1 2 3 4 5 6 7 8	Venkat Konda, Ph.D. 6278 Grand Oak Way San Jose, CA 95135 Telephone: (408) 472-3273 Email: vkonda@gmail.com Plaintiff <i>pro se</i> IN THE SUPERIOR COURT OI COUNTY OF S	Electronically Filed by Superior Court of CA, County of Santa Clara, on 11/7/2019 4:32 PM Reviewed By: Yuet Lai Case #19CV345846 Envelope: 3625402 F THE STATE OF CALIFORNIA SANTA CLARA
9 10	VENKAT KONDA, Ph.D., an individual, )	Case No. 19CV345846
10	Plaintiff,	Unlimited Civil Case
11	v. )	FIRST AMENDED COMPLAINT FOR:
13	<b>DEJAN MARKOVIC, Ph.D., an individual;</b>	<ol> <li>Unfair Business Practices</li> <li>Fraud - Intentional</li> </ol>
14 15	FLEX LOGIX TECHNOLOGIES, INC., a	Misrepresentation 3. Fraud – Concealment 4. Missemucariation of Tuo do Saconsta
16	Delaware Corporation;	<ol> <li>5. Ongoing Conspiracy</li> </ol>
17	THE REGENTS OF THE UNIVERSITY       )         OF CALIFORNIA;       )	
18	)	DEMAND FOR JURY TRIAL
19	MUNGER, TOLLES & OLSON LLP, a ) California limited liability partnership; and) DOES 1-20 inclusive	
20	DOES 1-20, inclusive,	
21 22		
$\begin{array}{c} 22\\ 23 \end{array}$		
24	)	
25		
26		
27		
28		
		1
	-	1-

Plaintiff Venkat Konda, Ph.D. (hereinafter referred to as "Dr. Konda" or "Plaintiff") alleges as follows:

1. This case involves a surreptitious scheme by a professor at the University of California Los Angeles (hereinafter referred to as "UCLA"), Dejan Markovic, Ph.D. (hereinafter referred to as "Defendant Markovic"), who conspired and coordinated with one of the graduate students he advised, Cheng C. Wang, Ph.D. (hereinafter referred to as "Defendant Wang"), (hereinafter collectively referred to as "Defendants" or Defendants Markovic and Wang) to misappropriate the intellectual property of a Silicon Valley company, Konda Technologies, Inc. (hereinafter referred to as "Konda Tech") through deception and manipulation under the cloak of legitimacy afforded by their association with UCLA which has benefitted by illicitly commercializing Konda Tech's intellectual property, first through a now-dissolved California Corporation Hierlogix, Inc. formed by Defendants Markovic and Wang with funding by UCLA's Institute for Technology Advancement (hereinafter referred to as "UCLA/ITA") and its successor Flex Logix Technologies, Inc. (hereinafter referred to as "Flex Logix").

2. Konda Tech's intellectual property (hereinafter referred to as "Konda Tech IP") relates to Field Programmable Gate Arrays (hereinafter referred to as "FPGAs"). FPGAs are semiconductor devices that are based around a matrix of configurable logic blocks (CLBs) comprising one or more Lookup Tables connected via programmable interconnects. After being fabricated, FPGAs can be reprogrammed to desired application or functionality requirements. They are used in many different applications from simple devices such as calculators to sophisticated artificial intelligence (AI) systems that require high-speed logic operations. FPGAs can perform these operations faster than a software application running on a computer's central processing unit (CPU).

3. Defendant Markovic was introduced to Dr. Konda by Flavio Bonomi, Ph.D. (hereinafter referred to as "Dr. Bonomi"), who was a Cisco Fellow, Vice President and the Head of the Advanced Architecture and Research Organization at Cisco Systems, Inc. in San Jose, California (hereinafter referred to as "Cisco"). After funding was orally offered by the Cisco Angel

1

2

3

-2-

Network, but later rescinded, Defendant Markovic reached out to Plaintiff to troll for information about Konda Tech IP beginning in or around March, 2009 and continuing through March, 2014.
4. On the pretense of obtaining funding for Konda Tech through UCLA/ITA, Defendant Markovic arranged a presentation by Dr. Konda on October 12, 2009, obtaining proprietary and confidential materials from Dr. Konda five days prior to the presentation. However, funding was not forthcoming because, as Defendant Markovic knew beforehand, the prerequisite nexus of a relationship between UCLA and Konda Tech did not exist. Nevertheless, within less than two months after the presentation, in December, 2009, unbeknownst to Dr. Konda, Defendant Wang with his fellow graduate students covertly implemented and fabricated FPGA devices comprising CLBs and interconnect based on Konda Tech IP through the graduate program at UCLA under the guidance of Defendant Markovic without Plaintiff's authorization.

5. Without disclosing that Defendants Markovic and Wang had implemented FPGA devices based on Konda Tech IP, Defendant Markovic eight to nine months later contacted Plaintiff to solicit submitting a confidential joint proposal to DARPA. When Dr. Konda learned about the covert work that had been carried out without his authorization at UCLA described in the draft DARPA proposal prepared by Defendant Markovic, Plaintiff told Defendant Markovic to cease any further work at UCLA. However, again on the pretense of obtaining funding for Konda Tech, but in actuality to obtain funding for Defendants' development of software tools to program the FPGA devices they had covertly implemented, Plaintiff agreed to the submission of the DARPA proposal with Konda Tech as the Principal Investigator. Defendant Markovic promised Plaintiff that if the DARPA proposal was granted, he would obtain a license for Konda Tech IP; otherwise he would cease any further implementation of Konda Tech IP at UCLA.

6. While the first DARPA proposal was still pending, Defendant Markovic solicited
Plaintiff to join in submitting a second confidential joint DARPA proposal approximately five
weeks later, but with UCLA as the Principal Investigator. Defendant Markovic again promised
Plaintiff that if the DARPA proposal was granted, he would agree to a license for Konda Tech
IP; otherwise he would cease any further implementation of Konda Tech IP at UCLA.

7. Both of the DARPA proposals were rejected in late 2010. At that point, Plaintiff believed that all of the FPGA device work incorporating Konda Tech IP at UCLA had ceased.
8. At or around that time, Dr. Konda contacted Defendant Markovic to inform him that Konda Tech had licensed a commercial FPGA supplier, QuickLogic Corporation (hereinafter referred to as "QuickLogic"), with whom Plaintiff had worked between late September, 2010 and mid-January, 2011 to prove the value of Konda Tech IP. Defendant Markovic trolled for confidential information regarding Plaintiff's work with QuickLogic. Dr. Konda informed Defendant Markovic in confidence regarding the licensee information and other confidential business information. Defendant Markovic was keenly interested in this information and further inquired what other FPGA suppliers Konda Tech had contacted.

9. Unbeknownst to Plaintiff, Defendants Markovic and Wang had formed Hierlogix Inc. ("Hierlogix") on January 4, 2011 with its principal place of business at Defendant Wang's apartment.

10. Unbeknownst to Plaintiff, Defendants Markovic and Wang also submitted a paper based on Konda Tech's IP in January 2011 to the VLSI Symposium without any attribution to Konda Tech IP, particularly Konda Tech 2D BFT layouts which is the cornerstone for achieving area, power, and performance improvements in FPGAs, and presented the paper in Japan in June 2011.

11. Also, unbeknownst to Plaintiff, Defendant Wang was completing his Ph.D. program under the guidance of Defendant Markovic and submitted his dissertation based on the implementation of Konda Tech IP, and he was awarded his Ph.D. in June 2013 and recognized for having submitted a distinguished Ph.D. dissertation.

12. Unbeknownst to Plaintiff, Defendants raised funding from UCLA/ITA for Hierlogix, Hierlogix believed to have been substantially based on the Konda Tech IP presentation given by Plaintiff to UCLA/ITA on October 12, 2009.

13. In the fall of 2013, Defendant Markovic invited Dr. Konda by email to meet him at
Stanford University while he was a Visiting Professor. When they met, Dr. Konda inquired
whether Defendant Markovic and his students had discontinued implementing Konda Tech IP as

-4-

part of the academic work. Defendant Markovic falsely replied "yes." During the conversation,
Defendant Markovic also asked Dr. Konda to inform him of the names of customers he was
currently working with to license Konda Tech IP.

14. In January, 2014, while Defendant Markovic was a visiting professor at Stanford University, Dr. Konda and Defendants Markovic and Wang met at Dr. Bonomi's residence. Dr. Bonomi had recently founded a startup company, Nebliolo Technologies, Inc. (hereinafter referred to as ("Nebliolo") and was interested in obtaining a supplier of FPGAs incorporating Konda Tech IP. He invited Defendants Markovic and Wang whom he understood had founded a semiconductor design company that he thought might implement FPGAs based on Konda Tech IP in an embedded FPGA block to supply Nebliolo. In that meeting, Defendant Markovic deceived Plaintiff by mentioning that he was in the process of raising funding for a startup company. When Plaintiff queried if their startup was in the area of wireless and digital signal processors (DSPs), Defendant Markovic said "yes," which was an intentional misrepresentation. Dr. Markovic concealed the fact that the technological focus of the startup was embedded FPGA ("eFPGA") blocks by covertly implementing Konda Tech IP without having a license from Konda Tech.

15. Unbeknownst to Plaintiff, Defendants Markovic and Wang were in involved in the founding Flex Logix on February 26, 2014 to continue the commercialization of eFPGA blocks implementing Konda Tech IP without having a license from Konda Tech.

16. In or about December, 2018 Plaintiff arranged to meet with Professor Vaughn Betz,
Ph.D. ("Dr. Betz") in the Department of Electrical and Computer Engineering at the University
of Toronto, Toronto, Canada, to discuss certain results Dr. Konda achieved with the Versatile
Place and Route ("VPR") tool suite developed by Dr. Betz using VPR to implement Konda Tech
IP. Plaintiff met with Dr. Betz in Toronto on or about December 18, 2015. During their
meeting, Dr. Betz asked Plaintiff if he knew of Flex Logix. Plaintiff responded that he was not
aware of Flex Logix. Nor was Plaintiff aware of Defendants' paper publications, Defendant
Wang's Ph.D. dissertation, Hierlogix, or Flex Logix at the time of his meeting with Dr. Betz on
December 18, 2015.

1

2

3

4

5

6

-5-

17. After returning to California after his meeting on December 18, 2015 with Dr. Betz, Dr. Konda sought facts regarding the activities of Defendants Markovic and Wang, Hierlogix, and Flex Logix. Plaintiff then prepared an email which he sent to Flex Logix and UCLA and others on March 27, 2016 requesting additional information from, and action by, those entities regarding wrongdoing that he first suspected had occurred on the part of Defendants Markovic and Wang when he prepared his email on the weekend of March 26-27, 2016, when he formed a belief that Flex Logix appeared to be implementing eFPGAs based on Konda Tech IP. Due to the intentional misrepresentations and concealment of Defendants Markovic and Wang, Dr. Konda was unsuspecting until that time of the illicit activities of Defendants Markovic and Wang until he was able to piece together the facts in his March 27, 2016 email. Until then, Dr. Konda was in disbelief that Defendant Markovic would have betrayed the confidences and trust of the relationship he believed he had with Dr. Markovic, who cloaked himself with and exploited the pretextual credibility of UCLA, heretofore promoted as a respected educational institution, but now exposed as a commonplace cutthroat competitor whose employees (*i.e.*, Defendant Markovic) deprive unsuspecting inventors of their innovations. In view of the intentional misrepresentations and concealment by Defendant Markovic and his co-conspirator Wang, the facts regarding their wrongdoing were concealed and thus not previously discoverable or known by Dr. Konda. At that time Dr. Konda realized for the first time that he has been harmed by the unauthorized commercialization of Konda Tech IP by the Defendants.

### **PARTIES**

18. Plaintiff Venkat Konda, Ph.D. is and at all times herein mentioned was a resident ofSanta Clara County, California. Konda Tech, a California Corporation, has assigned to Dr.Konda the right to bring this action in his individual capacity, as well as all right, title, andinterest to recover damages and injunctive relief.

19. Plaintiff is informed and believes, and thereupon alleges, that Defendant Markovic is an individual who is a resident of California and conducts business in Santa Clara County, California.

-6-

20. Plaintiff is informed and believes, and thereupon alleges, that Defendant Wang is an individual who is a resident of California and conducts business in Santa Clara County, California.

21. Plaintiff is informed and believes that Defendant The Regents of the University of California have their principal office in California and conduct business in Santa Clara County, California.

22. Plaintiff is informed and believes that Flex Logix has its principal place of business and conducts business in Santa County, California.

23. Plaintiff is informed that Munger, Tolles & Olson LLP has its principal place of business in California and conducts business in Santa Clara County, California.

24. Plaintiff is ignorant of the true names and capacities of Defendants sued herein as DOES 1 through 20, inclusive, and therefore sues these Defendants by such fictitious names. Plaintiff prays leave to amend this Complaint to allege their true names and capacities when the same have been ascertained.

25. Plaintiff is informed and believes, and thereupon alleges, that each of the Defendants sued herein is responsible in some manner for the occurrences herein alleged, and that Plaintiff's damages were proximately caused by such Defendants.

26. Plaintiff is informed and believes, and thereupon alleges, that at all times herein mentioned each of the Defendants, was and were, at all times, acting as principals or agents, employees, or representatives within the purpose and scope of such agency, employment, or representation as being responsible in some manner for the occurrences herein alleged.

### JURISDICTION AND VENUE

27. This Court has jurisdiction over this First Amended Complaint pursuant to California Code of Civil Procedure Section 395(a) as the transactions, occurrences, and omissions to act giving rise to the liability on the part of the Defendants occurred in Santa Clara County, California and/or they have directed their unlawful acts complained of herein in Santa Clara County, California. 28. This Court has personal jurisdiction over the Defendants for the additional reason that they have engaged in systematic and continuous contacts with Santa Clara County, California, *inter alia*, regularly conducting and soliciting business in Santa Clara County, and deriving substantial benefit from products and/or services provided to persons in Santa Clara County, California.

#### FACTUAL BACKGROUND

29. Dr. Konda founded Konda Tech, a California corporation, in 2007. Dr. Konda is a pioneer in FPGA routing fabric and interconnection networks technology. Konda Tech's business is based on Dr. Konda's work, and provides chip and system level interconnect technology solutions. Konda Tech has licensed FPGA interconnect architecture patent rights to two FPGA chip vendors, the first of which has made and sold three generations of chips. Dr. Konda has a Ph.D. in Computer Science and Engineering from the University of Louisville, Kentucky and has been granted twelve patents in the space.

30. In or around January 2009, Dr. Konda was introduced to Defendant Markovic by Dr. Bonomi. Defendant Markovic was and is a UCLA professor focused on circuits and embedded systems (which overlaps and compliments Konda Tech IP), and involved with UCLA/ITA. Defendant Markovic was not focused on FPGA work until he met Dr. Konda. Konda Tech was one of six startups that received an oral offer for funding from Cisco, that was later rescinded. Defendant Markovic became aware that Cisco's offer to Konda Tech had been rescinded, and that Konda Tech was still looking for funding. Defendant Markovic seized the opportunity to contact Plaintiff, claiming that Konda Tech could receive funding through UCLA/ITA. Defendant Markovic suggested that Dr. Konda present before UCLA/ITA. Dr. Konda provided Konda Tech's Business Presentation to Defendant Markovic on October 7, 2009 in confidence. However, after Dr. Konda arrived in Los Angeles on October 12, 2009 to present the Konda Tech business plan to UCLA/ITA, Defendant Markovic for the first time said to Dr. Konda that Dr. Konda should not expect UCLA/ITA to fund Konda Tech, because UCLA/ITA does not fund technologies built outside UCLA. Since the Konda Tech Business Presentation was also sent to UCLA/ITA on October 7, 2009, Dr. Konda made a presentation on October 12, 2009 to UCLA/ITA. Defendant Markovic started by presenting the Konda Tech Business Presentation to Dr. Les Lackman ("Dr. Lackman"), Deputy Director, Institute for Technology Advancement, UCLA and the other UCLA/ITA Directors in attendance, including Winn Hong. Dr. Lackman stopped Defendant Markovic after one slide and questioned Defendant Markovic "whose business plan is it?" or words to that effect. The Konda Tech Business Presentation was clearly marked "Konda Tech confidential and proprietary" on all of the slides. Defendant Markovic replied "It is Dr. Konda's." Dr. Lackman then said "Let Dr. Konda present it."

31. Dr. Konda's presentation on October 12, 2009 to UCLA/ITA was fruitless as had been made known by Defendant Markovic just prior to Dr. Konda presenting to UCLA/ITA.
UCLA/ITA rejected funding Konda Tech clearly stating that the complete Konda Tech Business Presentation was built outside UCLA and had nothing to do with UCLA or Defendant Markovic.
32. After presenting to UCLA/ITA, Defendant Markovic, enamored with Konda Tech IP, also asked Dr. Konda to give a seminar on Konda Tech IP to Defendant Markovic's students on October 12, 2009. Dr. Konda obliged Defendant Markovic by presenting an overview to Defendant Markovic's students only with respect to Konda Tech's published patent applications. Among those in attendance at the October 12, 2009 seminar was Defendant Wang, a graduate student and believed to be a Ph.D. candidate at the time. Defendant Wang subsequently grew similarly interested in Konda Tech IP.

33. When Dr. Konda presented to Defendant Markovic's students on October 12, 2009, Dr.
Konda clearly told them that what Dr. Konda was presenting to them was patent pending
technology by a commercial company, Konda Tech, and that the presented material was in the
published Konda Tech patent applications by then. So, no confidential material was presented to
the students.

34. Dr. Konda now believes that Defendant Markovic invited Dr. Konda to UCLA to get access to the Konda Tech Business Plan presentation and to give a seminar to Defendant Markovic's students including Defendant Wang. In this process Defendant Markovic, beginning at that time and continuing for the ensuing four years, continuously trolled Dr. Konda to learn

-9-

about all details of Konda Tech's technology including not only the disclosures in Konda Tech's patents, but also proprietary implementation details, technical know-how, and business knowhow, and the then customers and potential customers and Konda Tech's interaction with them which Plaintiff disclosed to Dr. Markovic in confidence in the belief that Defendant Markovic would maintain the information in confidence. As a result, Defendant Markovic learned about FPGA business models and the know-how of the FPGA industry with respect to interconnect technology and its evolution.

35. In June and July 2010, Defendant Markovic called Dr. Konda, and told him that he wanted to use Konda Tech IP in submitting two different applications for DARPA funding. Dr. Konda advised that he did not then have the time to work with Defendant Markovic. However, both times, Defendant Markovic assured Dr. Konda that he would not have to spend any time on the DARPA applications, and that he would incorporate the Konda Tech IP into the applications from the then published Konda Tech WIPO patent applications, as well as proprietary implementation details and technical know-how which Plaintiff had disclosed to Dr. Markovic in confidence, with the understanding that the DARPA applications were confidential. (But Defendant deceived concealing that Konda Tech proprietary information was revealed to Defendant Wang). Defendant Markovic assured Dr. Konda that he would secure a license from Konda Tech should a DARPA grant be approved for a DARPA project.

36. Attached hereto as Exhibits 1 and 2 are the June 23, 2010 and August 6, 2010 DARPA funding proposals (the "Two DARPA Proposals") that followed those conversations between Plaintiff and Defendant Markovic.

37. Both of the Two DARPA Proposals make clear that Konda Tech IP was at the heart of what Defendants Markovic and Wang were hoping to accomplish:

Konda Technologies inventions with regular VLSI layouts for Benes/BFT based hierarchical networks are seminal and subsumes all the other known network topologies such as Clos networks, hypercube networks, cube-connected cycles and pyramid networks, which makes these networks implementable in a FPGA devices with regular structures both interconnect distribution-wise and layout-wise which is the key to exploit improved area, power, and performance of FPGA devices. The regularity of Konda hierarchical layout is also the key for its commercializability in System-on-Chip interconnect devices, FPIC devices as well. Indeed, the Two DARPA Proposals state that they "will make use of hierarchically routed and proprietary Konda interconnect architecture." The first DARPA Proposal further estimated that Dr. Konda and Konda Tech would complete 620 task hours of the estimated 1020 task hours for key personnel.

38. Those Two DARPA Proposals, replete with references to Konda Tech IP, were rejected. However, Defendants Markovic and Wang were not dissuaded from continuing to work on implementation of Konda Tech IP without authorization from Konda Tech. In 2010, Defendant Markovic told Dr. Konda over the phone that his students, including Defendant Wang, were implementing Konda Tech IP as an "academic project," specifically the 2D layout, on an FPGA chip. When Defendant Markovic told Dr. Konda that his students had begun implementing Konda Tech's technology, Dr. Konda told him to stop. Defendant Markovic's answer was, as a university professor, he could implement any publicly available technology including any technology disclosed in patents or patent applications. Dr. Konda told Defendant Markovic that without a license from Konda Tech, Dr. Konda did not agree that he or UCLA had a right to implement Konda Tech's technology and to stop immediately.

39. Unbeknownst to Plaintiff, in defiance of Plaintiff's demand to Defendant Markovic to stop implementing Konda Tech IP, Defendants Markovic and Wang founded a Hierlogix on January 4, 2011 to commercialize Konda Tech IP without authorization. Notably, Hierlogix was incorporated approximately three months after the Two DARPA Proposals were rejected. Unbeknownst to Plaintiff, Hierlogix was funded by UCLA/ITA. Hierlogix even today on the UCLA/ITA website at https://www.ita.ucla.edu/companies/ describes "Hierlogix provides Energy-Efficient Hierarchical FPGA and Programming Tools. By developing a revolutionary new interconnect architecture, Hierlogix can provide hardware and software tools that are capable of greatly reducing FPGA power and size requirements, while producing higher speeds and performance". The "revolutionary new interconnect architecture" is in fact Konda Tech IP.

40. In June 2011, unbeknownst to Dr. Konda and without his authorization, Defendants Markovic and Wang presented a paper at the 2011 VLSI Circuits Symposium titled "A 1.1 GOPS/mQ FPGA Chip with Hierarchical Interconnect Fabric" (hereinafter referred to as the

-11-

"2011 VLSI Paper"), based on Konda Tech IP and disclosing trade secrets of Konda Tech that Dr. Konda had previously discussed with Defendant Markovic in confidence. Attached hereto as Exhibit 4 is the 2011 VLSI Circuits Symposium Paper.

41. Dr. Konda now believes that Defendants Markovic and Wang conspired so that Dr. Markovic provided Defendant Wang with access to the confidential Konda Tech Business Presentation to UCLA/ITA on October 12, 2009 which was clearly marked "Konda Tech confidential and proprietary."

42. Dr. Konda now believes that Defendant Wang illegally without any authorization from Plaintiff began implementing Konda Tech IP at least from the second quarter of 2009, i.e. before he was invited to UCLA/ITA on October 12, 2009.

43. Defendant Markovic did not follow the basic principles of respecting others' intellectual property provided to him in confidence and used his UCLA professorship as a shield for his illegal implementation of Konda Tech IP. Defendants Markovic and Wang brazenly incorporated Hierlogix on January 4, 2011 to implement Konda Tech IP without the authorization of Plaintiff. Defendants Markovic and Wang blatantly plagiarized Konda Tech IP and shamelessly published the 2011 VLSI Paper in which they intentionally misrepresented that Konda Tech's alternate vertical and horizontal layout of Benes/BFT layouts was their innovation in furtherance of their scheme of violating Konda Tech IP and unfairly competing against Konda Tech.

44. Subsequently, Defendant Markovic invited Dr. Konda by email in the fall of 2013 to
meet him at Stanford University while he was a Visiting Professor. When they met, Dr. Konda
inquired whether Defendant Markovic and his students had stopped implementing Konda Tech
IP as part of his academic work. Defendant Markovic replied "yes," Intentionally
misrepresenting that he and his students, including Defendant Wang, were no longer working on
implementing Konda Tech IP. During the conversation Defendant Markovic also asked Dr.
Konda to share the names of customers he was working with to license Konda Tech IP, and Dr.
Konda did so in confidence, because he was not aware that Dr. Markovic had betrayed him.

45. Between 2011 and 2014, Defendant Markovic and Dr. Konda had occasional phone calls, during which they spoke about the progress of their respective work, but Defendant Markovic never disclosed that Konda Tech IP was the subject of Defendant Wang's 2013 Ph.D. dissertation titled, "Building Efficient, Reconfigurable Hardware using Hierarchical Interconnects" (hereinafter referred to as "Wang's 2013 Ph.D. Dissertation"). Defendant Markovic never disclosed that he and Defendant Wang founded Hierlogix on January 4, 2011.

46. Dr. Konda was unaware of Wang's 2013 Ph.D. Dissertation until after December 18,
2015. Attached hereto as Exhibit 5 is a copy of Chapters II and III and portions of Chapters V and VI of the Wang's 2013 Ph.D. Dissertation. The disclosure in Chapters II and III and portions of Chapters V and VI of Wang's 2013 Ph.D. dissertation copies Konda Tech IP, especially the figures and layouts, as shown by the highlighted portions of Exhibit 5.

47. Dr. Konda met with Defendants Markovic and Wang at the home of Dr. Bonomi on January 28, 2014. Dr. Bonomi, who was no longer at Cisco, had invited them to his home because he wanted to share that he was in the process of forming his own startup, and wanted a supplier that would provide his company with FPGAs under license from Konda Tech and was looking for implementation help from Defendants Markovic and Wang. Over the course of their discussions, Defendants Markovic and Wang stated that they were looking for funding for a separate startup, but when queried if their startup was in the area of wireless and DSP, Defendant Markovic replied "yes" which was an intentional misrepresentation and a concealment of the fact that Defendants Markovic and Wang had already started up Hierlogix three years earlier with funding by UCLA/ITA. During the meeting, Dr. Konda gave an update of Konda Tech activities and details of Konda Tech FPGA Interconnect IP. Defendant Markovic later stated that he was potentially interested in working with Dr. Bonomi and cryptically stated that he "may need to license Konda Tech IP". Dr. Konda replied that most of the Konda Tech patents were published or granted and suggested that Defendant Markovic check them on the Web to see if a license was needed and if so to contact Dr. Konda to obtain a license.

48. While Dr. Bonomi was trying to setup the meeting on January 28, 2019, Dr. Bonomi and Dr. Konda were not aware that Defendants Markovic and Wang were building the FPGA startup

-13-

Hierlogix. Dr. Konda was not aware of the 2011 VLSI Paper and that Wang's 2013 Ph.D.
Dissertation was on FPGA interconnects. Defendants Markovic and Wang concealed from Dr.
Bonomi and Dr. Konda that Wang's 2013 Ph.D. Dissertation was on FPGA Interconnects.
Defendants Markovic and Wang concealed from Dr. Bonomi and Dr. Konda that they were
building an FPGA company called Hierlogix. Otherwise, Dr. Bonomi would not have set up the
meeting at his home for Dr. Bonomi and Dr. Konda to meet with Defendants Markovic and
Wang.

49. A couple of weeks later, unbeknownst to Plaintiff at the time, Defendants Markovic and Wang published a paper titled "A Multi-Granularity FPGA with Hierarchical Interconnects for Efficient and Flexible Mobile Computing" at the 2014 International Solid State Circuits Conference (the "2014 ISSCC Paper"). The 2014 ISSCC Paper is attached hereto as Exhibit 6. The 2014 ISSCC Paper was based on Konda Tech IP. The 2014 ISSCC Paper describes and demonstrates technologies that were invented by Dr. Konda, and monetized by Konda Tech and shows the scheme of Defendants Markovic and Wang of violating Konda Tech IP and unfairly competing against Konda Tech.

50. While publishing at circuit conferences, Defendants Markovic and Wang never attended or published any paper at the International Symposium on FPGAs held annually in Monterey, California. This is the primary FPGA conference, and one they know Dr. Konda attends every year.

51. On February 18, 2014, Dr. Bonomi set up a meeting for Dr. Konda and Defendants Markovic and Wang to meet with Sundar Iyer, Ph.D. ("Dr. Iyer"), Co-founder and Chief Executive Officer of Memoir Technologies Inc. ("Memoir") an IP company in the area of computer memory technologies. The objective of this meeting was for Dr. Iyer to share his experiences of building IP companies with Dr. Konda, Dr. Markovic, and Dr. Wang. Notably, Defendants Markovic and Wang did not use either Hierlogix or Flex Logix's email IDs in communicating to arrange the meeting with Dr. Iyer.

52. Subsequently, Dr. Konda and Defendants Markovic and Wang met Dr. Iyer at Memoir's offices on March 5, 2014. At that meeting Defendants Markovic and Wang again concealed

-14-

from Dr. Konda the fact that both Hierlogix and their new startup Flex Logix were building FPGA products based on Konda Tech IP relating to a revolutionary interconnect architecture. 53. While Dr. Bonomi was trying to setup the meeting with Dr. Iyer on March 5, 2014, Dr. Bonomi and Dr. Konda were not aware that Dr. Markovic and Dr. Wang were building the FPGA startup Hierlogix. Dr. Konda was not aware of the 2011 VLSI Paper and 2014 ISSCC Paper, or that Wang's 2013 Ph.D. Dissertation was on FPGA interconnects based on Konda Tech IP. Defendants Markovic and Wang concealed from Dr. Bonomi and Dr. Konda that Wang's 2013 Ph.D. Dissertation was on FPGA interconnects. Defendants Markovic and Wang did not tell Dr. Bonomi and Dr. Konda about the 2011 VLSI Paper or the 2014 ISSCC Paper. Defendants Markovic and Wang did not tell Dr. Bonomi and Dr. Konda that they were building an FPGA company called Hierlogix. Defendants Markovic and Wang did not tell Dr. Bonomi and Dr. Konda that they founded Flex Logix on February 26, 2014. Otherwise, Dr. Bonomi would not have set up the meeting with Dr. Iver at Dr. Iver's office for Dr. Konda and Defendants Markovic and Wang to meet with Dr. Iyer on March 5, 2014. And, Dr. Konda would not have attended the meeting on March 5, 2014. 54. Unbeknownst to Plaintiff at the time of the meeting at Dr. Iyer's office, Defendants

Markovic and Wang had already founded Flex Logix on February 26, 2014 as the successor to Hierlogix to compete against Konda Tech illicitly based on implementing Konda Tech IP with a license or authorization from Konda Tech.

