordinator position (APSA, Grade 10, Step 8, 1.0 FTE at \$98,848.53 per year including benefits) who is responsible for the development of employer relations and the generation of postings within the Program to help ensure that all students are placed in a position in their 4th and/or 5th semesters. A second position is a Co-op Student Career Advisor, (APSA, Grade 7, Step 6, FTE 1.0 at \$ 79,257.95 per year including benefits) required to assist students with resume writing, interview skills, etc. Finally, we anticipate needing a Program Assistant (CUPE, Gr. 7, Step 12 months, FTE 0.5 at \$ 26,744.96 per year including benefits). These expenses are based on workloads identified during the three special cohort offerings of the Professional Masters program and reflect the fact that our Program requires a placement for every student.

The budget also reflects support for students having English as an Additional Language. Offering a special section of co-op's new Job Search Success course for EAL curriculum facilitation will cost \$1500 for a section of 20 students. We expect half the students to sign up for this additional support. The Job Search Success course is a 4 week, self-paced, visually based, online course that is facilitated. The course focuses on the strengths of the multilingual and multicultural student and the curriculum goals include:

- improving students resume and cover letter within a Canadian context
- education on how to overcome barriers
- presentation of information regarding normalizing accents
- development of intercultural communication skills
- strategies for communicating across cultures.

The overall budget for the Program is provided below, with an assumption of a lab size of 20 students for an overall cohort size of 40 students. The budget assumes that 75% of the students are international, and that half of the students will do an extended (8 month co-op placement instead of a 4 month placement). Note that a minimum of 20 students is needed to break-even on the Program with a small contingency fund. Increasing the size of the program beyond 40 does not result in any dramatic revenue generation due to the high costs associated with delivery and overheads.

Professional Master of Science in Computer Science

Domestic Tuition	\$311,840	
International Tuition	\$1,218,234	
Total tuition income		\$1,530,074
Program Costs (including all lab sections)	\$774,776.00	
Program Costs (Fixed)	\$96,867.00	
University Overhead	\$566,127.39	
Total expenses		\$1,437,770
Carry forward / contingency	6%	\$92,304

The budget assumes the following overheads applied to the gross tuition for direction to other units in the university, including the following:

- SFU General University Revenue (GUR)
- Vice-President Academic
- Vice President Academic Service Units (percentage changes annually)
- Faculty of Applied Sciences

The budget assumes that the \$188,491 costs for the full time co-op coordinator, full time co-op student advisor and half time co-op program assistant are drawn from the university overhead, and are provided directly to SFU co-op.

• SFU General University Revenue (GUR)	8.30%
• Vice-President Academic (standard was 69.7%)	9.30%
• Vice President Academic Service Units (changes annually)	14.40%
• Faculty of Applied Sciences,	5.00%
TOTAL OVERHEADS	<u>37.00%</u>

PART D: Appendices

Appendix 1 Calendar entry

Professional Master of Science in Computer Science

MASTER OF SCIENCE

Admission Requirements

To qualify for admission to the Professional Master of Science in Computer Science, a student must satisfy the university admission requirements for a master's program as stated in Section 1.3.3 of the Graduate Admission section of the SFU calendar, and the student must hold a bachelor's degree, or equivalent in computer science or a related field, with a cumulative grade point average (GPA) of 3.00 (on a scale of 0.00 - 4.33) or the equivalent. Alternatively, a minimum GPA of 3.33/4.33 on the last 60 credits of undergraduate courses will also meet the GPA requirements for admission to the program.

The School's Graduate Admissions Committee may offer, at its discretion, M.Sc. admission to exceptional students without an undergraduate degree in computer science or a related field. Minimally we require demonstrated competence in computer science at the third year level equivalent to CMPT 300 (Operating Systems 1), CMPT 307 (Data Structures and Algorithms) and CMPT 354 (Database Systems).

Students who do not have the proper background in computer science may take the three courses listed above in the Summer semester before the Fall cohort begins and then join the program.

Program Requirements

Students will complete a minimum of 30 units of graduate work. These units are divided into three sections: a minimum of 15 units of discipline specific graduate

course work; 12 units of specialized lab work involving advanced industry relevant programming; 3 to 6 units of integrated learning via co-op.

Course work for Big Data Specialization (at least 15 credits)

CMPT 726 - Machine Learning (3)

One of

CMPT 705 - Design and Analysis of Algorithms (3)

CMPT 706 – Design and Analysis of Algorithms for Big Data (3) (recommended)

CMPT 711 - Bioinformatics Algorithms (3)

At least two of

CMPT 741 - Data Mining (3) (recommended)

CMPT 756 – Big Data Systems (3) (recommended)

CMPT 767 - Visualization (3)

CMPT 825 - Natural Language Processing (3)

IAT 814 - Knowledge, Visualization and Communication (3)

STAT 852 –Modern Methods in Applied Statistics (4)

One of

CMNS 815 Communication Theories in Technology and Society (3)

CMPT 829 - Special Topics in Bioinformatics (3)

CMPT 880 - Special Topics in Computing Science (3)

CMPT 881 - Special Topics in Theoretical Computing Science (3)

CMPT 882 - Special Topics in Artificial Intelligence (3)

CMPT 884 - Special Topics in Database Systems (3)

CMPT 885 - Special Topics in Computer Architecture (3)

CMPT 886 - Special Topics in Operating Systems (3)

CMPT 888 - Special Topics in Computer Graphics, HCI, Vision and Visualization (3)

CMPT 889 - Special Topics in Interdisciplinary Computing (3)

CMPT 894 - Directed Reading (3)

Other courses with permission of the School.

Note that STAT 652 – Statistical Learning (3) can be used in place of STAT 852 with permission of the School.

Lab Work for Big Data Specialization (12 credits)

Students will take the following two lab courses worth 6 credits each. Only students enrolled in the Professional Master of Science in Computer Science (Big Data) will be permitted to enrol in these courses:

CMPT 732 - Programming for Big Data 1 (6)

CMPT 733 - Programming for Big Data 2 (6)

Co-op (3 or 6 credits)

A co-op placement is an integral part of this program. Students will register for one or two co-op terms. With assistance from the co-op coordinator and co-op student advisor for this program, students will be expected to find a suitable industry

partner for the co-op placement. The student may instead choose to conduct research in Applied Computer Science at one of the various Computing Science research labs or elsewhere inside the University as a paid research assistant to satisfy the co-op requirement. In extenuating circumstances, a student may appeal to the program director to take an elective course from the list of electives for this program instead of a co-op. Students are required to enroll in at least one of the required courses in the term following the co-op term(s).

Academic Requirements within the Graduate General Regulations


All graduate students must satisfy the academic requirements that are specified in the graduate general regulations, as well as the specific requirements for the program in which they are enrolled, as listed above.

Appendix 2 New courses

CMPT 706 – Design and Analysis of Algorithms for Big Data (3)

CMPT 756 – Big Data Systems (3)

Both of these courses were previously developed and offered as special topics offerings (CMPT 881 and CMPT 886 respectively) with admission restricted to only Big Data students as part of the special cohort offering of the Professional Masters in Big Data. Based on feedback from the previous offerings, we are proposing permanent offerings of these courses as outlined in this appendix.



New Graduate Course Proposal

Please save the form before filling it out to ensure that the information will be saved properly.

Course Subject (eg. PSYC)	CMPT	Number (eg. 810)	706	Units (eg. 4)	3
Course title (max 100 characters including spaces and punctuation) Design and Analysis of Algorithms for Big Data					
Short title (for enrollment/transcript - max 30 characters) Algorithms for Big Data					
Course description for SFU Calendar * Concepts and problem-solving techniques that are used in the design and analysis of efficient algorithms. Special consideration and adaptations for big data applications will be emphasized.					
Rationale for introduction of this course This is a customized course for the Masters of Big Data program. It differs from CMPT 705 in its emphasis on big data problems and a more applications-oriented approach.					
Effective term and year Fall 2017			Course delivery (eg 3 hrs/week for 13 weeks) 3 hours/week for 13 weeks		
Frequency of offerings/year 1			Estimated enrollment/offering 40		
Equivalent courses (These are previously approved courses that replicate the content of this course to such an extent that students should not receive credit for both courses.) CMPT 705					
Prerequisite and/or Corequisite **					
Criminal record check required? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No If yes, then add this requirement as a prerequisite.					
Campus where course will be taught <input checked="" type="checkbox"/> Burnaby <input type="checkbox"/> Surrey <input type="checkbox"/> Vancouver <input type="checkbox"/> Great Northern Way <input type="checkbox"/> Off campus					
Course Components <input checked="" type="checkbox"/> Lecture <input type="checkbox"/> Seminar <input type="checkbox"/> Lab <input type="checkbox"/> Research <input type="checkbox"/> Practicum <input type="checkbox"/> Online <input type="checkbox"/>					
Grading Basis <input checked="" type="checkbox"/> Letter grades <input type="checkbox"/> Satisfactory/Unsatisfactory <input type="checkbox"/> In Progress/Complete				Capstone course? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No	
Repeat for credit? *** <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No		Total completions allowed? _____		Repeat within a term? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No	
Required course? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No		Final exam required? <input checked="" type="checkbox"/> Yes <input type="checkbox"/> No		Additional course fees? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No	
Combined with an undergrad course? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No If yes, identify which undergraduate course and what the additional course requirements are for graduate students:					

* Course descriptions should be brief and should never begin with phrases such as "This course will..." or "The purpose of this course is..." If the grading basis is satisfactory/unsatisfactory include this in the description.

** If a course is only available to students in a particular program, that should be stated in the prerequisite.

*** This mainly applies to a Special Topics or Directed Readings course.

RESOURCES

If additional resources are required to offer this course, the department proposing the course should be prepared to provide information on the source(s) of those additional resources.

Faculty member(s) who will normally teach this course Joseph Peters, Andrei Bulatov, Valentine Kabanets
Additional faculty members, space, and/or specialized equipment required in order to offer this course

CONTACT PERSON

Department / School / Program Computing Science	Contact name Fred Popowich	Contact email popowich@sfu.ca
--	-------------------------------	----------------------------------

DEPARTMENTAL APPROVAL

REMINDER: New courses must be identified on a cover memo and confirmed as approved when submitted to FGSC/SGSC. Remember to also include the course outline.

Non-departmentalized faculties need not sign

Department Graduate Program Committee Binay Bhattacharya	Signature <i>Binay Bhattacharya</i>	Date Oct. 31, 2016
Department Chair Greg Mori	Signature <i>Greg Mori</i>	Date Oct. 31, 2016

LIBRARY REVIEW

Library review done? YES

Course form, outline, and reading list must be sent by FGSC to lib-courseassessment@sfu.ca for a review of library resources.

OVERLAP CHECK

Overlap check done? YES N/A

The course form and outline must be sent by FGSC to the chairs of each FGSC (fgsc-list@sfu.ca) to check for an overlap in content. An overlap check is not required for some courses (ie. Special Topics, Capstone, etc.)

FACULTY APPROVAL

This approval indicates that all the necessary course content and overlap concerns have been resolved, and that the Faculty/Department commits to providing the required Library funds and any other necessary resources.

Faculty Graduate Studies Committee (FGSC) MIRZA FAISAL AREG	Signature <i>Mirza Faisal Areg</i>	Date Nov 1, '16
--	---------------------------------------	--------------------

SENATE GRADUATE STUDIES COMMITTEE APPROVAL

Senate Graduate Studies Committee (SGSC) PETER LILJEHOLM	Signature <i>Peter Liljeholm</i>	Date
---	-------------------------------------	------

ADMINISTRATIVE SECTION (for DGS office only)
 Course Attribute: _____
 Course Attribute Value: _____
 Instruction Mode: _____
 Attendance Type: _____

If different from regular units:
 Academic Progress Units: _____
 Financial Aid Progress Units: _____

CMPT 706 - Design and Analysis of Algorithms for Big Data

This course is primarily a course on problem solving - choosing good approaches and recognizing when a solution is a good solution. The objective of this course is to introduce concepts and problem-solving techniques that are used in the design and analysis of efficient algorithms for big data applications.

Course Outline

Algorithms with large numbers: Asymptotics, Cryptography, Hashing, Randomization

Divide-and-conquer: Recurrences, Sorting, Selection, Lower bounds

Graphs: Graph searches and applications, Trees, Shortest paths, Priority queues

Greedy algorithms: Spanning trees, Amortized analysis, Huffman encoding

Dynamic programming: Shortest paths, Longest subsequences, Knapsacks, Memoization

Algorithms for large-scale graphs: Vertex-centric and edge-centric approaches

Algorithm design for Map-Reduce: Complexity analysis and trade-offs

Consistency in large distributed systems: Paxos consensus, CAP theorem

Algorithms for large datasets: Time-accuracy trade-offs

Suggested Textbooks

Algorithms, S. Dasgupta, C. Papadimitriou, U. Vazirani, McGraw-Hill, 2008.


Algorithm Design, J. Kleinberg, E. Tardos, Addison Wesley, 2006.

Introduction to Algorithms (3rd Edition), T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein, MIT Press, 2009.

Assignments / Homework 30%

Quizzes 30%

Final Exam 40%



New Graduate Course Proposal

Please save the form before filling it out to ensure that the information will be saved properly.

