

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 29-05-2013		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 3-May-2005 - 2-Mar-2010	
4. TITLE AND SUBTITLE FANTOM : Algorithm-Architecture Codesign for High-Performance Embedded Signal and Image Processing Systems (post-final report)			5a. CONTRACT NUMBER W911NF-05-1-0248		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 5G10U4		
6. AUTHORS Xiaobai Sun, Nikos Pitsianis			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Duke University 130 Hudson Hall, Box 90271 2200 West Main Street Durham, NC 27705 -			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 48809-CS-DRP.1		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT This is the final report of project FANTOM, which was a part of the DESA program supported by the DSO office, DARPA, during the period between May 2005 and Aug. 2009, with a no-cost extension to Feb. 2010. The project team successfully designed, developed and delivered a novel algorithm-architecture co-design platform for semi-automatic or					
15. SUBJECT TERMS signal and image processing, embedded systems, architecture-algorithm codesign, optimization and automation, research translation					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Xiaobai Sun
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 919-684-3030

Report Title

FANTOM : Algorithm-Architecture Codesign for
High-Performance Embedded Signal and Image Processing Systems (post-final report)

ABSTRACT

This is the final report of project FANTOM, which was a part of the DESA program supported by the DSO office, DARPA, during the period between May 2005 and Aug. 2009, with a no-cost extension to Feb. 2010. The project team successfully designed, developed and delivered a novel algorithm-architecture co-design platform for semi-automatic or automatic production of signal and image application-on-a-chip (AOC), to be embedded in a custom system. This reports describes the technical gaps and problems the research project identified and addressed, summaries the major scientific results, inventions and engineering accomplishments, introduces successful employments of the FANTOM system as well as the methodology, and finally discuss on the impacts and implication of the research results. The report concludes with additional comments to relevant program mangers as well as tax payers.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
10	10
11	11
12	12
13	13
14	14
15	15
16	16
17	17
18	18
19	19
20	20
21	21
22	22
23	23
24	24
25	25
26	26
27	27
28	28
29	29
30	30
31	31
32	32
33	33
34	34
35	35
36	36
37	37
38	38
39	39
40	40
41	41
42	42
43	43
44	44
45	45
46	46
47	47
48	48
49	49
50	50
51	51
52	52
53	53
54	54
55	55
56	56
57	57
58	58
59	59
60	60
61	61
62	62
63	63
64	64
65	65
66	66
67	67
68	68
69	69
70	70
71	71
72	72
73	73
74	74
75	75
76	76
77	77
78	78
79	79
80	80
81	81
82	82
83	83
84	84
85	85
86	86
87	87
88	88
89	89
90	90
91	91
92	92
93	93
94	94
95	95
96	96
97	97
98	98
99	99
100	100

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Manuscripts:

Books

<u>Received</u>	<u>Paper</u>
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
10	10
11	11
12	12
13	13
14	14
15	15
16	16
17	17
18	18
19	19
20	20
21	21
22	22
23	23
24	24
25	25
26	26
27	27
28	28
29	29
30	30
31	31
32	32
33	33
34	34
35	35
36	36
37	37
38	38
39	39
40	40
41	41
42	42
43	43
44	44
45	45
46	46
47	47
48	48
49	49
50	50
51	51
52	52
53	53
54	54
55	55
56	56
57	57
58	58
59	59
60	60
61	61
62	62
63	63
64	64
65	65
66	66
67	67
68	68
69	69
70	70
71	71
72	72
73	73
74	74
75	75
76	76
77	77
78	78
79	79
80	80
81	81
82	82
83	83
84	84
85	85
86	86
87	87
88	88
89	89
90	90
91	91
92	92
93	93
94	94
95	95
96	96
97	97
98	98
99	99
100	100

TOTAL:

Patents Submitted

Patents Awarded

Awards

Co-PI M. Kandemir, and his students, won the best paper award
at International Parallel and Distributed Processing Symposium, 2008.