55. Until on or around March 2014, Defendant Markovic was pursuing Dr. Bonomi to serve as a reference for him, as he was applying to move as a Professor to Stanford University, Stanford, California ("Stanford") as well as California Institute of Technology, Pasadena, California ("Cal Tech").

56. Dr. Konda's expertise is not circuit design, so Dr. Konda never attended circuits conferences nor followed what is published in connection with such conferences.

57. Notably, the 2011 VLSI Paper and the 2014 ISSCC Paper list Defendants Markovic and Wang as affiliated with UCLA only. The 2011 VLSI Paper was published after Hierlogix was founded and the 2014 ISSCC Paper was published after Flex Logix was founded to conceal from

-15-

Dr. Konda that Defendants Markovic and Wang were associated in any way with those companies.

58. Dr. Konda was not aware of the publications by Defendants Markovic and Wang until after December 18, 2015 when Dr. Konda was told about Flex Logix by Dr. Betz during a visit to Dr. Betz's office. At that time, Dr. Konda first learned about Flex Logix and started investigating Flex Logix.

59. The conduct of Defendants Markovic and Wang makes clear that they employed subterfuge and deceit to gain access to Konda Tech IP, develop their fraudulent credibility in the technology through publications based on Konda Tech IP, and then used Konda Tech IP to launch their own startup company Hierlogix in competition with Konda Tech with funding by UCLA/ITA and its successor Flex Logix, covertly usurping Konda Tech IP as the cornerstone for establishing those companies.

60. Unbeknownst to Dr. Konda, Flex Logix was in stealth mode since its inception in February 2014 till on or about March 2015.

61. In the first part of December, 2015, Dr. Konda contacted Dr. Betz and they agreed to meet on December 18, 2015 in Toronto, Canada. They met on December 18, 2015 from approximately 1:00 – 2:30 PM at Dr. Betz's office at the University of Toronto to discuss certain results Dr. Konda achieved with the Versatile Place and Route tool suite ("VPR") by implementing Konda Tech IP using VPR. VPR was built by Dr. Betz and his students working for several years at the Department of Electrical and Computer Engineering at the University of Toronto, Toronto, Canada, and Dr. Betz was interested in discussing Dr. Konda's results.
62. During the meeting, Dr. Betz asked Dr. Konda whether Flex Logix was implementing Konda Tech IP. Dr. Konda was shocked when Dr. Betz asked him that question. Then, Dr. Konda replied that he hadn't heard about Flex Logix and inquired of Dr. Betz about Flex Logix. Dr. Betz told Dr. Konda that Flex Logix is an FPGA startup co-founded by Defendant Markovic and his student Defendant Wang.

63. From December 18, 2015, Dr. Konda investigated what Dr. Betz said to him by visiting the website for Flex Logix (www.flex-logix.com). Dr. Konda discovered that Flex Logix's

product purports to be embedded FPGA ("eFPGA") blocks and that Defendants Markovic and
Wang had intentionally misrepresented to Dr. Konda at the January 28, 2014 meeting with Dr.
Bonomi that they were involved with a startup company in the digital signal processor (DSP)
field for communications applications. As Dr. Konda continued his investigation he discovered
Wang's 2013 Ph.D. Dissertation and determined that it plagiarized the disclosure in Konda Tech
patents. Dr. Konda also found out for the first time about the 2014 ISSCC Paper and determined
that it described an implementation of the plagiarized Konda Tech patents.

64. Notably, Flex Logix touts the 2014 ISSCC paper on its website as describing Flex
Logix's "new, patented interconnect, XFLX<sup>TM</sup>." See, http://www.flex-logix.com/fpga-tutorial/.
65. On or about the fourth week of December 2015, Dr. Konda texted Dr. Bonomi to meet,
but Dr. Bonomi was in Italy at that time. On January 7, 2016, after Dr. Bonomi returned to
California, Dr. Konda met Dr. Bonomi at his office in Milpitas, California. Dr. Konda told him
about the meeting with Dr. Betz on December 18, 2015 and that he had subsequently searched on
the World Wide Web and only then found out that Defendants Markovic and Wang had founded
Flex Logix to manufacture eFPGAs based on Konda Tech IP, including Konda Tech trade
secrets that Dr. Konda had disclosed to Dr. Markovic in confidence. Dr. Konda asked if Dr.
Bonomi knew that Wang's 2013 Ph.D. Dissertation is on FPGA multi-stage interconnects and if
he knew about Hierlogix products or Flex Logix. Dr. Bonomi himself was shocked and said that

66. To the utmost shock of Dr. Konda, Defendants Markovic and Wang, beginning at time he first met Dr. Markovic and continuing for years, learned about all details of Konda Tech's technology including the proprietary implementation details, technical know-how, and business know-how and the then customers and potential customers of Konda Tech and Konda Tech's interaction with them. As a result, Defendants Markovic and Wang learned about FPGA business models and the know-how of the FPGA industry with respect to interconnect technology and its historical evolution and the revolutionary Konda Tech interconnect architecture. This provided Defendants Markovic and Wang a significant unfair head start in

-17-

launching Hierlogix and Flex Logix to compete with Konda Tech covertly using Konda Tech IP as the cornerstone for eFPGA products.

67. Defendant Markovic deceptively presented himself as an advisor to Dr. Konda. Dr. Konda believed that what he disclosed to Defendant Markovic was disclosed in confidence and believed that the intentions of Defendant Markovic were to help Konda Tech because he is a UCLA Professor and was not a competitor at the time. Dr. Konda always expected Defendant Markovic as a professor at a premier educational and research institution like UCLA, who teaches students and conducts research, to conduct himself in a fair and professional way. Dr. Konda was shocked to learn leading up to March 27, 2016 what Defendant Markovic did for years in such a deceitful manner.

68. Dr. Konda and Konda Tech made reasonable efforts to keep the trade secret information comprising the Konda Tech business plan, technical know-how, and business know-how secret. Any such information was and is disclosed to customers of Konda Tech under non-disclosure agreements. All written disclosures to Defendant Markovic were marked "proprietary and confidential." All such information which Dr. Konda verbally disclosed to Defendant Markovic likewise was considered to be in confidence and absolutely not intended to be used by Defendants Markovic and Wang to compete with Konda Tech.

69. After nearly exhausting his search for additional facts about Drs. Markovic and Wang and Flex Logix and consulting with an attorney, Dr. Konda formed his suspicion that Flex Logix was manufacturing eFPGAs based not only on the disclosures in Konda Tech patents, but also know-how trade secrets that Dr. Konda had communicated to Defendant Markovic and Defendant Wang in confidence, which enabled Hierlogix and its successor Flex Logix to start up its business. On March 27, 2016, Dr. Konda acted on his suspicion and prepared and sent an email with the subject "Flex-logix is Infringing Konda interconnect IP; Cheng Cheng Wang's UCLA PhD Dissertation is a blatant copy of Konda Technologies granted Patent(s) Flex-Logix Technologies" to Mr. Geoff Tate, Chief Executive Officer, Flex Logix ("Mr. Tate"), Mr. Peter Hebert, Co-founder/Managing Director of Lux Capital and board member of Flex Logix, and Mr. Shirish Sathaye, General Partner of Formation 8, Foundation Capital and board member of Flex

1

2

3

4

5

-18-

Logix, Dr. Gene David Block, Chancellor, UCLA, Dr. Jayathi Murthy, Dean of Henry Samuel School of Engineering and Applied Science, UCLA ("Dr. Murthy"), Dr. Les Lackman ("Dr. Lackman"), Deputy Director, Institute for Technology Advancement, UCLA, and Defendant Markovic with several of the facts that he had become aware of after he had visited Dr. Betz.

70. On March 28, 2016, Dr. Konda received a response from Mr. Tate saying that he would investigate the issue and get back to Dr. Konda in about a week. On March 28, 2016, Dr. Konda also received a response from Dr. Murthy saying that she would get back to him in a week.

71. Dr. Konda did not receive any response from Dr. Lackman whom Dr. Konda met during his confidential business presentation to UCLA/ITA on October 12, 2009. On April 7, 2016, Dr. Konda sent a follow-up email to Dr. Lackman. Not until that time did Dr. Konda discover on the UCLA/ITA website a company named Hierlogix and inquired of him whether that was an earlier name of Flex Logix. Dr. Konda pointed out that UCLA/ITA funding of Defendants Markovic and Wang for any work done at UCLA based on Konda Tech IP needed to be investigated by UCLA/ITA immediately. Dr. Konda also asked Dr. Lackman: "What are your policies and how soon do you take action against them?" Dr. Lackman has never responded to Dr. Konda's April 7, 2016 email leading to Dr. Konda's belief that UCLA/ITA was and is aware of the unauthorized use of Konda Tech IP.

72. Continuing with his investigation, in or about June 25, 2016, Dr. Konda discovered the 2011 VLSI Paper which plagiarized the disclosures in Konda Tech patents and discloses implementation details derived from the unauthorized use of Konda Tech IP including trade secrets at the time disclosed by Dr. Konda to Dr. Markovic. Dr. Konda was not aware of that paper until then.

73. On April 4, 2016, Dr. Konda received a response from Ann R. Karagozian, Ph.D. ("Dr. Karagozian"), Interim Vice Chancellor of Research at UCLA (and predecessor to Roger Wakimoto, Ph.D. ("Dr. Wakimoto"), Vice Chancellor of Research at UCLA), and Ann Pollack, Ph.D. ("Dr. Pollack"), Assistant Vice Chancellor of Research that they would investigate pursuant to UCLA Policies and Procedures. They categorized Dr. Konda's facts in Dr. Konda's March 27, 2016 email as research misconduct and patent infringement. Regarding the research

misconduct, Dr. Konda subsequently communicated with them and provided several documents. On May 7, 2018 Dr. Konda submitted a formal complaint of research misconduct. On July 19, 2018, Dr. Konda received a response from Dr. Pollack that the Two DARPA Proposals were not plagiarized, which is not what Dr. Konda had complained about. Instead, Dr. Konda claimed the Two DARPA Proposals were proof for the plagiarism by Defendants Markovic and Wang in their 2011 VLSI Paper, in Wang's 2013 Ph.D. Dissertation, and in their 2014 ISSCC Paper. Dr. 6 Konda requested a face to face meeting several times with Dr. Pollack but his requests were 8 declined. On April 30, 2019, Dr. Konda responded to Dr. Pollack that she still had not answered all the points he raised. Dr. Konda has received no further response from the Vice Chancellor of 10 Research at UCLA.

74. Regarding patent infringement, Dr. Konda submitted claim charts to Mr. Steven Drown ("Mr. Drown"), Senior Counsel, Educational Affairs, Office of the General Counsel, University of California and Mr. Robert Swerdlow ("Mr. Swerdlow"), Senior Counsel, UCLA on August 1, 2018. Dr. Konda has not received any response.

75. From March 30, 2016 until May 20, 2016, Dr. Konda provided additional information to Mr. Tate in at least three emails. To all of them his response was 1) We have reviewed your recent email, 2) We definitely do not agree with your analysis or your position(s), and 3) We will certainly consider any additional facts and listen to any new analysis you wish to provide us.

76. From on or about August, 2016 until on or about April, 2017, Konda Tech and Flex Logix had settlement negotiations represented by their attorneys. No resolutions were achieved. 77. Dr. Konda and Mr. Tate met on May 30, 2017 at a Starbucks in Mountain View, California. Dr. Konda explained at length regarding plagiarism by Defendants Markovic and Wang and patent infringement and requested Mr. Tate to act fairly and resolve these issues immediately. Mr. Tate then threatened Dr. Konda that "One of the senior board members of Flex Logix, alluding to Mr. Pierre Lamond, will ruin your career if you or Konda Tech files a lawsuit against Flex Logix."

78. Dr. Konda and Mr. Tate again met at the same place on June 1, 2017. During the discussions Mr. Tate again threatened Dr. Konda that "One of the senior board members of Flex

1

2

3

4

5

7

9

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

Logix, alluding Pierre Lamond, will ruin your career if you or Konda Tech files a lawsuit against Flex Logix."

79. From on or about September, 2018 until on or about December, 2018, Konda Tech and Flex Logix again had settlement negotiations represented by their attorneys. No resolutions were achieved.

80. Consistent with his threats to Dr. Konda on May 30, 2017 and June 1, 2017, Mr. Tate did the following. During the Design Automation Conference 2019 ("DAC 2019") held at the Las Vegas Convention Center, Las Vegas, Nevada from June 2 - 6, 2019, on June 3, 2019, Dr. Konda met with an executive of an FPGA manufacturer at the Center's food court regarding Konda FPGA interconnect technology from approximately 9 – 10 AM. Later on the same day while the executive was at his company's booth in the Las Vegas Convention Center, Mr. Tate, the Chief Executive Officer of Flex Logix which also had a booth at DAC 2019 approached the executive of his competitor and said to him "Can you provide me with the contact information of your lawyers, I think they shall talk together." At this, the executive replied to Mr. Tate "Geoff, we are grown up men, if you have something to tell me, tell me now." Mr. Tate said he saw Dr. Konda talking to him earlier in the day. Mr. Tate further threatened the executive that "I wanted to tell you that what we have is something totally different from Konda claims. I hope you are not helping him in any way."

#### FIRST CAUSE OF ACTION (Unfair Business Practices)

81. Plaintiff incorporates by reference every allegation contained in each and every one of the above paragraphs, as though set forth fully herein.

82. Defendants Markovic and Wang clearly knew that they did not have Dr. Konda's authorization to implement Konda Tech IP, or to publish technical papers by plagiarizing the Konda Tech Business Presentation and incorporating portions of Konda Tech IP, never citing Konda Tech's layouts in their publications or in Wang's 2013 Ph.D. Dissertation, in reckless disregard of Konda Tech's intellectual property rights.

83. Defendant Wang received a distinguished dissertation award (his dissertation advisor being Defendant Markovic), which dissertation included plagiarized and willfully misappropriated Konda Tech IP without attribution to Konda Tech. Additionally, Defendants Markovic and Wang won the best paper award at the 2014 ISSCC Conference for the 2014 ISSCC Paper, which plagiarized and willfully misappropriated intellectual property of Konda Tech IP. By doing so, Defendants Markovic and Wang usurped credit to the breakthrough technology developed by Konda Tech.

84. As a result, Defendants have systematically misappropriated Konda Tech IP. This has
substantially harmed Konda Tech by competing against Konda Tech usurping Konda Tech IP.
85. Plaintiff believes that Konda Tech has been deprived of customer licensees and revenue
by Defendant Markovic's and Defendant Wang's co-founding of Hierlogix and Flex Logix in
competition with Konda Tech using Konda Tech IP. Defendants Markovic and Wang and the
companies they co-founded, Hierlogix with funding by UCLA/ITA and its successor Flex Logix,
have caused severe harm in terms of Konda Tech's loss of business opportunities and taking
credit for the breakthroughs in technology that Konda Tech has made, which has negatively
impacted Konda Tech's ability to secure licenses from potential customers.

86. Furthermore, Mr. Tate, CEO of Flex Logix, has made threats, both directly to Dr. Konda and to the executive of a competitor of Flex Logix, to destroy Konda Tech.

87. Defendant Markovic's and Defendant Wang's tortious behavior and the behavior of Mr.
Tate, as described above, constitute unfair and unlawful business practices pursuant to Business
& Professions Code Section 17200, *et seq.*

88. The unlawful conduct described herein has resulted in economic harm to Konda Tech.89. As a direct and proximate result of their acts mentioned herein, Defendants have received and continue to receive ill-gotten gains belonging to Plaintiff.

90. Plaintiff is entitled to restitution for losses and has been damaged in an amount in excess of \$300,000 and to be established at trial.

91. Because the conduct alleged herein is ongoing, and there is no indication that Defendants will cease their unlawful conduct described herein, Plaintiff requests that this Court enjoin Defendants from further violations of California's laws.

## <u>SECOND CAUSE OF ACTION</u> (Fraud – Intentional Misrepresentation)

92. Plaintiff incorporates by reference every allegation contained in each and every one of the above paragraphs, as though set forth fully herein.

93. Defendant Markovic intentionally made false representations that harmed Plaintiff.

94. Defendant Markovic represented to Dr. Konda that he would assist Konda Tech to secure funding for Konda Tech to bring Konda Tech IP to the market. Defendant Markovic's representation was false.

95. Defendant Markovic knew that the representation was false when he made it because he knew that UCLA/ITA only funds technologies developed within UCLA. Neither Dr. Konda nor Konda Tech and Konda Tech IP have any affiliation or connection to UCLA and therefore do not qualify for funding by UCLA/ITA.

96. Dr. Konda did not become aware that funding by UCLA/ITA was not available until after he provided Konda Tech's confidential and proprietary business plan provided in confidence to Defendant Markovic prior to the scheduled presentation by Dr. Konda to UCLA/ITA. Had Dr. Konda been informed that UCLA/ITA does not fund technologies built outside UCLA, he would not have disclosed any of Konda Tech's confidential and proprietary information to Defendant Markovic.

97. Defendant Markovic intended that Dr. Konda and Konda Tech rely on the misrepresentation in providing proprietary and confidential information to Defendant Markovic.

98. But for that reliance, Dr. Konda and Konda Tech would not have disclosed their proprietary and confidential information with Defendant Markovic.

99. Konda Tech's reliance on the intentional misrepresentations by Defendant Markovic that he was helping Konda Tech was a substantial factor in the harm to Konda Tech. Defendant Markovic also contacted Dr. Konda pretending that he would help build Konda Tech by implementing Konda Tech's technology by submitting the Two DARPA Proposals with the promise that 1) if the Two DARPA proposals were granted, he would obtain a license from Konda Tech; and 2) otherwise if the proposals were rejected by DARPA, he would have his student Defendant Wang cease the chip implementations and any previous implementations would be used for academic purposes only.

100. Defendant Markovic intentionally made false representations to Plaintiff that the chip implementations by him and his students, including Defendant Wang, would be used for academic purposes when in fact the work by Defendants Markovic and Wang implementing Konda Tech IP were intended to form the groundwork and head start in Defendants Markovic and Wang obtaining funding from UCLA/ITA to start up Hierlogix and its successor Flex Logix in competition with Konda Tech based on Konda Tech IP.

101. Defendant Markovic intentionally made false representations to Dr. Konda in the fall of
2013 when he met with Dr. Konda at Stanford University while Defendant Markovic was a
Visiting Professor. When they met, Dr. Konda inquired whether Defendant Markovic and his
students had discontinued implementing Konda Tech IP as part of the academic work.
Defendant Markovic falsely replied "yes." During the conversation, Defendant Markovic also
asked Dr. Konda to inform him of the names of customers he was currently working with to
license Konda Tech IP, which Dr. Konda provided in confidence in the belief that Defendant
Markovic was still trying to help Konda Tech.

102. In January, 2014, while Defendant Markovic was a visiting professor at Stanford
University, Dr. Konda and Defendants Markovic and Wang met at Dr. Bonomi's residence.
When Plaintiff queried if their startup was in the area of wireless and digital signal processors
(DSPs), Defendant Markovic said "yes," which was an intentional misrepresentation because
Defendants Markovic and Wang had previously founded Hierlogix with funding from
UCLA/ITA to commercialize embedded FPGA blocks by covertly implementing Konda Tech IP
without having a license from Konda Tech.

103. Beginning at that time and continuing for four years, Defendant Markovic inquired about and collected all details of Konda Tech's technology that are not contained in the disclosures in Konda Tech's patents at the time, as well as proprietary implementation details, technical knowhow, and business know-how and the then customers and potential customers and Konda Tech's interaction with them that were disclosed to Defendant Markovic in confidence. As a result, Defendant Markovic learned about FPGA business models, competitive landscape, FPGA business opportunities and the know-how of the FPGA industry with respect to interconnect technology and its evolution, all based on the premise that he was helping to get Konda Tech funded or become one of Konda Tech's licensees.

104. Defendant Markovic intended that Konda Tech rely on his telling Dr. Konda that he was helping Konda Tech so that Konda Tech would provide him with Konda Tech's business plan and technical know-how and customer experience information over a period of years based on his intentionally misrepresenting his true intentions of helping Konda Tech when in truth he was extracting technical and business information from Dr. Konda for his own purposes in learning about the FPGA space to launch a commercial venture with Defendant Wang and exploit the information extracted from Dr. Konda using Konda Tech IP as the basis for Defendants commercial venture.

105. Defendants used the proprietary and confidential information provided by Dr. Konda and Konda Tech to develop FPGA chips which later became the basis for Defendants Markovic and Wang to found Hierlogix and its successor Flex Logix.

106. By Defendants Markovic and Wang using Konda Tech IP for the benefit of their later founded commercial ventures, they have deprived Konda Tech of revenue it would otherwise have received.

107. The unlawful conduct described herein has resulted in economic harm to Plaintiff.

108. As a direct and proximate result of their acts mentioned herein, Defendants have received and continue to receive ill-gotten gains belonging to Plaintiff and have unjustly enriched themselves at the expense of Plaintiff.

1

28

1

2

109. Monetary damages are not sufficient to compensate Plaintiff for the misappropriation of Konda Tech IP that enabled Defendants to found Hierlogix and its successor Flex Logix and illicitly receive equity in those companies. Plaintiff was damaged by being excluded as a founder notwithstanding the fact that the unauthorized use of Konda Tech IP by Defendants Markovic and Wang was instrumental to them being founders of Hierlogix and its successor Flex Logix. Without such unauthorized use of Konda Tech IP, Defendants Markovic and Wang would not have received the amount of founders shares they received in either of these companies. Plaintiff is entitled to equitable damages in an amount of Defendant Markovic's and Defendant Wang's shares of Hierlogix and its successor Flex Logix that should have been issued to Plaintiff to recognize the unauthorized and uncompensated use of Konda Tech IP by Defendants Markovic and Wang because they were unjustly enriched. Therefore, Defendants Markovic and Wang should disgorge to Plaintiff an amount of their founders shares in Hierlogix and its successor Flex Logix to be determined at trial.

110. Plaintiff is also entitled to punitive damages.

# THIRD CAUSE OF ACTION (Fraud – Concealment)

111. Plaintiff incorporates by reference every allegation contained in each and every one of the above paragraphs, as though set forth fully herein.

112. Defendants Markovic and Wang prevented Dr. Konda and Konda Tech from discovering certain facts and intended to deceive Dr. Konda and Konda Tech by concealing facts. Defendant Markovic intentionally hid the fact that his students had started to implement FPGA chips using Konda Tech IP without authorization from Konda Tech. Defendants Markovic and Wang published papers and received awards for those papers without acknowledging that their work was essentially based on the work by Dr. Konda and Konda Tech IP.

113. Defendant Markovic intentionally concealed the facts that his then student Defendant Wang intended to obtain a Ph.D. and that Defendants Markovic and Wang wanted to publish technical papers and then eventually co-found Hierlogix and its successor Flex Logix based on

proprietary and confidential technical and business information and other intellectual property of 2 Konda Tech.

114. Plaintiff was not aware of Defendant Wang's Ph.D. dissertation and technical paper publications by Defendants Markovic and Wang and eventual co-founding of Hierlogix and its successor Flex Logix. These facts were concealed from Plaintiff even though Defendant Markovic and Dr. Konda communicated occasionally over the period of October, 2009 through March, 2014. Dr. Konda did not discover they concealed facts until on or about December 18, 2015.

115. Defendant Markovic intentionally deceived Dr. Konda and collected substantially all of Konda Tech's details of technology including technology that is not contained in the disclosures in Konda Tech patents as well as implementation details, technical know-how, business knowhow, and the then customers and potential customers of Konda Tech and Konda Tech's interaction with them. Defendant Markovic also learned about FPGA business models and the know-how of the FPGA industry with respect to interconnect technology and its evolution. However, Defendant Markovic deceived Plaintiff by concealing how he was using the information that he obtained from Dr. Konda. Had Defendant Markovic told Dr. Konda that his intention, along with Defendant Wang, of how they were going to use the information they obtained from Dr. Konda, Plaintiff would have immediately stopped communicating with Defendants Markovic and Wang.

116. In January, 2014, while Defendant Markovic was a visiting professor at Stanford University, Dr. Konda and Defendants Markovic and Wang met at Dr. Bonomi's residence. During that meeting, Defendants Markovic and Wang concealed the fact that they had previously founded Hierlogix with funding from UCLA/ITA to commercialize embedded FPGA blocks by covertly implementing Konda Tech IP without having a license from Konda Tech.

117. Subsequently, Dr. Konda and Defendants Markovic and Wang met Dr. Iyer at Memoir's offices on March 5, 2014. At that meeting Defendants Markovic and Wang again concealed from Dr. Konda the fact that both Hierlogix and their new startup Flex Logix were building FPGA products based on Konda Tech IP relating to a revolutionary interconnect architecture.

1

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

-27-

118. As a result of the concealment of their use of Konda Tech IP, Defendants have caused severe harm in terms of Konda Tech's loss of business opportunities and taking credit for the breakthroughs in technology that Konda Tech has made, which has negatively impacted Konda Tech's ability to secure licenses from potential customers.

119. Defendants also concealed that they started a company, Hierlogix, and subsequently Flex Logix, using Konda Tech IP without disclosing these facts to Plaintiff.

120. Defendant Markovic and Defendant Wang concealed that their aim was to build their own company by misappropriating Konda Tech IP on the pretense of helping to obtain funding for Konda Tech.

121. Defendants concealed that their aim was to build their own company by misappropriating Konda Tech's IP.

122. Plaintiff did not become aware of the concealed facts until after December 18, 2015.

123. Had the omitted facts been disclosed to Dr. Konda or Konda Tech, Plaintiff reasonably would have behaved differently.

124. Plaintiff was harmed by the deprivation to Konda Tech of revenue it would otherwise have received as a result of the facts concealed by Defendants.

125. Defendants' concealment was a substantial factor in causing harm to Plaintiff.

126. Monetary damages are not sufficient to compensate Plaintiff for the misappropriation of Konda Tech IP that enabled Defendants to found Hierlogix and its successor Flex Logix and illicitly receive equity in those companies. Plaintiff was damaged by being excluded as a founder notwithstanding the fact that the unauthorized use of Konda Tech IP by Defendants was instrumental to them being founders of Hierlogix and its successor Flex Logix. Without such unauthorized use of Konda Tech IP, Defendants would not have received the amount of founders shares they received in either of these companies. Plaintiff is entitled to equitable damages in an amount of Defendants' shares of Hierlogix and its successor Flex Logix that should have been issued to Plaintiff to recognize the unauthorized and uncompensated use of Konda Tech IP by Defendants because Defendants were unjustly enriched. Therefore, Defendants should disgorge

to Plaintiff an amount of their founders shares in Hierlogix and its successor Flex Logix to be determined at trial.

127. Plaintiff is also entitled to punitive damages

## **FOURTH CAUSE OF ACTION** (Misappropriation of Trade Secrets)

128. Plaintiff incorporates by reference every allegation contained in each and every one of the above paragraphs, as though set forth fully herein.

129. Trade Secret is defined in California Civil Code Section 3426.1(d) as:

"Trade secret" means information, including a formula, pattern, compilation, program, device, method, technique, or process, that:

(1) Derives independent economic value, actual or potential, from not being generally known to the public or to other persons who can obtain economic value from its disclosure or use; and

(2) Is the subject of efforts that are reasonable under the circumstances to maintain its secrecy.

130. Plaintiff is the owner of Konda Tech's business plan, details of technology including all implementation details of the disclosures in Konda Tech's patents, layouts, specific architectural variations, Konda Tech's FPGA business models, and current and potential customer lists as well as Konda Tech's interaction with customers and potential customers and potential investors. All these are valuable trade secrets of Konda Tech (hereinafter referred to as "Konda Tech's Trade Secrets").

131. Dr. Konda provided a confidential list of Konda Tech's Trade Secrets to Munger, Tolles & Olson LLP on or about October 15, 2019 identifying Konda Tech's Trade Secrets, including trade secrets disclosed to Defendants Markovic and Wang in confidence during the period of October, 2009 through March, 2014.

132. Defendants Markovic and Wang systematically misappropriated Konda Tech's Trade Secrets and used them as the basis for Defendant Wang's 2013 Ph.D. Dissertation, for publishing technical papers, and as the cornerstone for the founding of Hierlogix and its successor Flex Logix. 133. The publication of trade secrets of Konda Tech by Defendants Markovic and Wang in the 2011 VLSI Paper misappropriated Konda Tech's Trade Secrets.

134. Plaintiff believes the misappropriation by Defendants Markovic and Wang of Konda Tech's Trade Secrets has caused harm to Plaintiff by depriving Konda Tech of customer licensees and revenue due to Defendants co-founding Hierlogix and its successor Flex Logix in competition with Konda Tech using Konda Techs Trade Secrets.

135. Defendants have caused severe harm in terms of Konda Tech's loss of business opportunities and taking credit for the breakthroughs in technology that Konda Tech has made, which has negatively impacted Konda Tech's ability to secure licenses from potential customers and has unjustly enriched Defendants.

136. The Konda Tech Trade Secrets had actual independent economic value at the time of disclosure to Defendants because said information was secret. Said trade secret information also had significant potential economic value at the time of disclosure as evidenced by the fact that Defendants used said information as the cornerstone to found Hierlogix and its successor Flex Logix.

137. Konda Tech made reasonable efforts to keep the Konda Trade Secrets secret. Any such information was and is disclosed to customers of Konda Tech under non-disclosure agreements.
All written disclosures to Defendant Markovic were marked "Proprietary and Confidential." All such information which Dr. Konda verbally disclosed to Defendant Markovic likewise was considered to be in confidence and not intended to be used by Defendants to compete with Konda Tech.

138. "Misappropriated" means the improper use of the trade secret. Defendants
misappropriated Konda Tech's Trade Secrets which were marked proprietary and confidential
when provided to Defendant Markovic and additional information disclosed to Defendant
Markovic in confidence by Dr. Konda. Defendant Markovic always presented himself as helping
Konda Tech get funded and be an advisor to Konda Tech. Dr. Konda believed that what Dr.
Konda disclosed to Defendant Markovic was disclosed in confidence and believed that the

intentions of Defendant Markovic were to help Konda Tech because he is a UCLA Professor and was not a competitor at that time.

139. On or about October 21, 2019, Munger, Tolles & Olson LLP made a public filing including the confidential list of Konda Tech's Trade Secrets. The public filing of the confidential list of Konda Tech's Trade Secrets was a publication of Konda Tech's Trade Secrets and constituted a misappropriation of Konda Tech's Trade Secrets.

