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 the burden, to the Department of Defense, Executive Service Directorate (0704-0188). 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 ORGANIZATION.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
22-03-2017	Final Technical Report	01/01/2016 - 12/31/2016

4. TITLE AND SUBTITLE
Heterogeneous IP Ecosystem enabling Reuse (HIER) Final Report

5a. CONTRACT NUMBER

5b. GRANT NUMBER
N00014-16-1-2039

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Dr. Jeffrey Draper

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Southern California (USC)
3720 S. Flower Street, Third Floor
Los Angeles, CA 90089-0001
USA

8. PERFORMING ORGANIZATION REPORT NUMBER
FINAL

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Office of Naval Research (ONR)
875 N. Randolph Street, Suite 1425
Arlington, VA 22203-1995

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for Public Release; distribution is Unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
In the Department of Defense (DoD) world, an IP re-use infrastructure has been lacking, even in the digital domain. Significant investments in custom chip designs have been made by the government, but the IP resulting from such efforts is not readily available for re-use, and even in cases where IP is available, porting to a common implementation platform for integration is often cost-prohibitive. Thus, an execution model and infrastructure to enable DoD-specific IP re-use is greatly needed. To address this critical gap, the University of Southern California conducted an exploratory effort to formulate the detailed requirements for accomplishing a successful Heterogeneous IP Ecosystem enabling Reuse (HIER). This effort explored chip fabrication process issues as well as tools and methodologies. The effort proposed and evaluated the concept of operations and a cost model for achieving a functioning infrastructure for heterogeneous IP reuse. The results of the study include an identification of where major investment is needed to make such a paradigm be as seamless as possible.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
					Jeff Draper
U	U	U	SAR	18	19b. TELEPHONE NUMBER *(Include area code)* 310-448-8750

Heterogeneous IP Ecosystem enabling Reuse (HIER)

Office of Naval Research (ONR) Award N00014-16-1-2039

Final Report

March 22, 2017

Submitted to:

ONR Elec Sensors & Networks Research Div
Attention: Dr. Brian Bennett
875 N Randolph Street
Code 31, Division 312
Arlington, VA 22203-1995
E-mail: brian.r.bennett@navy.mil

ONR REG Office San Diego
140 Sylvester Road, Bldg 140, Room 218
Code N66018
San Diego, CA 92106
Telephone: (619) 221-5490

Naval Research Laboratory
ATTN: CODE 5596
4555 Overlook Avenue SW
Washington, DC 20375-5320

Defense Technical Information Center
8725 John J Kingman Road Ste 0944
Fort Belvoir, VA 22060-6218

Office of Naval Research
ATTN: ONR BDCC, Code: N00014
One Liberty Center
875 North Randolph Street, Suite 1425
Arlington, VA 22203-1995

Submitted by:

University of Southern California – Information Sciences Institute
4676 Admiralty Way, Suite 1001, Marina del Rey, CA 90292-6601

Technical Point of Contact	Administrative Point of Contact
Co-PIs: Dr. Jeff Draper, draper@isi.edu Mr. Wes Hansford, hansford@isi.edu	Brigidann Cooper
University of Southern California Information Sciences Institute 4676 Admiralty Way, Suite 1001 Marina del Rey, CA 90292 Tel: (310) 448-8750, FAX: (310) 823-6714 draper@isi.edu	USC Information Sciences Institute 4676 Admiralty Way Suite 1001 Marina Del Rey, CA 90292 Tel: 310-448-9161 Email: brigidannc@usc.edu
Lead Organization and Reference Numbers: DUNS: 072933393; CAGE: 1B729	University of Southern California, Information Sciences Institute

Approved for Public Release; Distribution is unlimited.

Technical Approach and Justification

IP reuse is a cornerstone of the commercial electronics market particularly in the digital domain. Modern digital design involves billions of transistors, leveraging a large amount of commercially available and silicon proven IP because it is impossible to design at the individual device level for such scales. This IP typically consists of I/O's, high speed interfaces, memories, logic cores and even mixed-signal blocks. An end-to-end commercial infrastructure has been developed to support availability and access to such IP to enable designers to create complex designs with first pass success. In the DoD world, such an IP re-use infrastructure (of DoD-funded IP) has been lacking, even in the digital domain. Significant investments in custom ASIC designs have been made by the government, but the IP resulting from such efforts is not readily available for re-use, and even in cases where IP is available, porting to a common implementation platform for integration is often cost-prohibitive. Thus, an execution model and infrastructure to enable DoD-specific IP re-use is greatly needed. While such an effort is more of an infrastructure development rather than a research endeavor, it would pay handsome dividends to the DoD with respect to more efficient, lower cost chip design efforts in the future.