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
-------------	--------------------------

FTE Equivalent:

Total Number:

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
-------------	--------------------------

Tian Xiao	0.00
-----------	------

Paolo Bientinesi	0.00
------------------	------

FTE Equivalent:	0.00
------------------------	-------------

Total Number:	2
----------------------	----------

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
-------------	--------------------------

FTE Equivalent:

Total Number:

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
-------------	--------------------------

FTE Equivalent:

Total Number:

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale): 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

NAME

Kanwaldeep Sobti

Nihshanka Debroy

Total Number:

2

Names of personnel receiving PHDs

NAME

Prasanth Manglagiri

Kevi Irick

Yuanrui Zhang

Srinidhi Kestur

Chili Yu

Lanping Deng

Total Number:

6

Names of other research staff

NAME

PERCENT SUPPORTED

Xiaobai Sun

0.00

Nikos Pitsianis

0.00

Chaitali Chakrabart

0.00

Vijay Narayanan

0.00

Mahmut Kandemir

0.00

FTE Equivalent:

0.00

Total Number:

5

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

see the attached file.

Technology Transfer

FANTOM :
Algorithm-Architecture Codesign for
High-Performance Embedded Signal
and Image Processing Systems
(final report)

Sponsored by
Defense Advanced Research Projects Agency
Defense Science Office (DSO)
Program: DESA

Duration: 05/2005–08/2009
(with a no-cost extension to 02/2010)

May 25, 2013
(final report reconstruction)

Lead Organization

Duke University

Xiaobai Sun (PI, leading author) and Nikos P. Pitsianis (co-author)

Team Members

Arizona State University

Chaitali Chakrabarti

Pennsylvania State University

Mahmut Kandemir and Vijay Narayanan

In memory of Dr. Dennis Healy

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Contents

1	Project Objectives	4
1.1	Technical gaps	4
1.2	Milestone statements	5
1.3	The FANTOM system	6
2	Technical Problems	6
2.1	Mutual algorithm-architecture constraints	7
2.2	Cutting-edge application requirements	7
2.3	Hardware-Software (HS) partitioning on AoC	8
2.4	Optimization-Automation co-dependence	8
2.5	Leap forward from Phase-I to Phase-II	10
3	Major Research Contributions	10
3.1	Specification abstraction & expansion	11
3.2	Chiplet library & chip design hierarchy	12
3.3	Iterative forward-backward mapping	12
3.4	Novel instruction-function generation	12
3.5	Special-purpose mini-compilers	13
3.6	Processing irregular data on regular architecture	13
3.7	Shifting accuracy-efficiency trade-off boundary	14
3.8	System rendering	14
3.9	Accommodation of COTS products	15
3.10	Measure of success	15
4	Research Translation	16
4.1	Direct deployment and impact	16
4.2	Community recognition	16
5	Implications for Related/Future Research	17
6	Additional Comments (intended to program managers)	18
7	Bibliography	19

1 Project Objectives

1.1 Technical gaps

FANTOM represents a novel paradigm shift in design and development of embedded signal and image processing (ESIP) systems on chip, with high throughput and low power consumption (HL) in performance as well as with high turnaround rate and low overall cost (HL) in production.

There were a few large technical gaps, to come across, in design and development of ESIP systems, which are particularly important to defense research and applications.

- the gap between increasing demands for ESIP systems and the lack of means to meet the demands;
- the gaps in multiple aspects between general purpose systems with commodity software and hardware on the one side and embedded systems with customized hardware and software on the other side:
 - (+) performance: latency, accuracy, power consumption;
 - (+) form: size, weight;
 - (+) resource: area in particular
 - (−) development cycle (turnaround) for customized applications;
 - (−) adaptability for updating algorithms or applications or both;
 - (−) complexity in design and development;

where ESIP systems were known to have the potential in the aspects marked by (+), and face the challenges in other aspects marked by (−).

- the gap in knowledge and experience among the designers and developers, namely, the expert, the outmoded and the novice;

The long cycle in design and development, for example, had been a key factor, among others, responsible for the lack of, or slow, updates in many ESIP systems important or critical to defense applications.