140. Munger, Tolles & Olson LLP misappropriated Konda Tech's Trade Secrets by disclosure, because Munger, Tolles & Olson LLP disclosed them without Plaintiff's consent; and at the time of disclosure, knew that their knowledge of Konda Tech's Trade Secrets was acquired under circumstances giving rise to a duty to maintain confidentiality, which created a duty to keep the information secret, because the confidential list of Konda Tech's Trade Secrets was clearly marked confidential.

141. Munger, Tolles & Olson LLP's disclosure of Konda Tech's Trade Secrets was knowing, unnecessary, malicious, and vindictive.

142. Konda Tech's Trade Secrets had actual or potential independent economic value because they were secret.

143. Plaintiff made reasonable efforts to keep Konda Tech's Trade Secrets secret. All presentations made by Konda Tech included a "proprietary and confidential" statement.
Additionally, DARPA states that all information submitted by way of the BAA module is considered confidential. See, Exhibit 3 attached hereto.

144. Defendants' misappropriation of Konda Tech's Trade Secrets caused Defendants Markovic and Wang and their later formed commercial ventures Hierlogix and Flex Logix to be unjustly enriched.

145. Plaintiff is entitled to restitution for losses and has been damaged in an amount in excess of \$300,000 and to be established at trial.

146. Plaintiff is also entitled to punitive damages.

1	FIFTH CAUSE OF ACTION	
2	(Ongoing Conspiracy)	
3	147. Plaintiff incorporates by reference every allegation contained in each and every one of	
4	the above paragraphs, as though set forth fully herein.	
5	148. The Regents of The University of California joined the conspiracy to commit all the	
6	causes of action perpetrated by Defendants Markovic and Wang when UCLA/ITA funded	
7	Hierlogix.	
8	149. Flex Logix is the successor to Hierlogix and has committed additional wrongdoing	
9	constituting unfair business practices as a result of the actions by Mr. Tate, the CEO of Flex	
10	Logix.	
11	150. The Regents of The University of California and Flex Logix are responsible for all acts	
12	done as part of the conspiracy, whether the acts occurred before or after The Regents of The	
13	University of California and Flex Logix joined the conspiracy.	
14	151. Plaintiff is entitled to damages for losses and has been damaged in an amount in excess of	
15	\$300,000 and to be established at trial.	
16	152. Plaintiff is also entitled to punitive damages.	
17		
18	PRAYER AND RELIEF REQUESTED	
19		
20	WHEREFORE, Plaintiff respectfully prays for relief as follows:	
21	A. For judgment in Plaintiff's favor and against all Defendants as to each of the above	
22	causes of action;	
23	B. For damages in the amount to be determined at trial;	
24	C. Award Plaintiff for all damages legally and/or proximately caused by Defendants and	
25	equitable relief as set forth above, including costs and prejudgment interest and punitive	
26	damages; and	
27	D. Award Plaintiff such other or additional relief as the Court deems just and proper.	
28		
	-52-	

1	
2	
3	DATED this day of November, 2019.
4	
5	Respectfully submitted
6	
7	By: Venkah Kada
8	Venkat Konda, Ph.D.
9	
10	
11	
12	
13	
14	
15	
16	
17	
18	
19	
20	
21	
22	
23	
24	
25	
26	
27	
28	
	-33-

1		
2	VERIFICATION	
3		
4	I, VENKAT KONDA, Ph.D. declare:	
5	I have read the forgoing Complaint and know the contents thereof; that the same is true of	
6	my own knowledge, except as to the matters which are therein stated on my information and	
7	belief, and to those matters I believe it to be true.	
8	I declare under penalty of perjury under the laws of the State of California that the	
9	forgoing is true and correct.	
10	Executed on this day of November 2019, at Los Altos, California.	
11		
12	Jenhah Knok	
13	VENKAT KONDA, Ph.D.	
14		
15		
16		
17		
18		
19		
20		
21		
22		
23		
24		
25		
27		
28		
	-34-	

# **EXHIBIT 1**

# 1. Identification and Significance of the Problem or Opportunity

1.1 Objective: In this project, we plan to prepare a Phase I feasibility study of integrated circuit microcells based on regular geometry for use in our proprietary hierarchical FPGA interconnect architecture. The objective of the overall 3-phase FPGA project is to substantially reduce energy and cost of low-volume digital signal processing by using hierarchically routed interconnect, regulargeometry micro-cells, and associated tool-flow for routing and hardware mapping. Figure 1 illustrates our overall vision for the project. While the FPGA architecture and associated tool-flow for design and algorithm mapping reduces cost in design time and chip metrics (the focus of Phases II and III), enforcing regular layout geometries at the cell level provides additional reduction in the manufacturing cost, particularly in advanced technology nodes such as 32nm and below. This proposal will thus evaluate the design of regular layout cells for FPGA design and compare their circuit and cost metrics to standard-cell based CMOS design. Our team is formed from an industrial innovator (Dr. Venkat Konda) who has strong patent portfolio in routing networks (Phase II work), and academic leaders in the areas of regular geometry circuit design (Prof. Puneet Gupta), and energyefficient architecture design and associated tool-flows (Prof. Dejan Markovic). Our objective is to develop low-cost digital signal processing hardware and tool-flows for emerging markets such as wireless and sensor applications where cost and power consumption are key concerns.



**Figure 1:** Regular-fabric micro-cells and blocks (output of Phase I) will be used to route Konda's hierarchical interconnect architecture (output of Phase II) and further integrated on a demo FPGA chip with supporting mapping tool-flow (output of Phase III) to demonstrate significant improvements in chip size, performance, power, and also manufacturing cost.
*Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA* Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

- 1.2 Problem Addressed: With increasing cost of semiconductor design and manufacturing, enforcing regularity at all layers from device technology to hardware architecture is essential for future lowpower and low-cost digital signal processing hardware. At the architecture level, FPGA-like regularity is becoming an attractive solution particularly for low-volume applications, but the adoption of FPGAs will be greatly challenged with their excess power, area, and performance due to the massive FPGA interconnect. The complexity of the FPGA interconnect is a quadratic function,  $O(N^2)$ , of the number of processing elements, N. To mitigate the interconnect challenge, we will make use of hierarchically routed and proprietary Konda interconnect architecture which has greatly reduced complexity, O(N·log<sub>2</sub>N), which results in improved area, power, and performance of FPGA chips. Additionally, we must face unique challenges of scaled technology and enforce regularity in the layout cells (microcells). With the slow development of Extreme Ultra Violet (EUV) lithography, double patterning technology (DPT) appears as the most viable lithography solution for 32nm and later technology nodes [1]. DPT allows for more compact and better-yielding layout using mask decomposition to effectively increase pitch size. To find best DPT decompositions as applicable to FPGA micro-cells and building blocks, we propose to investigate micro-cell routing algorithms and characterize the cells in energy-area-performance space as compared to their standard-cell based CMOS counterparts.
- 1.3 Proposed Solution: Our approach will consist of:
  - (a) Research the state-of-the-art regular layout geometries and routing algorithms for micro-cells,
  - (b) Innovate and provide unique solutions to overcome challenges at the cell layout, circuit, and architecture levels,
  - (c) Develop modeling and simulation framework that will guide the final selection of regular-geometry micro-cells to be used in FPGA macro-blocks such as lookup tables (LUTs), DSP slices, block RAM (BRAM) modules, switch matrix (SM) elements that include switch boxes (SBs) and configuration memory, and
  - (d) Perform **energy**, **area**, **performance**, **yield**, **and variability evaluation** of the propose micro-cell and macro-block structures for use in hierarchical FPGA interconnect architecture.

As a quantitative measure of our Phase I study, we plan to provide an extensive list of circuit metrics as listed in Table 1. The metrics include area, energy, performance, variability, and yield estimates for standard-cell and proposed regular-geometry cells (both at the micro and macro levels). The outcome of Phase I will be to populate Table 1 with quantitative measures of functional-block metrics, and to provide associated solutions for layout cells. The layout cells from Phase I will be subsequently used in hierarchical FPGA interconnect architecture (Phase II), FPGA chip and hardware mapping tool-flow (Phase III) to provide over 10x improvement in power compared to the state-of-the-art FPGAs.

Matria /	Energy (fJ)		Delay (ps)		Variability (%)		Yield (%)	
Functional Block	Std- cell	Regular geometry	Std- cell	Regular geometry	Std- cell	Regular geometry	Std- cell	Regular geometry
NAND gate								
Flip-flop								
AOI gate								
Full adder								
4-input LUT								
DSP slice								
Switch box								
Switch matrix								

Table I: Feasibility study of quantitative figures of merit of layout cells for FPGA application.

*Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA* Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

- 1.4 Proposal Strength: The main strength of the proposed work is the multi-disciplinary approach that spans technology, circuits, architectures, and algorithms for exploiting regularity multiple hierarchical layers in the design of digital signal processing hardware. The combined effort in the aforementioned areas will lead to the development of low-cost FPGA platform based on regular-geometry layout cells, hierarchically routed interconnect architecture, and tool-flow for area-efficient hardware mapping of digital signal processing algorithms. Our strength in all aspects from layout to algorithms will allow for (a) layout cell development, (b) accurate development of device and circuit specifications, (c) allow for extensive analysis to predict yield, power consumption, chip area, and performance, (d) provide full hardware/software demonstration at the end of the 3-phase program. Furthermore, our team has extensive experience in energy-efficient integrated circuits and architectures, CAD algorithms and layout cells, as well as network architectures and supporting routing algorithms.
- **1.5 Market Opportunity:** We see a great market potential in broad area of digital signal processing hardware where the cost and power consumption are key challenges. Our final goal is to develop a simple and low-cost FPGA hardware/software technology based on regular layout cells, regular hierarchical interconnect architecture, and tools for block routing and hardware mapping. The potential markets include both commercial and defense segments. With greatly reduced power consumption and cost, the technology will particularly impact energy-starved applications such as embedded electronics and distributed sensors. The technology will also provide a solution to rapid prototyping and emulation for a variety of communications and imaging applications. To reduce design cost and ensure **scalability**, our approach will deliver **hierarchical methodology** from microcell layout to final chip architecture and supporting tool-flows.
- **1.6 Company Profile:** Konda Technologies, Inc. is a startup company based in San Jose, CA. The company was founded in 2007 to develop & commercialize interconnect IP applicable for various products including FPGA routing interconnect, System-on-Chip interconnects and warehouse-scale datacenter switch networks. Our main customer today is Tier Logic Inc, a 3D-FPGA startup. The company has been engaged with vendors such as Xilinx Corporation, Altera Corporation and Cisco Systems.

# 2. Phase I Technical Objectives

- **2.1 Objective 1: Development of reusable infrastructure of regularity evaluation at cell-level** We will develop a tool infrastructure to allow for evaluation and exploration of regular layout styles. This would include fast estimation-based methods as well as layout generation and simulation based methods.
- 2.2 Objective 2: Analysis of regularity tradeoffs at different layers and identification of layout styles suitable for the FPGA architecture

Using the regularity evaluation framework developed above, we will identify the optimal choice of regular layout styles on front-end layers (poly, active, M1, M2, contact). This will be applied to varying levels of design complexity ranging from standard cells to entire FPGA functional macros.

#### 2.3 Objective 3: Develop a comprehensive plan for Phase II

The outcome of Phase I will be a comprehensive study of micro-cells and macro-blocks that will be used in Phase II to implement hierarchical interconnect architecture. The objective is to significantly improve energy, area, and performance of the FPGA hardware. The output of Phase I will be guided by the metrics outlined in Table 1. Phase I solutions will be developed with tight interaction between architecture, circuit, and process parameters to ensure globally optimal solutions. We will propose an IP library and associated tools to reduce design cost and facilitate commercial adoption of our technology.

*Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA* Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

# 3. Phase I Work Plan

#### 3.1 Introduction and Prior Art

Generally, regularity makes patterning easier. Inserting dummy features to ensure uniform density or to "isolate" standard cells from surroundings has been commonly followed approach. For more regularity, set of layout constraints or restrictive design rules [2], can be enforced to guarantee a lithography-friendly regular layout. As an example, a unidirectional fixed-pitch poly layer is enforced in Intel's 45nm process. Because of the success of such gridded design rules in enhancing printability and reducing variations [4], such rules might be adopted to pattern other patterning layers such as metal and contacts/vias. This principle of restricting the layout is pushed to the extreme in [5] where layout is constructed out of pre-characterized regular fabrics (as opposed to design rules). A regular layout approach can be excessively conservative especially for layouts where patterning imperfections would otherwise be tolerable [2]. Nevertheless, increasing degree of regularity is expected to make patterning even feasible in the near term.

Another important point is that regularity need not imply 1D gratings. The basic "template" for regularity could be something else while still ensuring good, low-cost printability (e.g., see [6]). The template printability can be optimized, for example, using source-mask optimization (SMO) or using character projection in maskless E-beam direct-write. Part of our work will also investigate if regularity other than gratings can be useful.

#### 3.2 Our Approach

Our approach within this proposal, and in line with the SBIR call, is to examine routing tools for regular-geometry layout cells. The goal in Phase I will be to develop routing tools and layout cells for FPGA building blocks. The layout cells will vary in granularity from simple logic gates to complex blocks such as look-up tables, logic slices, switch boxes, and memory components. The regular cells will be characterized for density (area), yield, energy, and performance and compared to regular standard-cell based approach. The regular cells will be used in Phase II for interconnect architecture routing. The cells and routing tools will be made available as IP to facilitate rapid commercialization.



**Figure 2:** Regularity evaluation framework (left), DRE results on 45nm Nangate open-cell library (right).

Regularity is a continuum of possibilities and it has significant impact on area, delay, power as well as expected manufacturing yield. It therefore is very important to co-optimize design rules, regular layout styles as well as cell architectures. We have developed a Design Rule Evaluator (DRE) framework (see **Error! Reference source not found.**) which predicts the impact of layout style and design rule changes on important circuit metrics for standard cells as well as small custom blocks. DRE can run through a 100+ cell 45nm cell library in a few minutes with less than 2% average estimation error (see **Error! Reference source not found.**) making it perfectly suited for design space exploration of layout

#### Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

regularity, design rules and design styles. As an example preliminary study using DRE, consider increasing regularity on the polysilicon layer (see

Figure 1). We compare the cases of unrestricted 2D layout to 1D layout with arbitrary pitches and restricted 1D fixed pitch (i.e., grating-like) layout styles for a 45nm sequential benchmark design.



**Figure 1:** Comparing layout restrictions for a benchmark design on polysilicon layer in terms of area, catastrophic yield (POS or probability of survival) and current variability (change in W/L).

We will extend DRE to full-chip (FPGA) evaluation including local, intermediate and global metal/via layers. This will allow us to arrive at a principled choice of regularity and a layout style to enforce it for the FPGA interconnect fabric. The first phase will use DRE coupled with some layout design and simulation to identify optimal choice of regularity for basic building blocks of the FPGA.

#### 3.3 Task 1: Identification of candidate layout styles for regularity evaluations

For different candidate patterning technologies at 32nm, 22nm, 16nm nodes, we will identify what forms of regularity and on what layers will help the most. This set may be a large one, especially, for the 22nm and 16nm nodes where lithographic patterning choices are still unclear (double patterning, self-aligned double patterning, e-beam direct write, interference-assisted lithography, etc).

#### 3.4 Task 2: DRE-based exploration of regularity tradeoffs in FPGA building blocks

We will extend DRE framework to allow us to evaluate delay-power-yield-area-variability tradeoffs for regularity on polysilicon, contact, M1, M2 layers. The result will be a principled narrowing down of layout style choices with clear understanding of the tradeoffs.

#### 3.5 Task 3: Generation of layout, simulation, and comparison

Using the optimized design rules and regular layout stules derived above, we will draw layouts of FPGA micro-cells. These will then be analyzed for delay/power/variability/manufacturability using explicit lithography simulation using a projected 32nm lithography setup coupled with non-rectangular transistor models. This will allow us for a close-to-silicon comparison of different layout styles (e.g., irregular, 2D standard cells vs. regular layouts).



**Figure 4:** (left) Area-Energy-Delay space for comparing multiple circuit and micro-architectural options. (right) Energy-delay tradeoff in CMOS (solid line) indicating minimum-delay (MDP) and minimum-energy (MEP) points. Regular-geometry based designs marked in (X) are expected to provide better energy-delay tradeoff than standard-cell based CMOS.

#### Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

We will use our methodology for area-energy-delay optimization of CMOS circuits and architectures [7, 8]. The methodology is based on pareto curve analysis for various circuit and architecture realizations as indicated in Fig. 4(a). Each tradeoff curve is a result of optimization program that minimizes energy subject to a delay constraint for circuits. The optimal tradeoff for the circuit-level energy and delay is illustrated in solid line in Fig. 4(b) by tuning gate size, supply, and threshold voltage. The line is bounded by minimum-delay (MDP) and minimum-energy (MEP) points. All points above the line are suboptimal, all points below the line are infeasible.

The goal of regular-geometry explorations is to achieve better energy-delay tradeoff than regular standard-cell based CMOS approach as indicated by the (X) markers in Fig. 4(b). Points below MEP are the most desirable and it is expected that the micro-cell development will go mainly in this direction. Points above MEP but still below E-D plot of CMOS are also very desirable. We will use compact circuit models to formulate optimization problems, perform simulations and to populate the metrics in Table 1. This includes various FPGA datapath and storage functions.

#### 3.6 Task 4: A comprehensive Phase II development plan

Towards the end of Phase I, as outlined in Table 2, and based on the outcome of Phase I, we will create a comprehensive Phase II development plan. We aim to integrate hierarchical interconnect architecture in Phase II based on the micro-cells and macro-blocks from Phase I. Details of the proprietary hierarchical interconnect architecture will be available in our proposal at the conclusion of Phase I.

#### 3.7 Timeline

Month	1	2	3	4	5	6
Task 1						
Task 2						
Task 3						
Task 4						

Table II: Phase I project schedule

#### 3.8 Task Work Breakdown

Table III: Estimated task hours for key Personnel

Taak/Doroon	V. Konda	Scientist	D. Markovic	P. Gupta
Task/Person	(PI, Konda Tech)	(Konda Tech)	(UCLA)	(UCLA)
Task 1	100	140	20	80
Task 2	80	100	30	70
Task 3	60	60	50	30
Task 4	40	40	100	20
TOTAL	280	340	200	200

#### 4. Related Work

We have worked extensively on design-patterning interactions. We have developed methods for evaluation of regular layout styles through layout generation and simulation (DAC 2004) as well as through estimation and modeling (ICCAD 2009). Prof. Gupta has worked extensively on electrical modeling (SPIE'06, SPIE JM3'10, VLSID'10, ASPDAC'08, etc) and mitigation (TCAD'07, SPIE JM3'09, etc) of lithographic imperfections. Prof. Markovic has a strong track record in energy-efficient ASICs for digital signal processing.

*Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA* Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

**Relevant Business Relationships:** Since its founding in 2007, the company has attracted strong interest from a variety of companies. The company has been engaged with vendors such as Xilinx Corporation, Altera Corporation and Cisco Systems. Our main customer today is Tier Logic Inc, a 3D-FPGA startup.

**Related Work by Others:** Regular layouts have been under investigation to various extents in academia and industry for past few years. Commercial foundries enforce regularity to varying degrees using design rules (e.g., unidirectional, gridded poly is likely to be widely required at 32nm node). The origins of the approach lie in early work done by IBM on restricted design rules to be used for Alternating PSM patterning. Most cell libraries at 32nm node will use regular layouts, at least for the polysilicon layer. More regularity on other layers (contact, M1) has also been investigated in somewhat limited fashion by companies (e.g., PDBrix from PDF solutions and AreaTrim by Tela Innovations) but extensive tradeoff analyses between extent of regularity, area and yield is still an open problem.

# 5. <u>Relationship with Future Research or Research and Development</u>

#### 5.1 Phase I Results

The goal of Phase I is to develop a tool infrastructure to allow for evaluation and exploration of regular layout styles. This would include fast estimation-based methods as well as layout generation and simulation based methods. Using the regularity evaluation framework developed above, we will identify the optimal choice of regular layout styles on front-end layers (poly, active, M1, M2, contact). Our approach in Phase I will also focus the demonstration of library IP and the use of software to reduce cost and facilitate rapid commercialization. We will also study preliminary routing strategies for regular interconnect architecture for Phase II.

#### 5.2 Relationship to Phase II and its Objectives

The regularity layout cells will be extended to FPGA architecture. The FPGA devices have regularly placed LUTs (Look-up tables) in a 2D-plane on a silicon die. So far 2D-Mesh networks have been used in FPGA devices due to their regular structure, i.e., both interconnect distribution-wise as well as the horizontal and vertical routing tracks layout-wise. However the switch complexity of the 2D-Mesh based FPGA interconnect is a quadratic function,  $O(N^2)$ , of the number of processing elements, N. Even though Benes/Butterfly Fat Tree networks with switch complexity of  $O(N \cdot \log_2 N)$ , which results in improved area, power, and performance of FPGA chips, they are not implementable due to the lack of known regular VLSI layouts, till today. Konda Technologies inventions with regular VLSI layouts for Benes/BFT based hierarchical networks are seminal and subsumes all the other known network topologies such as Clos networks, hypercube networks, cube-connected cycles and pyramid networks, which makes these networks implementable in a FPGA devices with regular structures both interconnect distribution-wise and layout-wise which is the key to exploit improved area, power, and performance of FPGA devices. The regularity of Konda hierarchical layout is also the key for its commercializability in System-on-Chip interconnect devices, FPIC devices as well.

# 6. Commercialization Strategy

We believe that our fundamental intellectual property would help us to commercialize out IP by technology and tools licensing. We have already been successful with our current engagement with Tier Logic to incorporate our interconnect IP into Teir Logic's 3D-FPGA devices.

#### 6.1 General Commercial and Technology Landscape

In the regular layout cells space, there have several undertakings in both academia and industry. Commercial foundries also enforce regularity to varying degrees using design rules (for example, unidirectional, gridded poly is likely to be widely required at 32nm node). The origins of the approach lie in early work done by IBM on restricted design rules to be used for Alternating PSM patterning.

#### *Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA* Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

Most cell libraries at 32nm node will use regular layouts, at least for the polysilicon layer. More regularity on other layers (contact, M1) has also been investigated in somewhat limited fashion by companies (e.g., PDBrix from PDF solutions and AreaTrim by Tela Innovations) but extensive tradeoff analyses between extent of regularity, area and yield is still an open problem. Our advantage is system-wide visibility and consideration of regularity that starts from micro-cell level and goes up to our proprietary interconnect architecture.

#### 6.2 Market Opportunity

Our initial market focus will be in electronics for portable applications where energy consumption is limited and where cost is a key concern such that scalability can be achieved. The approach described in this proposal is to create library cells for programmable integrated circuits in advanced technology nodes such as 32nm and beyond. We expect this technology to complement existing patent portfolio at Konda Technologies, Inc. in the area of network routing algorithms for a variety of markets.

Commercialization of the technology is foreseen to be developed with close consultation with large semiconductor companies such as Qualcomm, Broadcom, ST Microelectronics, Novelics, Xilinx, and Altera where strategic partnerships have already been established. In addition, we expect large interest from defense companies such as Boeing and Northrop Grumman. We will certainly take inputs from both the civil and DoD companies to best tailor the technology platform to each market segment.

We foresee the opportunity to use the technology in application-specific integrated circuit (ASIC) markets as well as FPGA market. ASICs used in wireless devices are power-limited yet require large amounts of flexibility for multiple operation modes. FPGAs can provide the flexibility, but at a prohibitive cost in power and area. Our technology provides solution to both of these problems as we offer flexible yet low power FPGA technology. Our micro-cells and routing tools can be used as IP by communication and FPGA companies alike. In 2010, FPGA market is expected at \$4B, with projections of steady growth up to \$6B in 2015 [9]. The ASIC market, \$18B in 2009, is projected to exceed \$22B by 2010 [10].

We plan to expand our patent portfolio and issue soft IP (micro-cells and routing algorithms) on a non-exclusive license basis to ASIC and FPGA companies such as Qualcomm, Broadcom, Samsung, Xilinx, Altera, and Cisco.

# 7. Key Personnel

#### 7.1 Company Background

Based on a breakthrough and patent-pending layout for Benes/Butterfly Fat Tree network using horizontal and vertical tracks and with commercial potential for wide target applications such as FPGA devices, FPIC devices, logic emulation systems, Konda Technologies was founded in 2007 to commercialize the intellectual property into these markets. Our initial focus has been to commercialize interconnect IP into FPGA devices.

#### 7.2 Dr. Venkat Konda, Principal Investigator & CEO, Konda Technologies, Inc.

Venkat Konda is an inventor, experienced entrepreneur and the CEO of Konda Technologies which he founded in 2007 based on a breakthrough layout using only horizontal and vertical tracks for Benes/BFT hierarchical networks, seminal rearrangeably and strictly non-blocking multicast routing algorithms with an architecture optimum with switch cost, power and performance. Venkat is currently in the process of commercializing the IP in FPGA interconnects, System-on-Chip interconnects and warehouse-scale datacenter switches. Prior to it, Venkat invented seminal algorithms for rearrangeably and strictly non-blocking multicast routing algorithms for Clos Networks and founded a startup Teak Networks, to commercialize into packet switch fabrics which are also applicable to design cheaper optical cross connects. Venkat received PhD degree in Computer Science & Engg

*Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA* Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

from University of Lousiville, KY in 1992, and M.S in Electrical Engineering from Indian Institute of Technology, Kharagpur in 1988. Key patents/applications include:

- [1] Venkat Konda, "Fully connected generalized multi-stage networks", USPTO App# 12/530,207.
- [2] Venkat Konda, "Fully connected generalized Butterfly Fat Tree networks", USPTO App# 12/601,273.
- [3] Venkat Konda, "VLSI Layouts of Fully connected generalized networks", USPTO App# 12/601,275.
- [4] Venkat Konda, "Rearrangeably nonblocking multicast multi-stage networks ", US Patent # 6,885,669.
- [5] Venkat Konda, "Strictly nonblocking multicast multi-stage networks ", US Patent # 6,868,084.

#### 7.2 Prof. Dejan Markovic, UCLA, Electrical Engineering (Sub-contractor)

Dejan Markovic is an Assistant Professor of Electrical Engineering at UCLA. He completed the Ph.D. degree in 2006 at the University of California, Berkeley. In recognition of the impact of his Ph.D. work, he was awarded 2007 David J. Sakrison Memorial Prize at UC Berkeley. His current research is focused on integrated circuits for emerging radio and healthcare systems, design with post-CMOS devices, optimization methods and CAD flows. He will be contributing to the design and circuit demonstration tasks in this project. His responsibilities will include layout cell characterization, design and optimization of FPGA building blocks. Some relevant publications include:

- [1] D. Marković, C. C. Wang, L. Alarcon, T.-T. Liu, and J. M. Rabaey, "Ultralow-Power Design in Near-Threshold Region," Proceedings of the IEEE, vol. 98, no. 2, pp. 237-252, Feb. 2010.
- [2] R. Nanda, C.-H. Yang, and D. Marković, "DSP Architecture Optimization in Matlab/Simulink Environment," in Proc. Int. Symp. on VLSI Circuits (VLSI'08), June 2008, pp. 192-193.
- [3] D. Marković, B. Nikolić, and R.W. Brodersen, "Power and Area Minimization for Multidimensional Signal Processing," IEEE J. Solid-State Circuits, vol. 42, no. 4, pp. 922-934, April 2007.
- [4] D. Marković, R.W. Brodersen, and B. Nikolić, "A 70GOPS 34mW Multi-Carrier MIMO Chip in 3.5mm2," in Proc. Int. Symp. on VLSI Circuits (VLSI'06), June 2006, pp. 196-197.
- [5] D. Marković, V. Stojanović, B. Nikolić, M.A. Horowitz, and R.W. Brodersen, "Methods for True Energy-Performance Optimization," IEEE J. Solid-State Circuits, vol. 39, no. 8, pp. 1282-1293, Aug. 2004.

#### 7.3 Prof. Puneet Gupta, UCLA, Electrical Engineering (Sub-contractor)

Puneet Gupta (http://nanocad.ee.ucla.edu) is currently an Assistant Professor of Electrical Engineering at UCLA. He received the B.Tech degree in Electrical Engineering from Indian Institute of Technology, Delhi in 2000 and Ph.D. in 2007 from University of California, San Diego. He co-founded Blaze DFM Inc. (acquired by Tela Inc.) in 2004 and served as its product architect till 2007. He is a recipient of NSF CAREER award, ACM/SIGDA Outstanding New Faculty Award, IBM Ph.D. fellowship and European Design Automation Association Outstanding Dissertation Award. Dr. Gupta's research has focused on building high-value bridges between physical design and semiconductor manufacturing for lowered cost, increased yield and improved predictability of integrated circuits. He will be contributing to the design and circuit demonstration tasks in this project. His responsibilities will include optimization of regular layout styles, layout generation and characterization. Key relevant publications include:

- [1] P. Gupta, A. B. Kahng, D. Sylvester, and J. Yang, "Toward a Methodology for Manufacturability Driven Design Rule Exploration," in Proc. DAC, June 2004.
- [2] R. S. Ghaida and P. Gupta, ``A Framework for Early and Systematic Evaluation of Design Rules," in IEEE/ACM ICCAD, November 2009.
- [3] P. Gupta, A. B. Kahng, P. Sharma, and D. Sylvester, ``Gate-Length Biasing for Runtime Leakage Control," IEEE Transactions on CAD, June 2006.
- [4] T.-B. Chan, R. S. Ghaida, and P. Gupta, ``Electrical Modeling of Lithographic Imperfections," in Proc. IEEE/ACM VLSI Design Conference, 2010.
- [5] R. S. Ghaida and P. Gupta, ``Within-Layer Overlay Impact for Design in Metal Double Patterning," to appear in IEEE Transactions on Semiconductor Manufacturing, 2010.

#### 8. Facilities/Equipment

During Phase I of this project, no special facilities or equipment will be required to complete the proposed plan. Konda Technologies will only require the services of Profs. Gupta and Markovic from

**Proprietary** Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

Electrical Engineering at UCLA to provide design, modeling, and simulation capabilities. No equipment purchase will be necessary.

### 8.1 Government Equipment and Facilities

No government facilities or equipment will be used during this project.