IP re-use for heterogeneous integration is even more challenging. IP from widely disparate technologies including silicon CMOS/BiCMOS, compound semiconductors including InP/GaN/GaAs/InGaAs need to be properly modeled and simulated in an integrated environment. Such simulations need to also take into account the various heterogeneous packaging involved including 2.5D interposers and 3DIC integration. Developing the infrastructure to simulate and "silicon-prove" IP for heterogeneous integration is the most challenging (and highest payoff) aspect of such an effort.

The University of Southern California conducted an exploratory effort to formulate the detailed requirements for accomplishing a successful Heterogeneous IP Ecosystem enabling Reuse (HIER). The HIER project explored both fabrication process issues as well as tools issues. The results of the study identified where major investment is needed to make such a paradigm be as seamless as possible.

In the course of the HIER project, DARPA also established additional concepts in the formation of the Common Heterogeneous Integration and IP Reuse Strategies (CHIPS) program. In response to the request for information and broad agency announcement associated with the CHIPS program, the HIER project refined its approach to address requirements for that program.

Research Plan

The original vision for the HIER activity involved the evaluation of current process and tool barriers, development of research plans to fill these gaps, and analysis of the efforts involved in the research plan to determine which items are time-intensive, cost-intensive, or both to formulate an appropriate schedule for solving the problem. The key challenges in establishing the HIER paradigm are:

- Developing integration process technology that is broad enough to encompass integration of IP regardless of implementation technology yet adheres to a standard to facilitate seamless integration and

- Developing tool flows that lower the barrier to entry for even low-volume applications to cost-effectively take advantage of the HIER integration technology.

While activities addressing the original HIER vision were conducted, the HIER project adapted to address the needs of the upcoming DARPA CHIPS program. The research plan to support this program revolved around the concept of an infrastructure and business model for maintaining and distributing chiplet IP and associated documentation and simulation models that will be self-sustaining once DARPA investment concluded. A centralized repository of such IP is necessary for such a model to enjoy widespread adoption. As part of MOSIS' well-established experience as a successful non-profit service enterprise with both commercial and DoD customers, there is an extensive infrastructure that can be leveraged for accomplishing the chiplet IP management and distribution envisioned. An example of the current functioning of the MOSIS organization as a non-profit broker is shown in Figure 1. As part of its operation, MOSIS already manages distribution of commercial IP, process design kits (PDK), and physical chips, including the management of packaging with MOSIS vendors under MOSIS control. With minor adjustments, MOSIS can apply the same infrastructure to DoD IP management, in this case brokering NDA relationships between chiplet providers and customers, as well as managing the payment of chiplet IP providers when customers purchase some chiplets from the MOSIS-maintained chiplet inventory.

To enable this vision of chiplet reuse, MOSIS would extend infrastructure, as needed, to support all documentation and simulation models for the CHIPS chiplets. It would also integrate a system for maintaining and distributing the physical chiplets to customers into its existing chip distribution scheme. It would archive, update and distribute such IP to DoD-approved organizations in cooperation with the IP chiplet providers. Additionally, the MARINA research group at ISI would vet all chiplet IP in the inventory from a designer's perspective by conducting design experiments to validate that all models provided for chiplets can be integrated with other models using the same standard interface in a simulation environment.

In summary, the necessary components for a successful chiplet-based IP reuse model include:

1. A centralized infrastructure for brokering NDAs between chiplet providers and customers.
2. A centralized repository for maintaining and distributing chiplet documentation, simulation models, and other files related to integrating chiplets into design flows.
3. A centralized facility for distributing physical chiplets to customers.
4. A process for vetting associated design files for chiplet IP before the IP is officially added to the repository.
5. A process for tracking chiplet sales and forwarding payment to chiplet providers on a regular basis.
6. A business model for sustaining the infrastructure beyond initial DARPA investment.