Course Subject (eg. PSYC)	CMPT	Number (eg. 810)	756	Units (eg. 4)	3
Course title (max 100 characters including spaces and punctuation)					
Systems For Big Data					
Short title (for enrollment/transcript - max 30 characters)					
Systems For Big Data					
Course description for SFU Calendar *					
From health care to social media the world generates a tremendous amount of data every day, often too much to be processed on a single computer or even some-times a single data centre. In this graduate seminar we will learn about technologies and systems behind Big Data. In particular, we will discuss what challenges exist in processing and storing massive amounts of data. We will explore how these challenges are being solved in real-world systems as well as the limitations inherent in these designs. The evolution of these technologies will be explored by reading both current and historically significant research papers.					
Rationale for introduction of this course					
This course is currently being taught as specials topics in operating systems (cmpt-886) and is a course offered to students taking the Professional Master's Program in Big Data. With the conversion of the big data program to a permanent program and the importance of efficiently designed systems for processing big data, this course should be offered regularly.					
Effective term and year			Course delivery (eg 3 hrs/week for 13 weeks)		
Fall 2017 (1177)			3 hours per week for 13 weeks		
Frequency of offerings/year			Estimated enrollment/offering		
1			40		
Equivalent courses (These are previously approved courses that replicate the content of this course to such an extent that students should not receive credit for both courses.)					
CMPT 886 when offered as a special topics course in big data					
Prerequisite and/or Corequisite **					
Operating Systems (CMPT-300) and Data Base Systems (CMPT-354), or equivalents.					
Criminal record check required? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No If yes, then add this requirement as a prerequisite.					
Campus where course will be taught <input checked="" type="checkbox"/> Burnaby <input type="checkbox"/> Surrey <input type="checkbox"/> Vancouver <input type="checkbox"/> Great Northern Way <input type="checkbox"/> Off campus					
Course Components <input type="checkbox"/> Lecture <input checked="" type="checkbox"/> Seminar <input type="checkbox"/> Lab <input type="checkbox"/> Research <input type="checkbox"/> Practicum <input type="checkbox"/> Online <input type="checkbox"/> _____					
Grading Basis <input checked="" type="checkbox"/> Letter grades <input type="checkbox"/> Satisfactory/Unsatisfactory <input type="checkbox"/> In Progress/Complete				Capstone course? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No	
Repeat for credit? *** <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No		Total completions allowed? _____		Repeat within a term? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No	
Required course? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No		Final exam required? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No		Additional course fees? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No	
Combined with an undergrad course? <input type="checkbox"/> Yes <input checked="" type="checkbox"/> No If yes, identify which undergraduate course and what the additional course requirements are for graduate students:					

* Course descriptions should be brief and should never begin with phrases such as "This course will..." or "The purpose of this course is..." If the grading basis is satisfactory/unsatisfactory include this in the description.

** If a course is only available to students in a particular program, that should be stated in the prerequisite.

*** This mainly applies to a Special Topics or Directed Readings course.

RESOURCES

If additional resources are required to offer this course, the department proposing the course should be prepared to provide information on the source(s) of those additional resources.

Faculty member(s) who will normally teach this course Ryan Shea
Additional faculty members, space, and/or specialized equipment required in order to offer this course Arrvindh Shriraman can also teach the course Specialized equipment is available in labs within the School of Computing Science

CONTACT PERSON

Department / School / Program Computing Science	Contact name Fred Popowich	Contact email popowich@sfu.ca
--	-------------------------------	----------------------------------

DEPARTMENTAL APPROVAL

REMINDER: New courses must be identified on a cover memo and confirmed as approved when submitted to FGSC/SGSC. Remember to also include the course outline.

Non-departmentalized faculties need not sign

Department Graduate Program Committee Binay Bhattacharya	Signature <i>Binay K. Bhattacharya</i>	Date Oct. 31, 2016
Department Chair Greg Mori	Signature <i>Greg Mori</i>	Date Oct. 31, 2016

LIBRARY REVIEW

Library review done? YES

Course form, outline, and reading list must be sent by FGSC to lib-courseassessment@sfu.ca for a review of library resources.

OVERLAP CHECK

Overlap check done? YES N/A

The course form and outline must be sent by FGSC to the chairs of each FGSC (fgsc-list@sfu.ca) to check for an overlap in content. An overlap check is not required for some courses (ie. Special Topics, Capstone, etc.)

FACULTY APPROVAL

This approval indicates that all the necessary course content and overlap concerns have been resolved, and that the Faculty/Department commits to providing the required Library funds and any other necessary resources.

Faculty Graduate Studies Committee (FGSC) <i>M. Fleg</i>	Signature <i>M. Fleg</i>	Date Nov 1, '16
---	-----------------------------	--------------------

SENATE GRADUATE STUDIES COMMITTEE APPROVAL

Senate Graduate Studies Committee (SGSC) <i>DG76 LILCOAH</i>	Signature <i>[Signature]</i>	Date
---	---------------------------------	------

ADMINISTRATIVE SECTION (for D&S office only)

Course Attribute: _____
 Course Attribute Value: _____
 Instruction Mode: _____
 Attendance Type: _____

If different from regular units:
 Academic Progress Units: _____
 Financial Aid Progress Units: _____

Course Outline - CMPT 756 – Big Data Systems

Calendar Objective/Description

From health care to social media the world generates a tremendous amount of data every day, often too much to be processed on a single computer or even some-times a single data centre. In this graduate seminar we will learn about technologies and systems behind Big Data. In particular, we will discuss what challenges exist in processing and storing massive amounts of data. We will explore how these challenges are being solved in real-world systems as well as the limitations inherent in these designs. The evolution of these technologies will be explored by reading both current and historically significant research papers.

Instructor's Objectives

In this graduate seminar we will learn about technologies and systems behind Big Data.

1. What are the challenges in processing and storing massive amounts of data?
2. How are these challenges being solved in existing systems?
3. What are the limitations of existing systems?
4. We will learn by reading research papers, by watching technical talks and by discussing the material.

Prerequisites

Operating Systems (CMPT-300) and Data Base Systems (CMPT-354), or equivalents.

Topics

- Operating System Fundamentals
- Virtualization
- Big Data File Systems
- Big Data Bases (NoSQL)
- Cloud Computing
- GPGPU (CUDA and OpenCL)
- Big Data Processing Applications (Hadoop and Spark)

Grading

Project: 50%

Homework assignments: 25%

Class participation and Paper Summaries: 10%

Quizzes: 15%

Schedule CMPT 756

Introduction	
WEEK 1:	Lecture 1: <u>Course Introduction</u>
WEEK 2	Lecture 2/3: <u>O/S Review</u>
WEEK 3	Lecture 4/5: <u>O/S Fundamentals</u>
Topic: Multi-Core Programming and Intro Cloud	
WEEK 4	Lecture 4: <u>Multi-Core Revolution</u>
	Lecture 5: <u>Intro Cloud</u>
Topic: Multi-Core and Cloud Continued	
WEEK 5	Lecture 6: <u>Multi-Core Revolution(Part 2)</u>
	Lecture 7: <u>Cloud Computing – 2(Live Demo)</u>
Topic: Scheduling and Virtualization I	
WEEK 7	Lecture 8: <u>Scheduling</u>
	Lecture 9: <u>Virtualization</u>
Topic: Main Memory + Virtualization II	
WEEK 8	Lecture 10: <u>Main Memory -1</u>
	Lecture 11: <u>Multi-core quiz(+Virtualization II)</u>
Topic: Virtual Memory + GPGPU	
WEEK 9	Lecture 10: <u>Main Memory -2</u>
	Lecture 11: <u>Virtual Memory</u>

Topic: Virtual Memory and GPU Lab	
WEEK 10	Lecture 13: <u>Virtual Memory</u>
	Lecture 14: <u>GPGPU Lab (9840)</u>
Topic: File-System + Big-Databases	
WEEK 11	Lecture 15: <u>File-Systems</u>
	Lecture 16: <u>Cloud-Computing Lab (9840)</u>
Topic: Big Data Processing	
WEEK 12	Lecture 17: <u>Big Data Processing (Hadoop + Spark)</u>
	Lecture 18: <u>Big-Database</u>
Topic: Holiday + Final Quiz	
WEEK 13	Monday Holiday
	Lecture 20: <u>Quiz 3(Virtualization + Filesystem)</u>

Appendix 3 Market Analysis

Institutions Providing Similar Programs

The following table provides detailed information on similar programs at other Canadian institutions.

- http://www.sfu.ca/~popowich/private/appendix_bigdataprogams.htm

Interest in the Program

Through social media, we have been able to determine the amount of interest in the program while promoting it to potential students and employers. Below are some SFU Big Data Social Media Statistics and links, along with some comparison numbers for other SFU initiatives as of June 14, 2016.

Facebook

www.facebook.com/sfubigdata

- Launched account on May 25th, 2015.
- 1,569 likes.
- Compare to:
 - SFU Career Services: 1,563 likes
 - SFU Applied Sciences: 1,411 likes
 - SFU Graduate Studies: 779 likes

Twitter

www.twitter.com/bigdata_SFU

- 421 followers (compare to 18 followers as of November 2014)
- tweets : followers ratio = 0.80 (336 tweets : 421 followers)
- Compare to:
 - SFU Applied Sciences: 2.87
 - SFU Graduate Studies: 4.87
 - SFU Career Services: 2.22
- Note that we get the same number of followers as some of the comparable SFU accounts - with only 1/6 to 1/3 of the tweets

YouTube

www.youtube.com/user/SFUBigData

- 43 subscribers since launch on Mar. 7th, 2013

- 3 videos uploaded
 - compare to SFU Faculty of Applied Sciences: 69 subscribers for 41 videos uploaded

LinkedIn

www.linkedin.com/edu/school?id=187080

- 151 followers (since launch on Jan. 13th, 2016)
 - compare to 317 for Faculty of Applied Sciences page

Other Links

Convocation 2016 photo album (Facebook):

<https://www.facebook.com/media/set/?set=a.1058703020886023.1073741836.853205458102448&type=3>

News section from Big Data website:

<http://www.sfu.ca/computing/current-students/graduate-students/academic-programs/bigdata/news.html>

“Meet Our Students” Feature on Big Data website:

<http://www.sfu.ca/computing/current-students/graduate-students/academic-programs/bigdata/students2015.html>

Most recent Big Data Industry Newsletter issue:

<https://www.mailoutinteractive.com/Industry/View.aspx?id=788962&p=f310>

Global News features our students’ research on Vancouver housing crisis in a video:

<http://globalnews.ca/video/2640086/b-c-government-must-help-stop-soaring-house-prices-sfu-report/>

The Globe and Mail writes about our students’ findings on Vancouver’s housing problem:

<http://www.theglobeandmail.com/news/british-columbia/external-factors-setting-vancouvers-housing-costs-university-study-finds/article29640207/>

Metro News features a big data student’s project on chocolate:

<http://www.metronews.ca/news/vancouver/2016/02/15/sfu-student-using-big-data-to-find-chocolate-loving-cities.html>

SFU’s The Peak writes about our student’s project on chocolate:

<http://www.the-peak.ca/2016/02/sfu-student-tracks-chocolate-with-instagram-data/>

Additional employment information on the next several pages

Government
of CanadaGouvernement
du Canada

Canada

Job Bank

[Home](#) > [Explore Careers by Occupation](#) > [Job Market Report](#) > [Printing Options](#) > [Print Preview](#)

Software Engineers and Designers (NOC 2173)

Lower Mainland - Southwest Region, British Columbia

Description

Software engineers and designers research, design, evaluate, integrate and maintain software applications, technical environments, operating systems, embedded software, information warehouses and telecommunications software. They are employed in information technology consulting firms, information technology research and development firms, and information technology units throughout the private and public sectors, or they may be self-employed.

Included Job Titles

application architect, computer software engineer, embedded software engineer, software architect, software design engineer, software design verification engineer, software designer, software engineer, software testing engineer, systems integration engineer - software, technical architect - software, telecommunications software engineer.

+ [View more](#)

Wages - Canada

Wage Estimates

Location	Wage (\$/hr)			Note
	Low	Median	High	
Canada	25.48	41.03	58.65	Note
Alberta	26.92	45.00	75.00	Note
British Columbia	20.00	39.42	57.69	Note
Manitoba	18.87	34.69	51.19	Note
New Brunswick	18.86	39.29	49.63	Note
Newfoundland and Labrador	N/A	N/A	N/A	Note
Northwest Territories	N/A	N/A	N/A	Note
Nova Scotia	20.96	33.33	52.40	Note
Nunavut	N/A	N/A	N/A	Note
Ontario	26.92	43.27	59.13	Note
Prince Edward Island	N/A	N/A	N/A	Note
Québec	25.64	35.60	48.72	Note
Saskatchewan	14.99	37.28	54.21	Note

Yukon

N/A

N/A

N/A

Note

[Source: Labour Market Information - ESDC | Wage Methodology]

Please consult the Lower Mainland - Southwest Region and British Columbia tabs for more useful information related to wages.

Outlook & Prospects - Canada

National Outlook ' 10-Year Projection (2015-2024)

This section provides labour demand and labour supply projections for this occupation over the 2015-2024 period.

Occupational Outlook

SHORTAGE: This occupational group is expected to face labour shortage conditions over the period of 2015-2024 at the national level. The section below contains more detailed information regarding the outlook for this occupational group.

The data in the following table are derived from ESDC's Canadian Occupational Projection System (COPS). COPS uses a variety of models to produce a detailed 10-year labour market projection per broad skill level and per occupation at the national level, which focuses on the trends of labour supply and labour demand over the next ten years.