#### Proprietary Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

# 9. Subcontractors/Consultants

#### Subcontractors: Profs. Dejan Markovic and Puneet Gupta, UCLA Electrical Engineering

#### UNIVERSITY OF CALIFORNIA, LOS ANGELES

BERKELEY . DAVIS . IRVINE . LOB ANGELES . MERCED . RIVERSIDE . SAN DIEGO . SAN FRANCISCO

ANGISCO SANTA BARBARA • SANTA ERUZ

UCLA

Henry Samueli School of Engineering and Applied Science Electrical Engineering Department Engineering IV Building, 420 Westwood Plaza Los Angeles. California 90095-1594

June 22, 2010

To: Dr. Venkat Konda Konda Technologies, Inc. 6278 Grand Oak Way San Jose, CA 95135

Dear Dr. Konda:

We would like to express our interest in working with Konda Technologies, Inc. in support of the DoD SBIR project solicitation topic SB102-003 "Design Tools for Highly Regular Circuit Geometries." Our groups at UCLA would be willing to provide collaboration in accordance with the following statement of work:

Phase I tasks:

- 1. Expoloration of candidate layout styles for regular-geometry circuit building blocks.
- 2. Design rule evaluator (DRE) based evaluation of selected regular layout styles.
- 3. Layout generation, simulation, and characterization of energy, area, delay, variation, and yield metrics.

UCLA Budget Phase I: \$32,000.

We understand that the start date will be mid-late 2010 and the total duration will be six months.

Sincerely,

Dejan Marković, Assistant Professor UCLA Electrical Engineering Department 56-147E Eng-IV Bldg, 420 Westwood Plz Los Angeles, CA 90095-1594 Tel: (310) 825-8656, Email: dejan@ec.ucla.edu URL: http://www.ec.ucla.edu/~dejan

Puncet Gupta, Assistant Professor UCLA Electrical Engineering Department 6730C Boelter Hall, 420 Westwood Plaza Los Angeles, CA 90095-1594 Tel: (310) 825-1376, Email: puncet@ce.ucla.edu URL: http://www.ee.ucla.edu/~puncet

Regular-Geometry Micro-Cells and Design Tools for Butterfly FPGA Proposal Number: D102-0003-0305, Topic Number: SB102-003 (DARPA)

# 10. Prior, Current, or Pending Support of Similar Proposals or Awards

None.

### 11. References

- [1] G. E. Bailey, A. Tritchkov, J.-W. Park, L. Hong, V. Wiaux, E. Hendrickx, S. Verhaegen, P. Xie, and J. Versluijs, "Double pattern EDA solutions for 32nm HP and beyond," In Proc. SPIE 6521, 2007.
- [2] M. Lavin, F. Heng and G. Northrop, "Backend CAD flows for "restrictive design rules", in Proc. IEEE/ACM Intl. Conf. on Computer Aided Design, pp. 739-746, 2004.
- [3] R. S. Ghaida and P. Gupta, ``A Framework for Early and Systematic Evaluation of Design Rules," in IEEE/ACM ICCAD, November 2009.
- [4] M. C. Smayling, H. Liu, and L. Cai, "Low k1 logic design using gridded design rules," Proc. SPIE, p. 69250B, 2008.
- [5] T. Jhaveri, V. Rovner, L. Liebmann, L. Pileggi, A. J.Strojwas, J.D. Hibbeler, "Co-optimization of circuits, layout and lithography for predictive technology scaling beyond gratings,", IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems, vol. 29, no. 4, pp. 509-527, 2010.
- [6] R. S. Ghaida, G. Torres, and P. Gupta, ``Single-Mask Double-Patterning Lithography," in SPIE/BACUS Photomask Technology, September 2009.
- [7] D. Marković, C. C. Wang, L. Alarcon, T.-T. Liu, and J. M. Rabaey, "Ultralow-Power Design in Near-Threshold Region," Proceedings of the IEEE, vol. 98, no. 2, pp. 237-252, Feb. 2010.
- [8] D. Marković, V. Stojanović, B. Nikolić, M.A. Horowitz, and R.W. Brodersen, "Methods for True Energy-Performance Optimization," IEEE J. Solid-State Circuits, vol. 39, no. 8, pp. 1282-1293, Aug. 2004.
- [9] Electronics Weekly, [Online] Available: http://www.electronicsweekly.com/Articles/2010/05/19/48677/fpgamarket-soaring-to-4bn-in-2010-says-gavrielov.htm
- [10] BCC Research, "The Global Market for ASICs," Report Code: SMC067A, Published: June 2009 [Online] Available: http://www.bccresearch.com/report/SMC067A.html

# **EXHIBIT 2**

# Volume I Technical and Management Proposal

Title: Energy-Efficient Butterfly FPGA Hardware and Programming Tools

# A proposal submitted to Dr. William Harrod, DARPA/TCTO in response to

DARPA-BAA 10-78:	Omnipresent High Performance Computing (OHPC)
Technical Area:	Energy Efficient Computing
Lead Organization:	University of California, Los Angeles (UCLA) Department of Electrical Engineering Los Angeles, CA 90095-1594
Type of Business:	Other Educational
Team Members:	Dejan Markovic (PI) Venkat Konda (Consultant)

# **Technical Point of Contact:**

**Administrative Point of Contact:** 

Office of Contract and Grant Administration

UCLA Senior Grant Analyst

11000 Kinross Ave, Suite 102

Los Angeles, CA 90095-1406

Email: ocga5@research.ucla.edu

Tel: (310) 794-0155

Fax: (310) 943-1658

Ms. Julia Zhu

Dr. Dejan Markovic, PI UCLA Associate Professor Electrical Engineering Department 56-147D Engineering IV Building 420 Westwood Plaza Los Angeles, CA 90095-1594

Tel: (310) 825-8656 Fax: (310) 206-8495 Email: dejan@ee.ucla.edu

# Total funds requested: \$2.374.111

 	Year 1:	\$789,927
	Year 2:	\$792,100
	Year 3:	\$792,086



Date of proposal: August 4, 2010

#### UNIVERSITY OF CALIFORNIA, LOS ANGELES

BERKELEY • DAVIS • IRVINE • LOS ANGELES • MERCED • RIVERSIDE • SAN DIEGO • SAN FRANCISCO



SANTA BARBARA • SANTA CRUZ

OFFICE OF CONTRACT AND GRANT A DMINISTRATION BOX 951406 11000 KINROSS, SUITE 102 LOS ANGELES, CALIFORNIA 90095-1406

> PHONE: (310) 794-0102 FAX: (310) 794-0631

**UCLA** 

www.research.ucla.edu/ocga

August 5, 2010

DARPA/TCTO ATTN: DARPA-BAA-10-78 3701 N. Fairfax Drive Arlington, VA 22203-1714

The Regents of the University of California, Los Angeles, is pleased to submit the following proposal in response to solicitation DARPA-BAA-10-78.

Title: "	Energy-Efficient Butterfly FPGA Hardware and Programming Tools.
Requested Period of Performance	ce: September 15, 2010 – September 14, 2013
Amount Requested:	\$2,374,111
Principal Investigator:	Dr. Dejan Markovic Department of Electrical Engineering dejan@ee.ucla.edu 310-825-8656

This application is being submitted in contemplation of an agreement containing mutually agreeable terms and conditions applicable to educational institutions conducting unclassified fundamental research.

Since UCLA is a public/State institution, open dissemination of research results and information, commitment to students, accessibility for research purposes, and legal integrity and consistency are part of the University's Principles/Policy. The University does not discriminate and impose restrictions on any individual as a result of their nationalities.

If an award is made, please be advised that if it is funded by budget category 6.3(Advanced Research) and is considered Non-fundamental research, we will not be able to accept the award due to publication restrictions.

Your favorable consideration of this proposal would be appreciated. Technical questions should be directed to Dr. Markovic. Administrative and contractual questions, should be directed to me at (310) 794-0155 or via email at jzhu@research.ucla.edu.

Sincerely,

Julia Shy

Julia Zhu Senior Grant Analyst

# **Table of Contents**

# **Executive Summary**

UCLA offers to perform research on a revolutionary new FPGA technology consisting of FPGA hardware and supporting mapping tools. We will design, fabricate, and test hierarchical FPGA interconnect network to demonstrate FPGA technology that is 15x more energy-efficient than existing FPGAs. The new interconnect architecture allows for significant reduction in the number of switch points, buffers, and wire length in comparison to standard 2D-mesh architecture used by existing FPGAs. The proposed technology is a radical departure from 2D-mesh design, which for N logic blocks has complexity  $O(N^2)$ , incomplete and heuristic routing. The proposed technology has only  $O(N \cdot \log_2 N)$  complexity, complete and fully deterministic routing. The proposed technology has significant benefits: 15x lower power, 3x lower area, 2x higher performance compared to existing FPGA technology. The new FPGA technology will be used to demonstrate HPC benchmarks with a 15x higher power efficiency for DOD and commercial users. The PI has established interactions with industrial partners that will lead to the transition of ideas into the commercial space.

# Section II - Technical Details

# 2.1. PowerPoint Summary Chart



# 2.2. Innovative Claims for the Proposed Research

# **Problem Description**

Today's programmable FPGA devices are expensive in size, power, performance, scalability and flexibility. All of this is due to a fundamental problem in 2D-mesh interconnect architecture: it is large in size, has long latency, consumes lots of power, and is not scalable. Interconnect takes more than 75% of the FPGA chip area. Large number of inactive transistors also results in significant leakage power (about 50% of the total FPGA power). Due to inefficient interconnect architecture, there is a 30-50x energy-efficiency gap between FPGA and dedicated chips (Fig. 1).



**Figure 1:** Energy efficiency for various computing architectures: microprocessors, general purpose DSPs, FPGAs, and dedicated chips. The study is based on chips from the ISSCC conference (normalized to the same technology). FPGAs with DSP cores are 30-50x less energy efficient than dedicated chips.

# **Research Goals**

We will integrate hierarchical interconnect network to demonstrate significant improvements in speed, power, and area as compared to existing FGPAs technology. The hierarchical interconnect architecture requires at least 3x smaller number of active network elements, switch points and drivers. This is illustrated in Fig. 2 for a very simple 2x2 example.



Figure 2: 2D-Mesh and Konda networks for a design consisting of 4 CLB blocks.

For larger number N of configurable logic elements, the benefits of hierarchical network will be even more pronounced (Table 1). Such large cost of the 2D-mesh architecture forces designers to employ heuristics to reduce the number of switch points, which results in insufficient connectivity. The hierarchical network provides complete and deterministic routing.

Number of LUTs	2D-Mesh	Konda butterfly	Savings factor
1 k	1 M	9.97 k	<b>100x</b>
100 k	10 B	1.66 M	6,200x

 Table 1: Number of connections in 2D-Mesh and Konda networks.

# **Expected Impact**

The new FPGA platform will provide significant savings in power compared to today's FPGAs as shown in Fig. 3. Our FPGA technology, which includes hardware and supporting mapping tools, will provide an estimated 15x power reduction as compared to conventional FPGAs.



**Figure 3:** Power consumption for a range of applications. New FPGA will provide significant power reduction compared to typical Virtex-5 FPGA (normalized to the same technology).

We will provide new FPGA technology consisting of hardware and mapping tools. The hardware and mapping tools will provide significant impacts: 15x lower power, 3x lower area, 2x higher performance compared to existing FPGA technology. The new FPGA technology will be used to demonstrate HPC benchmarks with a 15x higher power efficiency for DOD and commercial users. Equivalently, our FPGA technology can provide >10x higher throughput for the same amount of power (as shown in Fig. 3). This technology will be of use for HPC applications and many other DOD applications which use FPGA technology.

# 2.3. Proposal Roadmap

**Main goals of the proposed research:** The main goal of the program is to develop energyefficient programmable hardware and supporting software mapping tools. The hardware is based on hierarchical interconnect architecture that provides significant reduction in interconnect complexity as compared to today's FPGA hardware. With a combination of new interconnect architecture and supporting toolflow, we project over a 15x improvement in energy efficiency while also considerably reducing chip area and improving performance. The proposed work builds on patent-protected network architecture and successful chip demonstrations. The work proposed here focuses on the investigation of needed level of connectivity for large-sale designs, and supporting mapping tools to make the technology accessible to end users.

**Tangible benefits to end users:** Over a 15x improvement in energy-efficiency, considerable reduction in chip area (3-4x), and considerable improvement in performance (> 2x) compared to today's FPGA chips. Mapping tools will be developed to automatically map algorithms into hardware and abstract away hardware-specific details from end users.

**Critical technical barriers:** Hierarchical interconnect networks have been known to the academic and industrial community for a long time, but physical realization of these networks precluded their successful deployment. **The critical difficulty associated with the hierarchical networks is routing congestion during chip synthesis.** Leopard Logic, Inc, is one example of a company that failed to deploy hierarchical interconnect architecture. FPGA startups today, most notably Abound Logic, Tier Logic, Blue Chip Designs, and Achronix, provide customized solutions for increased logic density or speed, but they still don't solve the problem of power inefficiency associated with FPGA chip interconnects.

Main elements of the proposed technical approach: Our approach is based on alternating vertical and horizontal routing. LUTs (or any other processing elements) are partitioned in a 2-D floorplan with switch-boxes placed to allow full routability. An *N*-LUT design requires  $log_2(N)$  levels of switch-boxes. Simple example of N = 4 is shown below to illustrate the concept.



Alternating Vert./Horiz. routing: N = 4 LUTs example (2 levels)

**Figure 4:** Hierarchical Konda interconnect architecture.  $O(N \cdot \log_2 N)$  interconnect switches are required for full connectivity. Routing is fully deterministic.

In the case shown in Fig. 4, 2 levels of switch-boxes are required for N = 4 LUTs. LUTs with indices from 0 N/2 - 1 are placed on the left, the remaining LUTs are placed on the right. Switch-boxes are placed next to the LUT columns. Routing between elements with adjacent index is provided as a vertical connection (1<sup>st</sup> level routing); routing between elements with 2 indices apart is provided with a horizontal connection (2<sup>nd</sup> level routing). The routing continues in vertical/horizontal fashion for larger N.

**Basis of confidence:** Konda network architecture is a patent-protected technology that is recognized by many semiconductor companies including Cisco, Xilinx, Altera, and LSI Logic. To demonstrate the network in hardware, UCLA team has taped out 3 chips and successfully implemented variants of Konda network and also variants of processor-block features.

*Chip 1 (90nm, LUT-slice FPGA, concept demo):* A 1024-LUT FPGA was made in 90nm 9SF technology (Dec 2009 run). Our synthesis estimates predict a 250 mW of power and a 600 MHz maximum performance. The chip occupies 2.6 x 2.5 mm<sup>2</sup> in 90nm. *Status:* lab testing.

*Chip 2 (65nm, LUT and DSP slices, small scale)*: A 256-LUT 240-DSP 8-BRAM FPGA was made in 65nm technology (June 2010 run). The chip is aimed to show asymmetric network and heterogeneous computing blocks. The chip occupies 2.1 x 3.1 mm<sup>2</sup> in 65nm. *Status*: taped out.

*Chip 3 (45nm, DSP-slice FPGA, small scale):* A 512-DSP slice FPGA is made in IBM 45 nm SOI technology (June 2010 LEAP run). We expect power consumption below 500 mW. This design will be applicable to small-scale applications such as micro UAVs. *Status:* taped out.

**Nature and description of end results to be delivered to DARPA:** We will provide several deliverables to DARPA and DOD community as listed below.

- Interconnect architecture and routing tools (software).
- Hardware library in 32nm IBM SOI process (compatible with Cadence software).
- Routing software for the new interconnect architecture and hardware library (software).
- Chip demos of varying scale to demonstrate algorithms of interest to DOD (hardware).
- Tool flow for mapping algorithms onto FPGA chips (software).
- Demonstrations of HPC benchmarks using commercial technology.

The first three items in the list are intermediate steps towards the final hardware demonstration that also includes user-friendly mapping tool interface.

**Cost and schedule of the proposed effort:** \$2,374,111 over 3 years.

# 2.4. Technical Approach

**Problem Description:** FPGAs are used in many signal processing and computing applications. DOD mission capability or computing performance can be greatly improved with more energy efficient hardware. FPGA based solutions are very attractive due to their flexibility, similar to that of CPUs. This flexibility comes at a very high energy cost, as shown in Fig. 1.

Looking at the energy efficiency (the amount of energy per unit operation) for a variety of chips from different categories, we observe a 1,000x gap in energy efficiency between microprocessors and dedicated designs. The root cause of this is architectural. Processors have general ALU-type processing unit(s) and large amounts of memory to support time-muliplexing of instructions and data into and out of the ALU(s). Dedicated chips have a variety of processing units, but are very expensive in low-volume and can't be programmed, so they can't be used for HPC applications. General-purpose DSPs are a viable compromise between microprocessors and dedicated designs. Recently, however, FPGA chips have started to gain attention with their increased computing capabilities. Look-up-table (LUT) based chips have energy efficiency similar to that of CPUs and are not very attractive alternatives to CPUs (CPUs are easier to program). Many today's FPGAs have dedicated kernels such as DSP slices, ARM cores, etc. These FPGAs have energy efficiency similar to DSP chips, but they are still 30-50x worse than dedicated chips. The root cause of energy inefficiency in these FPGAs is their interconnect architecture.

Today's FPGAs use 2D-mesh interconnect architecture shown in Fig. 5. Interconnect consists of switch boxes (shown as cross-points), connection points for the buses, and bus drivers (buffers). This architecture is not very scalable: it requires  $O(N^2)$  interconnect switches for *N* LUTs. For 1k processing units, this means 1M switches! To overcome this complexity issue, designers employ heuristics to reduce the number of switches. One of the ideas is to reduce connectivity around the edges, as shown in Fig. 5. Another idea is to reduce top-level connectivity in large designs and utilize local connections. These approaches are heuristic and lead to inefficient utilization of hardware resources. Readers may be have experienced that utilizing more than 80% of FPGA resources without sacrificing performance is a big challenge in commercial FPGA systems.



Figure 5: 2D-mesh interconnect architecture.  $O(N^2)$  interconnect switches are required for full connectivity. Heuristics are used to reduce the network complexity. These heuristics result in non-deterministic routing.

Even after reductions in network complexity, interconnect still occupies over 75% of area in today's FPGAs. For example, Xilinx Virtex-5 chip has 1.1B transistors; 275M are used for logic, 875M (80%) are used for interconnect. Most of FPGA power is dissipated by the interconnect, as shown in Fig. 6. Further simplifying interconnect (without sacrificing connectivity) would have multiple benefits. First, the interconnect power will decrease. Second, due to reduced interconnect area, overall chip area will also reduce. Third, since the chip area is reduced, the size of wires (and wire capacitance) also reduces. The reduction in wire length and complexity implies further reduction in power. It also implies



**Figure 6:** Power breakdown in a Virtex-5 FPGA.

improvements in performance. This excess performance can be traded for increased energy efficiency, or simply used to improve computational efficiency. Finally, we benefit from reduced clock power since the clock is now distributed over a smaller area. Therefore, reduction in interconnect complexity is crucially important for improved computing power and performance.

**Proposed Network Architecture:** In response to the interconnect challenge, we propose to use a proprietary Konda hierarchical interconnect architecture. This interconnect architecture has greatly reduced complexity,  $O(N \cdot \log_2 N)$ , and it is based on fully deterministic routing. The concept of Konda network is to use simple unidirectional switches and 2x1 multiplexers to hierarchically connect the computing resources (LUTs, DSP slices, ARM IP, etc.).

Eliminating routing congestions and making the 2D circuit layout possible are the key enabling features of the Konda network. An example of N = 8 LUT design with Konda network is shown in Fig. 7. For complete routing  $log_28 = 3$  levels of switch matrices are needed. First, vertical tracks connect nearest LUTs, then horizontal tracks are used to connect LUTs at the next level, and finally vertical tracks are used to connect the last level of switches. This structure has full connectivity and completely deterministic 2D routing.



Figure 7: Konda interconnect network architecture and routing tracks for N = 8 LUTs.

The benefits of this network architecture were evaluated using Toronto20 benchmarks. Toronto 20 benchmark suite originated from an FPGA place-and-route challenge that was set up by University of Toronto Researchers [1] to encourage FPGA researchers to benchmark their software design tool chains on large circuits These 20 benchmarks are from real designs and the placed netlists are provided - for a given FPGA logic block consists of a 4-input look-up table (LUT) and a flip flop - to experiment with different routing architectures and routing algorithms. The existing results are experimented with 2D-Mesh network based routing network by providing partial bandwidth i.e., with different switch-box flexibility, connection-box flexibility and a certain number of channels. Konda hierarchical network is also experimented with partial bandwidth provisioning and the results are compared on various dimensions such as 1) number of cross points, 2) route length (delay) 3) performance 4) speed of routing and 5) routability. Konda hierarchical network performed better in several easily-measureable ways and the results are presented in Tables 2 and 3.

				2D-Mesh Network		Konda Hierarchical Network		
Toror	nto20 Benc	hmark Info	ormation	Simulation		Simulation		
				(Bidirectio	onal wires)	(Unidirectional wires)		
			Number of	Max	Total	Total	Souings	Cross-
Name	Size	LUTs	Number of	channel	cross-	cross-	factor	points
			connections	width	points	points	Tactor	saved
alu4	40	1600	1514	9	177,174	58,737	3.02	118,437
apex2	44	1936	1875	10	237,660	83,180	2.86	154,480
apex4	36	1296	1243	11	175,890	52,482	3.35	123,408
bigkey	54	2916	1694	6	213,876	54,643	3.91	159,233
clma	92	8464	8302	10	1,026,780	359,846	2.85	666,934
des	63	3969	1347	7	338,730	57,044	5.94	281,686
diffeq	39	1521	1497	7	131,082	49,275	2.66	81,807
dsip	54	2916	1309	5	178,230	40,972	4,35	137,258
elliptic	61	3721	3604	9	408,510	129,507	3.15	279,003
ex5p	33	1089	1019	11	148,170	44,609	3.32	103,561
ex1010	68	4624	4588	9	506,790	192,391	2.63	314,399
frisk	60	3600	3556	11	483,186	134,686	3.59	348,500
misex3	38	1444	1383	10	177,900	58,866	3.02	119,034
pdc	68	4624	4535	15	844,650	239,484	3.52	605,166
s298	44	1936	1929	6	142,596	63,956	2.23	78 <i>,</i> 640
s38417	81	6561	6349	6	478,260	207,457	2.30	270,802
s38584.1	81	6561	6291	7	557,970	184,030	3.03	373,940
seq	42	1764	1717	10	216,780	73,880	2.93	142,900
spla	61	3721	3644	12	544,680	171,676	3.17	373,004
tseng	33	1089	975	6	80,820	31,599	2.56	49,221

 Table 2: Comparison of 2D-Mesh and Konda interconnect networks using Toronto20 behcmarks.

The benefits of Konda hierarchical network over 2D-Mesh network using Toronto20 Benchmarks are summarized in Table 4. Various configurations of Konda hierarchical network were tested for each benchmark and the results are verified as follows:

• All 20 benchmarks were routed by our algorithms in our network,

- Switches required to route was reduced significantly,
- Fundamental routing algorithms are proven,
- Speed of routing is proven,
- Benchmarks were profiled for Bandwidth requirements.

**Table 3:** Comparison of 2D-Mesh and Konda interconnect networks using Toronto20 behcmarks. Inaddition to considerable savings in the number of cross-points, Konda network uses has far betterpercentage utilization (fewer % is better) than the 2D-Mesh network.

Toroi Bench Inforn	Toronto20 Benchmark Information		2D-Mesh Network Simulation (Bidirectional wires)		Konda Hierarchical Network Simulation (Unidirectional wires)		Other Key the Sim	Results of ulation
Name	Size	Max Ch Width	Total Cross-pts	Total Cross-pts	Savings factor	Cross-pts saved	% Cross- pts used Konda	% Cross- pts used 2D-Mesh
alu4	40	9	177,174	58,737	3.02	118,437	7.9	66
apex2	44	10	237,660	83,180	2.86	154,480	9.3	70
apex4	36	11	175,890	52,482	3.35	123,408	9.4	60
bigkey	54	6	213,876	54,643	3.91	159,233	4.1	51
clma	92	10	1,026,780	359,846	2.85	666,934	8.1	71
des	63	7	338,730	57,044	5.94	281,686	2.9	33
diffeq	39	7	131,082	49,275	2.66	81,807	6.9	75
dsip	54	5	178,230	40,972	4,35	137,258	2.8	46
elliptic	61	9	408,510	129,507	3.15	279,003	7.0	63
ex5p	33	11	148,170	44,609	3.32	103,561	9.5	60
ex1010	68	9	506,790	192,391	2.63	314,399	8.4	75
frisc	60	11	483,186	134,686	3.59	348,500	7.5	56
misex3	38	10	177,900	58,866	3.02	119,034	8.7	66
pdc	68	15	844,650	239,484	3.52	605,166	10.5	56
s298	44	6	142,596	63,956	2.23	78,640	7.1	89
s38417	81	6	478,260	207,457	2.30	270,802	6.0	86
s38584.1	81	7	557,970	184,030	3.03	373,940	5.3	66
seq	42	10	216,780	73,880	2.93	142,900	9.0	68
spla	61	12	544,680	171,676	3.17	373,004	9.3	63
tseng	33	6	80,820	31,599	2.56	49,221	6.7	78

**Table 4:** Summary of the benefits of Konda hierarchical network. Analytical and empirical results are shown, the numbers are relative to 2D-Mesh network.

Criteria	Analytical	Empirical
Interconnect area	At most 1/3	At most 1/3
Connectivity	2-3x	2-3x
Interconnect Power	1/5 to 1/10	1/5 to 1/10
Interconnect Latency	1/5 to 1/10	1/5 to 1/10
Speed of compilation	Significantly faster	Significantly faster
Scalability across process generations	Close to linear	Close to linear

The conclusions of simulation of Toronto20 benchmarks using Konda hierarchical network matched the benefits derived in empirical analysis. The generic routing tool created for Konda hierarchical network delivers consistent and predictable results. Based on the Toronto20 benchmark results it can be projected that the gap between ASIC's and FPGA's can be closed as shown in Fig. 8, which would significantly improve performance and energy efficiency of HPC hardware. In the proposed work, we will explore further technology improvements.



**Figure 8:** Konda interconnect network architecture has substantial benefits over today's FPGAs. It is projected to have ASIC-like energy efficiency, power, and performance. Such energy-efficiency levels are more than 100x better than general purpose processors.

# 2.4.1. Network Architecture and Routing Tools

We will next work on homogeneous and heterogeneous networks featuring arbitrary level of connectivity. The decision about the connectivity level will be aided with feedback from the mapping tools (Task 6) in order to minimize hardware utilization.

**Task 1) Routing Architectures for Homogeneous Blocks:** Routing tool will be developed for the FPGA with homogeneous blocks. Routing algorithms need to be developed for uni-terminal nets and multi-terminal nets. The hierarchical routing network may be a symmetric network where the number of inputs and the number of outputs are the same. The routing network may also be asymmetric network where the number of inputs and the number of outputs are the same. The routing network may also be asymmetric network where the number of inputs and the number of outputs are not the same. Rearrangeably nonblocking and strictly nonblocking multi-terminal net algorithms will be implemented to demonstrate the routability and the speed of routing. Routing algorithms need to be implemented for configurations of Konda hierarchical network where some of the stages in the network may be partially connected and the other stages are fully connected. The LUT size of the network may be a perfect power of two or non-perfect power of two.

**Task 2) Routing Architectures for Heterogeneous Blocks:** We will also explore interconnect architectures suitable for heterogeneous blocks. The key architectural challenge is to adapt the Konda hierarchical network for FPGA architecture. A fully connected hierarchical network is an over-kill for FPGA applications. Our goal is to converge on the appropriate design of the routing network in three phases and also adopt it to many different applications end-user applications. Also we need to experiment with many varieties of hierarchical network designs such as Benes

network, butterfly fat-tree network and other optimizations related to properties of FPGA designs. One aspect is to analyze the locality typical in FPGA designs and optimizing or adopting Konda hierarchical network with optimum bandwidth for local connectivity and global connectivity. The typical LUT size that is known to be optimal in a 2D-Mesh routing network may not be optimal for Konda hierarchical network. This is because Konda hierarchical network provides richer connectivity with smaller switch and less number of tracks. Determining the appropriate length of the tracks is another aspect that will be addressed in this task.

# 2.4.2. Hardware Design

To fully demonstrate the benefits of the proposed interconnect architectures and routing tools, we will implement the network architectures on a series of chips. Hardware design tasks will concentrate on achieving two goals: 1) hardware demonstration of power, area, and performance benefits, 2) development of automated chip routers to facilitate technology transition.

**Task 3)** Chip Demonstrations: Multiple chip demonstrations are planned to further quantify the benefits of the interconnect technology, and to further optimize interconnects based on hardware experiments.

*Prior Work:* We have designed several chips prior to this program, as summarized in Table 5. This was a self-initiated self-supported work. The results of IBM-90 and TSMC-65 chips will be available in September 2010. The results will be made available to the OHPC community.

Chip ID	Features	Area	Power	Performance	Status
IBM-90	1k LUTs	6.5 mm <sup>2</sup>	250 mW	300-600 MHz	Lab testing
TSMC-65	256 LUTs 240 DSPs 1 BRAM	8 mm <sup>2</sup>	500 mW	400-700 MHz	Taped out 6/2010
IBM-45-SOI	512 DSPs	4.4 mm <sup>2</sup>	500 mW	500-800 MHz	Taped out 6/2010

**Table 5:** Summary of FPGA chips built prior to the OHPC program.

The chips summarized in Table 5 are an initial demonstration of hardware feasibility of the interconnect network. The chips also demonstrate the integration of heterogeneous blocks (LUT, DSP, BRAM) for small-scale examples. Before describing the features of proposed chips, below is the description of design methodology used in prior work.

Figure 9 illustrates hierarchical design approach that starts with switch-matrix design. The switch-matrix blocks are custom-designed to allow tiling and hierarchical expansion. Design techniques used here will become cornerstones for the automation (Task 4).



Figure 9: Herearchical design procedure starting with switch-matrix design, integration of a slice, a hierarchical macro, and chip-level integration.

The IBM-90 chip illustrates LUT-only design. Consistent with predictions from Tables 2 and 3, Konda network achieves at least 3x lower interconnect area. Conventional chips have over 75% of interconnect area. In our chip, we have 50% logic (LUTs) and 50% switches (wires), which confirms the 3x interconnect area reduction estimate. Even with nearly-full connectivity, our chip has only 50% of area of the interconnect (3x less than commercial).