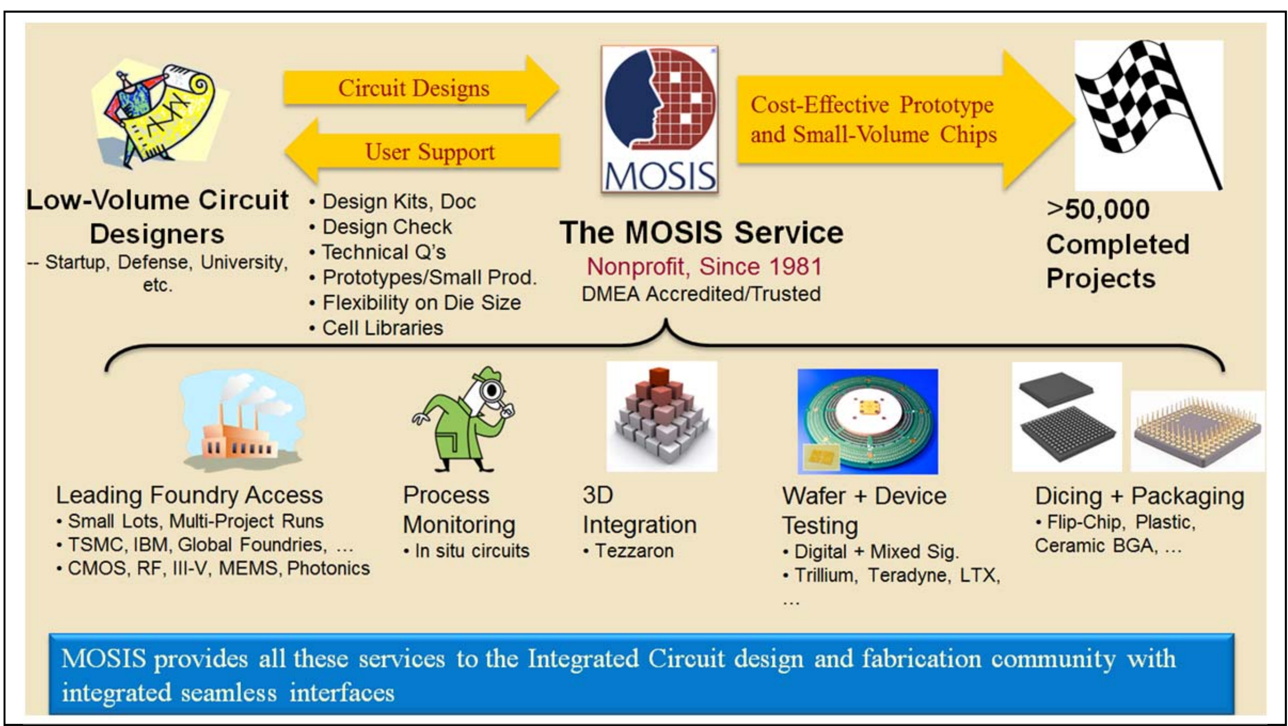

Figure 1: MOSIS Overview

Results of HIER Study

In addition to the HIER study results given in the outbrief presentation shown in Appendix A, the project also captured a high-level research plan that would ideally be executed to optimize an IP reuse model targeted at DoD. Much of this material was provided in a response to the RFI for the CHIPS program but is repeated here for the sake of completeness.

The example model for IP reuse as envisioned by the DARPA CHIPS program of IP instantiated as chiplets that can then be integrated into modular platforms promises to greatly reduce system implementation times, and therefore cost, while also delivering performance near what could be attained with system-on-chip (SoC) integration. An example of the impact of using such a model to implement a processor previously developed at USC is shown in Figure 2 below. Clearly, results will be very design dependent, but even if some custom design is involved, as was assumed for interfaces in the USC processor, we still expect design times and costs to reduce by at least a factor of 2. For more complex designs, where the baseline design times and costs are much greater than the simple USC processor, we expect improvements in design time and costs that could approach even 10x if the entire design can be composed of existing chiplets.