Our ultimate goal was to narrow down and bridge these gaps. This, we argued in our proposal, *can* be achieved by providing a semi-automatic or automatic platform to aid the design of ESIP systems. The term *codesign* used

in the project title implies also communication, connection and collaboration between the researchers in different expertise areas, namely, hardware design, software design, modeling, analysis, optimization and algorithms, and domain-SIP applications.

1.2 Milestone statements

PHASE-I.

Deliver an algorithm-architecture co-design platform for production of hardware accelerators, hosted by a personal computer, for geometrically structured matrix-vector products. The design will be competitive in time-power-area performance to existing ones. The turnaround time between a design specification and a customized design, by a senior EE undergraduate student, is (substantially) less than 10% of the time for manual tuning by an expert team.¹

PHASE-II.

Deliver an algorithm-architecture co-design platform for production of application-on-a-chip (AoC), to be embedded in a custom system, for geometrically structured matrix-vector products. The design will be competitive in time-power-area performance to existing ones. The turnaround time between a design specification and a customized design, by a senior EE undergraduate student, is (substantially) less than 10% of the time for manual tuning by an expert team.

The FANTOM-I system is an accelerator, driven by a host machine. Figure 1 is an initial and simple visual display of the FANTOM-I research mission. The FANTOM-II system is an autonomous system, which we describe in the next section. We will also describe in Section 2 the significant forward leap from Phase-I to Phase-II.

A couple of remarks are in order. We envisioned our mission beyond the stated milestones and we indeed far surpassed them. The milestone statements above understated the objectives in a sort of news reporting language, which reflect the main-stream management style at DARPA back then.

¹This milestone was revised by Dr. C. Schwartz when he took over the DESA program from Dr. D. Cochran.

Framework for Accelerating Numerical Transforms On Microchips




Figure 1: FANTOM I: high-level concept diagram.

1.3 The FANTOM system

The FANTOM project successfully rendered a *working* platform for AoC design, by the end of Phase-II. The platform system encodes and embodies our research results as a whole. It integrates vertically across layers of dependency, as much as possible, from the basic operations at the bottom, to the user interface at the top. It also integrates horizontally, as much as possible, between conceptually compartmentalized modules at each layer. In addition, the platform is designed to be adaptable, with regards and in response to the rapid advances in computer technology, computing techniques, and increasing demands or desires of new SIP applications.

We introduce in Section 3 the major research results that underlie the system, and how the accomplishments are measured.

2 Technical Problems

We describe in this section the major technical problems we were to overcome. Most of the problems lied across the boundaries between traditional disciplinary studies as well as those between stages and modules in traditional system design and development. The problems are fundamental and hard, many of them were

not previously addressed.

2.1 Mutual algorithm-architecture constraints

There are constraints imposed by architectures on algorithms and vice versa. For example, sparse and irregular data structures are a hallmark of 'smart' algorithms with low arithmetic complexity (linear or nearly linear complexity) in free space. Operations with such data structures were not well supported by modern architectures in hardware, which favor regularly strided data accesses and operations, such as array operations, and add tremendous extra cost in implementation and execution of the smart algorithms.

In general computational practice, we must deal with the following discrepancies

- in data locality : memory hierarchy vs. algorithm hierarchy

The memory hierarchy is in principle organized for the benefit of temporal data reuse and spatial data reuse. A tree-like algorithm hierarchy poses a great challenge on data reuse both temporally and spatially, because the temporal extension of a cluster of data requires the use of other data clusters. See Figure 2.

In addition, the memory hierarchy is regularly structured, while the algorithm hierarchy may be highly irregular depending on the data distribution.

- in parallelism: dependency-concurrency of an algorithm in free space vs. parallel operations and patterns supported or favored by architectures. Again, depending on the data distribution, the particular dependency-concurrency graph of algorithm execution may be highly irregular.

2.2 Cutting-edge application requirements

The application requirements may be described in terms of accuracy, latency, power consumption, and resource (area) consumption. Often, the requirements seem conflicting and infeasible with the existing techniques. These requirements drive the research to new frontiers.