Figure 10: Detail of switch-matrix block (left), Energy-delay optimization of LUT macros (right).

Figure 10 shows the detail of the switch-matrix block. Local I/O connections allow tiling of layout macros, while pins in the middle are being used for hierarchical routing. Plot on the right shows energy-delay optimization after gate sizing and supply voltage tuning. We designed for 0.9 V supply (corresponding to the nominal/slow corner). The design is based on the sensitivity optimization methodology [6] that balances impact of all tuning variables. At a solution point, sensitivities to all variables are equal. In the energy-delay space, this means that the energy-delay lines obtained by tuning individual variables around a design point should be tangent. This is shown in the final design, where the  $V_{DD}$  and sizing (W) lines have similar slopes at  $V_{DD} = 0.9$  V. With these optimizations being made at the circuit level, we ensure that power efficiency considerations are propagated from system level down to the technology level. Our 1k FPGA is estimated to consume total of 250 mW of power when fully utilized. The energy efficiency per CLB slice is 0.96 pJ at 0.9V. We also performed deep pipelining to maximize performance. Synthesis estimates show a 600 MHz performance. This performance is significantly better than 450 MHz achieved by Xilinx Virtex-5 parts (Virtex-5 is built in a faster 65 nm technology).



Figure 11: Switch-point in 2D-Mesh network highlighting bidirectional nets (left), Switch-matrix in Konda network (right) highlighting unidirectional nets.

A very important feature of Konda interconnect architecture is that it uses only unidirectional wires, as shown in Fig. 11. The 2D-Mesh architecture uses bidirectional nets as shown on the left. Going from bidirectional to unidirectional nets results in a lower switching net capacitance. Implementation of the Konda switch-matrix is done with simple 2x1 muxes.

After demonstrating the feasibility of Konda network architecture on chip, the next bit challenge is to support the integration of heterogeneous blocks (LUTs, DSPs, BRAMs, etc.). Also, one should explore irregular switch-matrix architecture to reduce top-level wiring. An irregular switch-matrix design shown in Fig. 12 is implemented on the TSMC-65 chip (Table 5).



**Figure 12:** Irregular switch-matrix architecture to reduce top-level wiring (left), Chip demonstration of LUT, DSP, and BRAM modules using the irregular SM architecture (right).

The idea implemented in the network architecture from Fig. 12 is to use full connectivity only near the center of the chip. In our previous FPGA designs, long wires are routed across the entire chip to connect the bordering switch matrices at the topmost level. The drawback of these long wires is the requirement of many buffer insertions, consuming excessive power and routing area. The irregular switch matrix architecture reduces the number of top-level buffers by 95% since the bordering switch matrices now route through the center switch matrices to connect to the other side.

*Proposed Work:* The chips summarized in Table 5 are just an initial effort towards optimized FPGA implementation. This proposal will focus on hardware designs with reduced-complexity irregular network architectures developed in Tasks 1 and 2. Together with mapping tools from Tasks 5 and 6, we will be able to minimize level of connectivity required for chip-level routing. This will be demonstrated in actual hardware designs as explained below.

We plan to make use of the DARPA LEAP program for chip fabrication. We will demonstrate medium-scale designs for which regularity of layout cells, in addition to architecture regularity, will play a key role in improving tolerance to manufacturing variations. We will thus explore designs with regular layout geometries in IBM's 32nm SOI process. The evaluation of circuit metrics as compared to standard-cell based CMOS design will result in a library of FPGA building blocks which will be used for chip demonstrations. In addition to the IBM library cores and routing tools (Task 4), we will also consider TSMC libraries due to their general use and to provide additional options for technology transition.

In CMOS run A (see Sec 2.8.1), we will test the use of library IP for the integration of mediumscale FPGA processor (5K LUTs). Since 32nm SOI process is new and not yet fully tested by designers, we will work with homogeneous LUT-only design to minimize the risk of potential manufacturing and design issues. This chip is a 10x larger than the IBM-90 and will allow us to confidently explore mapping algorithms for this level of design complexity.

In CMOS run B (see Sec 2.8.1), we will design a chip with 15K DSP slices. The chip will be compliant with FMC expansion modules (160 I/O pins). This design will be able to support DSP-centric applications such as signal and information processing. The chip will be tested using BEE4 module (coming out in Fall 2010) from BEEcube. Such setup will allow us to do side-by-side comparison with Virtex-6 Xilinx FPGAs. We will work with BEEcube on HPC application benchmarking and will also welcome inputs from the DOD community.

In CMOS run C (see Sec. 2.8.1), we will demonstrate LUT/DSP/BRAM based design with over 15K LUTs, over 15K DSP slices, and adequate BRAM memory. The chip will be also compliant with FMC for testing with the BEE4 module. Chips B and C will make use of irregular network architecture and optimized connectivity as described in Task 4. The chips from CMOS run C will be used for inter-module communication with multiple BEE4 boards to show expandability to large HPC benchmarks.

**Task 4)** Automated Chip Routing Tools: To facilitate the integration of medium- and largescale FPGA chips, and to enhance our technology transition capabilities, we will work on automated chip routing tools. The tools are intended to automate custom design strategies developed in our prior work. We will also automate design techniques further developed under Task 3, particularly CMOS runs A and B (see the scheduling chart in Sec. 2.8.1).

Advanced routing tools will need a library of switch-matrix blocks with varying degree of connectivity. Figure 13 shows example of full- and half-connectivity cells as well as full-to-half connectivity interface cells. The use of these simple cells, and others, will enable us to support network routers (developed in Tasks 1 and 2) with arbitrary level of connectivity.



Figure 13: Switch-matrix (SM) blocks include full-connectivity (left), full-to-half-connectivity (middle) and half-connectivity (right) features.

An example of cell design for future automated routing is shown in Fig. 14. In the chip shown in Fig. 12, DSP slices have to be designed with fixed width due to size constraints from configuration SRAM blocks. In our architecture, we use wordline (WL) to drive SRAM modules on both sides (left and right of the WL circuits). WL routing is done in metal 3 (M3) as shown in Fig. 14. We must allow M3 tracks for neighboring SM. The routing channel for 7 bits of SRAM



**Figure 14:** Switch matrix design showing detail of SRAM routing tracks. Upper SM (right) is a mirrored-version of the lower SM (left). Alternating 4/3 bits per row and custom muxes are used to facilitate 7-bit routing of SRAM configuration bits.

is made using alternating 4/3-bit horizontal tracks. We also make use of custom muxes to reduce redundant input inverters (details not shown on the figure). These concepts will be automated.

Automation of other routing tasks, in addition to the one illustrated in Fig. 14, will be performed. The outcome of Task 4 will be the router that can take arbitrary number of LUT, DSP, and BRAM cells and, for a given level of connectivity, provide routed chip that implements optimal network architecture developed in Task 2. This kind of routing capability will allow us to customize chip features and rapidly construct energy-efficient FPGAs for HPC applications (analogous to different Xilinx chip families, for example). The automated routing will also provide chip design community a tool for the utilization of our library macros. We will maintain library of macros in IBM 32nm SOI technology (for DOD community) and also TSMC 32/28nm technology (for DOD and commercial use).

# 2.4.3. Hardware Mapping

For an FPGA to be effectively used by its consumers, an automated mapping tool must be provided as well. Mapping tool for commercial FPGAs are provided to convert user-provided Verilog or VHDL into a gate-level design, and automatically place-and-route these gates onto the FPGA. As a result, the user has complete automation from Verilog/VHDL to a functional FPGA. The mapping and place-and-route software is a crucial component of this project, and major efforts ought to be allocated to provide an efficient tool.

**Gate-level Synthesis:** The first step of the mapping process is to create a mappable design from the Verilog or VHDL input. The process is called logic synthesis, an intricate procedure requiring complex algorithms.

To optimize our resource usage, we plan to adopt commercial synthesis tools such as Synopsys® Design Compiler or Cadence® RTL Compiler. Both these tools operates on a standard cell library that contains information regarding the timing and functional characteristics of each logic gate. The tool then converts the input design into a netlist consisting of gates from the standard cell library.

The main task here is to create a standard cell library mappable to our custom FPGAs. If the FPGA is constructed from 4-input LUTs, then the standard cell library ought to include different combinations of 2, 3, and 4-input logic gates. The FPGA mapper can then determine the appropriate values to program to each LUT based on the logic gate.

*Netlist Optimization:* Although the synthesized netlist is fully functional, it may not be optimal for our FPGA applications. The logic netlist should therefore be optimized by the mapper for area reduction and speed improvements. This is the second step of the mapping process.

In modern FPGAs, the majority of the area and delay are attributed to interconnects between logic gates. Therefore it is beneficial to maximize the logic function of each LUT instead of spreading the same function over many LUTs. To achieve such optimality, logic gates with less than 4 inputs are searched for logic recombination with its neighboring gates:



Figure 15: Illustration of gate-level synthesis.

Shown are two 2-input gates drive another 2-input gate. These 3 gates can be combined to a 4-input gate to fit into a single LUT instead of spreading over 3 LUTs, wasting both logic and routing resources. More strictly speaking, any sequential logic gates with a total unique input of 4 or less can be combined into a single 4-input LUT.

In our FPGAs, a 4-input LUT can be reprogrammed as two 3-input LUTs, where 2 of the inputs are shared. Two 3-input gates can potentially share a single LUT as long as the number of unique inputs is 4 or less. Two 2-input gates can always share a LUT as well. The mapper tool can exploit such feature during the placement process, as mentioned later.

**Task-5) Place & Route Algorithms and Tools:** Once the netlist has been optimized, place and route can begin. Each logic gate ought to be placed before it can be routed. Finding a suitable placement for each logic gate can greatly affect the performance of the final design, and poses a significant challenge to the tool.

The placement process is divided into coarse placement, which determines the gate partition, and fine placement, which determines the exact gate location in each partition. Gate partition takes place first, and the goal is to partition the gates into quadrants to minimize cross-partition

interconnects. This process can be modeled by an optimization problem, where the cost is total number of wires crossing horizontal, vertical, and diagonal boundaries. A penalty cost can be added to ensure even distribution of gates across partitions.

In most convex optimization problems, only the local minima can be determined based on the initial condition. Since an optimal initial condition cannot be determined, and the cost function may contain numerous local minima, a simulated-annealing based algorithm is used for gate partition. Based on the size of the design, numerous locally-optimal solutions can be found, and the lowest-cost solution is chosen.



Figure 16: Mapper for hierarchical-interconnect FPGAs.

Once the gate partition is determined, first level of fine placement can proceed. The longest interconnects are those requiring diagonal connections across the area. These gates are placed near the center of the area, where the diagonal connections are the shortest. Gates with vertical and horizontal connections are then placed near their respective edges to minimize their wire connections.

Hardware Routing: Since the routing resources in FPGA are limited, routing should occur at the same time as placement. This prevents a gate from being placed at a nonroutable location. When a gate is being considered for a location, all of its inputs and outputs should be located. The input and output gates that are already placed must be able to route to this location, else a new location must be determined. In cases where more than one possible routes exists, each routing candidate should be evaluated for interconnect length, and the shortest interconnect is chosen.



Figure 17: Recursive execution of sub-level PnR.

Automated Place and Route: The place-and-route tool can utilize recursive placement. Since all the partition-crossing gates are placed, all the non-placed gates in each partition are local to that partition, meaning they do not connect to any gates outside the partition. For each quadrant, the aforementioned place and route occurs again to place gates within the subquadrants, until all the quadrants (and subquadrants) are processed.

In some cases, a gate may not be placed inside the chosen partition, then other partitions at the level are considered for placement. If none of the partitions can accept the gate to be placed, a higher hierarchy is searched for placement.

**Task 6) Tools for Hardware Mapping:** The final step of the place-and-route process is the output the design. This step creates the exact bit sequence required to program the scan chain on the FPGA, which configures all the LUTs and interconnects to create the programmed function.

For each of our FPGA, we have knowledge of the exact bit location for each of the LUT configurations, as well as the switch-matrix configurations. For the logic block shown below, four LUTs are programmed as one configurable logic block. The scan-chain (SI) first controls internal configurations (such as 4-input/3-input mode, carry-chain propagation, register output, etc), and then controls each of the four LUT values. The corresponding configurations from the place-and-route output are then mapped to these bits on the scan-chain. The switch matrix bistream is determined in the similar fashion since all the interconnect configurations are already known from place-and-route.

**Task 7) Demonstrations and Technology Transition:** We will use chips from CMOS runs B and C (see Sec 2.8.1) to demonstrate the benefits of our hardware as compared to Xilinx Virtex-6 chips by using the BEE4 module from BEEcube.

*Collaboration with BEEcube:* We will work with BEEcube (support letter attached at the end of Volume I) to do technology demonstration and initial transition to HPC applications. We will make use of future BEE4 platform consisting of 4 Virtex-6 FPGA chips (LX240 family) and featuring FMC interface (160 pins/chip). A custom PCB board with our FPGA will comply to the FMC interface specifications. BEE4 FPGA chips will be used to execute computations on our custom chips, as shown in Fig. 19. This setup will also allow side-by-side benchmarking of Virtex-6 FPGAs and our chips.



**Figure 18:** Configuration bits for a Slice-L (4 LUTs) block.



**Figure 19:** Technology demonstration and initial transition using BEEcube HPC technology.

To further demonstrate inter-module communication and scalability to larger systems, we will connect 4 BEE4/Chip modules, as shown in Fig. 19. We will use BEEcube HPC benchmarks for initial evaluations. We welcome inputs from DOD community and other teams in the OHPC program about example algorithms.

In addition to BEEcube benchmarks, we will explore parallel execution of neural spike sorting algorithms from UCLA Medical School, Department of Neurology. We have data from our scientific collaborators (Prof. Rick Staba, and Prof. Chris Giza) who monitor neural action potentials in human patients with acute epilepsy. Data recording (64 channels, 10 bits, 28 kS/s) over 3 weeks aggregates 2 TB of data per patient and presents an extreme signal processing challenge. We have tried to process one hour of data (60 GB) using a CPU cluster with 40 3 GHz processors. Although we sample at 30 kHz, the total run-time was 68 minutes due to sequential nature of CPU processing. When the same algorithm was mapped to an FPGA running with a 300 MHz clock, we estimated processing time to be 0.4 seconds. Since we can execute complete algorithm iteration in one clock cycle, we get a 10,000x improvement in execution (300 MHz / 30 kHz). The bottleneck is how fast we can feed the data into the FPGA, not the speed of the FPGA itself. We would like to collaborate with other teams in the OHCP program who are exploring storage bandwidth issues.

Future Implications: Upon successful demonstration with BEEcube, the technology can be further scaled up to a rack system as shown in Fig. 20. A number of FPGA nodes will be used to execute common operations. Optionally, one could have a small number of general-purpose CPU blades for non-standard operations. This configuration will require software development to abstract away hardware details from the user and can be a topic addressed by other OHPC teams. Finally, Konda network can also be applied for efficient rack-to-rack connectivity to produce ultralow-power and ultra-highperformance HPC systems for extreme computations.



**Figure 20:** Future possibilities with our technology: FPGAs can be used in rack-scale systems (left), .Konda network technology can also provide superior rack-to-rack connectivity for server farms (right).

# 2.5. Statement of Work (SOW)

We propose to develop energy-efficient programmable hardware and supporting mapping tools. The hardware is based on proprietary Konda interconnect architecture that provides significant reduction in interconnect complexity as compared to existing FPGAs. The new architecture and supporting tools are projected to provide over a 15x improvement in energy efficiency while also considerably reducing chip area and improving performance.

# Task 1: Architectures for Homogeneous Blocks (Lead: Konda) \$200K, Q1-Q3

*Objective:* Define interconnect architecture and routing tools for designs with a given number of homogeneous blocks such as LUT or DSP slices.

*Approach:* Rearrangeably nonblocking and strictly nonblocking multi-terminal net algorithms will be implemented to demonstrate the routability and the speed of routing. Routing algorithms need to be implemented for configurations of Konda hierarchical network where some of the stages in the network may be partially connected and the other stages are fully connected. The LUT size of the network may be a perfect power of two or non-perfect power of two.

*Exit criteria:* When the chosen interconnect architecture demonstrates the optimal tradeoff between interconnect size and performance (as quantified by Toronto20 benchmarks).

Deliverables: Architectural diagram of the interconnect structure (block-level schematic).

# Task 2: Architectures for Heterogeneous Blocks (Lead: Konda) \$200K, Q3-Q6

*Objective:* Define interconnect architecture and routing tools for designs with a given number of heterogeneous blocks such as a combination of LUT, DSP slices, memory, and IP elements.

*Approach:* We will make use of the infrastructure developed for Task 1 and customize interconnects local to each type of heterogeneous block.

*Exit criteria:* Interconnect architecture with optimal size-performance tradeoff as quantified by Toronto20 benchmarks.

Deliverables: Architectural diagram of the interconnect structure (block-level schematic).

# Task 3: Chip Demonstrations (Lead: Markovic) \$900K, Q1-Q11

*Objective:* To demonstrate the feasibility of the interconnect architecture on different design complexity and compare performance and power with commercial FPGAs.

*Approach:* We will design and verify three FPGA chips with increasing level of complexity (5K, 15K, 45K LUTs). The chips will be constructed from custom-designed macros, including processing, memory, and interconnect blocks.

*Exit criteria:* Hardware demonstration of superior performance (>2x), energy (15x), and area metrics (>2x) as compared to commercial FPGAs.

Deliverables: Library of FPGA macros, chip demonstration results.

# Task 4: Automated Chip Routing Tools (Lead: Markovic) \$100K Q4-Q5

Objective: To automatically place-and-route an FPGA according to hardware requirements.

*Approach:* Library of FPGA macros for the technology of interest, automatic Verilog generation, supporting scripts for chip synthesis.
Exit criteria: Design and layout of an FPGA using the automated toolflow.

*Deliverables:* Scripts for synthesis and tools for automatic Verilog generation (that use the library of FPGA macros delivered in Task 3).

#### Task 5: Place and Route Algorithms and Tools (Lead: Markovic) \$300K, Q1-Q6

*Objective:* Determine the optimal physical location and routing for each LUT / macro to maximize resource utilization and chip performance.

*Approach:* Partition LUT / macro blocks into sets with coarse and fine levels of granularity. Place and route gates for minimum interconnect delay. The algorithm repeats hierarchically until all levels of interconnect are placed and routed.

*Exit criteria:* Successful place-and-route of Chip A (5K LUTs) to its full interconnect utilization. *Deliverables:* Software demonstration of an automated place-and-route flow.

#### Task 6: Tools for Hardware Mapping (Lead: Markovic) \$100K, Q3-Q4

Objective: To map the place-and-routed design to a bitstream format for FPGA programming.

*Approach:* Creates the exact bit sequence based on the scan chain configuration implemented on chip. This configures all the LUTs and interconnects to execute the programmed function.

Exit criteria: Successful programming of scan chain with demonstrated functionality.

Deliverables: An automated tool for bitstream programming.

#### Task 7: Demonstrations and Technology Transition (Lead: Markovic) \$500K, Q7-Q12

*Objective:* Demonstrate the benefits of our technology for HPC applications.

*Approach:* Use chips from CMOS runs B and C with the BEE4 platform (consisting of four Virtex-6 FPGA chips) to perform HPC benchmarking. The BEE4 FPGA chips will be used to execute computations on our custom chips and perform side-by-side comparison of Virtex-6 FPGAs against our chips.

Exit criteria: Functional HPC platform based on BEE4 and custom FPGA chips.

*Deliverables:* Results from HPC benchmarking, hardware demonstration of performance and energy efficiency on BEE4-based platform.

Task / Cost	Year 1	Year 2	Year 3	Total
Task 1: Homogeneous Architecture	\$200K	-	-	\$200K
Task 2: Heterogeneous Architecture	-	\$200K	-	\$200K
Task 3: Chip Demonstrations	\$350K	\$350K	\$200K	\$900K
Task 4: Automated Chip Routing	\$50K	\$50K	-	\$100K
Task 5: Place-and-Route Algorithms and Tools	\$200K	\$100K	-	\$300K
Task 6: Tools for Hardware Mapping	\$100K	-	-	\$100K
Task 7: Demonstrations and Technology Transition	-	\$100K	\$400K	\$500K

Table 6: Task Cost and Schedule
---------------------------------

#### **2.6. Intellectual Property**

Venkat Konda has filed 10 patent applications and assigned them to Konda Technologies Inc. More patent applications are in the pipeline. The following is the complete list of patent applications:

- [1] Large Scale Crosspoint Reduction with Nonblocking Unicast & Multicast in Arbitrarily Large Multi-stage Networks
  - US Provisional Patent Application Number: 60/905,526
  - Date Filed: March 6, 2007
- [2] Fully Connected Generalized Multi-stage Networks
  - US Provisional Patent Application Number: 60/940, 383
  - Date Filed: May 25, 2007
  - [2a] Fully Connected Generalized Multi-stage Networks
    - PCT Patent Application Serial Number: PCT/US08/56064
    - Date Filed: March 6, 2008
  - [2b] Fully Connected Generalized Multi-stage Networks
    - US Patent Application Serial Number: 12/530,207
    - Date Filed: September 6, 2009
- [3] Fully Connected Generalized Butterfly Fat Tree Networks
  - US Provisional Patent Application Number: 60/940, 387
  - Date Filed: May 25, 2007
- [4] Fully Connected Generalized Multi-Link Butterfly Fat Tree Networks
  - US Provisional Patent Application Number: 60/940, 390
  - Date Filed: May 25, 2007
  - [4a] Fully Connected Generalized Butterfly Fat Tree Networks
    - PCT Patent Application Number: PCT/US08/64603
    - Date Filed: May 22, 2008
  - [4b] Fully Connected Generalized Butterfly Fat Tree Networks
    - US Patent Application Number: 12/601,273
    - Date Filed: November 22, 2009
- [5] Fully Connected Generalized Rearrangeably Nonblocking Multi-Link Multi-stage Networks
  - US Provisional Patent Application Number: 60/940, 389
  - Date Filed: May 25, 2007
- [6] Fully Connected Generalized Strictly Nonblocking Multi-Link Multi-stage Networks
  - US Provisional Patent Application Number: 60/940, 392
  - Date Filed: May 25, 2007

- [7] Fully Connected Generalized Folded Multi-stage Networks
  - US Provisional Patent Application Number: 60/940, 391
  - Date Filed: May 25, 2007
  - [7a] Fully Connected Generalized Multi-link Multi-stage Networks
    - PCT Patent Application Serial Number: PCT/US08/64604
    - Date Filed: May 22, 2008
  - [7b] Fully Connected Generalized Multi-link Multi-stage Networks
    - US Patent Application Serial Number: 12/601,274
    - Date Filed: November 22, 2009
- [8] VLSI Layouts of Fully Connected Generalized Networks
  - US Provisional Patent Application Number: 60/940, 394
  - Date Filed: May 25, 2007
  - [8a] VLSI Layouts of Fully Connected Generalized Networks
    - PCT Patent Application Number: PCT/US08/64605
    - Date Filed: May 22, 2008
  - [8b] VLSI Layouts of Fully Connected Generalized Networks
    - US Patent Application Number: 12/601,275
    - Date Filed: November 22, 2009
- [9] VLSI Layouts of Fully Connected Generalized Networks with Locality Exploitation
  - US Provisional Patent Application Number: 61/252, 603
  - Date Filed: October 16, 2009
- [10] VLSI Layouts of Fully Connected Generalized and Pyramid Networks
  - US Provisional Patent Application Number: 61/252, 609
  - Date Filed: October 16, 2009

#### 2.7. Management Plan

Program management plan is shown in the figure below.



UCLA and KondaTech will closely work together to provide demonstrations and technology transfer to the DoD community. Below is a description of various tasks and their interaction.

**Routing Architectures:** KondaTech will be responsible for the development of interconnect architectures and supporting routing tools. This includes both homogeneous-block (Task 1) and heterogeneous-blocks (Task 2) architectures. UCLA will provide information about building blocks (e.g., LUT, DSP, memory) for Task 2. The developed architectures and routing tools will be benchmarked using standard Toronto20 FPGA benchmarks.

UCLA will be responsible for hardware design and hardware mapping efforts.

**Hardware Design:** Chip demonstrations (Task 3) will be initially made using theoretical routing architecture concepts developed at KondaTech. In this phase, we will investigate types of chip routing procedures that need to be automated during chip design. Custom routing of low-level interconnects will be enabled by the regularity of the interconnect architecture. KondaTech will provide input into automated chip routers (Task 4) to ensure that routing tools are properly transferred to chip design. We will then make use of the chip routers for the final chip design.

**Hardware Mapping:** UCLA will lead the hardware mapping effort, which goes in conjunction with chip design. Here, we will focus on developing algorithms (Task 5) that would optimally map algorithms to the newly developed interconnect architectures. We also have the capability to do wordlength optimization and high-level architecture transformations prior to mapping. The algorithm mapping will attempt to minimize resource utilization and power, and also maximize performance. The details of the mapping algorithms will be abstracted away from the user by developing mapping tools (Task 7) to provide automated algorithm mapping onto hardware.

**Demonstrations and Technology Transition:** We will demonstrate the execution of complete algorithms in order to validate power and performance gains of our technology. We will actively solicit and welcome inputs from the DOD community about the algorithmic examples that would best serve the demonstration of hardware and mapping tools.

**Demonstration Platform:** We will work with **BEEcube** to execute initial technology transition plan. BEEcube is now a well recognized provider of high-performance processing technology. The technology is based on FPGA hardware, a library of IP cores, and software development tools for high-performance computing and other applications. We will use BEEcube technology for our chip demonstrations. In particular, we will make use of the FMC expansion modules on their hardware platforms to control our chips. This includes programming of the chips and program execution. The 160-pin connection slots based on VITA Standard 57.1 provide BEE-to-chip interface. With this setup, we will be able to boost the performance of the existing hardware and quantify the benefits of our technology in actual computing environment. We will use four hardware units from BEEcube and setup inter-module communication in order to demonstrate scalability of our design (as described in Section 2.4).

#### Team size and composition:

UCLA team:

Dejan Markovic, PI

Four full-time graduate students to work on hardware design (Tasks 3 and 4), hardware mapping (Tasks 5 and 6), and technology demonstrations. They will also collaborate with Dr. Konda on Task 2. Each task will require the effort of two full-time students.

*KondaTech:* Venkat Konda, Consultant

UCLA and KondaTech will work closely to maximize the impact of our technology.

#### 2.8. Schedule and Milestones

Project schedule, task description, and management plan are described in this section.

#### 2.8.1. Schedule Graphic



Key deliverables from the program will include (due dates are referenced to the program start):

- (1) Layout library of FPGA building blocks in IBM's 32nm SOI technology (available through the DARPA LEAP program). *Due: after 4 months.*
- (2) Routing architectures for homogeneous blocks (e.g. LUTs) with varying degree of connectivity. Deliverable: software benchmarking results. *Due: after 8 months.*
- (3) Routing architectures for heterogeneous blocks (LUTs, DSP slices, BRAMs, other IP) with varying degree of connectivity. Deliverable: software benchmarking results. *Due: after 17 months.*
- (4) Test results from CMOS run A to demonstrate small-scale integration (< 5k LUTs). Deliverable: hardware measurements. *Due: after 15 months*.
- (5) Test results from CMOS run B to demonstrate medium-scale integration (< 10k LUTs). Deliverable: hardware measurements. *Due: after 26 months*.
- (6) Test results from CMOS run C to demonstrate large-scale integration (< 20k LUTs). Deliverable: hardware measurements. *Due: after 32 months*.
- (7) Tools for generating a bitstream for a mapped design. *Due: after 12 months*.
- (8) Tools for performing place-and-route of the optimized netlist for FPGA programming. *Due: after 15 months.*
- (9) Tools for integrating place-and-route tool with existing testing solutions. *Due: after 24 months.*
- (10) Hardware and tool flow demonstrations for relevant DOD algorithms; commercialization of technology. *Due: after 36 months.*

#### 2.8.2. Detailed Task Description

#### **Part I: Network Architectures and Routing Tools**

We will next work on homogeneous and heterogeneous networks featuring arbitrary level of connectivity. The decision about the connectivity level will be aided with feedback from the mapping tools (Task 6) in order to minimize hardware utilization.

#### **Task 1: Routing Architectures for Homogeneous Blocks**

Routing tool will be developed for the FPGA with homogeneous blocks. Routing algorithms need to be developed for uni-terminal nets and multi-terminal nets. The hierarchical routing network may be a symmetric network where the number of inputs and the number of outputs are the same. The routing network may also be asymmetric network where the number of inputs and the number of outputs are not the same. Rearrangeably nonblocking and strictly nonblocking multi-terminal net algorithms will be implemented to demonstrate the routability and the speed of routing. Routing algorithms need to be implemented for configurations of Konda hierarchical network where some of the stages in the network may be partially connected and the other stages are fully connected. The LUT size of the network may be a perfect power of two or non-perfect power of two.

#### **Task 2: Routing Architectures for Heterogeneous Blocks**

The key architectural challenge is to adapt the Konda hierarchical network for FPGA architecture. A fully connected hierarchical network is an over-kill for FPGA applications. Our goal is to converge on the appropriate design of the routing network in three phases and also adopt it to many different applications end-user applications. We will experiment with many varieties of hierarchical network designs. One aspect is to analyze the locality typical in FPGA designs and optimizing or adopting Konda hierarchical network with optimum bandwidth for local connectivity and global connectivity. The typical LUT size that is known to be optimal in a 2D-Mesh routing network may not be optimal for Konda hierarchical network. This is because Konda hierarchical network provides richer connectivity with smaller switch and less number of tracks. Determining the appropriate length of the tracks is another aspect that will be addressed.

#### Part II: Hardware Design

To fully demonstrate the benefits of the proposed interconnect architectures and routing tools, we will implement the network architectures on a series of chips. Hardware design tasks will concentrate on achieving two goals: 1) hardware demonstration of power, area, and performance benefits, 2) development of automated chip routers to facilitate technology transition.

#### **Task 3: Chip Demonstrations**

In CMOS run A, we will test the use of library IP for the integration of medium-scale FPGA processor (5K LUTs). This chip is 10x larger than the IBM-90 and will allow us to confidently explore mapping algorithms for this level of design complexity.