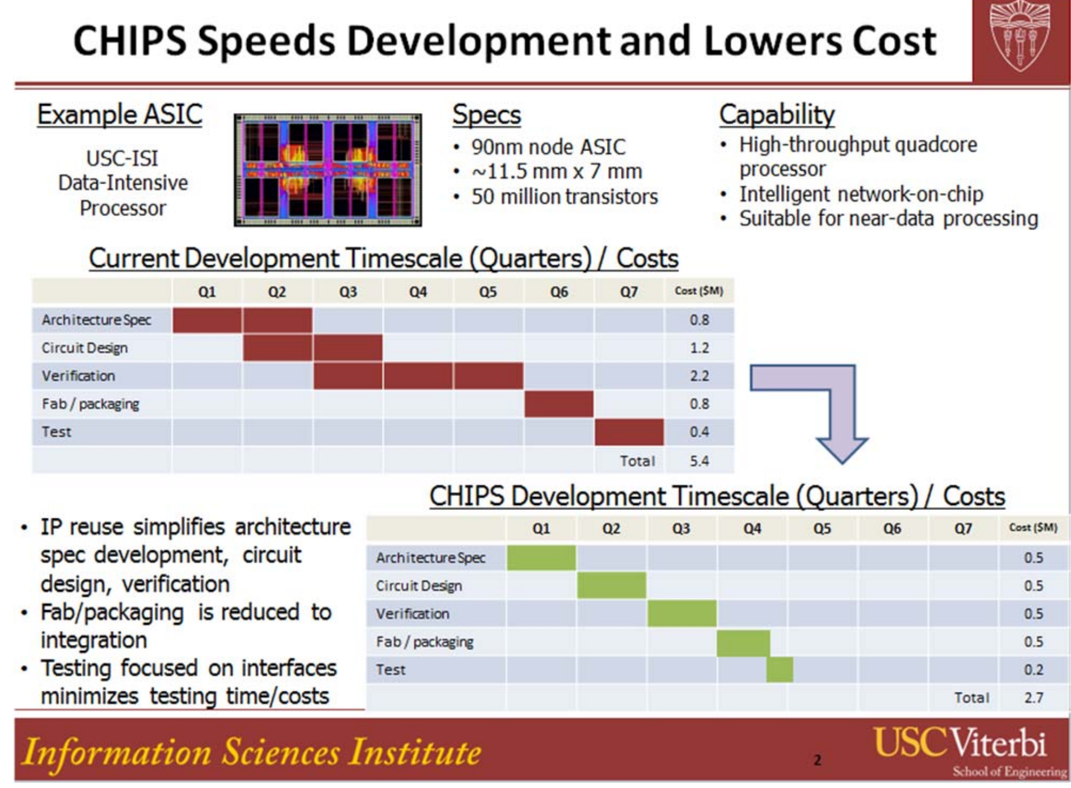

Figure 2: Potential Impact of IP Reuse on USC Processor

The challenges for the CHIPS chiplet ecosystem can be largely grouped into categories that mirror the design and fabrication of a chip itself: architecture and implementation. While the challenges associated with the implementation may be more numerous, mostly due to the detail involved at that level, the architecture challenges are far more important. What small set of IP block chiplets should be developed so that a vast majority of future DoD systems could be implemented simply by interconnecting chiplets from this set? Clearly there are broad common functionality categories that are prevalent in DoD electronic systems, such as processors and sensors, but ascertaining how generic versus parameterizable (or configurable) to make even these types of chiplets attractive to a wide set of applications is very challenging.

A research program that truly addresses this ecosystem architecture part of the problem will need to:

1. Identify and detail a set of DoD applications that represents a large majority of all DoD applications and where the CHIPS concept is likely to have the most impact. Examples may be (Radar, EW, Radio, etc.).
2. Define common functionality among the applications that leads to a small set of chiplet types in the CHIPS IP inventory. Define chiplet types in a manner such that chiplet granularity boundaries can be optimized in subsequent evaluations.
3. Develop and define cost functions (or metrics) for multi-objective optimization experiments. Cost functions include chiplet granularities and boundaries, typical system performance metrics like throughput, application execution times, energy, size, etc. In addition to these typical system metrics, perhaps of most importance to the CHIPS concept are the metrics of system design/implementation time and

cost, as these metrics are where the CHIPS paradigm are expected to have the greatest impact.

4. Conduct architecture simulations/evaluations to quantify design parameters that yield optimal metrics.

The sequential listing of the steps above by no means implies a serial sequence of steps. Like many of the steps in chip implementation itself, a two-way information flow is expected between the CHIPS ecosystem architecture exploration activities described above.

The ecosystem implementation issues are many, but most can be addressed by ascertaining best standard practices. For example, all chiplet IP must be accompanied by datasheet documentation and simulation models for easy integration into system designs. Constraints on physical implementation will also need to be specified depending on modular platforms supported. For example, some 3DIC or 2.5D silicon interposer approaches may impose die height restrictions, thermal envelopes, power budgets, etc. One of the most difficult challenges in the CHIPS ecosystem implementation is how IP providers will assure functionality of the chiplets. Chiplet IP providers must develop testing methodologies that ensure inventory chiplets function according to spec. For the CHIPS concept to succeed, it must surpass the known-good-die (KGD) hurdles from the multichip module (MCM) era. Some methods for ensuring die-level functionality without necessitating costly and time-intensive numerous steppings of probecards across a wafer involve embedding sacrificial tester chips in a wafer that are connected to neighboring die for the mere purpose of running acceptance tests. The wafer test/probe process therefore accesses only the tester chips to determine which chiplets are functional.