Employment in 2014	48,400
Median Age of workers in 2014	41
Average Retirement Age in 2014	62

Occupation Projection for Canada

In order to determine the expected outlook of an occupation, the magnitude of the difference between the projected total numbers of new job seekers and job openings over the whole projection period (2015-2024) is analyzed in conjunction with an assessment of labour market conditions in recent years. The intention is to determine if recent labour market conditions (surplus, balance or shortage) are expected to persist or change over the period 2015-2024. For instance, if the analysis of key labour market indicators suggests that the number of job seekers was insufficient to fill the job openings (a shortage of workers) in an occupational group in recent years, the projections are used to assess if this situation will continue over the projection period or if the occupation will move towards balanced conditions.

The analysis of key labour market indicators such as employment and wage growth as well as the unemployment rate suggests that the number of job openings exceeded substantially the number of job seekers in this occupational group over the 2012-2014 period.

For **Software engineers and designers**, over the period 2015-2024, new job openings (arising from expansion demand and replacement demand) are expected to total **18,600**, while **18,200**

new job seekers (arising from school leavers, immigration and mobility) are expected to be available to fill them.

As job openings and job seekers are projected to be relatively similar over the 2015-2024 period, the labour shortage conditions seen in recent years are expected to continue over the projection period. A more detailed analysis of the outlook of this occupation will be released in the coming weeks.

This Chart contains data for Projection of Job Openings vs. Job Seekers for Canada. Information is available in the following tables.

Category	Openings	%
Expansion Demand	9,600	52%
Retirements	7,100	38%
Other Replacement Demand	1,200	6%
Emigration	800	4%
Projected Job Openings	18,600	100%

Category	Seekers	%
School Leavers	18,400	101%
Immigration	7,200	40%
Other	-7,500	-41%
Projected Job Seekers	18,200	100%

[Source: Canadian Occupational Projection System (COPS)]

In which industry or sector do people in this occupation find jobs in Canada?

This table shows the industry and sectors employing the highest number of people in this occupation.

Industry / Sector	%
Professional, scientific and technical services	50.03
Information and cultural industries	13.96
Manufacturing	10.39
Wholesale trade	9.22
Finance and insurance	5.79
Public administration	3.15

[Source: National Household Survey 2011 - Statistics Canada | Legal note]

What percentage of people in this occupation are self-employed?

According to the Labour Force Survey (2014), in Canada, 11% of workers in this occupation were self-employed, while the average for all occupations was 14%.

[Source: Labour Force Survey - Statistics Canada | Legal note]

What proportion of people in this occupation work full-time?

According to the Labour Force Survey (2014), in Canada, 99% of workers in this occupation worked full-time, compared to the average of 81% for all occupations.

[Source: Labour Force Survey - Statistics Canada | Legal note]

What is the proportion of women working in this occupation?

According to the National Household Survey (2011), in Canada, women represented 17% of workers in this occupation compared to the average of 48% for all occupations.

[Source: National Household Survey 2011 - Statistics Canada | Legal note]

What percentage of people in this occupation are members of a union?

This occupation (**Software Engineers and Designers**) is part of a larger group called **Computer and Information Systems Professionals (NOC 217)**. According to the Labour Force Survey (2014), in Canada, the unionization rate for this group was 20%, while the unionization rate for all occupations was 30%.

[Source: Labour Force Survey - Statistics Canada | Legal note]

Please consult the Lower Mainland - Southwest Region and British Columbia tabs for more useful information related to outlook.

Education & Job Requirements - Canada

Employment Requirements

Employment requirements are prerequisites generally needed to enter an occupation.

- A bachelor's degree, usually in computer science, computer systems engineering, software engineering or mathematics
or
Completion of a college program in computer science is usually required.
- A master's or doctoral degree in a related discipline may be required.
- Licensing by a provincial or territorial association of professional engineers is required to approve engineering drawings and reports and to practise as a Professional Engineer (P.Eng.).
- Engineers are eligible for registration following graduation from an accredited educational program, three or four years of supervised work experience in engineering and passing a professional practice examination.
- Experience as a computer programmer is usually required.

[Source: National Occupational Classification 2006 - ESDC]

Regulation by Province/Territory

Some provinces and territories regulate certain professions and trades while others do not. If you have a licence to work in one province, your licence may not be accepted in other provinces or

territories. Consult the table below to determine in which province or territory your occupation/trade is regulated.

Location	Regulation
Alberta	Regulated
British Columbia	Regulated
Manitoba	Regulated
New Brunswick	Regulated
Newfoundland and Labrador	Regulated
Northwest Territories	Regulated
Nova Scotia	Regulated
Nunavut	Regulated
Ontario	Regulated
Prince Edward Island	Regulated
Québec	Regulated
Saskatchewan	Regulated
Yukon	Regulated

[Source: Labour Market Information - ESDC]

Education Programs

Programs in the order in which they are most likely to supply graduates to this occupation (Software Engineers and Designers):

- Computer Science
- Electrical, Electronics and Communications Engineering
- Computer Engineering
- Computer and Information Sciences and Support Services, General
- Engineering, General

[Source: National Household Survey 2011 - Statistics Canada]

Apprenticeship Grants

There are two types of Apprenticeship Grants available from the Government of Canada:

- The **Apprenticeship Incentive Grant (AIG)** is a taxable cash grant of \$1,000 per year, up to a maximum of \$2,000 per person. This grant helps registered apprentices in designated Red Seal trades get started.
- The **Apprenticeship Completion Grant (ACG)** is a taxable cash grant of \$2,000. This grant helps registered apprentices who have completed their training become certified journeypersons in designated Red Seal trades.

[Source: CanLearn - HRSDC]

Information for Newcomers

Please consult the Lower Mainland - Southwest Region and British Columbia tabs for more useful

Information related to education and job requirements.

Date Modified: 2013-09-20

Appendix 4 Consultation comments

Before launching the special cohort in 2014, we commissioned a survey to reach potential applicants to the proposed program. We received close to 500 responses to this survey and there was overwhelming interest in this program. The summary of the survey results is attached in this section.

The main group of people we were able to reach in our survey who were enthusiastic about the special cohort program were SFU students. The majority of these students, from the Faculty of Applied Sciences and from other units across campus, stated that they would be interested in obtaining a professional master's degree. The survey clearly shows that there is a perceived need for such a fulltime professional master's program among our own students.

Prior to the launch of the special cohort program, we had also received more than 50 emails expressing interest in this program from international students who had joined a mailing list that would provide details about the proposed program when it was approved. Since then, we have received well over 3000 inquiries from prospective applicants for this program.

We have consulted with other universities in the development of the Program, and have requested letters of support from a selection of these universities, as summarized below.

- Dalhousie University
- Ryerson University
- University of British Columbia

We also had presentations and feedback on our special cohort program at the following refereed conference.

Anoop Sarkar, Fred Popowich, Alexandra Fedorova (2014). A Professional Big Data Master's Program to train Computational Specialists, BIG DATA AND ANALYTICS EDCON 2014, October 25th and 26th, Mandalay Bay: Las Vegas, Nevada.

Based on feedback obtained from co-op employers, students, and advisors during our first offering of the curriculum in 2014-2015 as a cohort special arrangements master's program in Big Data through the Office of Graduate Studies at SFU, the curriculum was restructured to have key courses moved to earlier in the program (before the co-op placement). Although this results in a very high credit load during the first two semesters of the program, the students from cohort two (2015-16) and from cohort three (2016-17) have generally been able to manage the load, plus we have made arrangements for a reduced course-load alternative that allows for a more even distribution of work over the duration of the program. The 35% attrition rate we saw for the first cohort became less than 10% for cohorts two and three. Additionally, we have made provisions for additional non-computing science courses to be available to students.

Fred Popowich, PhD
Professor
School of Computing Science
Simon Fraser University

August 16, 2016

Dear Dr. Popowich,

I would like to express my full support to your proposed professional MSc program in Computer Science. The program focuses on Big Data, and computational skills to build models, design of algorithms for data science and analytics.

This program would be highly relevant to the needs of society by promoting business growth, economic development and employment generation in Canada and globally. The increased flow of digital information, which is characterized by high volume and variety, provides opportunities for transforming this big data into business intelligence. Demand is growing for data analytics expertise within all sectors, and across a variety of business domains. Tens of thousands of individuals have already been hired globally by organizations mandated by statute to use their technical and analytical skills to run and maintain information technology systems and databases. However, only a small subset of these professionals has been trained in the creation, management, integration, organization and delivery of data analytics methods and processes applied to data science and analytics implementation. Therefore the technical focus of the program combined with the co-op opportunities makes it unique and very relevant for Canada and globally.

School of Computing at Simon Fraser University is well positioned to deliver this program since the faculty members possess the diverse skill sets needed for this emerging discipline. Core technical skills such as math, statistics, machine learning, programming, artificial intelligence, and domain knowledge are the core competencies that are in greatest demand in the marketplace. The masters program is well designed to meet the requirements of the marketplace. One other strength of the program is that it aims to engage industry partners from diverse domains to access extensive data as well as opportunities for the students to work on industry problems.

Based on the above, I strongly support the proposed professional MSc Program in Computer Science.

Yours sincerely,



Ayse Bener, PhD
Professor and Director
Data Science Lab

Director, Big Data, Office of the
Provost and Vice President Academic



a place of mind
THE UNIVERSITY OF BRITISH COLUMBIA

Department of Statistics
The University of British Columbia
3182 Earth Sciences Building
2207 Main Mall, Vancouver, BC
Canada V6T 1Z4
Website : www.stat.ubc.ca

Paul Gustafson, Ph.D.
Professor

Tel: (604) 822-1300
Fax: (604) 822-6860
Email: gustaf@stat.ubc.ca
Web: www.stat.ubc.ca/~gustaf

June 24, 2016

Professor Fred Popowich
Director, SFU Professional Master's Program in Big Data
Simon Fraser University
8888 University Drive
Burnaby BC
V5A1S6

Re: Support of your new program

Dear Dr. Popowich:

We are writing to provide our strong support for the creation of the Professional Master of Science in Computer Science degree program at Simon Fraser University.

The initial offering of the proposed program in the form of the SFU Professional Masters' program in Big Data has already established itself as providing knowledge and training needed by B.C. industry. This demand is evident from the strong placements of graduates from the Big Data Masters' program.

Studies in Canada and the U.S. both call out the critical need for graduates with skills in big data and data science. The value chain related to big data and data science spans from obtaining and handling large amounts of data through to the extraction of knowledge from the data to taking action on extracted knowledge to effect change in organizations. At UBC, the Faculty of Science is in the midst of launching a new Master of Data Science (MDS) program with the first students starting in September 2016. The 10-month MDS program is offered jointly by the Department of Computer Science and the Department of Statistics and provides students with the skills to apply descriptive and predictive techniques to extract and analyze a wide variety of data. The program also focuses on providing graduates the ability to communicate knowledge gained from data analyses to guide change in organizations. Despite a short period of time available for recruitment for the September 2016 launch, the program has attracted a very strong applicant pool, with an anticipated acceptance rate of approximately 25%.

The UBC MDS program complements the SFU Professional Masters on Big Data. The UBC program targets students without degrees in Computer Science or Statistics, whereas the SFU program is designed for students already having a degree in Computer Science. As a result, the

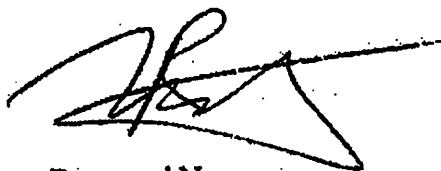
two programs do not overlap in the target population of applicants. Graduates from the SFU program are finding jobs primarily as data architects and software engineers. The target jobs for UBC MDS graduates are more likely to be as data analysts or data scientists.

Together, the UBC MDS program and the SFU program help position B.C. to be the leader in Canada in data science and big data education and help position B.C. to be a leader within North America. Both programs will provide graduates that will help grow B.C.'s technology sector as well as enhance other sectors that are increasingly dependent on data, such as health and resources.

Yours Sincerely,



Paul Gustafson
Professor of Statistics
Co-Director, Master of Data Science Program



Raymond Ng
Professor of Computer Science
Co-Director, Master of Data Science Program

August 18, 2016



Prof. Dr. Fred Popowich
Professor, School of Computing Science
Simon Fraser University
tel: (778) 782-4193
popowich@sfu.ca

Dear Prof. Popowich,

I am writing in support of the Professional Master of Science in Computer Science, and in particular its instantiation in Big Data. The proposed program intends to formalize its offering as a cohort special arrangements program over the last three years. Strengths of the proposed program include:

1. A solid research foundation. The proposed program will capitalize on the research excellence in the area of Big Data of the School of Computing Science, and its leadership in collaborative research with industry through initiatives such as VIVA and the VARDEC consortium. The proposed program will offer first-class education and training by eminent researchers.
2. The focus on programming for Big Data. This is an emerging technical skill that is in high demand by companies that are active in the Big Data space, such as IBM, Google, Microsoft, Amazon, Facebook, SAP and LinkedIn. It is a necessary part of the training helping bridge the gap between theory and practice, that is weaker in other similar programs.
3. The highly structured nature of the program and the tight timeline of 16-18 months for completing the program. Prospective students with a career objective to work in industry want to acquire the necessary skills efficiently and then join the work force in a timely manner. The proposed program fulfills this requirement.
4. The mandatory co-op component. Graduates of this program will be significantly more marketable after having acquired Canadian work experience through the co-op term.
5. Success as a special arrangements program. The proposed program has been running for the last three years, and it has demonstrated its popularity, as the impressive growth in the number of applicants clearly demonstrates. Such growth, combined with keeping the enrollment number fixed at its current level ensures a very strong student and graduate pool. Placement rate for internships has been 100%, which suggests that the demand by industry is quite strong.