In CMOS run B (see Sec 2.8.1), we will design a chip with 15K DSP slices. The chip will be tested using BEE4 module (coming out in Fall 2010) from BEEcube. Such setup will allow us to do side-by-side comparison with Virtex-6 Xilinx FPGAs. We will work with BEEcube on HPC application benchmarking and will also welcome inputs from the DOD community.

In CMOS run C (see Sec. 2.8.1), we will demonstrate LUT/DSP/BRAM based design with over 15K LUTs, over 15K DSP slices, and adequate BRAM memory. The chips from CMOS run C will be also used for inter-module communication with multiple BEE4 boards to show expandability to large HPC benchmarks.

#### **Task 4: Automated Chip Routing Tools**

To facilitate the integration of medium- and large-scale FPGA chips, and to enhance our technology transition capabilities, we will work on automated chip routing tools. The tools are intended to automate custom design strategies developed in our prior work. We will also automate design techniques further developed under Task 3, particularly CMOS runs A and B.

#### Part III: Hardware Mapping

For an FPGA to be effectively used by its consumers, an automated mapping tool must be provided as well. The mapping and place-and-route software is a crucial component of this project, and major efforts ought to be allocated to provide an efficient tool.

#### **Task 5: Place and Route Algorithms and Tools**

Finding a suitable placement for each logic gate can greatly affect the performance of the final design, and poses a significant challenge to the tool. The placement process is divided into coarse placement, which determines the gate partition, and fine placement, which determines the exact gate location in each partition. Gate partition takes place first, and the goal is to partition the gates into quadrants to minimize cross-partition interconnects. This process can be modeled by an optimization problem.

Since the routing resources in FPGA are limited, routing should occur at the same time as placement. In cases where more than one possible routes exists, each routing candidate should be evaluated for interconnect length, and the shortest interconnect is chosen. We will hierarchically extend the partition and place-and-route within an automated flow.

#### **Task 6: Tools for Hardware Mapping**

The final step of the place-and-route process is the output the design. This step creates the exact bit sequence required to program the scan chain on the FPGA, which configures all the LUTs and interconnects to execute the programmed function. We will make the mapping tool compliant with existing FPGA design environments such as Xilinx XSG/EDK toolset.

#### Task 7: Demonstrations and Technology Transfer

We will use chips from CMOS runs B and C to demonstrate the benefits of our hardware as compared to Xilinx Virtex-6 chips by using the BEE4 module from BEEcube as a test platform for HPC benchmarking. A custom PCB board with our FPGA will comply to the FMC interface specifications. BEE4 FPGA chips will be used to execute computations on our custom chips. This setup will also allow side-by-side benchmarking of Virtex-6 FPGAs and our chips. To further demonstrate inter-module communication and scalability to larger systems, we will connect 4 BEE4/Chip modules in an evaluation platform.

#### 2.8.3. Project Management and Interaction Plan

UCLA and KondaTech will closely interact during the program. Our interaction plan consists of several means of communication:

- We will conduct weekly teleconference meetings using desktop sharing software
- Due to geographic proximity, Dr. Konda will visit UCLA once a month to meet with UCLA team and conduct detailed discussions about the project
- The PI will additionally interact with Dr. Konda regarding project management

Project management is illustrated in Section 2.7 and also summarized in Table 7.

Person / Task	Task 1	Task 2	Task 3	Task 4	Task 5	Task 6	Task 7
D. Markovic		+	+	+	+	+	+
V. Konda	+	+		+		+	+

KondaTech (V. Konda) will be responsible for the development of interconnect architectures (Tasks 1 and 2). UCLA will provide information about building blocks (e.g., LUT, DSP, memory) for Task 2 and assist in verification of interconnect architectures.

UCLA (D. Markovic) will be responsible for hardware design and hardware mapping efforts (Tasks 3 to 6). KondaTech will assist in automating chip routers (Task 4) to ensure integration of the interconnect architectures from Tasks 1 and 2. The expertise of KondaTech will also be used to transition mapping software to commercial tool environments.

UCLA and KondaTech will work together on HPC demonstrations described in Section 2.4.

#### 2.9. Personnel, Qualifications, and Commitments

**Dejan Markovic** is an Assistant Professor of Electrical Engineering at the University of California, Los Angeles. He completed the Ph.D. degree in 2006 at the University of California, Berkeley. His current research is focused on integrated circuits for emerging radio and healthcare systems, design with post-CMOS devices, optimization methods and CAD flows.

Prof. Markovic's research accomplishments include sensitivity-based circuit optimization [6] and DSP architecture optimization [11] for reduced power and area. As a demonstration of these concepts, his group has designed a number of complex digital chips for parallel data processing. Representative chips shown in Fig. 21 show performance range by 5 orders of magnitude (kHz to multi-GHz) and power density range of 3 orders of magnitude ( $\mu$ W/mm<sup>2</sup> to mW/mm<sup>2</sup>).



Figure 21: Sample chips designed by PI's group showing broad range in performance and ultra-low power density.

The PI's research in low power has been recognized by multiple awards:

- **2007 David J. Sakrison Memorial Prize** (Best Ph.D. Dissertation at UC Berkeley), awarded to the PI for his contributions to low-power digital circuit design.
- **2008 Outstanding MS Thesis Award**, UCLA EE Dept, received by R. Nanda, for her M.S. Thesis titled: "DSP Architecture Optimization in Matlab/Simulink Environment," June 2008.
- **2010 Outstanding MS Thesis Award**, UCLA EE Dept, received by V. Karkare, for his M.S. Thesis titled: "A 130 uW, 64-Channel, Spike-Sorting DSP Chip," Dec. 2009.
- 2010 DAC/ISSCC Student Design Contest Winner, awarded to Chia-Hsiang Yang for his paper titled "A 2.89mW 50GOPS 16x16 16-Core MIMO Sphere Decoder."

PI's group has unparalleled infrastructure for the design and optimization of digital chips, based on a number of tools developed and maintained over the past decade. PI's key publications in the area of low-power design are provided in [2-16].

**Venkat Konda** is an inventor, experienced entrepreneur and the CEO of Konda Technologies which he founded in 2007 based on a breakthrough layout using only horizontal and vertical tracks for Benes/BFT hierarchical networks, seminal rearrangeably and strictly non-blocking multicast routing algorithms with an architecture optimum with switch cost, power and performance. Venkat is currently in the process of commercializing the IP in FPGA interconnects, System-on-Chip interconnects and warehouse-scale datacenter switches. Prior to it, Venkat invented seminal algorithms for rearrangeably and strictly non-blocking multicast routing algorithms for Clos Networks and founded a startup Teak Networks, to commercialize into packet switch fabrics which are also applicable to design cheaper optical cross connects. Venkat received PhD degree in Computer Science & Engineering from the University of Lousiville, KY in 1992, and M.S in Electrical Engineering from the Indian Institute of Technology, Kharagpur in 1988.

Key Individual	Project	Pending/Current	2010	2011	2012
Dejan Markovic	HEALICs	Current (Co-PI)	176 hours	176 hours	176 hours
	STT-RAM	Current (Co-PI)	176 hours	176 hours	176 hours
	NEMS	Current (Co-PI)	176 hours	176 hours	176 hours
	NSF CAREER	Current (PI)	88 hours	88 hours	88 hours
	C2S2	Current (PI)	88 hours	88 hours	88 hours
	NVL	Pending (Co-PI)	176 hours	176 hours	176 hours
	FPGA	Proposed	176 hours	176 hours	176 hours
Venkat Konda	FPGA	Proposed	1000	1000	1000
			hours	hours	hours

Table of individual time commitments is provided below:

*Note:* Dejan Markovic is a Co-PI on three DARPA projects, two of which can directly benefit to the OHPC program. Namely, STT-RAM can be used for the realization of FPGA memory blocks. Also, zero-leakage NEM relay switches can be used for effective realization of FPGA switch-matrix blocks. PI's commitment to the proposed work is already evidenced by self-initiated and self-supported work described in Sec 2.4.1. The OHPC program will, therefore, not add a significant workload to the PI.

### 2.10. Organizational Conflict of Interest Affirmations and Disclosure

#### 2.11. Human Use

#### 2.12. Animal Use

### 2.13. Statement of Unique Capability Provided by Government or Government-funded Team Member

Not applicable.

### 2.14. Government or Government-funded Team Member Eligibility

#### 2.15. Facilities

The description of UCLA and KondaTech facilities is provided below.

#### Dejan Markovic, PI

Office: UCLA Engineering IV Building, Room 56-147E (approx. 140 sq. ft).

*Graduate student offices:* UCLA Engineering department allocates the required office space for graduate students from a common pool. The PI has 14 graduate students.

*Computing resources:* The PI and his students have access to a 10-node high-performance linuxbased compute cluster, 4 windows-based remote desktop servers. Hardware resources are complemented with software tools for chip design (Cadence, Synopsys, Mentor), algorithm design (Matlab, Simulink), and FPGA prototyping tools (Xilinx, Synplicity).

*Laboratory space:* The PI has access to several labs equipped with chip instrumentation for the testing of digital circuits. He is primarily using Integrated Circuits and Systems Lab (ICSL) at UCLA EE department. The equipment includes signal generators, spectrum analyzers, logic analyzers, high-speed oscilloscopes, and a high-speed probe station.

#### Venkat Konda, Consultant

Office: Konda Technologies, 6278, Grand Oak Way, San Jose, CA (approx. 100 sq. ft).

*Computing resources:* Dr. Konda has a server with two quad-core AMD Athlon 64x2 processors and fedora Linux operating system, two windows-based desktop/laptop machines.

No Government Furnished Property is required for conduct of the proposed research.

#### Referenes

- [1] [Online], available: http://www.eecg.toronto.edu/~vaughn/challenge/challenge.html
- [2] D. Marković, C. C. Wang, L. Alarcon, T.-T. Liu, and J. M. Rabaey, "Ultralow-Power Design in Near-Threshold Region," Proceedings of the IEEE, vol. 98, no. 2, pp. 237-252, Feb. 2010.
- [3] C.-H. Yang and D. Marković, "A Flexible DSP Architecture for MIMO Sphere Decoding," IEEE Trans. Circuits & Systems I, vol. 56, no. 10, pp. 2301-2314, Oct. 2009.
- [4] V. Wang, K. Agarwal, S.R. Nassif, K.J. Nowka, and D. Marković, "A Simplified Design Model for Random Process Variability," IEEE Trans. Semiconductor Manufacturing, vol. 22, no. 1, pp. 12-21, Feb. 2009.
- [5] D. Marković, B. Nikolić, and R.W. Brodersen, "Power and Area Minimization for Multidimensional Signal Processing," IEEE Journal of Solid-State Circuits, vol. 42, no. 4, pp. 922-934, Apr. 2007.
- [6] D. Marković, V. Stojanović, B. Nikolić, M.A. Horowitz, and R.W. Brodersen, "Methods for true energy-performance optimization," IEEE Journal of Solid-State Circuits, vol. 39, no. 8, pp. 1282-93, Aug. 2004.
- [7] F. Chen, et al., "Demonstration of Integrated Micro-Electro-Mechanical Switch Circuits for VLSI Applications," in Proc. IEEE Int. Solid-State Conference (ISSCC'10), Feb. 2010, pp. 26-27.
- [8] V. Karkare, S. Gibson, and D. Marković, "A 130-uW, 64-Channel Spike-Sorting DSP Chip," in Proc. IEEE Asian Solid-State Circuits Conference (A-SSCC'09), Nov. 2009, pp. 289-292.
- [9] C.-H. Yang and D. Marković, "A 2.89mW 50GOPS 16×16 16-Core MIMO Sphere Decoder in 90nm CMOS," in Proc. IEEE European Solid-State Circuits Conference (ESSCIRC'09), Sep. 2009, pp. 344-348.
- [10] F. Chen, H. Kam, D. Markovic, T.-J. King Liu, V. Stojanovic, and E. Alon, "Integrated Circuit Design with NEM Relays," in Proc. Int. Conf. on Computer Aided Design (ICCAD'08), Nov. 2008, pp. 750-757.
- [11] R. Nanda, C.-H. Yang, and D. Markovic, "DSP Architecture Optimization in Matlab/Simulink Environment," in Proc. Int. Symp. on VLSI Circuits (VLSI'08), June 2008, pp. 192-193.
- [12] D. Markovic, C. Cheng, B. Richards, H. So, B. Nikolic, and R.W. Brodersen, "ASIC Design and Verification in an FPGA Environment," in Proc. Custom Integrated Circuits Conference (CICC), Sep. 2007.
- [13] D. Markovic, B. Richards, and R.W. Brodersen, "Technology Driven DSP Architecture Optimization within a High-Level Block Diagram Based Design Flow," in Proc. 40th Asilomar Conf. on Signals, Systems and Computers, Nov. 2006. (Invited)
- [14] D. Markovic, B. Nikolic, and R.W. Brodersen, "A 70 GOPS Multi-Carrier MIMO Chip in 3.5mm2 and 34mW," in Proc. IEEE Int. Symposium on VLSI Circuits (VLSI), June 2006, pp. 196-197.
- [15] R. W. Brodersen, M.A. Horowitz, D. Markovic, B. Nikolic, and V. Stojanovic, "Methods for true power minimization," in Proc. Int'l Conf. on Computer Aided Design (ICCAD), Nov. 2002, pp. 35-42. (Invited)
- [16] V. Stojanovic, D. Markovic, B. Nikolic, M.A. Horowitz, and R.W. Brodersen, "Energy-delay tradeoffs in combinational logic using gate sizing and supply voltage optimization," in Proc. 28th European Solid-State Circuits Conference (ESSCIRC), Sept. 2002, pp. 211-14.



39465 Paseo Padre Parkway Suite 3700 Fremont, CA 94538 510.252.1136 (P) 888.700.8917 (F)

August 3, 2010

Dr. William Harrod DARPA TCTO

Dear Dr. Harrod:

I have reviewed the proposal entitled "Energy-Efficient Butterfly FPGA Hardware and Programming Tools," to be submitted by Dejan Markovic and his colleagues to DARPA TCTO.

I find this work of significant value to the development of our future HPC products and am interested to have my engineers interact with UCLA team during the course of this project to ensure that the final outcome can be smoothly transferred to a product.

I look forward to working with Dejan Markovic and his colleagues on this research.

Sincerely,

NO

Dr. Chen Chang Founder, CEO BEEcube, Inc. 39465 Paseo Padre Parkway Suite 3700 Fremont, CA 94538 Phone: (510) 252-1136

## **EXHIBIT 3**

## **EXHIBIT 3**



## DARPA Guide to Broad Agency Announcements and Research Announcements

November 2016

RELEASABILITY: UNLIMITED. This Instruction is authorized for public release.

Office ADPM when there have been material changes to the content of the briefing. Any Review Team Member who does not attend the ethics briefing will be required to document and self-certify the date of his or her last ethics briefing in the COI Self-Certification form.

Prior to proposal review, all Review Team Members shall be required to complete and submit a written self-certification, for the record, to document any known or apparent COIs or stating that they have none relevant to reviewing BAA proposals, as well as any other requirements regarding information access during the Scientific Review Process. Review Team Members complete this form after receipt of proposals. The Technical Office will retain the self-certification forms as part of the documentation in accordance with paragraph 2.E. below. The briefing charts and the self-certification form are available on the DARPA portal on the GC home page.

The PM is responsible for ensuring that each Review Team Member has access to or receives a copy of both the briefing charts and the self-certification form. After verifying that each member of the Review Team has sufficiently completed the self-certifications forms, the PM will review the forms with the CO and GC regarding potential COIs and appearance issues in the self-certifications, as necessary. The PM will brief all support contractor personnel having access to the proposals and ensure that no support contractor personnel have any COIs. Support contractor personnel with COIs participating in the Scientific Review Process must work out their participation in the process with GC, the CO, and the PM. The PM must also ensure that support contractor personnel have a nondisclosure agreement on file signed when they began their duties with DARPA. The PM shall remind them of the restrictions and requirements that are contained in that agreement as they relate to the handling and review of proposal material in accordance with section 2.E. below. A sample of a nondisclosure agreement is available in DI 70, "Contractor Relationships: Inherently Governmental Functions, Prohibited Personal Services, and Organizational Conflicts of Interest."

2.F.2. <u>Scientific Review Training</u>. The CO will attend the Scientific Review Team Kick-off Meeting and provide training on how to sufficiently document proposal reviews.

2.G. Protection of Sensitive Data. All participants in the Scientific Review Process (including SMEs and SETAs) are prohibited from, unless permitted by law, knowingly disclosing contractor bid, or proposal information, or source selection information in accordance with FAR 2.101, and the Procurement Integrity Act, 41 U.S.C. §§ 2101-2107 (implemented in FAR 3.104). Unauthorized disclosure of proprietary or confidential information, either before or after the award, is prohibited by the Trade Secrets Act, 18 U.S.C. § 1905, the Privacy Act, 5 U.S.C. § 552a, and by other laws and regulations. Prior written authorization from DIRO, or the CO must be obtained prior to releasing protected information outside the Scientific Review Team. The requirement for prior written authorization does not apply to the personnel associated with standard operational support activities such as preparing/processing/reviewing funding requests for selected proposals by Financial/Comptroller personnel, or archiving solicitation documentation on the Agency server or SharePoint sites by information technology or SETA support personnel.

The PM shall monitor and maintain all source selection information (as defined by FAR 2.101) within a secured physical and network area. This includes ensuring that information

## **EXHIBIT 4**

## **EXHIBIT 4**

#### A 1.1 GOPS/mW FPGA Chip with Hierarchical Interconnect Fabric

Cheng C. Wang, Fang-Li Yuan, Henry Chen, Dejan Marković

Electrical Engineering Department, University of California, Los Angeles, CA

#### Abstract

A 2048 look-up-table FPGA with a radix-2 hierarchical interconnect network is realized in 3.94mm<sup>2</sup> in 65-nm CMOS. It has an interconnect-to-logic area ratio of 1:1, which is a 3-4x reduction from modern FPGAs while allowing up to 100% resource utilization. As a proof of concept, it is designed with standard cells, achieving 16.4 GOPS/mm<sup>2</sup> at 370MHz. Peak energy efficiency of 1.1 GOPS/mW is measured at 0.5V.

#### Introduction

Field-programmable gate arrays (FPGAs) are effective for rapid verification and prototyping of VLSI designs. They are also used in products that require periodic hardware changes and short time to market. However, FPGAs incur penalties in area (17–54x), speed (2.5–6.7x), and power (5.7–62x) over standard-cell ASICs [1], hindering their expansion into ASIC markets. The overhead is primarily due to interconnects, which account for over 75% of area and delay.

For over 20 years, FPGAs have used 2D-mesh interconnects, where look-up tables (LUTs) are placed in configurable logic blocks (CLBs), and arrays of switch boxes are placed at interconnect crossings (Fig. 1). Since a full array requires too much area, various heuristics are used to simplify switch-box arrays at the cost of resource utilization. Yet 80% of the 1.1B transistors on Virtex-5 are used for interconnects [2]. This paper demonstrates an FPGA with hierarchical interconnects where interconnect area is 51%, a 3–4x reduction from commercial FPGAs while preserving connectivity. An energy efficiency of 1.1 GOPS/mW is the highest among reported FPGAs. The chip is tested up to 400MHz.

#### **Hierarchical Interconnect Architecture**

The key issue with 2D-mesh is scalability; the number of switch boxes grows as  $O(N^2)$  with the number of LUTs. Using Rent's rule, interconnect complexity is still  $O(N^{1.75})$  for random logic, requiring FPGA size to scale much faster than Moore's Law. In the proposed hierarchical interconnect, a folded Beneš network is employed to reduce the complexity to  $O(N \cdot log N)$  [3]: 4 LUTs are connected via 2 stages of switch matrices (SMs), and another 4 LUTs are connected with a 3<sup>rd</sup> SM stage (Fig. 2a). Each SM has 4 unidirectional connections per direction. Although this architecture reduces interconnect complexity, each SM stage doubles the routing congestion. This O(N) congestion makes physical design difficult.

To alleviate congestion, routing is alternated between x-y directions to reduce congestion to  $O(N^{0.5})$  (Fig. 2b). At every hierarchy, the LUTs near the center are interconnected to create shorter routes, and the edge routes are longer. This gives routing tools options for faster paths on timing-critical routes.

The test chip has 2048 4-input LUTs: 1024 LUTs form 256 Logic CLBs, 896 LUTs form 224 DSP CLBs, and 128 LUTs form 16 Block RAMs (BRAMs) of 1kb each. In practice, the majority of the logic connections are local, requiring fewer connections on upper hierarchies. Therefore full connectivity is preserved up to 6 SM stages (Fig. 3a), then half-connectivity SMs are used to reduce the complexity of upper hierarchies. This partitions the interconnect into 3 sub-networks:  $N_{8:2}$ ,  $N_{6:2}$ , and  $N_{6:1}$ . The chip is divided into 16 macros (Fig. 3b). Macros  $N_{8:2}$  are centered for shorter top-level routing, branching into  $N_{6:2}$  and  $N_{6:1}$ . Each of the macros contains 32 CLBs—a combination of Logic, DSP, and BRAM (Fig. 3c).

#### **Circuit Implementation**

The CLBs include four 4-input LUTs with selectable asynchronous/synchronous output stages (Fig. 4a). Each LUT

is configurable as one 4-input LUT or two 3-input LUTs with up to 4 unique inputs. A Logic CLB includes a carry chain to support 4b additions where Propagate and Generate are driven from LUTs. The Logic CLB is especially useful when two outputs per bit are required, such as in 3:2 compressors.

The DSP CLB (Fig. 4b) has a LUT combiner to support 5/6-input LUTs, and a carry chain that is configurable as one 8b or two 4b adders. The adder cells are shared with a  $4b\times4b$  Wallace-tree multiplier. Based on the configuration, the appropriate outputs are sent to the output stage. Due to the level of configurability, the synthesized CLB has 50 logic gates on its critical path (shaded), amounting to a 1.1ns delay.

Configuration bits are required to control CLBs and SMs, but traditional SRAM arrays are not suitable because all bits cannot be accessed simultaneously. A scan chain is adopted in [4] to control 6 CLBs, but it is not scalable to larger designs. Therefore an SRAM-based bit cell (BC) is designed where the output of each BC is directly routed to the configuration inputs of CLBs and SMs (Fig. 5a). The BC area is 5x smaller than a DFF-based scan cell. The bit-line (BL) and word-line (WL) controls are implemented as scan chains to write one row of BCs at a time. The BC arrays are local to each CLB, so only the BL and WL controls are propagated to top level. Overall, the memory area is reduced (Fig. 5b), and total interconnect area is 51%, a 3–4x reduction over 2D-mesh [5] for a fixed logic area.

#### **Automated Mapper**

An automated mapper is developed to map RTL onto this FPGA. A standard-cell library of LUT functions is created to enable logic synthesis using commercial tools. The LUT netlist is imported into an automated, custom place-and-route tool that generates the bitstream for FPGA programming. This tool is also used during architecture design to evaluate interconnect connectivities by mapping Toronto20 benchmarks.

#### **Measurement Results**

Our chip achieves 16.4 GOPS/mm<sup>2</sup> when all Logic and DSP CLBs are utilized, executing 175 16b accumulators at 370MHz. Since a 16b adder uses 2 DSP CLBs or 4 Logic CLBs, the DSP adders are faster, reaching 400MHz. Performance is hindered by equipment limitations due to a 0.25ns input-clock jitter at 400MHz. The energy-delay curve and the power breakdowns for minimum delay and minimum energy are shown in Fig. 6.

In comparison, [4] has no interconnects, the full-custom CLB in 32-nm LVT is 2.5x faster, but achieves 2.6 GOPS/mW at 0.34V for 8b operations, which is 0.65 GOPS/mW for 16b (2 CLBs per operation at half the speed). With interconnects, our 65-nm chip reaches 1.1 GOPS/mW at 0.5V.

Leakage is well-controlled even without power gating. A 1.08 GOPS/mW is attainable with only 112 DSP accumulators active and most of the Logic CLBs idle (Table I). The FIR filter achieves 274MHz due to longer routing, but interconnect delay is still under 50%. The  $2\times 2$  MIMO FFT uses 10 BRAMs to implement various delay lines. With many control signals and a critical path of 11 CLBs, the FFT achieves 83MHz.

Figure 7 shows the die photo. The top 3 metal layers (out of 9) are sparsely used, leaving ample room for larger designs.

#### Acknowledgments

We thank STMicroelectronics and C. Yang for helpful discussions. References

- [1] I. Kuon et al., Found. Trends in Elec. Design, 2008.
- [2] I. Bolsens, *MPSOC*, 2006.
- [3] V. Konda, U.S. Patent 2010/0172349.
- [4] A. Agarwal et al., ISSCC Dig. Tech. Papers, 2010.
- [5] M. Lin et al., FPGA '06.



Figure 4: a) CLB block diagram and b) DSP CLB schematic.

Figure 7: Die micrograph and chip summary.

## **EXHIBIT** 5

# **EXHIBIT** 5

#### **ABSTRACT OF THE DISSERTATION**

## Building Efficient, Reconfigurable Hardware using Hierarchical Interconnects

by

Chengcheng Wang Doctor of Philosophy in Electrical Engineering University of California, Los Angeles, 2013 Professor Dejan Marković, Chair

In the semiconductor industry today, ASICs are able to offer 10x-1000x higher energy and area efficiencies than non-dedicated chips, such as programmable DSP processers, fieldprogrammable gate arrays (FPGAs), and microprocessors. Not surprisingly, SoCs today have become an integration of many ASIC blocks, each performing a few dedicated tasks. The growing size of modern SoC chips, accelerated by the increasing demands for functionalities, has exposed the major drawback of ASIC: design cost. These large SoCs are re-designed a few times a year to rectify hardware-bugs and to support new features. Because ASICs are not reconfigurable, even the smallest hardware change would require a re-design. Additionally, design cost is rising exponentially with every technology generation.

The rising design cost of ASICs has exposed a huge need today: efficiency and flexibility must co-exist. But among flexible hardware candidates, microprocessors and programmable DSP

processors are far too slow to meet the throughput requirements of ASICs. FPGAs do come close in terms of performance, but are extremely inefficient due to its high energy and large area overhead. We must bridge the huge gap in efficiency for FPGA to become a viable contender to ASICs.

The primary culprit for FPGA inefficiency is interconnect, which accounts for over 75% of area and delay. For over 20 years, 2D-mesh network has been the back-bone of FPGA interconnects, but full connectivity in a 2D-mesh require O(N2) switches, requiring interconnects to grow much faster than Moore"s Law. As a result, various heuristics are used to simplify switch-box arrays at the cost of resource utilization, but interconnect area of modern FPGA is still around 80%. This work builds FPGA using hierarchical interconnects based on Beneš networks, requiring O(N·log·N) switches. Although Beneš is commonly used in telecommunication, this work is its first silicon realization of a FPGA. To realize a highly efficient interconnect architecture, significant pruning of the network is required. Novel techniques such as fast-path U-turns and unbalanced branching are also implemented. A custom place-and-route software is developed to map benchmark designs on a variety of interconnect candidates. From mapping results, the architecture is updated based on network utilization until an optimized design is converged. The large area of FPGA chip requires aggressive power gating (PG), but interconnect signals often lack spatial locality, make it block-level PG difficult. A novel PG circuit technique is developed to power-gate individual interconnect switches with very small overhead in area and performance. Such technique requires fundamental circuit changes, even modifying the CMOS inverter.

With innovations in chip architecture, circuit design, and extensive software development, this work has demonstrated 5 user-mappable FPGAs (from 1K-16K LUTs) all

iv

with around 50% interconnect area: a 3–4x reduction from commercial FPGAs while preserving connectivity. An energy efficiency of 1.1 GOPS/mW is the highest among reported FPGAs, and is 22x more efficient than the most efficient commercial FPGA today, significantly bridging the efficiency gap between FPGA and ASIC.

### TABLE OF CONTENTS

Ι	Intro	Introduction1				
	1.1	The Drive Towards Efficiency1				
	1.2	What is Efficiency?				
	1.3	The Efficiency Tradeoff				
	1.4	Efficiency and Flexibility – Current Solutions				
	1.5	Keeping Up with the Standards7				
	1.6	The Cost of Chip Design				
	1.7	Candidates for Reconfigurable Hardware9				
	1.8	Thesis Outline11				
Π	FPGA Interconnects: the Source of its Inefficiency1					
	2.1	Brief History of FPGAs				
	2.2	Interconnects: the Backbone of an FPGA18				
	2.3	Scaling a 2D-mesh Network				
	<u>2.4</u>	Hierarchical Network – A Scalable Solution				
	2.5	Prior Attempts at Hierarchical FPGAs				
	2.6	Our Challenges				

III	Arch	Architecture Design of Hierarchical FPGAs				
	<u>3.1</u>	Realizing Large-Scale Beneš Networks				
	<u>3.2</u>	Implementing a 2048-LUT FPGA Interconnect				
	<u>3.3</u>	Radix-3 Boundary-less Interconnect				
	<u>3.4</u>	Fast-Path Interconnect				
	<u>3.5</u>	Interconnect Cost vs. Gate Cost47				
	<u>3.6</u>	Local Interconnect vs. Branch Interconnect				
	<u>3.7</u>	Micro-architecture of a Switch Matrix				
	<u>3.8</u>	Implementing a 16K-LUT FPGA Interconnect				
IV	Inter	connect Circuit Design58				
	4.1	Key Building Blocks in Interconnect Circuits				
	4.2	Static Multiplexers and Area-Performance Tradeoff59				
	4.3	Strategies for Interconnect Buffering63				
	4.4	Designing Configuration Bit-Cells66				
	4.5	Power-gating Switch Matrices				
	4.6	Power-On Sequence of the Interconnect Network73				

V	Conf		
	5.1	Configurable Logic Blocks for the 2048-LUT FPGA	79
	5.2	Macro-based Chip Integration for the 2048-LUT FPGA	86
	5.3	Fine-Grained CLBs for the 16K-LUT FPGA	91
	5.4	Medium-Grained CLBs for the 16K-LUT FPGA	
	<u>5.5</u>	Coarse-Grained CLBs for the 16K-LUT FPGA	<u></u>
	5.6	Macro-based Chip Integration for the 16K-LUT FPGA	
VI	Softw	vare Flow and Design Mapping	113
	6.1	Overview of FPGA Software Mapping Flow	
	6.2	FPGA Synthesis and LUT Packing	116
	6.3	FPGA Partitioning and Placement	
	<u>6.4</u>	FPGA Routing	124
	6.5	Bitstream Generation	
VII	Test	Infrastructure and Measurement Results	130
	6.1	Matlab Simulink-based Testing Infrastructure	
	6.2	Measurement Results of our 2048-LUT FPGA	134
	6.3	Updated Testing Infrastructure	
	6.4	Measurement Results of our 16K-LUT FPGA	140
	6.5	Chips Summary and Die Photos	
VIII	Conc	lusion and Future Outlook	147

### **CHAPTER II**

### **FPGA Interconnects: the Source of its Inefficiency**

#### 2.1 Brief History of FPGAs

The concept of a reconfigurable hardware started over 30 years ago, but it was regarded as prohibitively expensive because of its large overhead in area over ASICs. Transistors were expensive, and no one wanted to pay the huge area penalty for reconfigurability. Fortunately, the semiconductor industry rapidly expanded at the pace of Moore''s law, and such large area overhead became more tolerable, finally leading to a first FPGA by Xilinx Corporation in 1985. The original FPGA, XC2000 series, had 64 or 128 look-up-tables (LUTs). As shown in Figure 2.1 a), each configurable-logic block (CLB) contains just one LUT and one selectable flip-flop. With so few CLBs, the interconnect network is also simple. The interconnects run in x- and ydirection around the CLBs, twisting with every segment, and some of the intersections have switch matrices placed diagonally, consisting of 6 pass-transistors per switch (Figure 2.1 b) [Brown92].

