Request for Technology Fee Funds: FY20

NOTE: A separate request should be made for each initiative.

I. Department Number/Department Name: 360 College of Computing

Title of Request *(please be brief):* Reconfigurable Cluster Initiative

Amount of Request *(formula from detailed budget below):* $74,905

Type of Proposal: Atlanta or Dist Lrng/Non-Atl Atlanta

Was this project request funded in FY19? No (Yes or No)

Are there installation/renovation costs associated with this request? No (Yes or No)

If "Yes" then indicate the source of approved funding:
(Note: Tech Fees are not allowed for installation/renovation)

Executive Summary of Request (100 words or less):

This request will develop a common reconfigurable computing infrastructure that can be used for multiple undergraduate classes. We propose to acquire several FPGA boards that will be used for both undergraduate labs as well as VIP-focused research courses.

Specific class and/or lab initiative(s) if applicable: Center for Research into Novel Computing Hierarchies

Contact person for this request (incl. phone #): Jeffrey Young (404-385-1513)

Responsible faculty for this request (incl. phone #) Hyesoon Kim, Jason Riedy, Lee Lerner

Indicate priority per department if applicable: Number ____ of ____

Indicate priority per college or unit: Number 1 of 9

II. Impact on Students - Provide course title, course number, and anticipated enrollments:

Titles/Numbers of Course(s) CS 3220, 4803/7643, 7290; ECE 2031, 2601, 3601, 4601

Anticipated Enrollments

Graduate:	*165*	*(per* Year) sem or yr
Undergraduate:	1,104	(per Year) sem or yr
Total:	1,269	

The estimated percent use of the resources in the item by:

Students	90%
Faculty	10%
Other	0%
Total:	100%

Brief explanation of how estimate was achieved.

Estimates are based on making the FPGA resources available primarily for students during the school year and opening these resources up for general usage during summer. The cluster-based scheduler will prioritize student usage.

NOTE: Other impacts on students should be described in narrative to include benefits to the students affected.

III. Detailed Budget - Requested Items by Category *List separately any equipment, software, and other allowable expenses (see Tech Fee Guidelines). There is a formula in the "total column" that multiplies the number of items times the unit price. You may enter a figure into the total column if the unit pricing is not applicable. If you need additional rows, contact the Budget Office to receive a modified form. Software or data license proposals should indicate how many years the item has been funded through student tech fees in narrative.*

Supporting documentation is required *- Include price justification in some form, such as quotations, published price lists, etc. as a separate PDF attachment. All supporting information should be in a single PDF.*

	Proposed Number of Items	Estimated Price per Unit	Total ($)
Xilinx VCU1525 PCIe FPGA Accelerator	2	$4,000	$8,000
Xilinx ZCU106	4	$1,995	$7,980
Xilinx Alveo U250	1	$12,995	$12,995
Xilinx XUPVVH-0020	2	$12,995	$25,990
4U Host Servers Capable of Hosting 4-5 FPGAs	2	$9,970	$19,940
Total *(linked to the total amount of request line above)*			$74,905

Please return form via e-mail in Excel format to: techfees@business.gatech.edu. Supporting information only in a PDF file.

IV. Narrative - Provide narrative justification for your intended use of the technology fee funds. Include narrative on how the education or research of the students will be enhanced. To include curricular, co-curricular, and extracurricular benefits expected to accrue to students through provision of this resource, including students outside the unit. Briefly state how information regarding similar technology use elsewhere on campus to benefit from lessons learned, to standardize, or differentiate, and to avoid duplication. Also include how the request aligns with the Strategic Plan of Georgia Tech.

We propose to acquire nine Field Programmable Reconfigurable Array (FPGA) devices and two host servers to seed a new reconfigurable cluster initiative. We anticipate supporting undergraduate and graduate students using this technology for coursework, but will also make the resource available to all interested students via our existing TSO-supported testbed, the "Rogues Gallery". This effort will provide easier access to novel hardware for students as well as reduce the need for individual professors to acquire and maintain this type of hardware for their classes. This cluster will initially be used by Atlanta-based students; however, if this proposal is successful, we will likely seek additional funding from OMSCS in the future and make these resources accessible via remote access and cluster-based scheduling.

FPGAs are a useful piece of technology because they excel at performing irregular and high-speed computations in programmable hardware. That is, the hardware can be customized to solve problems that are normally might be run on low-power embedded systems (e.g, the Raspberry Pi), high-performance GPU cards, or even more specialized accelerators like Digital Signal Processing chips or machine learning and deep learning accelerators like Google's Tensor Processing Unit (TPU). The FPGA device is flexible enough to allow us to implement custom accelerators that can replace all these specialized pieces of hardware.

To promote a diverse environment for students, we propose to acquire four different types of FPGA accelerator devices and two host servers to host remote access for the requested devices. We propose to acquire 2 Xilinx VCU1525 FPGA accelerators and 4 ZCU106 Xilinx+ARM development kits for general-purpose use in classes like CS 3220 - Processor Design with FPGAs, CS 7290 – Advanced Microarchitecture, and ECE 2601 – Digital Design Laboratory, ECE 3056 - Architecture, Concurrency, and Energy in Computation. We will also support students in ECE (2,3,4)601/(2,3,4)602/4603 - Vertical Integrated Project (VIP) classes run by Drs. Young and Riedy (Rogues Gallery) and by Dr. Lerner and Dr. Wills (Configurable Computing and Embedded Systems.

We also propose to acquire one Xilinx Alveo U250 board that will be provided for special projects in classes like CS 4803 and CS 7643 (Deep Learning) via our scheduling interface and hosted server. We will host the Xilinx machine learning framework, which will allow users to run common deep learning frameworks like Caffe2, Tensorflow, and MxNet. Finally, we have requested two Bittware XUPVVH boards which include high-bandwidth memory, a new type of DRAM that allows for greatly speeding up applications with high-performance data requirements. This particular card will be made available to our traditional architecture classes in addition to our high-performance parallel computing courses, CSE 6230, where it will be made available for students to implement and evaluate high-performance algorithms written in languages like OpenCL and OpenMP. Most interestingly, we have also been in discussion with colleagues at University of Tennessee to use this type of board to host their neuromorphic or brain-inspired DANNA processor and related high-level framework that can be used by students in computing courses and can be extended to support new courses related to GTNeuro.

We will host the proposed infrastructure as part of the TSO-supported Rogues Gallery testbed. This testbed currently has 30 users and supports student researchers in addition to faculty and graduate students for a variety of hardware, including one Intel FPGA. The Rogues Gallery infrastructure will allow for scheduling of FPGA resources and hosting of FPGA tools in a common network-accessible location. These two capabilities will remove one of the key barriers for student usage of these devices – typically student projects are limited by the need to acquire new accounts on a machine that may sit in a professor's lab. This typically requires both physical access to a lab on GT's campus as well as a custom login to be created for each new student. The Rogues Gallery uses LDAP-based user authentication to give students access and techniques like USB-IP (a technique for accessing USB-based development boards over a network), remote power toggling, and a common set of programming and debugging tools. A robust set of wiki pages are currently available for this infrastructure (https://github.gatech.edu/crnch-rg/rogues-docs/wiki), and we provide a community mailing list and discussion group for students to request assistance.

Due to the nature of this request, and the modularity of this resource, we are able to receive partial funding for specific components from this request and will request any remaining funding next year. The Xilinx VCU, ZCU, and the 2 Host Servers are the most pressing portion of this request.

Please return form via e-mail in Excel format to: techfees@business.gatech.edu. Supporting information only in a PDF file.