.../2

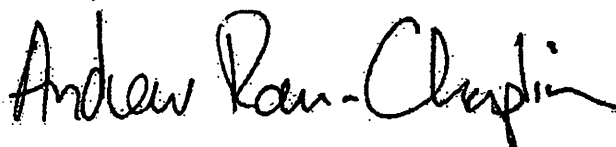
6. Flexibility in admission requirements. The flexibility of the program in requiring computer science competence at the 3rd year level on the basis of three core courses, which can be taken in the summer before joining the program.
7. Complementary nature. The discussion of the main competitors, both within and outside BC, makes a strong case that the proposed program offers a degree that complements well existing offerings in Canada. I expect that the growth in the demand of big data skills allows room for several more programs like the proposed one.

Some questions regarding the resource requirements of the proposed program include:

1. What will be the level of utilization of the VIVA lab by students in the program? Is there a plan for equipment renewal?
2. Amazon compute cloud and Google Compute Engine are not free. Has the cost of these resources been factored into the financial planning for the proposed program?
3. Although the program appears to generate significant surplus, a detailed financial analysis would be helpful in planning.

Overall, I am pleased to lend my strong support to the proposed degree. It aims to satisfy the significant growth of the big data job market, and it builds on research strengths of the School, and a very positive experience with a special offering of the program over the last three years.

Yours sincerely,

A handwritten signature in black ink that reads "Andrew Rau-Chaplin". The signature is written in a cursive, flowing style.


Dr. Andrew Rau-Chaplin
Dean

Appendix 5 Letters of support

In addition to the resources within the School of Computing Science described in part C, the attached letters and memos indicating that sufficient space and other resources are available for the Program. Also included are letters of support from other units within Simon Fraser University.

The letters and memos are from

- The Vancouver Institute for Visual Analytics
- The Department of Statistics
- The Faculty of Education
- The Faculty of Health Sciences



Simon Fraser University
VIVA (ASB 10905)
8888 University Drive
Burnaby, British Columbia
Canada V5A 1S6
Contact: info@viva-viva.ca
Web: www.viva-viva.ca

Vancouver Institute for Visual Analytics

26 June 2016

Dr. Greg Mori, Director
School of Computing Science
Simon Fraser University
8888 University Drive
Burnaby, BC
V5A 1S6

Re: Support for Professional Masters Program in Big Data

Dear Dr. Mori,

As Director of the Vancouver Institute for Visual Analytics (VIVA), I am writing to express our enthusiastic support for the School of Computing Science's proposed Professional Masters Program in Big Data.

VIVA is a joint UBC-SFU institute dedicated to developing Visual Analytics (VA) expertise in industry, government and institutions across Canada, as well as fostering collaborative VA research. VA is the science of analytical reasoning facilitated by interactive visual interfaces. VA tools and techniques enable analysts to synthesize and explore information and derive insight from massive, dynamic and possibly conflicting data collections. VA constantly deals with issues in big data, so graduates from the proposed program would be well suited for careers with the industrial partners of VIVA, and in careers dealing with visual analytics research, development and applications.

VIVA is committed to assisting the School of Computing Science on the proposed program in a number of ways. First, we will provide students in new program with full access to our Visual Analytics Research and Instructional Labs (VARI Labs). The laboratory contains a wide selection of state of the art visual analytics software, along with the capacity to develop and deploy new software for use on a wide range of platforms. Note that Western Economic Diversification Canada has provided \$513,141 in funding for the labs, with another \$1.5 million in-kind contribution from IBM, along with other big data software provided by VIVA partners. Second, we will assist in finding internships for the graduate students through our partner organizations and through our Andrew Wade Visual Analytics Challenge Program. Third, faculty members affiliated with VIVA will be willing to act in advisory roles in curriculum development activities. Finally, VIVA will also work at promoting the new program through its partnerships, including the newly established national not-for-profit Canadian Network for Visual Analytics (CANVAC), involving researchers from Dalhousie University, OCAD University, Simon Fraser University, University of British Columbia, the University of Alberta, the University of Calgary, and York University.

We sincerely look forward to the opportunity of working with the School of Computing Science on this initiative.

Sincerely,

Dr. Ted Kirkpatrick
Director
Vancouver Institute for Visual Analytics
ted@sfu.ca +1.778.782.8639



FACULTY of SCIENCE
Statistics & Actuarial Science

SC K10545
8888 University Drive, Burnaby, BC
Canada V5A 1S6

TEL 778.782.3803
FAX 778.782.4368

www.stat.sfu.ca

Date: 30 June 2016

To: Fred Popowich, School of Computing Science
Re: Support for Professional Masters in Big Data

Our department has reviewed the proposal for a Professional Master of Science in Computer Science degree, which offers students specialized training in Big Data. We see the need for Big Data training as a large and growing trend, and heartily support this program.

We offer two courses that could be used to supplement the program's offerings: STAT 852 Modern Methods in Applied Statistics, and a forthcoming course STAT 652 Statistical Learning. The latter course is an introductory course that will offer a statistical perspective on machine learning, accessible to an audience with a relatively low background in statistics and mathematics, while the former teaches expansive topics in the same area at a level assuming advanced regression and some intermediate statistical theory.

As the Professional Master of Science in Computer Science program expands its areas of specialization, there may be need for topics where there is considerably more need for education in Statistics. We stand ready to collaborate with and support our colleagues in Computing Science if/when such a need arises.

Sincerely,

Thomas M. Loughin
Professor and Chair
tloughin@sfu.ca
778-782-8037

MEMO

Office of the Associate Dean,
Graduate Studies & Research

8888 University Drive
Burnaby BC V5A 1S6
Canada

T: 778.782.4858
F: 778.782.4203

www.educ.sfu.ca

ATTENTION Fred Popowich | TEL
FROM Phil Winne, Associate Dean, Graduate Studies & Research
RE Professional Master of Science in Computer Science: Big Data
DATE 2016 June 27 | TIME 11:10

Fred,

Thanks very much for sharing your School's proposal. While some may perceive a weak connection between this program and the education enterprise, they would be overlooking the vast and varied data about more than 5M students enrolled in Canadian public elementary and secondary programs and 2M Canadian post-secondary students.

It's apparent the objectives, design and intellectual resources of the program are first class. I anticipate opportunities for students enrolled in your program to begin to realize benefits of big data science for one of Canada's most core social enterprises.

P.H.W.

Phil Winne, PhD
Professor & Associate Dean, Graduate Studies & Research
Office of the Dean | Faculty of Education |



SIMON FRASER UNIVERSITY
ENGAGING THE WORLD

FACULTY OF HEALTH SCIENCES

June 28, 2016

Dr. Fred Popowich
Professor and Associate Director (Research and Industry Relations)
School of Computing Science

Dear Fred

I am writing in support of your proposal to introduce a Big Data Professional Master's Degree in Computing Science. As the proposal notes, there is an urgent and growing need for computing science professionals who can make sense of large data sets from a wide variety of sources. In health Sciences, our scientists and health service managers are increasingly challenged by the scale and complexity of administrative data sets that can be utilized to understand patterns in health care and health needs of the populations.

I would like to communicate to you strong support for your proposal from the Faculty of Health Sciences.

Yours sincerely,

A handwritten signature in black ink, appearing to read "John O'Neil", is written in a cursive style.

John D. O'Neil, PhD, FCAHS
Dean of Health Sciences

28-4

Appendix 6 Details of program steering committee (if applicable)

The following faculty members from Computing Science are the members of the steering committee for the Program:

- Binay Bhattacharya, Professor and Director of Graduate Programs, binay@sfu.ca
- Martin Ester, Professor, ester@sfu.ca
- Katie Knorr, Big Data Academic Program Coordinator, bigdata@sfu.ca
- Greg Mori, Professor and School Director, mori@sfu.ca
- Fred Popowich, Professor and Program Director, popowich@sfu.ca

The Program also has an advisory board. Please see appendix 8.

Appendix 7 Abbreviated curriculum vitae for faculty

Attached are CVs for the instructors listed in Part C of the proposal.

- **Greg Baker**
- **Andrei Bulatov**
- **Martin Ester**
- **Apala Guha**
- **Valentine Kabanets**
- **Greg Mori**
- **Oliver Schulte**
- **Ryan Shea**
- **Jiannan Wang**
- **Ke Wang**

Greg Baker

School of Computing Science, SFU
Burnaby, BC, Canada, V5A 1S6
Telephone: 778-782-5755
Fax: 778-782-3045
E-mail: ggbaker@sfu.ca
Canadian citizen

Academic Employment

- 09/2006–present Senior Lecturer, School of Computing Science, Simon Fraser University.
09/2000–08/2006 Lecturer, School of Computing Science, Simon Fraser University.

Education

- 08/1998–08/2000 Simon Fraser University, M.Sc. in Computing Science with senior supervisor was Arthur Liestman. The title of the thesis was “Caches in a Theoretical Model of Multicasting”.
08/1994–04/1998 Queen’s University. Honours Bachelor of Science with a major on mathematics and minor in Computer Science.

Honours and Awards

- 2014 SFU Faculty of Applied Sciences Excellence in Teaching Award
2005 Community Involvement Award, from Computing Science Students Society

Teaching Experience

- 09/2000–present Courses taught include:
- CMPT 120, Introduction to Computing Science and Programming I
 - CMPT 165, Introduction to the Internet and WWW
 - CMPT 376, Technical Writing and Group Dynamics
 - CMPT 383, Comparative Programming Languages
 - CMPT 470, Web-based Information Systems
 - CMPT 732, Programming for Big Data I
- 07/2015–12/2015 Revision of CMPT 732 (Programming for Big Data I) for the Professional Masters Program in Big Data. Updated course to cover big data technologies which were not available during the previous offering of the course. Updated curriculum and delivery to integrate into modified program curriculum.

- 09/2015–08/2016 Complete revision of CMPT 165 (Introduction to the Internet and WWW) for use in distance education and on-campus offerings. The goal of the revision was to update the course so it gives a comprehensive academic overview of modern front-end web development.
- 05/2015 Invited lecture, University of the Philippines Cebu. “Why ‘slow’ languages don’t have to be slow”.
- 02/2013–06/2013 Taught Discrete Mathematics and Web Development at Zhejiang University in Hangzhou, China as part of the SFU CS Dual Degree Program.
- 05/2004–09/2004 Design and creation of course materials for CMPT 120 (Introduction to Computing Science and Programming) in preparation for initial offering in the fall of 2004. This is be the new first course for Computing Science students at SFU. It includes a broad survey of computing science topics and an introduction to programming using Python.

Selected Service Activities

- 01/2010–present Design, development, supervision of student development, and maintenance of the course management system (CourSys) used by the School of Computing Science and other departments. System has expanded to include support for administrative workflow within the University.
- 09/2010–present Supervision of 40 undergraduate project students. This included Dual-Degree students’ Capstone projects, as well as other software development projects, and several directed studies courses.
- 09/2006–08/2008 Director of Undergraduate Studies, School of Computing Science. Responsible for chairing undergraduate curriculum committee; curriculum and procedural oversight; external relations with respect to undergraduate issues. Key initiatives: increased outreach activities to high school teachers and students; defining learning outcomes for courses to standardize offerings across instructors and campuses; academic enhancement program introducing learning skills.

June 16, 2016

Curriculum Vitae

Dr Andrei Bulatov / Professor / Computing Science

Educational Background

- 1995 Ph.D. Mathematics, Institute of Mathematics and Mechanics, Ural Branch, Russia
Algebraic properties of clone lattices
- 1991 M.Sc. Mathematics, Ural State University, Russia
Identities in clone lattices

Employment History at Academic Institutions

- September 2012 - Current Professor, Computer Science, SFU
- September 2008 - August 2012 Associate Professor, Computer Science, SFU
- August 2004 - August 2008 Assistant Professor, Computer Science, SFU
- July 2002 - August 2004 Research Officer, Computer Science, Computing Laboratory, The University of Oxford
- September 1997 - July 2002 Associate Professor, Mathematics and Computer Science, Ural State University, Ekaterinburg, Russia
- September 1993 - August 1997 Assistant Professor, Mathematics, Ural State University

Semesterly Activity at Simon Fraser University

In the past 6 years I taught the following courses: CMPT 705 (Design and analysis of algorithms), MACM 101 (Discrete Mathematics I), CMPT 404 (Cryptography and protocols), CMPT 404 (Design and analysis of algorithms)

Senior Supervisory Duties of a Thesis/Dissertation

2 PhD and 7 MSc students graduated under my supervision: Skvortsov, Evgeny (2010), Cong Wang (2014), Liyue Wang (2016), Hedayat Zarkoob (2014), Omelchenko, Oleksii (2014), Peng, Jian (2013), Shariaty, Shabnam (2012), Bolourani, Siavash (2011), Wang, Cong (2009). I currently supervise 3 PhD and 1 MSc students: Hayati, Maryam, Lu, Zhenjian, Omelchenko, Oleksii, Chang Xu.

I have supervised the following postdoc fellows: Egri, Laszlo, Arseny Shur, Catarina Carvalho, Gabor Kun, Victor Dalmau.

Selected Recent Publications

Graphs of relational structures: restricted colours. In Proceedings of LICS 2016.