Standard interfaces for the chiplets are crucial for a system to be implemented primarily through composition of chiplet components. Given the building blocks of this chiplet model are in bare die form and are likely to be assembled/integrated with 3DIC and 2.5D technology, it is probable that existing interface standards developed for PCB-based systems, such as QPI or PCI-express, will not be ideal for the chiplet model. Since the interconnect pitch and distances envisioned for the chiplet model are much smaller than those of PCBs, these characteristics can be exploited to simplify interfaces. More likely candidates to serve as at least starting points for such interfaces are SoC standards like AMBA or one-off solutions. Once the architecture explorations described in the previously are conducted and more detail can be ascertained about preferred chiplet granularities and boundaries, existing interfaces can be evaluated for suitability in such systems. In parallel and speculatively, alternative interfaces specifically targeted to expected integration platforms should be explored.

Metrics for interfaces will include not only the typical interconnect measures of throughput (bandwidth), latency, and energy, but there also needs to be some measure of applicability to chiplet types. This is specifically important with regard to analog/mixed-signal (AMS) chiplets. For integrating AMS chiplets, it is probably best to consider external system interfaces to the outside world, which may be custom, versus internal system interfaces, which couple to other chiplet components. For AMS internal system interfaces, the interfaces should look more digital in nature to be able to couple with any other generic data generating/receiving chiplet. For the external system interfaces, there is no need for standardization from the chiplet perspective. External system interfaces for AMS components are typically defined by the input/output driving point impedances and voltage/current levels. For instance, operational amplifiers may be specified for a given load capacitance. In discrete realization of radio-frequency circuits, the interface is typically defined by an impedance value that corresponds to the characteristic impedance of the transmission line that is used

at the interface; in most cases, this is 50 Ohms. In on-chip RF circuits, the interface between various blocks may or may not be 50 Ohms. For instance, in typical RFIC receivers, the low-noise amplifier (LNA) is directly connected to the frequency down-conversion mixer without any 50-Ohm transmission line. In such a case, the LNA is designed to be able to "drive" the mixer input. For analog, mixed-signal, and RF building blocks, the driving point impedance and/or drive capability (in units of current or voltage) are an important metric at the interface. In the case of analog and mixed-signal IPs, it may be appropriate to specify a range of impedances for which such IPs are expected to operate under a given performance specification. Insertion loss and bandwidth, while somewhat appropriate for RF interfaces, may not be meaningful for analog and mixed-signal interfaces. In the latter interfaces, power delivery, for which insertion loss is defined, is typically not an objective.

Regardless of chiplet type, the interface will need to be scalable and configurable to support a broad range of chiplet types, implementation technologies, speeds, etc. In fact, given the likelihood of disparate chiplet operating speeds, some type of globally-asynchronous locally-synchronous (GALS) interconnect scheme is advisable. Therefore it may be necessary to design a polymorphic interconnect framework, which should be customizable through a software defined methodology and reprogrammable over the lifetime of the integrated system, as applications change with missions. Similarly, one could imagine that a grossly configurable FPGA-like chiplet to serve as an interface between commodity devices that do not adhere to the chiplet interface scheme and the rest of the chiplet-based system will be needed as part of the chiplet inventory.

In summary, a research program to address interface requirements for the CHIPS chiplet IP reuse model will need to include the following activities:

1. Define metrics for characterizing data transport needs of interactions between chiplets. Likely candidates include throughput (bandwidth), latency, and energy.
2. Characterize the aggregate data transport needs of chiplets using defined metrics.
3. Analyze existing interfaces for suitability or adaptation to provide a solution for systems composed of chiplets.
4. Define a minimal set of adaptable, scalable interfaces that satisfy not only the data transport needs of existing chiplets but likely future chiplets.
5. Explore the use of globally asynchronous schemes to enable easier integration of chiplets across technology node types and generations.