Among other important and influential applications, we used the SAR image formation as a case study. Figure 3 depicts a key processing component for SAR image formation. High-efficiency, high-resolution image formation remains a




Figure 2: Mapping between Algorithm Hierarchy and Memory Hierarchy

significant challenge for modern power/volume constrained weapons systems. It requires broad-band frequency ranges and high image formation speed but with data sampled on non-Cartesian grids.

2.3 Hardware-Software (HS) partitioning on AoC

An AoC system includes FPGAs and CPUs. The most challenging task in mapping an SIP application, such as an SAR application, lies in partitioning the computation tasks between the reconfigurable fabric and multiple CPU cores. The CPUs may be internal or external to the FPGA. Furthermore, when internal, the CPUs may be software cores or hard cores. See Figure 4.

2.4 Optimization-Automation co-dependence

Modeling, estimation and optimization were conventionally restricted to each operation module and often carried out manually, by expert designers. Joint




Figure 3: 2D Translation and Transformation for SAR image formation, provided by research personnel at the Radar Signal Processing Center at Raytheon Missile Systems

optimization at each integration stage, and across the stages, is hard and often skipped. The lack of automation in practice was partially due to the lack of modeling, analysis and optimization. On the other hand, over a complex design space, optimization must be assisted by automation.

This includes, for example, the joint optimization in accuracy and efficiency. Accuracy is affected by numerical ranges in input, output and intermediate data, and data representation on an AoC (format and precision). In conventional ESIP system design, the fixed-point representation with few bits was much preferred for efficiency.

Conventional methods confined the accuracy and throughput, for instance, in a tight trade off space. A higher throughput was to be traded off with a




Figure 4: Hardware and Software Partitioning and Codesign.

lowered accuracy.

2.5 Leap forward from Phase-I to Phase-II

We highlight in Figure 5 the major changes from the first to the second generation of FANTOM system, broken down in terms of the system components. In the rest of the section we elaborate on the challenging issues in FANTOM-II and our approaches to the solution.

3 Major Research Contributions

We have identified in Section 2 the major gaps and the fundamental technical problems to overcome. In this section, we describe briefly a few major scientific results, inventions and engineering accomplishments by the research project. In other words, we describe the conceptual and engineering components the FANTOM system is based upon and made of. We annotate each item with the development Phase, I or II, in which it was completed. Most of the details can




Figure 5: Advances from FANTOM-I to FANTOM-II with Major/Medium/Minor Changes Color-Coded in Blue/Green/Yellow, Respectively.

be found in the FANTOM bibliography, where the most closely related works by other researchers are cited.

3.1 Specification abstraction & expansion

Conventionally, an AoC design specification is provided in code in a programming language such as C. We refer to such method as the code specification. First, a code specification is subject to the expression scope of the language, the programmer’s understanding of the computation task as well as the programmer’s capability to describe the task with his or her command of the language usage. Next, the code specification is *instantiated* in hardware instruction by instruction, basically. A translation word by word, phrase by phrase, is not even a good approach for natural language without regard to the context and structure. In other words, the code specification and instantiation approach is the first spot that not only narrows down the design space but also deforms it

most likely.

We broke this conventional barrier right at the beginning. We specify the application algorithm in a high-level description, in particular, in MATLAB with FANTOM annotations. No architectural constraints are imposed unnecessarily in task specification. We then expand the specification at different levels of detail.

3.2 Chiplet library & chip design hierarchy

Exploiting the very same idea behind software libraries, we developed the chiplet library, which consists of parametric chip models for the basic and primitive operations. The importance of the library becomes evident during the project period, in which the reconfigurable fabric hardware was changed and updated multiple times. These changes at the very bottom of the system design did not crash and collapse our system thanks to the chiplet library we had abstracted and established first. Most of the changes will be located at the chiplet levels. There are two types of chiplets, basic operations (bricks) and basic compositions (mortar). In Phase-I.