Conservative constraint satisfaction re-revisited. *J. of Comp. and System Sci.*, 82(2), 2016, 347-356

(with E.Skvortsov) Phase Transition for Local Search on Planted SAT. *MFCs (2) 2015*: 175-186

Boolean max-co-clones. *Algebra Universalis*, 74(1-2), 2015

Counting Constraint Satisfaction Problems. In the Proceedings of the International Congress of Mathematicians, Seoul, Korea, August 13-21, 2014, vol 4, p. 321-342

(with C.Wang) Approximating Highly Satisfiable Random 2-SAT. *SAT 2014*: 384-398

Curriculum Vitae

Dr Andrei Bulatov / Professor / Computing Science

- (with D.Marx) Constraint Satisfaction Parameterized by Solution Size. *SIAM J. Comput.* 43(2): 573-616 (2014)
- (with C.Wang) Inferring Attitude in Online Social Networks Based on Quadratic Correlation. *PAKDD (1) 2014*: 139-150
- (with V. Dalmau and M. Thurley) Descriptive complexity of approximate counting CSPs. *CSL 2013*: 149-164
- (with M. Dyer, L. Goldberg, M. Jerrum) Log-supermodular functions, functional clones and counting CSPs. *STACS 2012*: 302-313
- The complexity of the counting constraint satisfaction problem. *J. ACM* 60(5): 34 (2013)
- (with M.E. Dyer, L.A.Goldberg, M.Jerrum, C.McQuillan) The expressibility of functions on the boolean domain, with applications to counting CSPs. *J. ACM* 60(5): 32 (2013)
- (with M. Dyer, L. Goldberg, M. Jalsenius, M. Jerrum, D. Richerby) The complexity of weighted and unweighted #CSP. *J. Comput. Syst. Sci.* 78(2): 681-688 (2012)
- (with V. Dalmau, M. Grohe, D. Mar): Enumerating homomorphisms. *J. Comput. Syst. Sci.* 78(2): 638-650 (2012)

Selected Invited Talks

- May 2016 International conference on Universal Algebra and Computing. Boulder, USA, Gibbs measures and finite duality.
- August 2014 International Congress of Mathematicians, Seoul, Korea, Republic of. The Counting Constraint Satisfaction Problems.
- May 2014 2014 North American Annual Meeting of the Association for Symbolic Logic, Boulder, United States. Descriptive Complexity of Counting Constraint Satisfaction Problems.
- February 2014 87 Arbeitstagung Allgemeine Algebra (AAA87), Linz, Austria. Counting Approximation: a hunt for Galois connection.
- August 2013 General Algebra and its Applications, Melbourne, Australia. Counting Approximation: a hunt for Galois connection.
- September 2012 2012 ACP Summer School in Constraint Programming. Constraints: Counting and Approximation

Recent Funding

Contract/Grant: Operating Grant **Awarded:** 2015 **Period:** 2015 – 2019
Project Title: NSERC Discovery Grant **Funding:** NSERC **Type:** External
Annual: \$43,000 **Total:** \$215,000

Awards, Honors and Scholarships

- 2013 **Title:** Best Paper Award **Type:** Research
Organization: 43rd Annual International Symposium on Multiple-Valued Logic (ISMVL 2013)
- 2004 **Title:** Best paper award **Type:** Research
Organization: Algorithmic Learning Theory, 15th International Conference (ALT 2004)
- 2002 **Title:** Best Paper award **Type:** Research
Organization: 43rd Annual IEEE Symposium on Foundations on Computer Science (FOCS)



CURRICULUM VITAE

Family name: Ester
 Position: Professor
 Office phone: (778) 782-4411
 Email: ester@cs.sfu.ca

Given name & initials: Martin M.
 Department/School: Computing Science
 Fax: (778) 782-3045
 Personal URL: <http://www.cs.sfu.ca/~ester/>

ACADEMIC BACKGROUND

PhD	1989	Computer Science	Swiss Federal Institute of Technology (ETH Zurich), Switzerland
MSc	1984	Computer Science	University of Dortmund, Germany

AREAS OF RESEARCH EXPERTISE

Data mining, clustering, graph mining, social network analysis, recommender systems, trust-based recommendation, social network analysis, opinion mining, bioinformatics, biological network analysis, efficient and scalable algorithms, probabilistic graphical models

AWARDS, CITATIONS AND HONOURS

- H-index of 50 (according to Google Scholar, June 20, 2016)
- Total number of citations 19731 (according to Google Scholar, June 20, 2016)
- Paper “A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise”, published at the 2nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD ‘96), received a KDD 2014 Test-of-Time Award
- Paper “A Matrix Factorization Technique with Trust Propagation for Recommendation in Social Networks” received the Best Paper award of the 4th ACM Conference on Recommender Systems (RecSys 2010)

PUBLICATIONS (selection)

Journal papers (total 5 in last 6 years)

1. Gao B.J., Ester M., Xiong H., Cai J., and Schulte O.: “The Minimum Consistent Subset Cover Problem: A Minimization View of Data Mining”, IEEE Transactions on Knowledge and Data Engineering (TKDE): 25(3): 690-703, 2013.
2. Gao B., Griffith O.L., Ester M., Xiong H., Zhao Q., and Jones S.J.M.: On the Deep Order-preserving Submatrix Problem: A Best-Effort Approach. IEEE Transactions on Knowledge and Data Engineering (TKDE), 24(2): 309-325 (2012).

Conference papers (total 41 in last 6 years)

- 3 Yao Wu , Xudong Liu, Min Xie, Martin Ester, Qing Yang:
 CCCF: Improving Collaborative Filtering via Scalable User-Item Co-Clustering. WSDM 2016: 73-82.

CV Martin Ester

- 4 Beidou Wang, Martin Ester, Jiajun Bu, Yu Zhu, Ziyu Guan, Deng Cai:
Which to View: Personalized Prioritization for Broadcast Emails. WWW 2016: 1181-1190.
- 5 Yu Liu, Martin Ester, Bo Hu, David W. Cheung:
Spatio-Temporal Topic Models for Check-in Data. ICDM 2015: 889-894
- 6 Yao Wu, Martin Ester:
FLAME: A Probabilistic Model Combining Aspect Based Opinion Mining and Collaborative Filtering. WSDM 2015: 199-208
- 7 Mohammad A. Tayebi, Martin Ester, Uwe Glässer, Patricia L. Brantingham:
Spatially embedded co-offence prediction using supervised learning. KDD 2014: 1789-1798
- 8 Bo Hu, Mohsen Jamali, Martin Ester:
Spatio-Temporal Topic Modeling in Mobile Social Media for Location Recommendation. ICDM 2013: 1073-1078

GRANTS (only those on whom I am the PI)

Computational infrastructure for cloud computing, computational biology and data mining	NSERC RTI	\$110,711	2014	2014
Computational Methods for the Analysis of the Dynamics and Diversity of Genomes	NSERC Create	\$300,000 per year	2013	2017
Probabilistic Graphical Models for Data Mining and Recommendation in Social Media	NSERC Discovery	\$34,000 per year	2012	2016

GRADUATE STUDENT SUPERVISION

(last 6 years: 4 PhD graduated, 2 MSc graduated, current students listed below)

Mansouri, Mehrdad	PhD	2015	Detecting causal patterns of adverse drug reactions
Tong, He	MSc	2014	Drug target interaction prediction
Heidemeyer, Marten	MSc	2014	Drug target interaction prediction
Nabaei, Boshra	MSc	2014	Online Activity Monitoring
Khakabi, Sahand	PhD	2014	Data mining for precision cancer medicine
Wang, Xin	PhD	2013	Recommendation in social networks
Wang, Beidou	PhD	2013	Prioritization of broadcast emails
Wu, Yao	PhD	2012	Efficient and scalable recommendation

APALA GUHA

aguha@sfu.ca

Postdoctoral Fellow ◊ School of Computing Science ◊ Simon Fraser University

<https://sites.google.com/site/apalaguha/home>

My research interests are in the intersection of energy-efficient architecture and high-level languages. Due to the end of Moore's Law, specialized architecture design is becoming standard, while high-level languages such as Python are becoming standard, for example in big data programming. My goal is to exploit the richer semantic information exposed by high-level languages to design custom hardware.

EDUCATION

University of Virginia

Ph.D. in Computer Engineering

Dissertation: Memory Optimization of

Dynamic Binary Translators for Embedded Systems

Advisors: Prof. Mary Lou Soffa and Prof. Kim Hazelwood

August 2010

Charlottesville, USA

POSITIONS

School of Computing Science, Simon Fraser University

Postdoctoral Fellow

Feb 2016 - Present

Vancouver, Canada

Department of Computer Science and Engineering, IIT-Delhi

Adjunct Professor

Jun 2016 - Present

New Delhi, India

SELECTED PUBLICATIONS (STUDENT NAMES IN BOLD)

Rakhi Hemani, Subhasis Banerjee and Apala Guha. Easy and expressive LLC contention model, *The 2016 International Conference on High Performance Computing Simulation (HPCS 2016)*, Innsbruck, Austria, July 2016.

Rakhi Hemani, Subhasis Banerjee and Apala Guha. On the applicability of simple cache models for modern processors, *2nd International Conference on Green High Performance Computing (ICGHPC)*, Nagercoil, India, February 2016.

Rakhi Hemani, Subhasis Banerjee and Apala Guha. ACCORD: An Analytical Cache Contention Model using Reuse Distances for Modern MultiProcessors, *21st Annual International Conference on High Performance Computing Student Research Symposium (HiPC SRS)*, Goa, India, December 2014.

Apala Guha, Karan Kalra, Sandip Aine. Instruction Set Architecture Customization for Heuristic Search Applications, *Iridian Symposium on Computer Systems (IndoSys)*, Hyderabad, India, June 2014.

Apala Guha, Yao Zhang, Raihan ur Rasool, Andrew A. Chien. Systematic Evaluation of Workload Clustering for Extremely Energy-Efficient Architectures, *ACM SIGARCH Computer Architecture News*, Vol. 41, Issue 2, Pages 22-29, May 2013.

Prasanna Balaprakash, Darius Buntinas, Anthony Chan, Apala Guha, Rinku Gupta, Sri Hari Krishna Narayanan, Andrew Chien, Paul Hovland, Boyana Norris. Exascale Workload Characterization and Architecture Implications, *21st High Performance Computing Symposia (HPC)*, San Diego, USA, April 2013.

Apala Guha, Kim Hazelwood, Mary Lou Soffa. Memory Optimization of Dynamic Binary Translators for Embedded Systems, *8th International Conference on High-Performance and Embedded Architectures and Computing (HiPEAC)*, Berlin, Germany, January 2013.

Prasanna Balaprakash, Darius Buntinas, Anthony Chan, Apala Guha, Rinku Gupta, Sri Hari Krishna Narayanan, Andrew Chien, Paul Hovland, Boyana Norris. Exascale Workload Characterization and Architecture Implications, *Preprint ANL/MCS-P9019-0712*, July 2012.

Apala Guha, Andrew Chien. Systematic Evaluation of Workload Clustering for Designing 10x10 Architectures, *University of Chicago Technical Report TR-2012-05*, June 2012.

Apala Guha, Pietro Cicotti, Allan Snaveley, Andrew Chien. The 10x10 Foundation for Heterogeneity: Clustering Applications by Computation and Memory Behavior, *University of Chicago Technical Report TR-2012-01*, February 2012.

Apala Guha, Kim Hazelwood, Mary Lou Soffa. Memory Optimization of Dynamic Binary Translators for Embedded Systems, *Transactions on Architecture and Code Optimization (TACO)*, Volume 9 Issue 3, September 2012.

Apala Guha. Memory Optimization of Dynamic Binary Translators for Embedded Systems, Ph.D. Thesis, Department of Computer Science, University of Virginia, August 2010.

Apala Guha, Kim Hazelwood, Mary Lou Soffa. Balancing Memory and Performance through Selective Flushing of Software Code Caches, *International Conference on Compilers Architecture and Synthesis for Embedded Systems (CASES)*, Scottsdale, USA, October 2010.

Apala Guha, Kim Hazelwood, Mary Lou Soffa. DBT Path Selection for Holistic Memory Efficiency and Performance, *ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE)*, Pittsburgh, USA, March 2010.

Apala Guha, Kim Hazelwood, Mary Lou Soffa. Code Lifetime-Based Memory Reduction for Virtual Execution Environments, *6th Workshop on Optimizations for DSP and Embedded Systems (ODES) held in conjunction with the 2008 International Symposium on Code Generation and Optimization (CGO)* Boston, USA, April 2008.

Apala Guha, Kim Hazelwood, Mary Lou Soffa. Reducing Exit Stub Memory Consumption in Code Caches, *International Conference on High Performance Embedded Architectures and Compilers (HiPEAC)*, Ghent, Belgium, January 2007.

ADVISING EXPERIENCE

PhD students

- Rakhi Hemani. Computer Science Department, IIT-Delhi. Summer 2014-.
- Manideepa Mukherjee. ECE Department, IIT-Delhi. Spring 2015-Spring 2016.

Masters Students

- Bernard Lin. Big Data Masters Program, SFU. Summer 2016.
- Denis Zubo. Big Data Masters Program, SFU. Summer 2016.
- Jay Naidu. Big Data Masters Program, SFU. Summer 2016.
- Wenzhen Gong. Big Data Masters Program, SFU. Summer 2016.
- Zelin Tian. Big Data Masters Program, SFU. Summer 2016.
- Saeed Soltani. Big Data Masters Program, SFU. Summer 2016.
- Antara Ganguly. ECE Department, IIT-Delhi. Spring 2015-Fall 2015.
- Sachin Negi. Computer Science Department, IIT-Delhi. Summer 2015-Spring 2016.
- Rahul Sharma. Computer Science Department, IIT-Delhi. Fall 2015.
- Pooja Gupta. Computer Science Department, IIT-Delhi. Fall 2014-Summer 2015.
- Nishant Adhikari. Computer Science Department, IIT-Delhi. Summer 2015.
- Rohit Jain. Computer Science Department, IIT-Delhi. 2014.
Now at MathWorks.