The initial perceptions of the XC2000 were "small, slow, expensive, and "different" [Alfke07], but the XC3000 introduced in 1987 became very successful even with very rudimentary software support. Fast-forward to today, FPGAs can support up to 500,000 LUTs per die, and the largest Xilinx Virtex-7 even supports 2 million LUTs using Stacked-Silicon Technology (Figure 2.2) [Saban12].



Figure 2.1: Schematic diagram from a Xilinx XC2000 of a) CLB and b) interconnects.



Figure 2.2: Illustration of Stacked-Silicon Technology in Xilinx Virtex-7.

Due to yield and fabrication constraints, each die is limited to around 500,000 LUTs, occupying 529 mm<sup>2</sup> in 28nm. "Stitching" the 4 chips together requires a very large interconnect bandwidth, far greater than that offered by standard packaging solutions. Therefore, a 65-nm passive silicon interposer is mounted onto the 4 FPGA dies to create a high-bandwidth interconnect, providing more than 10,000 connections between each adjacent die. For communication with external I/Os, the interposer uses through-silicon vias (TSVs) to connect the FPGA die to the C4 bumps on the package. Although the stacked silicon technology is not monolithic, many of the performance and cost benefits of a 3-D monolithic FPGA from [Lin07] still apply.

Of course, FPGA progressions are more than just area expansion, the CLB core of the FPGAs has also evolved over the years (Figure 2.3) [Rose93]. Many features are added to implement commonly-used ASIC features very effectively, such as multiple flip-flops with clock-enables (XC3000), a dedicated ripple-carry chain (XC4000), and LUT-combining multiplexers (XC5200).





### XC4000


# XC5200



Figure 2.3: CLB diagram of Xilinx a) XC3000, b) XC4000, and c) XC5200.

Over the past ten years, CLB sizes grew even more. Xilinx has transitioned to four 4input LUTs per CLB in its Virtex-4 [XlinxV408], then to four 6-input LUTs per CLB in Virtex-5 [XlinxV512]. The newer Virtex-6 and 7 even have dual flip-flops mated to each of the 6-input LUTs (Figure 2.4) [XilinxV6CLB12].

The newer CLBs place an even greater emphasis on software design. The performance of the FPGA depends heavily on the mapping algorithm – packing critical-path gates within a CLB would provide much faster performance than spreading the critical path across multiple CLBs. Since the interconnect network cannot provide full connectivity across all CLBs (Chapter III), packing LUTs locally into CLBs can reduce the number of I/Os required by the CLB [Betz98], and the software tool also needs to provide quality place-and-route results to ensure feasible design mapping.

# Virtex 6/7



Figure 2.4: CLB diagram of Xilinx a Virtex-6 and 7 series FPGA.

Over the years, the FPGA software support has developed into a complete design suite. With extensive support for automated design mapping from HDL into bitstream, very little effort is required by the end-user. High-level synthesis tools even support mapping software programs (such as C or Matlab models) directly onto the FPGA. This many layers of abstraction provide a simple user experience, but it also shields us from seeing the intricate details of a FPGA design, especially interconnect routing.

### 2.2 Interconnects: the Backbone of an FPGA

In FPGA design, great emphasis is placed on the CLBs and other programmable blocks, and documentations are widely available. On the other hand, interconnects have mostly remained in the dark. Although FPGAs have grown enormously in size since the XC2000, the fundamental interconnect architecture still remains (Figure 2.5). In 2D-mesh interconnects, LUTs are placed in configurable logic blocks (CLBs), and interconnects run in the x- and y- direction surrounding the CLBs. I/O connection switches tie the CLB I/O to the interconnect network. Arrays of switch boxes are placed at interconnect crossings to select and buffer the programmed path. Each switch-box contains pass-transistors programmable by the configuration memory. Since a full switch-box array at every interconnect crossing requires too much area, various heuristics are used to simplify the arrays at the cost of interconnect connectivity [DeHon99, Tessier00, Lin09]. In Figure 2.5, the example network only implements switch boxes along one main diagonal and two sub-diagonals of the switch-box array. In this simplistic case, each interconnect trace enters a switch-box at every interconnect crossing, the selected path is then buffered to drive the next trace.



Figure 2.5: A sample 2D-mesh architecture with I/O connections and switch boxes.

To improve routing performance and add path diversity, each interconnect trace can be heuristically designed to travel for 1, 2, 4, 6, or even more CLBs before reaching the next switch. A path from one switch to the next is called a "hop". From an illustration in Xilinx XC4000 interconnects (Figure 2.6) [XilinxXC99], we see different interconnects labeled as "single", "double", "quad", "long", or even "global" based on the distance of each hop. Coming out of a CLB, a signal can be connected to a selection of hop lengths, giving the router freedom to choose a longer or shorter hop based on its routing requirements. Modern FPGAs have also migrated towards uni-directional routing, thus removing bi-directional buffers and significantly reducing interconnect loading [Lemieux04, Lee06].



Figure 2.6: Interconnect architecture of a Xilinx XC4000 FPGA [XilinxXC99].

With extensive techniques in interconnect pruning, along with ever more complex CLBs, one may expect the FPGA area to be dominated by CLBs. It is called a "gate-array" after all. Surprisingly, even with such heuristics, 80% or more of the area on modern FPGAs are occupied by interconnects [Bolsens06]. The interconnect area is actually 4 times the logic area! In addition, interconnect also accounts for the majority of the delay and power in today"s FPGAs (Figure 2.7). The reality could be even worse: if we were to remove the larger IP blocks and accelerators from the FPGA, and compare the area of interconnect versus the area of CLBs, the ratio could be closer to 10:1.



Figure 2.7: Area, delay and power breakdown of a modern 2D-mesh FPGA.

### 2.3 Scaling a 2D-mesh Network

The key cause for interconnect overhead is the scalability of 2D-mesh interconnects. In the worst case, the number of switch boxes grows as  $O(N^2)$  with the number of LUTs. Although heuristics are able to reduce the number of switches, there is a limit. Rent"s rule ( $T = t \cdot g^p$ ) can be used to model interconnects, where g is the number of gates, exponent p is the Rent"s coefficient for modeling the number of I/Os, and t is a constant of proportionality. In typical cases, the interconnect complexity per logic block is  $O(N^{0.75})$  for random logic, which is still  $O(N^{1.75})$  for a chip of N logic blocks [Landman71].

For very regular designs, such as a memory banks, the complexity per logic is  $O(N^{0.5})$ . Since FPGA mapping software employs intelligent gate placements, the logic is not completely random, but it is certainly not as regular as memory banks. We therefore expect the actual Rent"s exponent *p* to be between 0.5 and 0.75 [Tessier00]. But for very large designs (large *N*),  $O(N^{0.5})$  to  $O(N^{0.75})$  provides too large of a range for this model to be useful. Nevertheless, it provides us theoretical lower and upper bounds on interconnect complexity.

Even using an optimistic exponent of p = 0.5, the total complexity of O(N<sup>1.5</sup>) still

requires FPGA sizes to scale much faster than Moore''s Law. Figure 2.8 shows the interconnect expansion from Xilinx Virtex-4 to Virtex-5 [XlinxV506, Minev09]. Adding 50% of interconnect logic per CLB poses a significant area increase even for just 1 product generation. Scaling *N* from 64 in XC2000 to 500,000 in modern FPGAs, it becomes clear why interconnect area is a key concern today.



Figure 2.8: Interconnect resources per CLB for Xilinx Virtex-4 vs. Virtex-5 [XlinxV506].

In more recent years, many have proposed asynchronous architecture for FPGAs, aiming to improve its performance [Teifel04, Teifel204, Manohar06]. Such techniques have claimed to achieve > 1 GHz performance from FPGAs by using asynchronous hand-shake and token-based heavy pipelining. However, such technique failed to recognize the root cause of FPGA overhead, which is the scalability of the interconnect area. In contrast, asynchronous FPGAs require a 3x overhead in interconnect area: replacing 1 signal with 3 asynchronous hand-shake signals, further exacerbating the effect of interconnect overhead. Whenever signal fan-outs are required, complex acknowledgement circuitry is required to wait for the slowest path to return the token before passing it on. More recent work by [LaFrieda10] acknowledged the large area and power overhead required by asynchronous FPGAs, and proposed a two-phase logic and voltage-scaling in the acknowledge signals to reduce the power consumption, but the large overhead in area remains. Although asynchronous FPGAs claims to run up to 3x faster than their synchronous counterparts, the 3x penalty in interconnect area will quickly nullify any performance advantages on large designs. Recent work in [Devlin11] uses dual pipeline (separate pipelines for precharge and evaluation phases) to further improve asynchronous performance, but requires 5 physical wires for 1 interconnect signal. Clearly, these approaches are not scalable to larger designs. For efficient, high-performance FPGAs, what we need is an interconnect architecture that is scalable in area and performance, and not brute-force circuit implementations.

#### 2.4 Hierarchical Network – A Scalable Solution

<u>To address the non-scalability of 2D-mesh, we adopted a hierarchical interconnect</u> <u>architecture based on a Beneš network.</u> In telecommunication, Clos,Beneš, and similar hierarchical networks are well-known to be rearrangeably non-blocking network for point-topoint connections, and are commonly used in communications [Clos53, Benes62, Kleinrock77, Yang99, Dally04]. <u>There has not been a silicon realization of a Beneš network for FPGAs until</u> this work. To demonstrate its feasibility, the original Beneš network is first modified into a <u>realizable FPGA architecture</u>.

As a demonstration, we start with 2 LUTs, each with just 2 inputs and 2 outputs (Figure

2.9). This network requires 3 stages, and each stage uses 2x2 switch matrices (SMs) for signal routing. Each SM can support both uni-cast and multi-cast of incoming signals, as shown. This network is rearrangeably non-blocking for uni-cast, meaning the signal routing can be rearranged to support arbitrary LUT-to-LUT connections.



Figure 2.9: A simple 3-stage Beneš network connecting 2 LUTs.

In FPGA applications, it is common to use 4 to 6 input LUTs with 2 outputs. To illustrate a 4-input, 2-output LUT network, the 3 stage network is recursively extended to a 5-stage network (Figure 2.10), and can be further extended to larger networks. This network remains non-blocking for uni-cast, and because there are only half as many LUT outputs as inputs, it is virtually non-blocking even for multi-cast based on our simulations. Since each LUT only has 2 outputs, the red SMs can always multi-cast the signals, and can be removed. In addition, the 4 inputs to a LUT may arrive in any order, therefore the gray SMs can be removed as well. Note that for some CLBs, such as DSP accelerators or control signals, the inputs may not arrive in any order, and in those cases the grey SMs must remain. For simplicity, the center 3 stages are abstracted as a single 4-input, 4-output SM, which is essentially a 2-bit 2x2 switch because it propagates two paths in each direction. The simplified diagram is shown on the bottom of Figure 2.10.



Figure 2.10: A 5-stage Beneš network merged into a 3-stage using 2-bit 2x2 switches.

Scaling to a larger network, we observe one key problem with the original Beneš network. Figure 2.11 shows an 8-LUT network using 5 SM stages. The downside is that *all* paths are required to traverse on all 5 stages regardless of the physical distance between the source and destination. As shown in Figure 2.11, LUT 7 and 8 are physically adjacent to each other, but the network requires the signal to traverse through all hierarchies while a simple switch in the first stage would suffice. Another issue with this network is input/output locality. In an FPGA, the input and output of a LUT is coming from one hardware block, but in this network, the inputs and outputs are split across two sides of the network. Since this diagram is not representing physical implementation, it can be misleading to the FPGA designer.



Figure 2.11: A 5-stage Beneš network connecting 8 LUTs.

To avoid traversing unnecessary hierarchies to speed up interconnect routing, and to provide an interconnect that closely resembles the physical implementation, we employ a folded Beneš network (Figure 2.12), also called a fat-tree network by [Leiserson85]. This similar architecture has been employed in supercomputing machines, such as the Connection Machine CM-5 with 256, 544, and even over 1000 processing nodes [Leiserson96].

As shown, 4 LUTs are connected via 2 stages of SM, and another 4 LUTs are to be connected with a 3rd SM stage. This effectively leads to an interconnect complexity of  $O(N \cdot \log N)$ , which scales much better than  $O(N^2)$  in 2D-mesh interconnects.



Figure 2.12: A 3-stage folded Beneš network connecting 8 LUTs (4 LUTs shown).

Although drawn with 2 arrows, each trace is actually 2 uni-directional signals. Each switch matrix then performs 4 uni-directional connections both upwards and downwards. Signals will come from the LUT output, traverse up to the required hierarchy, and traverse back down to the LUT input. Because the network is still rearrangeably non-blocking, full connectivity can be obtained.

Although this architecture reduces interconnect complexity by reducing the number of switches, routing congestion remains an issue. In Figure 2.12, the first SM stage has 2x2 wires crossing each other, but the second stage has 4x4 wires crossing, and the  $3^{rd}$  stages has 8x8. Each additional SM stage doubles the routing congestion. This O(N) congestion requires much larger area for higher level SMs, making physical design more difficult and less area-efficient.

Fortunately, implementing a Beneš network on silicon gives us freedom in both x- and ydirections. Although [Manuel 07] illustrated a manual layout method for a Beneš layout on a 1dimensional array, most silicon implementations allow for a 2-dimensional layout. To alleviate congestion, routing is alternated between the x-y directions, doubling the routing congestion for every 2 stages. The routing congestion is reduced from O(N) to  $O(N^{0.5})$  (Figure 2.13), and the fully symmetrical implementation also eases physical design.



Figure 2.13: A hierarchical Beneš interconnect architecture using alternated x-y routing.

Another change from the original Beneš network is unequal wire lengths. At every hierarchy, the LUTs near the center are connected to create shorter routes, and the LUTs near the edges have longer routes. In terms of logic connectivity, this wiring difference is an isomorphic transformation from the original network, thus the interconnect connectivity remains unchanged [Wu80, Duato02, Konda08]. Yet this difference in wire lengths gives routing tools options for faster paths on timing-critical routes. In physical design, this also allows the center routes to remain at the lower metal layers without crossing over the longer routes on the upper metal layers, further avoiding congestion.

### 2.5 **Prior Attempts at Hierarchical FPGAs**

Numerous publications have discussed hierarchical FPGA implemented as tree-ofmeshes (Figure 2.14) [Greenberg88, Lai97, Tsu99, Wong04, DeHon04]. It is a limited bisection network, where the mesh connectivity decreases for upper hierarchies. In some implementations [Tsu 99], even connectivity at local levels is limited. Additionally, a centralized routing network is required at every hierarchy, which increases routing congestion, and central switches are still based on 2D-mesh. The layout in [Greenberg88] intelligently distributes the meshes across the layout into "cubies", but the complexity of every hierarchy remains that of a mesh-based switch.



Figure 2.14: A hierarchical interconnect architecture using alternated x-y routing [DeHon04].

Unlike tree-of-mesh interconnects, our Beneš interconnect architecture evenly distributes routing across all LUTs instead of crowding them into centralized "hubs," easing routing congestion and shortening the wire length significantly. This is different from the butterfly layout in [DeHon00, Wong04] where centralized hubs are used, but hubs are distributed across different "cubies," thus requiring each signal to traverse across different hubs in different cubies just to switch hierarchy, significantly increasing interconnect delay.

There is one known silicon implementation of a tree-of-mesh FPGA, the hierarchical, synchronous reconfigurable array (HRSA) [Tsu99]. The architecture uses a Radix-4 topology with centralized switches and bi-directional routing. Rent's exponent of 0.5 is used, so every hierarchy prunes the interconnect connectivity by 50%. Due to the centralized hubs used in this architecture, processing elements (PEs, equivalent to LUTs) that are physically close to each other may be required to use a detour routing. A heuristic is then employed to add "shortcuts" to

connect these PEs using additional wiring (Figure 2.15).



Figure 2.15: The HSRA architecture without (left) and with (right) wiring shortcuts.

The HSRA architecture was able to maintain good operation frequency due to its heavy pipelining, but the interconnect network with a Rent's exponent of 0.5 offered "very limited" connectivity. There has not been a follow-up chip after the original HSRA in 1999.

A multilevel hierarchical FPGA was published by [Mrabet06], although no silicon realization is attempted. The architectures use a Radix-4 topology with a Rent<sup>\*\*</sup>s exponent of 1, but only on the downward paths. The upward path, on the other hand, provides no path diversity (Figure 2.16). Therefore, the overall path diversity of this architecture is very limited, and the interconnect connectivity when mapping real-world designs is about 30-50%, often requiring a 2K-LUT FPGA to map 1K-LUT designs.



Figure 2.16: The multilevel hierarchical FPGA architecture.

## 2.6 Our Challenges

Although hierarchical FPGA has great appeal on paper, it has not received much attention in practice. The main reason is that it has yet to demonstrate *any* advantage over 2D-mesh: its 30-50% logic utilization is significantly lower than the 85% utilization achievable by commercial FPGAs, and it has yet to demonstrate any notable performance, power, or area advantage. The speed improvement in HRSA is due to heavy pipelining, not interconnect improvements.

On the other hand, commercial FPGAs today are already very mature products, often made as full-custom designs with state-of-the-art processes (and needing more than 10 layers of metal). The CAD tools are also capable of delivering very high quality-of-results (QoR) within a easy-to-use framework.

For our work to be considered worthwhile, we need to demonstrate and realize a

hierarchical FPGA with significant benefits in performance and efficiency. To demonstrate its practical values, software development is also needed to allow users to map their own designs. Overall, this project requires innovation and extensive work in creating an interconnect architecture, realizing it in silicon, and developing software tools to demonstrate its advantages. These details are covered in the following chapters of this thesis.

# **CHAPTER III**

# **Architecture Design of Hierarchical FPGAs**

### 3.1 Realizing Large-Scale Beneš Networks

To illustrate the silicon realization of the Beneš network, we start with the architecture design applied to our two FPGAs. The two chips shown in this dissertation have approximately  $10 \times$  difference in logic capacity, and have different interconnect architecture as well. The first chip is a more straight-forward implementation, while the second chip utilized extensive architectural optimization techniques illustrated in Section 3.3 through 3.7.

The first test chip we published in [Wang11] contained 2048 look-up-tables (LUTs), each with 4 inputs and 2 outputs. Built on a Radix-2 architecture, it requires 11 levels of interconnects. Since every level translates to one SM stage, 11 levels of SMs are required. To ensure 100% connectivity in all cases, every LUT would need to have 11 levels of SM to preserve the full Beneš network. Using the 2D-layout method illustrated in Figure 2.13, expanding from 4 stages for 16 LUTs to 11 stages for 2048 LUTs would still be feasible to route, but it would occupy a significant amount of area. According Rent"s rule, this brute-force implementation represents a Rent"s exponent of p = 1. Realistically, there is no need to implement an interconnect network with more than p = 0.75 connectivity, as the area penalty associated with building larger interconnects far outweighs the benefits from chip utilization [Tessier00].

Mathematically speaking, implementing a network with p < 1 requires interconnect pruning at every stage. For example, when p = 0.75, every additional stage should implement 25% fewer wires than the previous stage. For FPGA realizations, there are three key reasons that make this exact implementation impractical.

First, mapping FPGA design is a very non-deterministic process that depends heavily on the design to be mapped and the algorithms used by the place-and-route (P&R) software. The design to be mapped can have a Rent<sup>\*\*</sup>s exponent p anywhere between 0.5 and 0.75, which is a very wide range for interconnect routing. A very regular design, such as a feed-forward finiteimpulse-response (FIR) filter, combined with a high-quality P&R tool, could be easily mapped onto an architecture with p = 0.5. On the other hand, a more complex design such as fast Fourier transform (FFT) will consume significantly more interconnect resources. There is no single exponent that can accurately represent all design complexities.

Second, the interconnect utilization is uneven across the SM stages. An effective P&R software would attempt to keep most of the signals local, thus shortening the critical path and reducing the active wire lengths. As a result, it is important to have sufficient routing resources for the lower levels to provide sufficient path diversity for the P&R tool. It can be worthwhile to use a Rent<sup>\*\*</sup>s exponent of p = 1 for the lower hierarchies, and use a more aggressive pruning (e.g. p = 0.5) for the upper hierarchies. From our architecture evaluations, pruning the lower hierarchies, even with p = 0.75, can lead to sever routing problems and performance degradation.

Lastly, and most importantly, the FPGA architecture needs to be realized in a 2-dimensional layout, and its large size can lead to a very complex physical design if not planned carefully. As shown in Figure 3.1, an efficient physical implementation can allow the FPGA chip designer to start with creating just one LUT macro and its SMs. Although the interconnect wire length between the macros can be different, the hardware logic and the I/O port for each macro are identical. The fully symmetrical architecture allows the LUT macro to be replicated throughout the entire chip, drastically improving design time. The designer can also add more hierarchies to the physical design flow, such as creating a 4-LUT macro out of the 1-LUT macro, then creating a 16-LUT macro from the 4-LUT macro. However, if the interconnect is to be pruned at every stage, the regularity of the layout can no longer be preserved: assuming all LUTs have SMs at stage 1, using p = 0.75, only 75% of the LUTs will have stage-2 SM, and only 56% of the LUTs will have stage-3 SM, and so on. Without regularity in the layout, not only will the interconnect take much longer to design, the reduced SM does not necessarily lead to reduced area. In Figure 3.1, if SMs are reduced for LUT 4, 8, 12, and 16, it would leave a gap in the middle of the layout because the surround macros are larger. This results in a worst-case situation of lost interconnect connectivity *and* lower layout density due to wasted area. When pruning SMs, the designer needs to make sure the reduced SM actually leads to reduced area, and must not over-complicate the layout process. This requires very judicious SM pruning at very strategic locations.



Figure 3.1: A hierarchical macro-based implementation of a 2D-Beneš network.

Overall, realizing a large Beneš network in FPGAs requires 3 things to keep in mind: interconnect connectivity, layout regularity, and layout density.

### **3.2 Implementing a 2048-LUT FPGA Interconnect**

The 2048-LUT test chip requires 11 levels of interconnects. To preserve interconnect connectivity for lower levels, we maintained connectivity (Rent"s p = 1) until SM stage 7, followed by 2 stages of p = 0.5, and full connectivity for the top 2 stages. One quadrant of the FPGA architecture is shown in Figure 3.2: the quadrant is divided into 4 macros, each containing 128 LUTs. Inside each 128-LUT macro, all the LUT macros are identical; they are implemented similarly to Figure 3.1, but with 7 stages of SM per LUT. The half-SMs shown in yellow allow 2 out of 4 inputs to propagate upwards, realizing Rent"s p = 0.5. Two concatenated half-SMs leads to a top-level connectivity of 25%.



Figure 3.2: Interconnect architecture for our 2048-LUT FPGA, one quadrant shown.

The interconnect network is partitioned into three sub-networks:  $N_{8:2}$ ,  $N_{6:2}$ , and  $N_{6:1}$ , where  $N_{P:Q}$  represent a network of P full-SMs and Q half-SMs. Intelligent SM-pruning also requires the pruned SM to translate to an area reduction. From the architecture in Figure 3.2, it is clear that the 3 types of SM macros,  $N_{8:2}$ ,  $N_{6:2}$ , and  $N_{6:1}$ , will each occupy a different area, because they each contain different number of SM stages.  $N_{8:2}$  is the largest macro, followed by  $N_{6:2}$ , with  $N_{6:1}$  being the smallest. To avoid gaps in the layout area, all SMs have the same width. Therefore  $N_{6:1}$  macros are shorter.  $N_{6:2}$  is also shorter than  $N_{8:2}$ , leading to some open space. In Chapter V, we will see that the opened space is used by Block RAMs. Because BRAM CLBs are larger than regular CLBs, the area pieces together very densely.

The top level of the chip is shown in Figure 3.3 with the 4 hierarchies of top-level wires shown in colors corresponding to those in Figure 3.2. The top-level layout is symmetrical in the x- and y- direction, allowing the single 512-LUT quadrant to be replicated to form the other 3 quadrants. The chip is divided into 16 macros of 128 LUTs each: macros with  $N_{8:2}$  interconnects are placed near the center for shorter top-level routing, branching into  $N_{6:2}$  on the left and right.  $N_{8:2}$  and  $N_{6:2}$  then both branch into  $N_{6:1}$  on the top and bottom. This physical placement avoids long wires at the top level, and therefore minimizes interconnect buffers and further reduces area.



Figure 3.3: Interconnect architecture for our 2048-LUT FPGA, one quadrant shown.

This 2048-LUT architecture is relatively straightforward, using only 2 types of SMs to form 3 types of LUT macros. Scaling into larger designs with even more hierarchies, more advanced architectural techniques are used to further optimize the design. They are highlighted in the following sections (3.3 - 3.6).

#### 3.3 Radix-3 Boundary-less Interconnect

Although hierarchical routing "s  $O(N \cdot \log N)$  complexity is much better than  $O(N^2)$  from 2D-mesh, it is sometimes inefficient for local routing if the leaves are crossing a high-radix boundary. For example, In Figure 3.4a), LUT 8 and 9 are neighbors, but signals have to traverse up 4 stages of network, and then zig-zag their way down the hierarchy to for LUTs to communicate with each other. Such lack of spatial locality is not desirable.

One method to shorten the nearest-neighbor routing lengths is an isomorphic transformation, as shown in Figure 3.4b). Connections from LUT 8 to LUT 9 can now traverse directly up to stage 4, make a U-turn, and traverse straight down. In terms of connectivity, it is well known that isomorphic butterfly structures maintain the same logic connectivity [Wu80]. Although the wire length travelled has reduced, the number of switches has not: the signal still needs to traverse up and down 4 hierarchies for communication between LUT 8 and 9.



nearest-neighbor lengths, and c) with boundary-less radix-3 switches in stage 1.

In this section, we propose a method of applying higher radix switches on the lower SM levels to utilize spatial locality in routing, allowing efficient interconnect routing for direct neighbors. We call such network a boundary-less radix-3 network [Wang13].

To convert a radix-2 network to a boundary-less radix-3 network, we first identify the center 2x2 routing of each stage, shown in the dashed circle in Figure 3.4b). It is noted that such center 2x2 routing only connects across an interconnect length of 1 ( $2^0$ ). The first stage transformation into a radix-3 boundary-less interconnect is shown in Figure 3.4c). All center 2x2 routing in the dashed circles are moved to stage 1. This converts stage 1 into a radix-3 interconnect, and all stage-1 switches are capable of communicating with their immediate neighbors, both up and down the SM stages.

With stage 1 completes, we now convert stage 2 to a boundary-less radix-3 switch. We first identify the remaining center  $2x^2$  routing above stage 2 (Figure 3.5a), shown in dashed circles. Note that these  $2x^2$  routings only connect across an interconnect length of 3 ( $2^{1}+1$ ). These  $2x^2$  routings are then moved down to between stages 1 and 2 (Figure 3.5b), converting the second stage into a radix-3 boundary-less interconnect.



Figure 3.5: A 16-LUT Beneš network with a) boundary-less radix-3 switches in stage 1, and b) with boundary-less radix-3 switches in stages 1 and 2.

The same transformation continues for stage 3-4: we first identify the remaining center 2x2 switches above stage 3, shown in dashed circle (Figure 3.6a). For stage 2-3, we can note the remaining 2x2 switches are actually double pairs, one for LUTs 6–11, and one for LUTs 5–12. The inner 2x2 of the double pair connects across a distance of 5 ( $2^{2}$ +1), while the outer 2x2 connects across a distance of 7 ( $2^{2}$ +3). To maintain consistency, we then move the center double

pair from stage 3-4 (dashed circle) down to stage 2-3 (Figure 3.6b), transforming stages 2-3 into a boundary-less interconnect. It is clear that this stage-by-stage transformation can be continued to the top of the hierarchy. Alternatively, the designer may also choose to stop the transformation at any hierarchy, and preserve the remaining upper hierarchies as traditional radix-2 network.



b) with boundary-less radix-3 switches in stage 1-3, and c) rearranged for distributed routing.

From the intermediate result in Figure 3.6b), we have shown that 50% of the wires branching out above stage 1 have been removed, and the wires on the bottom-most stage have doubled. Since the upper-stage wires are long, and the bottom-stage wires are very short, such tradeoff results in significant wire-length reduction for the architecture. Though shown for a 16-LUT example, this methodology can be extended to a network of arbitrary size.

From this illustration, we see that all stages above stage 1 have unevenly distributed

routing: some switches have to connect more routing than others. This scenario occurs because the wires above stage 1 have been reduced by 50%. To form a regular routing pattern, one method is to evenly re-distribute the interconnect routing: the dual routes branching out of stages 1-4 are re-distributed across all switches, resulting in the final routing architecture shown in Figure 3.6c). We see that the re-distributed routes for stages 1-4 use only single 2x2 butterflies, as opposed to the double 2x2 butterflies used below stage 1.