3.3 Iterative forward-backward mapping

We established a formal description of the algorithm-architecture codesign space in terms of forward and backward mappings, a systematic modeling framework for the two-way mappings, and an efficient and effective approach for efficient and effective performance estimation and adaptive search for optimal design(s), in Phase-I.

For instance, the mapping between the potential concurrency in the fast multipole method (FMM) and the parallel patterns favored by hardware structures (Figure 2) takes a few iterative steps.

3.4 Novel instruction-function generation

We invented and developed a novel and unique FPGA design system that enables *recursive* and *nested* generation of functions modules with increasing complexity, within the restricted resources on FPGAs. Designed and developed in Phase-II.

This recursion idea and technique allows model functions constructed, generated and encapsulated shell by shell, with the same fixed and very limited amount of resource on an FPGA. Roughly speaking, a function at the inner




Figure 6: FANTOM special-purpose (SP) compiler to automate the generation of lookup-table based implementations for convolution kernel functions, under accuracy and resource constraints and two architectural-aware geometric-tiling-based parallelization strategies for convolutions, targeting both uniformly-distributed and non-uniformly distributed input samples.

shell can be used as if it were a native instruction to the composition at the next level. If this might sound similar to a nested software design, think about putting a strong memory limitation on the entire instruction set.

3.5 Special-purpose mini-compilers

Most of the above is enabled and automated with our compiler techniques. We developed special-purpose compilation techniques in multiple stages for translating, interpreting, transforming, and mapping from application-specific algorithm specification at a graphical user interface, all the way, to a high-performance instantiation on FPGAs. For accelerators in Phase-I, for systems-on-chip in Phase-II.

3.6 Processing irregular data on regular architecture

We developed the very first system design platform that permits modern SIP applications with irregular sampling, a significant leap from simulated or post-

processing studies. As mentioned in Section 2.2, this is important to SAR image formation. See Figure 7. The initial module is for accelerators in phase I, it is extended to systems-on-chip in Phase II.




Figure 7: Multi-resolution image formation by adaptive formation from non-uniformly sampled data (simulation result)

3.7 Shifting accuracy-efficiency trade-off boundary

We established a distinct methodology for co-optimization in accuracy and speed performance (subject to resource constraints), which shifts the trade-off boundary well beyond the traditional ones, in Phase-I.

For example, we may represent a filter kernel and filtered data in two parts. We use economic representation for the first part. Specifically, we make a combined use of low-bit data representation and lookup tables. The low-bit representation includes fixed-point representation, or blockwise fixed-point representation, or customized low-bit floating-point representation. We then use fast processing during execution to get the second part and reach the required accuracy.

3.8 System rendering

We produced a video screencast to introduce, demonstrate and give a tutorial on the system. S. Kestur, a graduate student then, narrated in his gentle voice

	GRAPE-6	FANTOM-I
Device type	ASIC	FPGA(XC2VP100-6)
Device technology (μ m)	0.25	0.13
PEs/chip	6	3
Frequency(MHz)	90	125
Real Peak (GFLOPS)	17.2	22.2
Power Consumption (W)	12	6.5

Table 1: Performance comparison with GRAPE-6

step by step how to use the system. The late manager, Dr. Healy, was very pleased with it.

3.9 Accommodation of COTS products

In addition to the above efforts and accomplishments, we have also envisioned and engaged in FANTOM research on accommodation of emerging commercial-off-the-shelf (COTS) products. The COTS products include game or graphics processors by IBM, AMD, and nVIDIA. They are between general-purpose processors and special-purpose processors. In Phase II.

3.10 Measure of success

FANTOM accomplishments have been measured in three different ways.

First, as promised in the milestone statements, we carefully set up a comparison environment at the Department of Electrical and Computer Engineering, PSU. The comparison results were as expected and reported in DARPA reviews, in Phase-I and Phase-II.

Next, we took the conventional benchmarking approach and made a comparison to GRAPE-6, in Phase-I, to the best special-purpose computer, made in Japan and supported fully by the Japanese government, for molecular dynamics simulation, See Table 1.