CV

Dr. Valentine Kabanets / Associate Professor / Computing Science

Educational Background

2000 Ph.D. Computer Science, University of Toronto, Canada
Nonuniformly Hard Boolean Functions and Uniform Complexity Classes (*Supervisor: Stephen A. Cook*)

Employment History at Academic Institutions

September 2009 - Current Associate Professor, Computing Science, Simon Fraser University
July 2003 - August 2009 Assistant Professor, Computing Science, Simon Fraser University

Senior Supervisory Duties of a Thesis/Dissertation/or Major Project

Name	Degree	Project/Thesis Title	Status	Began	Completed
Anton Cherniavskyi	M.Sc.		Active	2015-3	
Lu, Zhenjian	Ph.D.		Active	2014-3	
Chen, Ruiwen	Ph.D.	Meta-algorithms versus circuit lower bounds	Completed	2011-3	2014-2
Gakkhar, Sitanshu	M.Sc.	Hardcore measures, dense models and low complexity approximations	Completed	2010-3	2012-2
Amiri, Ehsan	Ph.D.	Fingerprinting Codes: Higher Rates, Quick Accusation co-supervised with Gabor Tardos	Completed	2004-3	2010-3

Current Research Interests

Meta-algorithms versus Circuit lower bounds (2012-current). The main goal of computational complexity is to understand what is *efficiently* computable. Given a computational problem, there are two obvious directions to try: design an efficient algorithm (*upper bound*), or prove that no such efficient algorithm exists (*lower bound*). Traditionally, designing efficient algorithms is the subject of the theory of algorithms, while lower bounds are sought in complexity theory. It turns out, however, that there is a deep connection between the two directions: better algorithms (for a certain class of problems) also yield strong lower bounds (for related problems), and vice versa, strong lower bounds translate into more efficient algorithms.

CV

Dr. Valentine Kabanets / Associate Professor / Computing Science

Completed Works (2010-2016)

Papers in Refereed Journals

- R. Impagliazzo and V. Kabanets, Fourier concentration from shrinkage, *Computational Complexity*, pages 1–47, 2016.
- E. Allender, D. Holden, and V. Kabanets, The Minimum Oracle Circuit Size Problem, *Computational Complexity*, pages 1–28, 2016.
- R. Chen, V. Kabanets, and N. Saurabh, An Improved Deterministic #SAT Algorithm for Small De Morgan Formulas, *Algorithmica*, pages 1–20, 2015.
- R. Chen, V. Kabanets, A. Kolokolova, R. Shaltiel, and D. Zuckerman, Mining circuit lower bounds for meta-algorithms, *Computational Complexity* (the special issue for CCC'14), 24(2): 333–392, 2015.
- R. Chen, V. Kabanets, and J. Kinne, Lower bounds against weakly-uniform threshold circuits, *Algorithmica* (the special issue for COCOON'12), 70(1), pages 47–75, 2014.
- H. Dell, D. van Melkebeek, V. Kabanets, and O. Watanabe, Is Valiant-Vazirani's isolation probability improvable?, *Computational Complexity* (special issue for CCC'12), 22(2):345–383, 2013.
- R. Impagliazzo, V. Kabanets, and A. Wigderson, New Direct-Product Testers and 2-query PCPs, *SIAM Journal on Computing*, 41(6):1722–1768, 2012 (special issue for STOC'09).
- R. Impagliazzo, R. Jaiswal, V. Kabanets, and A. Wigderson, Uniform Direct-Product Theorems: Simplified, Optimized, and Derandomized, *SIAM Journal on Computing*, 39(4):1637–1665, 2010.

Papers in Refereed Conferences

- M. Carmosino, R. Impagliazzo, V. Kabanets, and A. Kolokolova, Learning algorithms from natural proofs, *CCC 2016*.
- S. Artemenko, R. Impagliazzo, V. Kabanets, and R. Shaltiel, Pseudorandomness when the odds are against you, *CCC 2016*.
- M. Carmosino, R. Impagliazzo, V. Kabanets, and A. Kolokolova, Tighter connections between derandomization and circuit lower bounds, *RANDOM-APPROX 2015*.
- R. Chen and V. Kabanets, Correlation bounds and #SAT algorithms for small linear-size circuits, *COCOON 2015*.
- E. Allender, D. Holden, and V. Kabanets, The Minimum Oracle Circuit Size Problem, *STACS 2015*.
- R. Chen, V. Kabanets, and N. Saurabh, An Improved Deterministic #SAT Algorithm for Small De Morgan Formulas, *MFCS 2014*.
- R. Chen, V. Kabanets, A. Kolokolova, R. Shaltiel, and D. Zuckerman, Mining circuit lower bounds for meta-algorithms, *CCC 2014*.
- R. Impagliazzo and V. Kabanets, Fourier concentration from shrinkage, *CCC 2014*.
- R. Chen and V. Kabanets, Lower Bounds against Weakly Uniform Circuits, *COCOON 2012*.
- H. Dell, V. Kabanets, D. van Melkebeek, and O. Watanabe, Is Valiant-Vazirani's Isolation Probability Improvable?, *CCC 2012*.
- R. Impagliazzo and V. Kabanets, Constructive proofs of concentration bounds, *RANDOM-APPROX 2010*.
-

Research/Project Funding - Received

Contract/Grant: Research Grant Awarded: 2012 Period: 2012 - 2017
Project Title: Meta-Algorithms versus Circuit Lower Bounds
Funding: NSERC Type: External Annual: \$34,000 Total: \$170,000
Involvement: Principal Investigator

Greg Mori

CONTACT INFORMATION

Associate Professor
School of Computing Science
8888 University Drive
Burnaby, BC V5A 1S6, Canada

Email: mori@cs.sfu.ca
Web: <http://www.cs.sfu.ca/~mori>
Tel.: +1 (778) 782 7111
Fax: +1 (778) 782 3045

RESEARCH INTERESTS

Computer Vision, Machine Learning, Video Analysis, Human Action Recognition, Human Pose Estimation, Pedestrian Detection and Tracking, Object Recognition

EDUCATION

Ph.D. in Computer Science 2004
University of California at Berkeley

Hon. B.Sc. in Computer Science and Mathematics with High Distinction 1999
University of Toronto

PROFESSIONAL EXPERIENCE

Associate Professor Sept 2010 - Present
School of Computing Science, Simon Fraser University

Assistant Professor Aug 2004 - Aug 2010
School of Computing Science, Simon Fraser University

Graduate Student Researcher Aug 1999 - July 2004
Department of Electrical Engineering and Computer Sciences,
UC Berkeley

Student Intern, Computer Vision Group June 2000 - Aug 2000
Intel Corporation, Santa Clara, CA

PUBLICATIONS

Most recent publications in refereed journal papers (names in bold face are my students).

J. Li, H. Hajimirsadeghi, M. Zaki, G. Mori, and T. Sayed. Cyclist's Helmet Recognition Using Computer Vision Techniques. *Transportation Research Record: Journal of the Transportation Research Board*, 2014 (in press)

H. Hedyeh, T. Sayed, M. Zaki, and G. Mori. Pedestrian Gait Analysis Using Automated Computer Vision Techniques. *Transportmetrica A: Transport Science*, 10(3), pp.214-232, 2014.

O. Aziz, E. Park, G. Mori, and S. Robinovitch. Distinguishing the Causes of Falls in Humans Using an Array of Wearable Tri-Axial Accelerometers. *Gait & Posture*, 39(1), pp.506-512, 2014.

S. Oh, S. McCloskey, I. Kim, A. Vahdat, K. Cannons, H. Hajimirsadeghi, G. Mori, A. G. Perera, M. Pandey, J. J. Corso. Multimedia Event Detection and Recounting with Multimodal Feature Fusion and Temporal Concept Localization. *Machine Vision and Applications, MVA* 25(1) pp.49-69, 2014.

M. Ranjbar, T. Lan, Y. Wang, S. Robinovitch, Z. Li, and G. Mori. Optimizing Non-Decomposable Loss Functions in Structured Prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI 35(4) pp.911-924, 2013.

TEACHING

CMPT 225: Data Structures and Programming 2010, 2012, 2013
CMPT 419/726: Machine Learning 2008, 2009, 2011, 2013, 2014
CMPT 415/416: Special Research Projects 2005, 2007, 2014
CMPT 310: Introduction to Artificial Intelligence 2004, 2005, 2006, 2007, 2008
CMPT 882: Recognition Problems in Computer Vision 2004, 2005, 2006, 2007, 2009
CMPT 888: Human Activity Recognition 2010

SUPERVISION

Summary of graduate student senior supervisory duties:

Ph.D.		M.Sc.		Total
Active	Complete	Active	Complete	
4	4	5	19	32

My graduate students who have completed their degrees all did so in a timely fashion (avg. 6 trimesters for MSc students, 10 trimesters for PhD students).

Most recent graduated students:

Jinling Li. M.Sc. Road User Detection and Analysis in Traffic Surveillance Videos.	2012 - 2014
Yasaman Sefidgar. M.Sc. Discriminative Key-Segment Model for Interaction Detection.	2012 - 2014
Amir Hossein Bakhtiari. M.Sc. Detecting Pedestrians Using Motion Patterns: A Latent Tracking Approach. Next Vidigami.	2011 - 2013
Tian Lan. Ph.D. From Flat to Hierarchical: Modeling Structures in Visual Recognition. Next postdoc at Stanford.	2010 - 2013

GRANTS

The most recent grants in which I have played a major role. Total amounts received from 2004 to present: \$1,797,934 Operating + \$360,124 Equipment

NSERC Engage Plus Grants (EPG). Improved face analysis for school photo collections (with Vidigami). G. Mori: \$25,000 (2014)

Disney Research. Hierarchical Max-margin Clustering and Applications. G. Mori: \$12,500 (2014)

NSERC Engage Grants (EG). Matching of social media profile photographs (with ThinkCX). G. Mori: \$25,000 (2014)

Accelerate BC (MITACS) Internship. Using Machine Learning Techniques to Improve Automatic Keyword Extraction from Textual Web Content H. Hajimirsadeghi and G. Mori: \$15,000 (2014)

NSERC Engage Grants (EG). Photo Quality and Content Detection for Image Collections (with Vidigami). G. Mori: \$25,000 (2014)

PROFESSIONAL ACTIVITIES

Associate Director, Research and Industrial Relations, School of Computing Science, SFU AY 2011 - present

Associate Editor, IEEE Transactions on Pattern Analysis and Machine Intelligence, T-PAMI 2010 - present

Associate Editor, IPSJ Transactions on Computer Vision and Applications, CVA 2013 - present

Editorial Board, International Journal of Computer Vision, IJCV 2013 - present

Tenure and Promotion Committee, School of Computing Science, SFU AY 2012-13, 2013 - 2014

Faculty Search Committee, School of Computing Science, SFU AY 2005-06, 2011-12, 2012-13, 2013-14

Education	<p>Ph.D. 1997 (Logic and Computation) Carnegie Mellon University, Pittsburgh, PA, USA, Department of Philosophy Thesis title: "Hard Choices in Scientific Inquiry". Supervisor: Kevin Kelly.</p> <p>M.S. 1994 (Logic and Computation) Carnegie Mellon University, Pittsburgh, PA, USA, Department of Philosophy Thesis title: "The Computable Testability of Uncomputable Theories"</p> <p>B.Sc. 1992 (Cognitive Science) with high distinction University of Toronto, Ontario, Canada</p> <p>Abitur (grade 13, final grade average 1.1) 1988 Hamburg, Germany.</p>
University Appointments	<p>Associate Director. 2007-2009, School of Computing Science, Simon Fraser University. Associate Professor. September 2006-current. School of Computing Science, Simon Fraser University.</p> <p>Associate Professor. August 2004–September 2006. Assistant Professor July 2001–August 2004. Department of Philosophy and School of Computing Science, Simon Fraser University.</p> <p>Assistant Professor. (Tenure-Track) July 1997–June 2001. Department of Philosophy, University of Alberta.</p> <p>Adjunct Professor. Department of Computing Science, University of Alberta. October 1998–present.</p>
Honors and Awards	<p>Distinguished Junior Scholar in Residence, Peter Wall Institute for Advanced Studies; July 1999, University of British Columbia, Vancouver, Canada.</p> <p>Phi Kappa Phi Honor Society; inducted 1996.</p> <p>Werner von Siemens Scholarship; 1992-94 sponsored by Siemens AG (covered tuition and living expenses for 2 years during Master's Degree).</p> <p>Daniel Berlin Scholarship; awarded 1991 (an award for Cognitive Science Specialists at the University of Toronto).</p> <p>Undergraduate Research Award; awarded 1991 (by the National Sciences and Engineering Research Council of Canada).</p> <p>University College Scholarship; awarded 1988 and 1989 (by University College, University of Toronto).</p>
Professional Activities	<p>Program Committee Member: ICML 06 (Senior), ICML (04,06–Senior Comm, 08), UAI (02,04), FLAIRS (06), AAI (08), IJCAI 09.</p> <p>Action Editor for <i>Computational Intelligence</i></p> <p>Associate Editor for <i>Philosophy in Review</i> (Epistemology, Philosophy of Science, Decision Theory), 1998–2001.</p> <p>Referee for <i>Algorithms</i>, <i>J. of AI Research</i>, <i>J. of Machine Learning Research</i>, <i>Computational Intelligence</i>, <i>Annals of Mathematics and Artificial Intelligence</i>, <i>IJCAI</i>, <i>Erkenntnis</i>, <i>Synthese</i>, <i>British Journal for the Philosophy of Science</i>, <i>Journal of Symbolic Logic</i>, <i>Mathematical Social Sciences</i>, <i>Dialogue</i>, <i>Canadian Journal of Philosophy</i>, <i>Canadian Philosophical Association</i>, <i>Philosophy of Science</i>, <i>Logic and Computation</i>.</p> <p>Member of <i>Philosophy of Science Association</i>, <i>Canadian Philosophical Association</i>.</p> <p>Organizer of 2000 Meeting of the Western Canadian Philosophical Association (with Jeff Pelletier)</p>
Research Experience	<p>Research Assistant for Office of Naval Research Project (August 1995 - May 1996) Project Title: "Coordination and Cooperation among Tactical Picture Agents". ONR contract N00014-95-1-1161: Principal Investigator: C. Bicchieri, Carnegie Mellon University.</p> <p>Research Intern with Siemens Corporate Research (Summer 1994) Internship with the Learning Systems Department. Princeton, New Jersey.</p> <p>Research Assistant in Computational Linguistics (Summer 1991) Supervisor: Graeme Hirst, Department of Computer Science, University of Toronto. Funded by an Undergraduate Research Award from the National Sciences and Engineering Research Council of Canada.</p>
Languages	<p>English and German with native competence. Reading knowledge of French.</p>
Citizenship	<p>Canadian and German.</p>
Interests	<p>I hold the title of a FIDE chess master. During high school I played for the strongest professional team in Germany, together with world champions and contenders such as Boris Spasski, Nigel Short, John Nunn and Lubomir Kavalek.</p>

Abbreviations: NSERC = Natural Sciences and Engineering Research Council Canada, SSHRC = Social Science and Humanities Research Council Canada, MITACS: Mathematics, Information Technology and Complex Systems.