Given the 50% wiring reduction above stage 1, an alternative method of wire redistribution is to prune the number of switches above a certain hierarchy. As shown in Figure 3.7a), one method is to prune the switches in stage 3 by moving some wires to a double wire, reducing the number of stage-3 switches by half. Since the remaining stage-3 switches are centered, this results in shorter interconnect length for stage 3-4, and reduces the number of switches in stage 4 by 50%.

Another method is to prune the switches in stage 4 by moving some wires to a double wire, reducing the number of stage-4 switches by half (Figure 3.7b). This can allow the stage-4 switches to reside on 1 half of the network, which can be useful in reducing the wire length of upper hierarchies. For example, for the 2048-LUT FPGA in Figure 3.2, SM stages 7 and 8 can benefit from this technique because the wires are merged toward the center, where the  $N_{8:2}$  interconnects reside.



Figure 3.7: A boundary-less radix-3 network with switches pruned at a) stage 3 and b) stage 4.

Although the illustrations above use a radix-3 boundary-less architecture as an expansion to radix-2, it is not limited to this case. For example, a radix-6 architecture can be used as an expansion to radix-4; a radix-12 architecture can be used as an expansion to radix-8; and so on.

For the sake of completeness, Figure 3.8a) illustrates a radix-4 Fat Tree using 4x4 switches. Two stages of radix-4 switches are required to implement a 16-LUT network. To construct a boundary-less network, we first identify the wires in stage 1-2 that have a distance of 4  $(4^0)$ : these wires are bolded in Figure 3.8a). These selected wires are then moved down to below stage 1 (Figure 3.8b) to form a boundary-less network in the first stage. The center switches for LUTs 5-12 are radix-6, while LUTs 1-4 and LUTs 13-16 are only radix-5 in this illustration because they rest on the boundary of the network.



Figure 3.8: a) An original radix-4 16-LUT Beneš network and b) with boundary-less radix-6 switches in stage 1.

### **3.4 Fast-Path Interconnect**

In VLSI designs, there usually exists a critical path, that is, a path that is more difficult to meet timing constraints. In most VLSI designs, the vast majority of the paths do not reside on the critical path, but those that are on the critical path usually determine the performance of the entire design. We therefore propose an addition to the interconnect SMs to allow faster performance for critical-path gate: fast path.

In the example in Figure 3.9a), we see an example routing from LUT 2 to LUT 16. One possible route is highlighted. Beneš network offers many path diversity (thus it is rearrangeably

non-blocking; without path diversity, the network offers very limited connectivity (such as [Mrabet06] from Section 2.5)), and we are simply choosing one path as an example. The signal needs to traverse up to stage 4 before U-turning back down. With the addition of fast-path, signals are allowed to travel from the LUT output directly to all SMs within its macro (Figure 3.9b). Therefore, the signal is able to travel directly from the output of LUT 2 to the SM on stage 4, and then U-turning back down. Following the macro-based design methodology highlighted in Figure 3.1, a LUT is placed with all its SMs in one macro during physical design, so adding fath-path routing within the macro does not add any interconnect routing outside the macro.



Figure 3.9: A routing example from LUT 2 to 16 a) without fast path and b) with fast path.

For each point-to-point connection, there is always at least one fast-path available, but other routes that conflict with the fast-path routes must take the slower route. In a timing-driven place-and-route flow, this gives the software tool freedom to choose a faster path for more timing-critical routes.

In cases where routing obstructions occur, it is sometimes still possible to utilize portions of a fast-path, and use regular routing for the remainder of the routes. One such example is illustrated in Figure 3.10a), although it would be ideal to have fast-path directly connected to SM stage 4, the router can still connect fast-path to SM stage 3, and use regular routing to complete the route. In other cases, it is sometime impossible to use any fast-path, and regular routing must be used entirely (Figure 3.10b). Even under such cases, path-diversity allows for many routing choices, and the boundary-less radix-3 network sometimes even allows for fewer SM stages. In Figure 3.10b), one example route requires 4 SM stages, while another requires just 3 SM stages. It is up to the timing-driven P&R tool to select the faster path for timing-critical nets.



Figure 3.10: A routing example with routing obstruction that a) still allows a slower fast-path and b) allowing no fast-path.

#### **3.5** Interconnect Cost vs. Gate Cost

In an FPGA, upper-level interconnects are often required to travel long distances, and it would be beneficial to reduce the number of these nets. On the other hand, interconnect switches are also a dominating factor for chip area, and it would be beneficial to reduce the number of these gates as well. Although it is ideal to reduce both, there also exists a trade-off between these two factors.

From the simple example in Figure 3.11, the two types of SMs have the same gate cost. Actually the 4-input muxes in Figure 3.11b) cost more when implemented as a traditional mux, as it takes three 2-input muxes to implement. As a static parallel mux (Chapter IV), a 4-input mux occupies as much area as two 2-input muxes. The muxes in Figure 3.11a) only allow for odd-to-odd and even-to-even switching, but the SM has double the number of muxes.



Figure 3.11: Two SM design with same gate cost, but a) with more wiring than b).

In terms of connectivity, the design in Figure 3.11a) is superior. For example, if input (1) travels to output (1), the design in Figure 3.11b) will not be able to send another signal in the horizontal direction. But the design in a) is still able to send another signal through output (2) as long as it does not need to route input (3). Overall, design in a) provides more path diversity for

routing.

A different scenario arises when the wire lengths are long, and signals (3) and (4) would need to be buffered (sometimes more than once). When the wires are long and the buffers are large, the signal buffering area can easily outweigh the mux area. In this case, the design in a) is clearly inferior: it requires double the number of buffers but does not provide double the connectivity of b).

For lower-level SMs, where the wiring is short and does not require additional buffers, it is beneficial to use limited-input muxes, but implement more of them to improve path diversity. For upper-level SMs with high wiring cost, it is beneficial to reduce to number of wires, in which case full-input muxes should be used, but fewer should be implemented to save wiring cost.

#### 3.6 Local Interconnect vs. Branch Interconnect

In FPGA, interconnect wiring is expensive, because it contributes to routing congestion and buffer gate area. But local interconnects are much cheaper to implement. In traditional Beneš networks, each SM provides just as much local interconnects as branch interconnects (Figure 3.12), even though interconnects that branch to long wires cost significantly more hardware area. To reduce hardware, it is more effective to prune branch interconnects before pruning local interconnects. Local interconnects alone can also contribute to path diversity. In the example in Figure 3.12 (right), the fastest route from LUT 2 to LUT 14 is using the fast path, but let us assume two downward paths between SM stage 1 and 2 are blocked by other timing-critical signals. In this case, a design with traditional SM switches would be required to take a longer route, but a SM design with more local interconnects (4 in this example) can still provide a downward path for this route.



Figure 3.12: An example where traditional-Beneš based SM experiences local interconnect congestion, whereas a SM design with more local interconnects can utilize the fast path.

An example SM design with 4 local interconnect and 1 branch interconnect is shown in Figure 3.13. When implemented as a SM macro, the local interconnects are contained inside the macro. Compared to the traditional-Beneš based SM design, the new design reduces the interconnect wiring in and out of the macro by 50%, but doubles the local interconnects. Such SM design is very effective for upper-level SMs where the branch interconnects are expensive. This essentially follows the same optimization strategy from Section 3.5: it adds more wires and uses simpler muxes when the wire cost is low, but use larger muxes and fewer wires when the wire costs are high.



Figure 3.13: A switch-matrix example with more local interconnects than branch interconnects.

# 3.7 Micro-architecture of a Switch Matrix

A switch matrix (SM) is the most commonly used building block in the hierarchical FPGA – our FPGA has more than 10x as much SMs as LUTs. It is therefore important to have an SM design that is as small as possible, yet provides sufficient connectivity. Figure 3.14 shows an example SM micro-architecture of a radix-3 SM used in our most recent FPGA (details in Section 3.8). Not surprisingly, a SM consists of simply of a collection of muxes. The number of SM outputs determines the number of muxes it needs, but we need to carefully decide how much connectivity to build into each mux, for that has a large impact on the SM area, which has a significant impact on the final chip area.



Figure 3.14: Internal mux interconnect of an example radix-3 switch matrix.

In Figure 3.14, signals 1–4 are upstream signals. Signals 1 and 2 travel internally inside the SM macro, and signals 3 and 4 are branches. From the mux design of 1 and 2, we see the first pruning heuristic: muxes 1 and 2 are allowed to propagate only signals 1 and 2 upwards, respectively, and both 3 and 4 are allowed. This is because outputs 1 and 2 travel in the same path. Not allowing switching between paths 1 and 2 has minimal impact on routing results, but reduces the mux complexity for 1 and 2. Using this simplification, the incoming signal from branch 3 and 4 will be assigned to path 1 or 2 (or both if decided by the router), and remain in the assigned path until it branches out again.

Similar approach applies for the downward paths: incoming signals can be assigned downward paths 5, 6, or 7, and are not allowed to switch between these paths until the signal branches out again. For U-turns, another simplification can be made. For example, there is no need to U-turn from input 1 and 2 back down to output 5, 6, or 7, because they come from the
same SM. There is never a need to ascend one hierarchy and U-turn back to the same SM. Similarly, output 8 travels back to the same SM that input 3 is coming from, so there is no need to performance that U-turn either; the same case applies to output 9 and input 4.

These micro-architectural techniques are effective in reducing SM complexity, thus reducing area and improving mux performance. But even with these techniques, the muxes still poses a large overhead on area and performance. Many circuit-level techniques are applied to implement these SMs efficiently, as discussed in Chapter IV.

# **3.8** Implementing a 16K-LUT FPGA Interconnect

In the previous 2048-LUT FPGA chip (Section 3.2), the architecture was optimized manually, and two types of SMs are utilized. To fully demonstrate the scalability of hierarchical interconnects, the new FPGA has expanded the interconnect architecture by 10x. Since there is no theoretical method to calculate the optimal connectivity at every level of the hierarchy, we have also developed a software tool to map designs onto our architecture (Chapter VI), which allows us to explore the optimal interconnect architecture using an iterative, closed-loop design process: we explore different architectures by mapping benchmarks and commonly-used designs, then examine the interconnect usage across different SM stages and locations, then refine the architecture accordingly and perform the mapping process again.

This FPGA consists of 16K "LUTs" arranged on a 64×320 array. Because it is a heterogeneous FPGA (Chapter V), each "LUT" is limited not to a look-up table, but is more like a SM macro that provides I/O capabilities: in this case, each SM macro provides 5 inputs and 2 outputs to any CLBs, logic, memory, DSP, or others. For example, a SLICE L CLB contains 30 inputs and 12 outputs, it therefore requires 6 SM macros to implement its interconnect; on the

other hand a DSP CLB requires 165 inputs and 66 outputs, requiring 33 SM macros in a 3×11 array.

The SM architecture of the 16K-LUT FPGA is shown in Figure 3.15 and 3.16. Figure 3.15 illustrates the lower 10 SM levels on a 1-dimensional drawing, although physical implementation is 2-D. Figure 3.16 illustrates the top-level physical architecture, highlighting wiring for the top 5 SM stages. The SM architecture is symmetrical across horizontal bisection, and is composed of 7 types SM macros, ranging from 10 to 14 stages of SM. The bottom 10 stages of SM are common across all SM macros, and are illustrated in Figure 3.15.

The CLB-input requirements in this chip ranges from 30, 35, 165, or 180 inputs, therefore the switch matrix in this architecture is chosen to contain 5 inputs and 2 outputs as a common denominator. From Figure 3.15, it shows each LUT to provide 5 inputs and 4 outputs, that is because each output is multi-casted to both local and branch interconnects, similar to the multicast concept from Section 2.4. The bottom 5 stages of the SM utilizes boundary-less radix-3 interconnect, providing short routing distance to neighbors and providing extra path diversity for the network. Above stage 5, we transition back to a radix-2 network to save interconnect area. Additionally, having all radix-3 network in all hierarchies would make the entire architecture boundary-less, which drastically increases place-and-route time. The current timing-driving routing algorithm is based on breadth-first search, and by having radix-2 in the upper hierarchies, the P&R tool is able to converge more quickly due to its reduced search radius. From our P&R evaluations, a radix-3 to radix-2 transition at SM stage 5 provides sufficient path diversity and routing performance.

This SM architecture uses 2 local interconnects per SM on the upward path, but 3 local interconnects per SM on the downward path. This is due to the assistance of fast-path, which

allows many signals to travel directly from the LUT output to the upper-level SM without occupying local interconnects along the way. This alleviates the routing congestion upwards, but does not alleviate the congestion downwards (the fast-path signals still need to traverse downwards on regular interconnects).

Another key distinction between the lower 5 SM stages and upper stages are the upward branch interconnects. From Figure 3.15, we the lower 5 SM stages to have branching on the upward path, but above stage 5, upward branching has been pruned, and only local upward interconnects remain. The exception is for SM stages 10, 11, and 12, for those stages allow the SM to branch upwards upon the termination of the SM macro. As shown in Figure 3.15, the SMs on the bottom half only have 9 stages, and therefore must branch into the SMs on the top half at stage 10 to continue signal propagation, else the signal would reach a "dead-end". This pruning methodology trades off local vs. branch interconnects: it allows branching when the wire costs are low, therefore providing more path diversity, but for the upper hierarchies, path diversity is sacrificed to reduce interconnect congestion and gate area. However, local interconnects are not pruned even for upper hierarchies, because those wire costs remain low, and having 3 local interconnects on the downward paths provides additional path diversity without increasing the area significantly.



Figure 3.15: 1-D SM architecture of the 16K-LUT FPGA, showing the lower 10 SM stages.



Figure 3.16: 2-D SM architecture of the 16K-LUT FPGA, showing the top 5 stages of wiring.

In the top level, the SM architecture is divided into 40 macros, each containing 512 SM macros. From the iterative interconnect optimization process, we"ve converged to an architecture shown in Figure 3.16. There are 7 types of SM macros, shown in 7 different colors. The most commonly-used SM macro has 9 stages, spanning across rows 2, 4, 7, and 9. The remaining SM macros have 11, 13, or 14 stages of SM (labeled in Figure 3.16). The largest SM has 14 stages, shown in the inset of Figure 3.16. These SMs reside in the center of the top and bottom halves of

the network.

Figure 3.16 also illustrates the mixed-radix implementation in the top level. SM stages 10 and 11 are actually radix 3, but are not boundary-less (with the exception of some stage-10 routing that crosses the horizontal bi-section). This is partially because the number of rows (320) is not a radix-2 number. Without utilizing radix-3 SM, another stage of SM would be required. However, since 320 is not much larger than 256, adding a SM stage appears wasteful. The other reason is due to the wiring cost of stage-14 routing, which needs to span half the height of the FPGA. This results in very long wires, and are very expensive to buffer. To reduce the requirements on the number of stage-14 routing, boundary-less stage-10 routes are implemented along the horizontal bisection. This addition allows gates that are placed near the horizontal bisection to use the shorter, and faster, stage-10 routing. Only the gates that are required to communicate across the entire chip need to occupy stage-14 routing.

The final architecture in Figure 3.16 is arrived through extensive iterative improvements to the architecture. The automated P&R flow (Chapter VI) greatly expedited the evaluation process, and gives us confidence in the routability and performance of the optimized design. The architecture techniques discussed in 3.3–3.6 have greatly improved the routing quality of the interconnect network, and reduced interconnect area. Although we have expanded from 3 types of SM macros to 7, it remains a feasible implementation. The circuit-level implementation of the interconnect are detailed in Chapter IV, and the physical integration details are discussed in Chapter V.

# 5.5 Coarse-Grained CLBs for the 16K-LUT FPGA

Since this chip primarily targets high-throughput communication applications, we have integrated two coarse-grain accelerators. The first block is a 16-core, highly-efficient communications DSP accelerator, reconfigurable to perform many common communications algorithms very efficiently. The 16-core architecture is illustrated in Figure 5.15. Core-to-core communications utilize both local, fast-path interconnects running vertically and horizontally, as well as a 4-stage hierarchical interconnect network spanning the 16 cores.



Figure 5.15: Core schematic and interconnect architecture of a 16-core DSP processor.

<u>Each core is realized using radix-2 butterfly architecture</u>, performing  $2\times2$  matrix computations, called a butterfly-computation element (BCE). This provides the versatility for various fundamental  $2\times2$  operations, e.g. permutation, CORDIC, multiplication-and-accumulation (MAC), unitary transformation, etc. Higher level of integration such as multi-stage pipeline is achievable with multiple cores. Each BCE is designed to be run-time reconfigurable

reached the bottom hierarchy, resulting in a partition size of 1 CLB. The corresponding CLB is then placed in the current partition.

## 6.4 FPGA Routing

Due to the large overhead of interconnect area, FPGA routing is performed on very limited routing resources. In our hierarchical FPGA design, the interconnect architecture is also designed to provide just sufficient routing resources to avoid area waste. As a result, FPGA routing places large emphasis on the quality of the software router. The router need to not only resolve all routing congestions, minimize critical-path, and complete the task in a reasonable (hours of less) run-time even for large designs.

<u>As shown in Figure 6.5, the hierarchical interconnect architecture</u> was implemented to have many path diversities, therefore improving connectivity. However, not all paths result in the same timing performance, as illustrated by the routing preferences. It is generally preferred to travel the shortest routings, using fast-path whenever possible, to reduce overall interconnect capacitance. But in the case of routing congestion, re-routing must be done, and some nets may be required to take non-preferred routes.

Modern routers generally employ global routing before detailed routing. The purpose of global routing is to provide a best-case timing performance of the design, and to estimate routing congestion. Being agnostic to routing congestion, the router is able to perform global routing very quickly, such as using the shortest-path algorithm [Nair82] and [Nair87]. In our hierarchical interconnects, the hierarchical architectures allows for very deterministic global routing. The router may utilize fast-path to perform no branching on the upward path, make a U-turn at the required hierarchy, and the downward path is very deterministic (computed by the radix-2

boundaries).



Figure 6.5: A routing-preference example for a point-to-point connection, LUT (S) to LUT (E).

Global routing gives the router valuable information, such as timing feasibility and routing congestion, but all congestions must be resolved for the design to be realizable. The initial version of our router employs rip-up-and-reroute detailed routing to resolve routing congestions [Dees81]. However, the algorithm we implemented was not timing-driven, and is dependent on the routing order of the nets. Therefore the routing results often have inconsistent timing, and sometimes fail to converge. Unsatisfied with our routing results, we implemented a new routing algorithm to the PathFinder router [McMurchie95, Ebeling95].

The PathFinder is a negotiation-based router that iteratively improves routing congestion by de-touring the lesser-performance-critical gates. It is able to incorporate global routing and detailed routing into a unified algorithm. The first iteration of the router is performed only based on interconnect delay, and not routing congestion, resulting in a minimum-delay design with many routing conflicts. However, the router does not attempt to rip-up the conflicting nets, instead it reroutes the design iteratively, but each successive iteration places a higher cost on routing conflicts. Eventually, the cost of routing through a faster, congested net will outweigh the

# **EXHIBIT 6**

## 27.5 A Multi-Granularity FPGA with Hierarchical Interconnects for Efficient and Flexible Mobile Computing

Cheng C. Wang<sup>1</sup>, Fang-Li Yuan<sup>1</sup>, Tsung-Han Yu<sup>2</sup>, Dejan Markovic<sup>1</sup>

<sup>1</sup>University of California, Los Angeles, CA, <sup>2</sup>Qualcomm, Irvine, CA

Following the rapid expansion of mobile computing in the past decade, mobile system-on-a-chip (SoC) designs have off-loaded most compute-intensive tasks to dedicated accelerators to improve energy efficiency. An increasing number of accelerators in power-limited SoCs results in large regions of "dark silicon." Such accelerators lack flexibility, thus any design change requires a SoC re-spin, significantly impacting cost and timeline. To address the need for efficiency and flexibility, this work presents a multi-granularity FPGA suitable for mobile computing. Occupying 20.5mm<sup>2</sup> in 40nm CMOS, the chip incorporates 2,760 fine-grained configurable logic blocks (CLBs) with 11,040 6-input look-up-tables (LUTs) for random logic, basic arithmetic, shift registers, and distributed memories, 42 medium-grained 48b DSP processors for MAC and SIMD operations, 16 32K×1b to 512×72b reconfigurable block RAMs, and 2 coarsegrained kernels: a 64-8192-point fast Fourier transform (FFT) processor and a 16-core universal DSP (UDSP) for software-defined radio (SDR). Using a mixradix hierarchical interconnect, the chip achieves a  $4\times$  interconnect area reduction over commercial FPGAs for comparable connectivity, reducing overall area and leakage by 2.5×, and delivering a 10-50% lower active power. With coarse-grained kernels, the chip's energy efficiency reaches within 4-5× of ASIC designs.

Although commercial FPGAs can come close to ASICs in performance, they are highly inefficient due to their high energy and a large area overhead. This is mainly due to the programmable interconnect. For over 20 years, a 2D-mesh network has been the backbone of FPGA interconnect, but full connectivity in a 2D mesh requires  $O(N^2)$  switches, requiring interconnects to grow faster than Moore's Law O(N). As a result, various heuristics are used to simplify switches at the cost of resource utilization, but the interconnect area is still ~4× the logic area in modern FPGAs. By effectively pruning a Beneš network, a hierarchical interconnect network is realized where the number of switches is less than  $O(N \cdot logN)$ , allowing us to maintain an interconnect-to-logic-area ratio of 1:1.

The O(N·logN) complexity of Beneš network is well-known in telecommunications, but such a network is seldom used in hardware primarily due to its implementation complexity. In a traditional Beneš, wirelength doubles for every stage. With an equal number of wires for all stages, this leads to long, congested wires in the upper hierarchies. An efficient implementation requires pruning the upper hierarchies, and we alternate the routing in the x- and ydirections so wirelength doubles every two stages [1]. Another drawback is the delay across radix boundaries. As shown in Fig. 27.5.1, communication between neighboring computing elements (CE) 4 and 5 requires 3 hierarchies. A boundary-less radix-3 network is created to restore spatial locality by shifting all local connections to the lower switch matrices (SMs). In the simplified illustration, radix-3 SMs are used in the lower stages to increase local bandwidth, allowing even fewer radix-2 SMs in the upper hierarchies. For improved timing and reduced power, fast-path routing allows hops directly to the required hierarchy level, routing only half the network on the return path. Our router automatically assigns fast-path interconnect based on congestion and timing.

Boundary-less radix-3 SMs are used in the lower 5 hierarchies (Fig. 27.5.2), and pruned radix-2 switches are used from stage 6 to 14, except stage 10 and 11. Stage 10 employs boundary-less radix-3 across the horizontal bisection to improve bisection bandwidth. The top-level connectivity (stage 14) is pruned to only 5%. This is a result of closed-loop optimization by mapping various FPGA benchmarks, then pruning or expanding each stage based on congestion and performance. To ease physical design, the chip is divided into 40 interconnect regions, each with 512 SM macros, with 9 to 14 stages per SM macro.

The fine-grained and medium-grained CLBs offer behavior identical to commercial FPGAs, allowing for a direct comparison of interconnects by executing identical netlists. To target common communications designs, two coarse-grained kernels were implemented. A 64-8192-point reconfigurable FFT is beneficial for digital baseband processing. It has a small dedicated memory, and interconnects to the FPGA memory to realize the long delay lines for 2048-8192-point FFTs. A 16-core UDSP targets a variety of SDR algorithms, where each core is reconfigurable for arbitrary  $2\times 2$  matrix operations using a flexible instruction-set architecture. Unlike the medium-grained DSP processor, the  $2\times 2$ 

butterfly core in the UDSP is very efficient for complex arithmetic, capable of many SDR functions, such as filtering, equalization, CORDIC, and sphere decoding by simply concatenating multiple butterfly stages. FFT and UDSP both connect to the interconnect network.

Power gating (PG) is desirable for large chips, but each interconnect signal often traverses many blocks, making block-level PG ineffective. A fine-grained PG is needed for individual switches. Traditional PG becomes very inefficient because the footer PG transistor is no longer shared by the entire block, so it cannot be made very large (Fig. 27.5.3), but a smaller footer can degrade performance by 30-50%. To power gate without a footer, a PG branch is added to the mux, and the pass-gate is separated into NMOS and PMOS segments, where enabling PG leaves the output floating, reducing the coupling capacitance on neighboring wires. When conducting, the NMOS segment is driven by PMOS pass-gates, thus it can rise much faster than the PMOS segment driven by NMOS passgates, which settles to  $V_{\text{DD}}$ -V<sub>t</sub> (and vice versa). This results in larger transient leakage, but does not degrade performance significantly, because the output current is the *difference* of the pull-up and pull-down branches. A small high-V<sub>t</sub> keeper pulls together the NMOS and PMOS voltages to overcome the V<sub>t</sub> drop. This results in a 5-10% performance penalty, but reduces leakage by more than 50% (now gate-leakage dominated). The output floats during PG, so it cannot drive a CMOS gate, but can only enter a pass-gate that can be disabled during PG. Over 90% of the switches utilized this PG scheme, except those driving long wires that require buffer insertion.

With over 9 million configuration bits, an automated mapping tool is developed. The tool supports two modes (Fig. 27.5.4). Mode 1 maps an identical netlist as used by commercial FPGAs for a direct comparison of performance, power, and area utilization: the user design is first synthesized using commercial tools, then the output netlist is parsed into our custom tool, which performs timing analysis, floorplan, placement, routing, and bitstream generation for our FPGA. Mode 2 incorporates our coarse-grained kernels into the P&R flow. Although the configuration SRAM cells are distributed throughout our FPGA, their word-lines (WL) and bitlines (BL) are organized as one large memory for easy initialization. The FPGA core can only be powered on after configuration finishes.

Measurement results of our FPGA with CLBs, and with coarse-grained kernels are compared against processors, a commercial FPGA, and an ASIC (Fig. 27.5.5). Although the CLBs alone achieve over 1.5GOP/mW, an energy efficiency of 0.86GOPS/mW is achievable when mapping an FIR filter, which is 4× more efficient than commercial FPGA (both in 40nm). An 8× efficiency gain can be achieved by using UDSP kernels. FFT operations, which are dominated by memory and control, are 13× more energy efficient when mapped to the FFT kernel instead of CLBs. A 2-2.5× reduction in leakage is attained from smaller chip area and fine-grained PG, even with the disadvantage of dual-oxide transistors. Our chip is built with standard-cells, yet we are often within 20% of the performance of high-end FPGAs, though our software is still improving.

With efficient interconnect, our FPGA is within  $20 \times$  of ASIC efficiency for most designs (Fig. 27.5.6). Coarse-grained kernels further improve the efficiency, bringing it within 4 to  $5 \times$  of ASICs. The key to coarse-grained efficiency is to identify compact, reconfigurable kernels that improve efficiency, apply to a variety of applications, and leverage existing FPGA resources where possible. Our chip (Fig. 27.5.7) is able to attain the energy efficiency suitable for mobile applications while maintaining the full flexibility of an FPGA.

#### Acknowledgments:

The authors thank Dr. Sanjay Raman and DARPA for funding support.

### References:

[1] C.C. Wang, *et al.*, "A 1.1 GOPS/mW FPGA Chip with Hierarchical Interconnect Fabric," *IEEE Symp. VLSI Circuits*, pp. 136-137, 2011.

[2] Z. Yu, *et al.*, "An 800 MHz 320 mW 16-Core Processor with Message-Passing and Shared Memory Inter-Core Communication Mechanisms," *ISSCC Dig. Tech. Papers*, pp. 64-65, 2012.

[3] "FFT Implementation on the TMS320C5535 DSP," *TI Technical Reference Manual*, pp. 111-134, 2012.

[4] T-H. Yu, *et al.*, "A 7.4 mW 200 MS/s Wideband Spectrum Sensing Digital Baseband Processor for Cognitive Radios," *IEEE J. Solid-State Circuits*, vol. 47, no. 9, pp. 2235-2245, 2012.

[5] F-L. Yuan, *et al.*, "A 256-Point Dataflow Scheduling 2x2 MIMO FFT/IFFT Processor for IEEE 802.16 WMAN," *Asian Solid-State Circuits Conf.*, pp. 309-312, 2008.

[6] J. Thompson, *et al.*, "An Integrated 802.11a Baseband and MAC Processor," *ISSCC Dig. Tech. Papers*, pp. 126-127, 2002.





1	PROOF OF SERVICE
2	<i>Konda v. Markovic, et al</i> Santa Clara Superior Court Case No. 19CV345846
3	I am a resident of the State of California, over the age of eighteen years, and not a party
4	to the within action. My business address is: 184 13 <sup>th</sup> Street, Suite 2, Oakland, CA 94612. On the below-mentioned date, I caused to be served the within documents:
5	FIRST AMENDED COMPLAINT
6	By hand service upon Defendant's counsel of record at the address set forth below during regular business hours.
7 8	By electronically serving the document(s) described above via email to the recipients registered with the e-filing service for Santa Clara County Superior Court pursuant to the Court Order establishing the case website and authorizing
9	<ul> <li>by transmitting via facsimile the document(s) listed above to the fax number(s) set</li> <li>forth below on this date before 5:00 n m</li> </ul>
10 11	<ul> <li>by placing the document(s) listed above in a sealed envelope with postage thereon fully prepaid, in United States mail in the State of California at Oakland addressed</li> </ul>
12	as set forth below.
12	□ by placing a true copy thereof enclosed in a sealed envelope, at a station designated for collection and processing of envelopes and packages for overnight delivery by the United States Post Office as part of the ordinary business practices.
14	of Law Offices of James Farinaro described below, addressed as follows:
15	Steven M. Perry MUNCED TOLLES & OLSON LLD
16	MUNGER TOLLES & OLSON LLPGregory.stone(a)mto.com350 South Grand Avenue, 50th FloorSteven.perry(a)mto.comLos Angeles, CA 90071Steven.perry(a)mto.com
17	I am readily familiar with the firm's practice of collection and processing correspondence for mailing. Under that practice it would be deposited with the U.S. Postal
18	Service on that same day with postage thereon fully prepaid in the ordinary course of business. I am aware that on motion of the party served, service is presumed invalid if postal cancellation
19 20	date or postage meter date is more than one day after the date of deposit for mailing in affidavit.
21	I declare under penalty of perjury under the laws of the State of California that the above is true and correct.
22	Executed on November 7, 2019 at Oakland, California.
23	
24	Venkat KONDA Ph.D.
25	VENKAT KONDA, TU.D.
26	
27	
28	
	PROOF OF SERVICE

Law Offices of James Farinaro 184 13<sup>th</sup> Street, Suite 2 Oakland, CA 94612