Finally, more excitingly, directly and remarkably, the FANTOM system and the methodology have been employed, and hence tested, since 2010 by other research and development projects, at DARPA and elsewhere. In particular, three research projects under the NeoVision2 program at DARPA employed the FANTOM system and methodology. Each project team had their own ESIP design experts. With the aid of the FANTOM system and methodology, each

of the teams reached a reduction in power consumption by a factor of $O(10^3)$ and a faster processing speed, without compromising accuracy. We describe in Section 4 more of FANTOM research translation and impact.

4 Research Translation

4.1 Direct deployment and impact

More directly, the FANTOM platform and methodology have been employed and applied to multiple projects at DARPA and elsewhere, in particular, by the following three research teams under the NeoVision2 program at DARPA

- M. Peot’s team at Teledyne Scientific Imaging, Inc.
- L. Itti’s team at the Univ. of Southern California
- D. Khosla’s team at HRL Laboratories, LLC.

Other research groups at the Office of Naval Research and at Intel Co. have been interested in exploring with and exploiting FANTOM methodologies. Three of the co-PIs have been approached by Intel researchers for potential collaborations on commercial applications.

4.2 Community recognition

In addition to the publications listed in the Bibliography section, all co-PIs have given FANTOM talks at various conferences. In particular,

- Co-PI V. Narayanan has been highly visible and influential for his major contributions in FANTOM project, among his other projects. He was invited as a keynote speaker at many conferences, including
 - International Symposium on High Performance Computing Architecture, Jan 2010
 - FETCH 2012, Alpe dHuez, France,
 - VLSI Design Conference, January 2013, Pune India
 - 4th Workshop on SoCs, Heterogeneous Architectures and Workloads, (SHAW-4). February 24th 2013, Shenzhen, China.

Duke	Nihshanka Debroy Paolo Bientinesi Tian Xiao	MS postdoc postdoc	Deloitte Consulting RWTH Aachen University Wave Computation Technologies
ASU	Lanping Deng Chi-Li Yu Kanwaldeep Sobti	PhD PhD MS	Oppo Digital Marvell Semiconductors AMD
PSU	Jungsub Kim Prasanth Mangalagiri Kevin Irick Yuanrui Zhang Srinidhi Kestur Sungho Park	PhD PhD PhD PhD PhD PhD candidate	Samsung Intel Co. Silicon Scapes (founder & CEO) Intel Co. Intel Co.

Table 2: Placement of former students and postdocs partially or fully supported by FANTOM

- Workshop on Neuromorphic and Brain-Based Computing Systems (Neu-Comp 2013), Grenoble , March 2013
- Co-PI M. Kandemir, and his students, won the best paper award at International Parallel and Distributed Processing Symposium, 2008.
- Two conference papers on project FANTOM got special invitations as contributed papers in special issues of high-influence journals.

All FANTOM students and postdocs were recruited by highly competitive research institutes or groups, see Table 2 for the placement. In particular, Dr. Kevin Irick started a company Silicon Scapes and he has been the CEO.

5 Implications for Related/Future Research

We speculate that more FANTOM-like methodologies are needed in the next decade, at least, for developing small, smart and special-purpose ESIP systems, considering the following factors.

- At the time project FANTOM started, there was no iPhone. The emerge of iPhone in June 2009 was in the finishing days of the FANTOM project. Developers and researchers at large, began to realize, not much ahead of

the populace, the importance and great potential with small, smart and special-purpose ESIP systems.

- The approach to the limitation of Moore’s law implies that we must not continue the rich man’s way to use hardware resources with ever swelling software packages.
- While mobile computing devices are enabled by AoC techniques, mobile communicating techniques will expand and advance the AoC techniques for larger and larger scale processing.

6 Additional Comments (intended to program managers)

FANTOM research results are remarkable, especially, under the following conditions. A small team (5 co-PIs, 2 postdoctoral years, 10 graduate students, in total), a short time span (05/2005-02/2010, including no-cost extension of 6 months), and a modest budget. Most of the FANTOM publication was done in the late stage of Phase-II, not only because of the fruitful results but also thanks to a welcome change at DARPA from the extreme product driven, deadline driven style during the earlier years of FANTOM project.