Type	Role	Period	Project Title	Funding Agency	Total Amount	Notes
Research Grant	PI	2008 - 2013	Machine learning for entity-relationship databases	NSERC	\$85,000	
Industrial Internship	PI	2008 - 2009	Data Mining for Distributed Database with Encrypted Information	MITACS	\$15,000	Student: Z.Lu with Bits Republic Technologies
Research Grant	PI	2004-2007	The Epistemology of Rational Choice and Its Applications	SSHRC	\$79,000	
Research Grant	PI	2003 - 2007	A Learning-Theoretic Approach to Discovering Causal Models from Large Datasets	NSERC	\$72,000	
Research Grant	PI	2001-2003	The Epistemology of Rational Choice in Social Interactions	President's Research Grant	\$10,000	
Research Grant	PI	1999 - 2003	Automated Inference of Conservation Principles in Particle Physics	NSERC	\$61,200	
Research Grant	PI	1999 - 2002	The Epistemology of Rational Choice in Social Interactions	SSHRC	\$40,420	
Conference Grant	Organizer	2000	WCPA meeting	SSHRC	\$8,000	

Presentations and Talks: Summary

Conference Presentations (Total: 21, 1994-2008)

Invited Lectures (Total: 24, 1998-2008)

Local Seminars (Total: 16, 2000-2008)

In Summary: 60 presentations, lectures, and seminars given at various institutions. This includes the following, sorted by location.

- **USA:** Stanford University, ; Carnegie Mellon University; California Institute of Technology; University of Michigan Ann Arbor; University of California at San Diego; University of Washington.
- **Canada:** Simon Fraser University, University of Alberta, University of British Columbia, University of Victoria, University of Lethbridge.
- **Europe:** University of Maastricht, Netherlands; Paul Sabatier University, Toulouse, France; University of Konstanz, Germany.
- **Australasia:** University of Tsukuba, Japan; Japanese Advanced Institute for Science and Technology; Australian National University; Australian National Logic Summer School.

Ryan Shea
University Research Associate
Simon Fraser University
Burnaby, British Columbia, Canada

June 2016
ryan.shea@sfu.ca
604-762-4826

Research Area

- Big Data, Cloud Computing, Cloud Gaming, Computer and Network Virtualization.

Education

- **Simon Fraser University** Burnaby, B.C., Canada
Doctor of Philosophy 2010 - 2016
 - Major: Computing Science
 - Thesis: Performance and energy efficiency of virtual machine based clouds
 - Supervisor: Professor Jiangchuan Liu
- **Simon Fraser University** Burnaby, B.C., Canada
Graduate Certificate 2012
 - Certificate in University Teaching and Learning for Graduate Students
- **Simon Fraser University** Burnaby, B.C., Canada
Bachelor of Science 2005 - 2010
 - Major: Computing Science
 - Obtained President's Roll in Fall 2009 and Summer 2010 (4.0+ GPA)
 - Obtained Honour Roll status for Spring 2010 (3.5+ GPA)

Research, Teachings and Work Experience

- **University Research Associate** Simon Fraser University
School of Computing Science Jan 2016 - Current

Researcher and Instructor associated with the Professional Master's Program in Big Data. Designed and facilitated the graduate seminar course Big Data System (CMPT-886). Currently researching topics related to big-data processing and future data-center design.
- **Cloud and Big Data Architect** Simon Fraser University
School of Computing Science Jan 2015 - Current

Designed the architecture for SFUcloud.ca, an advanced IaaS cloud based on Xen Server and Apache's Cloud stack. The clouds primary purpose is to process Big Data workloads and collect data through advanced instrumentation. This project is funded through a NSERC Research Tools and Instruments Grant.
- **Sessional lecturer** Simon Fraser University
School of Computing Science Jan 2015 - April 2015
- **Teaching and Teaching Assistant** Simon Fraser University
School of Computing Science Fall 2010 - Fall 2016

Selected Awards, Grants & Honours

Postgraduate Major Awards (SFU, 2012-Present)

NSERC Alexander Graham Bell Canada Graduate Scholarship (\$105,000 Over 3 Years) . . . 2013

Postgraduate (SFU, 2010-Present)

SFU Graduate Fellowship (\$6500) per year 2012-2015
Best Student Paper Award IEEE/ACM IWQoS'12 2012
Nokia Research Fellowship (\$3500) 2012

Selected Publications

Peer-Reviewed Articles

1. W. Cai and R. Shea and C. Y. Huang and K. T. Chen and J. Liu and V. C. M. Leung and C. H. Hsu. "The Future of Cloud Gaming [Point of View]" *The Proceedings of the IEEE - Special Issue: Control Challenges in Microgrids and the Role of Energy Sustainable Buildings*, 2016
2. Ryan Shea, Di Fu, and Jiangchuan Liu. Cloud Gaming: Understanding the Support from Advanced Virtualization and Hardware *IEEE Transactions on Circuits and Systems for Video Technology - Special Issue on Visual Computing in the Cloud: Cloud Gaming and Virtualization*, 2015
3. Ryan Shea, Jiangchuan Liu. "Cloud gaming: architecture and performance" *IEEE Network Special Issue on Cloud and Data Center Performance July-August 2013*, 2013.
4. Ryan Shea, Jiangchuan Liu. "Performance of Virtual Machine Under Networked Denial of Service Attacks: Experiments and Analysis". *IEEE Systems Journal Special Issue - Security and Privacy in Complex Systems (2012)*, 2012.
5. Ryan Shea, Jiangchuan Liu. "Network Interface Virtualization: Challenges and Solutions" *IEEE Network Special Issue on Network Virtualization September/October 2012*, 2012.
6. Ryan Shea, Di Fu, and Jiangchuan Liu. "Towards bridging online game playing and live broadcasting: design and optimization." *Proceedings of the 25th ACM Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV'15)* 2015.
7. Ryan Shea, Di Fu, and Jiangchuan Liu. "Rhizome: utilizing the public cloud to provide 3D gaming infrastructure." *Proceedings of the 6th ACM Multimedia Systems Conference (MMSYS'15)*, 2015.
8. Ryan Shea, Feng Wang, Haiyang Wang, Jiangchuan Liu. "A Deep Investigation into Network Performance in Virtual Machine based Cloud Environment". *Proceedings of the 33rd Annual IEEE International Conference on Computer Communications (INFOCOM'14)*, 2014.
9. Ryan Shea, Haiyang Wang, Jiangchuan Liu. "Power Consumption of Virtual Machines with Network Transactions: Measurement and Improvement". *Proceedings of the 33rd Annual IEEE International Conference on Computer Communications (INFOCOM'14)*, 2014.
10. Ryan Shea, Jiangchuan Liu. "Understanding the Impact of Denial of Service Attacks on Virtual Machines" In *Proceedings of the IEEE/ACM 20th International Workshop on Quality of Service '12*, 2012
- Best Student Paper Award

Jiannan Wang

CONTACT INFORMATION	TASC 1 9237 School of Computing Science Simon Fraser University	<i>Voice:</i> 1-778-782-4288 <i>E-mail:</i> jnwang@sfu.ca <i>Homepage:</i> http://www.cs.sfu.ca/~jnwang
RESEARCH INTERESTS	Data Management, Database Systems, Big Data, Crowdsourcing, Data Science, Data Cleaning, Data Integration, Approximate Query Processing, Information Retrieval	
CURRENT POSITION	Simon Fraser University (SFU), Burnaby, BC, Canada Assistant Professor, School of Computing Science	Jan 2016 - Present
EDUCATION	Tsinghua University (THU), Beijing, China Ph.D., Computer Science and Technology	Sept 2008 - Jul 2013
	Harbin Institute of Technology (HIT), Harbin, China B.Sc., Computer Science and Technology	Sept 2004 - Jun 2008
PROFESSIONAL EXPERIENCE	University of California Berkeley (UCB), Berkeley, CA, USA Postdoctoral Associate, AMPLab at EECS	Aug 2013 - Jan 2016
	Qatar Computing Research Institute, Doha, Qatar Research Associate, Data Analytics Group	Nov 2012 - Feb 2013
	University of California Berkeley, Berkeley, CA, USA Visiting Scholar: Database Group and AMPLab	Aug 2011 - Mar 2012
	Chinese University of Hong Kong, Hong Kong Visiting Scholar: Database Group	Mar 2010 - Sept 2010
PUBLICATIONS	Selected Recent Publications	
	<ul style="list-style-type: none">• Sanjay Krishnan, Jiannan Wang, Michael J. Franklin, Ken Goldberg, Tim Kraska. PrivateClean: Data Cleaning and Differential Privacy. SIGMOD 2016, San Francisco, CA, USA.• Daniel Haas, Jiannan Wang, Eugene Wu, Michael J. Franklin. CLAMShell: Speeding up Crowds for Low-latency Data Labeling. PVLDB 9(4) (2015)• Sanjay Krishnan, Jiannan Wang, Michael Franklin, Ken Goldberg, Tim Kraska. Stale View Cleaning: Getting Fresh Answers from Stale Materialized Views. PVLDB 8(12) (2015)• Yudian Zheng, Jiannan Wang, Guoliang Li, Reynold Cheng, Jianhua Feng. QASCA: Quality-aware Task Assignment System for Crowdsourcing Applications. ACM SIGMOD 2015, Melbourne, VIC, Australia• Jiannan Wang, Sanjay Krishnan, Michael Franklin, Ken Goldberg, Tova Milo, Tim Kraska. A Sample-and-Clean Framework for Fast and Accurate Query Processing on Dirty Data. ACM SIGMOD 2014, Snowbird, UT, USA• Jiannan Wang, Nan Tang. Towards Dependable Data Repairing with Fixing Rules. ACM SIGMOD 2014, Snowbird. UT, USA	

- Jiannan Wang, Guoliang Li, Tim Kraska, Michael J. Franklin, Jianhua Feng.
Leveraging Transitive Relations for Crowdsourced Joins.
ACM SIGMOD 2013, New York, NY, USA
- Jiannan Wang, Tim Kraska, Michael J. Franklin, Jianhua Feng.
CrowdER: Crowdsourcing Entity Resolution.
PVLDB 5(11) (2012).

TEACHING CMPT 733: Big Data Programming Spring 2016

- MENTORSHIP**
- Ruochen Zhang, Undergraduate, Simon Fraser University Feb 2016 - Present
 - Dongxiang Zhang, Undergraduate, Simon Fraser University Feb 2016 - Present
 - Sanjay Krishnan, PhD, UC Berkeley Fall 2013 - Present
 - Liwen Sun, PhD, UC Berkeley Spring 2014 - Present
 - Daniel Haas, PhD, UC Berkeley Summer 2014 - Jan 2016
 - Juan Sanchez, MS, UC Berkeley Summer 2014 - Aug 2015
 - Wenbo Tao, Undergraduate, Tsinghua University Summer 2014
 - Yudian Zheng, Undergraduate, Nanjing University Fall 2012 - Summer 2013

- GRANTS**
- [G2] NSERC Discovery Grant. Crowdsourced Data Cleaning. J. Wang: \$170,000 (PI) (2016-2021)
 - [G1] NSERC RTI Grant. Computational Infrastructure for Online Big Data Analytics. J. Pei, A. Shriraman, and J. Wang: \$148,408 (Co-PI) (2016 - 2017)

- PROFESSIONAL ACTIVITIES**
- **Program Committee**
 - SIGMOD (2016, 2017), SIGMOD Demo (2016)
 - HCOMP (2016)
 - ICDE/TKDE poster (2016)
 - WAIM (2014, 2015, 2016)
 - APWeb (2016)
 - **Chairing**
 - SIGMOD 2017 Registration Chair

Ke Wang

Education

- Ph.D, Information and Computer Science, Georgia Institute of Technology, Atlanta, USA, December 1986; MSc., Information and Computer Science, Georgia Institute of Technology, Atlanta, USA, August 1984.

Employment

- 2003 – Present: Professor, School of Computing Science, Simon Fraser University.
- 2000 – 2003: Associate Professor, School of Computing Science, Simon Fraser University.
- 1999 — 2000: Associate Professor, School of Computing, National University of Singapore.
- 1994 — 1998: Senior Lecturer, School of Computing, National University of Singapore.
- 1992 — 1994: Research Fellow, National University of Singapore.
- 1990 — 1991: Research Scientist, University of Lethbridge and University of Alberta
- 1987 — 1990: Associate Professor, Chongqing University

Ke Wang's research interests include database technology, data mining and knowledge discovery, machine learning, and emerging applications. He has published the book on the topic of privacy preservation: "Introduction to Privacy-Preserving Data Publishing: Concepts and Techniques" (B. Fung, K. Wang, A. Fu, and P. Yu, Data Mining and Knowledge Discovery Series. Chapman & Hall/CRC, August 2010) and over 100 papers in database, information retrieval, and data mining conferences, including SIGMOD, SIGIR, PODS, VLDB, ICDE, EDBT, SIGKDD, SDM and ICDM. Recent publications can be found at <http://www.cs.sfu.ca/~wangk>. 30 of his papers received over 100 citations each. His research is supported by NSERC Discovery Grants and several CRD grants and RTI grants. He has worked extensively with industries including BC Hydro, MDA, and CIBC. 15 MSc students and 10 PhD students have graduated, with another 5 students currently pursuing PhD, under his supervision. One of his students is a CRC II chair at University of McGill. Ke Wang is an associate editor for ACM Transactions on Knowledge Discovery from Data, and was an editorial board member for Journal of Data Mining and Knowledge Discovery from 2007 to 2011 and an associate editor of the IEEE Transactions on Knowledge and Data Engineering from 2003 to 2007, the PC co-chair for SIAM Conference on Data Mining 2008, the general co-chair for SIAM Conference on Data Mining 2015 and 2016.

Recent Publication

Ke Wang, Peng Wang, Ada Fu, Raymond Wong. Generalized Bucketization Scheme for Flexible Privacy Settings. Information Sciences, Volume 348, 20 June 2016, Pages 377–393

Hongwei Liang, Ke Wang, Feida Zhu. Mining Social Ties Beyond Homophily. ICDE 2016

Junqiang Liu, Ke Wang, Benjamin Fung. Mining High Utility Patterns in One Phase without Generating Candidates . TKDE, 201

Yue Wang, Ke Wang, Ada Wai-Chee Fu, and Raymond Chi-Wing Wong. KeyLabel Algorithm for Keyword Search in Large Graphs. IEEE Big Data 2015

Wei Xie, Feida Zhu, Siyuan Liu, and Ke Wang. Modelling Cascades Over Time in Microblogs. IEEE Big Data 2015

Aungon Nag Radon, Ke Wang, Uwe Glaesser, Hans Wehn, and Andrew Westwell-Roper. Contextual Verification for False Alarm Reduction in Maritime Anomaly Detection. IEEE Big Data 2015

Xiaoning Xu; Chuancong Gao; Jian Pei; Ke Wang; Abdullah Al-Barakati. Continuous Similarity Search for Evolving Queries. KAIS, pp 1-30, 15 October 2015

Chenyi Zhang, Hongwei Liang, Ke Wang, Jianling Sun. Personalized Trip Recommendation with POI Availability and Uncertain Traveling Time. CIKM 2015. The runner-up for the best student paper award.

Shuaiqiang Wang, Yun Wu, Byron J. Gao, Ke Wang, Hady W. Lauw, and Jun Ma. A Cooperative Coevolution Framework for Parallel Learning to Rank. TKDE, Issue No.12 - Dec. (2015 vol.27), pp: 3152-3165

Lei Dong, Xuan Chen, Jianxiang Zhu, Hong Chen, Ke Wang, Cuiping Li. A Secure Collusion-aware and Probability-aware Range Query Processing in Tiered Sensor Networks. The 34th IEEE Symposium on Reliable Distributed Systems, 2015.

Chao Han and Ke Wang. Sensitive Disclosures under Differential Privacy Guarantees. IEEE BigData Congress 2015.

Chenyi Zhang and Ke Wang. POI Recommendation through Cross-Region Collaborative Filtering. KAIS, pp 1-19, February 2015.

Ke Wang, Chao Han, Ada Waichee Fu, Raymond Chi Wing, Philip S. Yu. Reconstruction Privacy: Enabling Statistical Learning. EDBT 2015.

Loc Do, Hady W. Lauw, Ke Wang. Mining Revenue-Maximizing Bundling Configuration. PVLDB 2015.

Chenyi Zhang, Ke Wang, Ee-peng Lim, Qinneng Xu, Jianling Sun, and Hongkun Yu. Are Features Equally Representative? A Feature-Centric Recommendation. AAAI 2015.

Appendix 8 Advisory Board

Dennis Buttera (IBM, Ottawa)
Jesse Calderon (Tableau)
George Chow (Simba)
Chris Neumann (founder of Data Hero)
Paul Terry (PHEMI)

INVITATION LETTER TEMPLATE:

[DATE]

[NAME]

[ADDRESS]

[CITY], [STATE/PROVINCE] [POSTAL CODE]

[COUNTRY]

Dear [NAME],

I am writing to invite you to join the Big Data Professional Master's Program External Advisory Council at Simon Fraser University. The council plays a valuable role in providing advice on the activities and future directions of the program.

The Professional Master's Program in Big Data was launched in 2014 and has seen tremendous growth in student numbers since then. We will see the first cohort of 13 students convocate this June. The second cohort of 39 students began the program in September 2015 and will finish in June 2017. And this year, we will be welcoming our third cohort of over 50 students.

To ensure the continued success of the program, we must be responsive to changes in its external environment due to technological advances and industry needs, government priorities, community and professional expectations, and a challenging fiscal climate. In this context, we rely on advice from the advisory council to stay well-informed of expectations of industry, related professions, and society.

The advisory council is currently made up of 5 external members, all senior leaders from industry. In terms of commitment involved, the advisory council meets once a year via video conference to discuss strategic priorities for the program. At times, individual members may be called upon for feedback on initiatives that take place during the year. Should you accept, your term on the council would be three years starting [DATE].

If you are willing to serve on the advisory council, please let me know via e-mail. If you have questions, please don't hesitate to contact me via phone at +1 (778)782-4193 or via email at popowich@sfu.ca.

I sincerely hope that you become a member of the Big Data Professional Master's Program External Advisory Council to help shape the future of the program and ensure its continued success.

Sincerely,

Dr. Fred Popowich,
Big Data Program Director, Professor, and Associate Director, Research & Industry Relations

Appendix 9 Supplementary financial information

Tuition levels

The following table provides details on past and projected tuition levels for domestic and for international students. For the fall of 2017, we are projecting a 6.5% general increase in addition to the 2% annual SFU increase for domestic students plus a 16.5% general increase for international students in order cover the expenses associated with the program as identified during the special cohort offerings.

Domestic	Base	Tuition Increase	Change to amount	Total new tuition
1147	\$26,000			
1157	\$26,000	2.00%	\$520	\$26,520
1167	\$26,520	2.00%	\$530	\$27,050
1177	\$27,050	8.50%	\$2,299	\$29,350
1187	\$29,350	2.00%	\$587	\$29,937
1197	\$29,937	2.00%	\$599	\$30,535

International	Base	Tuition Increase	Change to amount	Total new tuition
1147	\$31,000			
1157	\$31,000	2.00%	\$620	\$31,620
1167	\$31,620	2.00%	\$632	\$32,252
1177	\$32,252	18.50%	\$5,967	\$38,219
1187	\$38,219	2.00%	\$764	\$38,983
1197	\$38,983	2.00%	\$780	\$39,763

Appendix 10 Support Letters from employers

In response to the Senate Graduate Studies Committee request from October 27th, we have contacted a few organizations that are employing our graduates.

- A letter from BuildDirect Technologies Inc, which is described on the Government of Canada's Innovation, Science and Economic Development website as "the world's leading online wholesaler of building materials", is attached.
- Also included is a letter from ScotiaBank's Vice President Decision Sciences in Toronto.
- Finally, we have provided a letter from Samsung R&D Canada.
- A video filmed at RBC in Toronto includes comments from Chris Laver, Chief Data Scientist, about our Big Data Masters program and students.
 - <https://www.sfu.ca/computing/current-students/graduate-students/academic-programs/bigdata/coop.html>
- A web page providing more detail on some of their graduates and their current positions is available at
 - <http://www.sfu.ca/computing/current-students/graduate-students/academic-programs/bigdata/students/classof2016.html>



November 2, 2016

To: Fred Popowich, Professor, SFU

Re: SFU Big Data PMP

I am writing this letter to share our experience with the SFU Big Data program to date. We have successfully hired 3 interns out of the Big Data program and are currently interviewing for our 4th intern. We have been impressed with the level of knowledge these candidates bring as well as prior work experience that is very relevant for us.

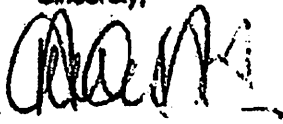
One of our interns from Cohort 1 is now working for us permanently as a Data Scientist and he has brought tremendous value and knowledge to the team already. In a short period of time he was able to demonstrate that he could own the delivery of a project by himself. This is in large part due to the exposure of a wide range of relevant technologies and techniques that he received while attending the SFU Big Data program. This exposure armed him with the right tools to tackle some of the most complex problems that the Analytics team is asked to solve at BuildDirect. This same statement can be applied to all the interns from this program. It is refreshing to be in a position to assign individual projects to interns with little supervision.

BuildDirect is currently building a data-driven eCommerce platform. One of the goals of the platform is to be largely self-correcting, i.e. the user experience will be modified based on the performance of previous user sessions. This is done to maximize all portions of the conversion funnel. The data scientists at BuildDirect have a large part in the development of this functionality and the need for this role will only increase as we tackle ever more complex scenarios.

This SFU Big Data program is very valuable for BuildDirect and our Analytics Department because it has proven to be a great source of well-trained data scientists in the Vancouver area. In fact, we have had great difficulty in finding suitably trained resources outside of SFU. I can not recommend more highly that the program should transition to a permanent program.

Please contact me should you have any questions.

Sincerely,

per 

Hilan Amin
VP Analytics
BuildDirect.com Technologies Inc.
hilanamin@bulldirect.com

Tel

Fax



November 20, 2016

Dr. Fred Popowich
Professor and Associate Director Research and Industry Relations, School of Computing Science,
Director, SFU Professional Master's Program in Big Data

Simon Fraser University
Applied Science Building 9971
8888 University Drive
Burnaby, British Columbia
Canada V5A 1S6

Professor Fred Popowich,

I am writing you to let you know about our continued interest in SFU's Professional MSc in Computer Science - Big Data program. As you know we have recently hired several of your graduating students onto full time Data Scientist roles in the Decision Sciences area at Scotiabank. Additionally, we have also hired for several students from the program for the upcoming co-op term. They will be joining the team of about sixty data scientists; which will likely grow to around one hundred over two groups.

During our recent recruiting trip to SFU, we found the graduating students to be well prepared for entry into Decision Sciences group in the bank. The Professional MSc in Computer Science - Big Data program has a strong curriculum and we like the direction that you are taking it in. I would also like to thank you for soliciting our input into future enhancements to the program.

As you know, people with the both the applied skills and strong academic foundation for data scientist work are in high demand. We are glad that you are prudently expanding the class to both fill some of the demand but also maintain a high caliber of graduates.

We are looking forward to onboarding our first group of SFU Big Data program graduates and co-ops. We hope to keep in contact with you and the department and look forward to the opportunity to come out and talk with the students again next year.

Thank you,

A handwritten signature in black ink, appearing to read 'Andrew Storey', written in a cursive style.

Andrew Storey
Vice President, Decision Sciences
Scotiabank
647 919 0369

SAMSUNG

PRIVATE & CONFIDENTIAL

Dr. Fred Popowich
Program Director, Professional Master's Program, Big Data
School of Computing Science
Simon Fraser University
8888 University, Drive
Burnaby, BC V5A 1S6

Re: Letter of Support, Master's of Big Data Program

Please accept this letter as support for converting the SFU Master's of Big Data program into a permanent offering for students.

In May 2016, Samsung R&D Canada hired our first Data Analytics Intern from the Master's program. We were very pleased with the knowledge and level of expertise she brought with her. The student showcased abilities in both software development and data science; this resulted in Samsung offering her an extension to her internship placement. With the conclusion of the internship placement, Samsung has offered the student a full time permanent position with the company.

Samsung R&D Canada has been impressed with the caliber of students we have seen come through the program. For the upcoming May 2017 term, the Samsung Analytics team has once again hired another Data Analytics Intern from this Master's program.

We would like to encourage Simon Fraser University to continue offering this program on a permanent basis. The students' experiences with Data Mining and Machine Learning algorithms, Hadoop and Spark for processing Big Data, Python and Java for development puts them in line with the latest trends in the Big Data technologies. The skillsets gained from the program will place the graduates in high-demand within the technology sector. We hope to hire more talented students who come out of this program.

Regards,



Joti Dhaliwal

On Behalf of Cristina Doroftei
Technical Software Manager, B2B Analytics
Samsung R&D Canada
200 - 8500 Glenlyon Parkway
Burnaby, BC, V5J 0C6