In the 4-year project duration, the project management of FANTOM was passed among the hands of 4 program managers at DARPA. We the co-PIs thank Dr. D. Cochran for having the initial vision and putting the researchers in different research areas together. We thank Dr. C. Schwartz for managing the project with professional appreciation and intense passion. We remember late Dr. D. Healy, who passed away in Sept. 2009, with absolute respect and admiration for his insightful and gentle guidance and inspiration, for his dedication to scientific research, to DARPA and to researchers on DARPA projects. We thank Dr. A. Kane for taking the position left behind by Healy a year after reviewing our final project report in Jan. 2011.

At the test and validation stage of the project. FANTOM’s partner at Raytheon quit from his company and hence from this project. Fortunately, FANTOM was tested in other and perhaps more effective ways, see Section 3.

Many end-products of DARPA projects fall into the category of embedded systems, The FANTOM research results can be applied to more and on-going DARPA research projects, if the results are broadly introduced via program managers. We have seen and heard of some projects struggling with FPGA implementation by learning from scratch.

We have developed the FANTOM system and methodology and provided important and critical assistance to broader application system designs. We also emphasize the importance of maintaining the leading researchers and active research in ESIP system design and development. Systems must be updated or transformed or replaced one way or another, sooner or later. Through continuously active and advanced research activities, we make new advances and foster a new generation of researchers. In this aspect, the most important FANTOM result is the FANTOM students, who are hot recruit targets and new entrepreneurs.

This report is requested by ARO for the official closure in paper work of the research project. Because of the loss of some FANTOM data and files due to a failed workstation in 2011, the authors took extra time and efforts to reconstruct this report from the remaining and published material.²

7 Bibliography

1. G. Chen, L. Xue, J. Kim, K. Sobti, L. Deng, X. Sun, N. Pitsianis, C. Chakrabarti, M. Kandemir and N. Vijaykrishnan, "Geometric Tiling for Reducing Power Consumption in Structured Matrix Operations," in the Proceedings of IEEE International SOC Conference, pp 113-114, 2006.
2. K. Sobti, L. Deng, C. Chakrabarti, N. Pitsianis, X. Sun, J. Kim, P. Mangalagiri, K. Irick, M. Kandemir and V. Narayanan, "Efficient Function Evaluations with Lookup Tables for Structured Matrix Operations," Proceedings of IEEE Workshop on Signal Processing Systems, pp 463-468, 2007.
3. J. S. Kim, P. Mangalagiri, K. Irick, M. Kandemir, V. Narayanan, K. Sobti, L. Deng, C. Chakrabarti, N. Pitsianis and X. Sun, "TANOR: A Tool For Accelerating N-Body Simulations On Reconfigurable Platform," Proceedings of FPL 2007, 17th International Conference On Field Programmable Logic And Applications, pp 68-73, 2007.
4. X. Sun and N. Pitsianis, "FFTs of Arbitrary Dimensions on GPUs," in the Proceedings of 11th Annual Workshop on High Performance Embedded Computing 2007.

²If anyone at ARO or DARPA reads this report, please drop a line or a note at xiaobai@cs.duke.edu. The authors would highly appreciate that.

5. P. Bientinesi, N. Pitsianis and X. Sun, "Multi-dimensional Array Operations for Signal Processing Algorithms," Proceedings of PARA 2008: 9-th International Workshop on State-of-the-Art in Scientific and Parallel Computing, May 2008.
6. N. Debroy, N. P. Pitsianis and X. Sun, "Accelerating nonuniform fast Fourier transform via reduction in memory access latency," Proc. SPIE 7074, 707404 (2008).
7. J. S. Kim, C.-L. Yu, L. Deng, S. Kestur, V. Narayanan and C. Chakrabarti, "FPGA Architecture for 2D Discrete Fourier Transform based on 2D Decomposition for Large-Sized Data," IEEE Workshop on Signal Processing Systems, pp 121-126, Spring 2009.
8. L. Deng, C. Chakrabarti, N. Pitsianis, and X. Sun, "Automated optimization of look-up table implementation for function evaluation on FPGAs," Proc. SPIE 7444, 744413 (2009).
9. K. M. Irick, M. DeBole, S. Park, A. Al Maashri, S. Kestur, C-L. Yu, N. Vijaykrishnan, "A Scalable Multi-FPGA Framework for real-time Digital Signal Processing", Proc. of SPIE, Vol 7444, Aug 2009.
10. Y. Zhang, M. T. Kandemir, N. Pitsianis, X. Sun, "Exploring parallelization strategies for NUFFT data translation," EMSOFT 2009, pp 187-196 (2009).
11. J. Kim, L. Deng, P. Mangalagiri, K. Irick, K. Sobti, M. Kandemir, V. Narayanan, C. Chakrabarti, N. Pitsianis, X. Sun, "An Automated Framework for Accelerating Numerical Algorithms on Reconfigurable Platforms Using Algorithmic/Architectural Optimization", IEEE Trans. on Computers Vol. 58 no. 12, pp 1654-1667, (December 2009).
12. Y. Zhang and M. Kandemir, "A Hardware-software Codesign Strategy for Loop Intensive Applications," In Proceedings of IEEE Symposium on Application Specific Processors. (SASP 2009) pp 107-113, 2009.
13. Y. Zhang, J. Liu, E. Kultursay, M. Kandemir, N. Pitsianis and X. Sun, "Scalable parallelization strategies to accelerate NuFFT data translation on multicores," Euro-Par'10 Proceedings of the 16th international Euro-Par conference on Parallel processing: Part II, LNCS 6272, pp 125-136, 2010.

14. Y. Zhang, P. Yedlapalli, S. Muralidhara, H. Zhao, M. Kandemir, L. Deng, C. Chakrabarti, N. Pitsianis and X. Sun, "A Special-Purpose Compiler for Look-Up Table and Code Generation for Function Evaluation", Design Automation and Test in Europe (DATE 2010), Dresden Germany, March 8-12, 2010.
15. S. Kestur, S. Park, K. M. Irick, V. Narayanan, "Accelerating the Nonuniform Fast Fourier Transform using FPGAs", Proc. of Intl. IEEE Symposium on Field-Programmable Custom Computing Machines, Charlotte, NC, May 2010.
16. V. Narayanan, A. Al Maashri, K. M. Irick, M. DeBole, S. Park, "AutoFLEX: A Framework for Image Processing Applications on Multiple-FPGA Systems", Proc. of Intl. Conf. on Engineering of Reconfigurable Systems and Algorithms, pp 59-66, July 2010.
17. C.-L. Yu, C. Chakrabarti, S. Park and V. Narayanan, "Bandwidth-intensive FPGA Architecture for Multi-dimensional DFT," Proc of the Int Conf on Acoustics, Speech and Signal Processing, pp 1486-1489, March 2010.
18. S. Kestur, K. M. Irick, S. Park, A. Al Maashri, V. Narayanan, C. Chakrabarti, "An Algorithm-Architecture Co-design Framework for Gridding Reconstruction using FPGAs", pp 585-590, Design Automation Conference 2011.
19. C.-L. Yu, K. Irick, C. Chakrabarti and V. Narayanan, "Multidimensional DFT IP Generator for FPGA Platforms," IEEE Trans on Circuits and Systems I, pp 755-764, April 2011.
20. L. Deng, K. Sobti, Y. Zhang and C. Chakrabarti, "Accurate Area, Time and Power Models for FPGA-based Implementations," Signal Processing Systems, 63(1), pp 39-50, 2011.
21. C.-L. Yu and C. Chakrabarti, "Transpose-free SAR imaging on FPGA Platforms," Int. Symp. on Circuits and Systems, pp 762-765, 2012.