Appendix A: Outbrief Presentation

Heterogeneous IP Ecosystem enabling Reuse (HIER) January 23, 2017 Update

Jeff Draper
Deputy Director of Computational Sciences and Technology, ISI
Research Associate Professor, Ming Hsieh Department of EE
draper@isi.edu

Wes Hansford
Director, MOSIS
hansford@mosis.com

University of Southern California

4676 Admiralty Way
Marina del Rey, CA

01/23/2017

3811 N Fairfax Drive
Arlington, VA

Information Sciences Institute

HIER Primary Objectives

- **Develop integration process technology that is broad enough to encompass integration of IP regardless of implementation technology yet adheres to a standard to facilitate seamless integration**
- **Develop tool flows that lower the barrier to entry for even low-volume applications to cost-effectively take advantage of the HIER integration technology**

Information Sciences Institute 2

HIER Tasks

- Develop preliminary cost model for heterogeneous IP integration, interconnect, and foundry services
- Characterize the scope of tool flow activities needed for making heterogeneous IP integration as seamless as possible
- Describe methodology needed to enable a holistic heterogeneous IP reuse ecosystem

Heterogeneous IP Integration and Foundry Services

IP Management - 1

- **Currently in discussion with a limited set of IP providers to hash out a model for providing broker-like access**
  - Envision a cost-sharing model for IP much like the MOSIS model for fabrication cost-sharing
- **Challenges**
  - How to accommodate the various types of typical IP pricing
    - Initial license fee, annual maintenance fee, per-design use fee, production royalty fee, etc.
  - What level of effort is needed for broker to provide IP support or will individual users still need to set up individual arrangements with IP provider for excessive support?

IP Management - 2

- **MOSIS has discussed with several leading DoD Contractors , e.g. Boeing, NGC, etc.**
  - These firms have a 'Fabless', not a IP business model
- **MOSIS also works closely with a number of highly successful IP firms, e.g. ARM, Cadence, Synopsys**
  - Dedicated business model for IP development, marketing, pricing – and (especially) verification, ongoing support
- **DoD-Contractor firms would need a major change to business model to offer Hard or Soft IP**
  - So CHIPS program can be a better fit to these firms
- **DoD-Contractor IP pricing model likely to be very different from traditional IP providers**
  - DoD-Contractor NRE often paid for by previous contracts, whereas IP providers amortize NRE over IP pricing

IP Management - 3

- **IP from DoD-Contractors, e.g. Boeing, NGC**
  - While there are pricing models for fabrication, e.g.
    - MPW – cost per sq mm, varies by process (mask set cost)
    - Dedicated – cost of masks (number, complexity); wafers (volume, complexity)
  - There are not industry-wide pricing models for IP, e.g.
    - CHIPS users will need to determine fabrication costs
    - But also will need to determine 'value'
    - Companies (fab, e.g. Intel – fabless, e.g. Qualcomm) have determined these costs
      - What the market will bear, what the competition charges
  - We can't further advise of CHIPS pricing without detailed information of what will be offered
    - But the price will likely reflect fabrication, testing, and incidental management costs plus some markup to incentivize DoD contractors to participate
      - The markup value is key to the model maintaining success and is probably best determined by collaboration among potential IP providers in what will be a relatively small market

IP Fabrication Costing Example

- **Assume 3mm x 3mm design is fabricated on a dedicated 28nm run with a minimum-sized lot of six 300mm wafers**
  - Assume run cost is ~ $2M
  - Wafer sort testing can be in the $200K range depending on the level of testing to be done, even for this small lot
  - Reticle sizes for this technology node tend to be in the range of 25mm x 30mm, and there are roughly 100 reticle steppings
    - Roughly 8,000 3x3 chips per wafer available, so 48,000 per lot
- **For even a 90% yield, per-chip costs would at least be $2.2M / (48,000 * 0.90) = $51**
  - Chiplet price would be based on this cost plus some markup
- **Per-chip cost could scale with chip size in some fashion**
  - Yield and testing costs can have a significant impact on cost function

Heterogeneous Foundry Services

- **ISI/MOSIS has experience in working with Heterogeneous fabrication**
  - DARPA COSMOS, DAHI; also IARPA TIC
- **Challenges**
  - How to accommodate the various combinations of CMOS and III-V (et al) processes
  - Requires multiple ‘flavors’ of PDKs, et al
- **Opportunity**
  - MOSIS can support access to scheduled MPW runs for various CMOS processes, enable post processing
  - Allows flexible use, without dedicated run costs

Heterogeneous Integration Services

- **ISI/MOSIS has experience in working with a wide range and number of:**
  - Fabrication: CMOS, SiGe, Photonics, etc.
  - Assembly: Flip-chip, wirebond, interposers, etc.
  - IP sources: Foundry, Commercial, Users
- **Challenge**
  - User IP/Chips require careful vetting, plus neutral third part for inventory, fee payment, etc.
- **Opportunity**
  - ISI/MOSIS has unique experience and established systems for many of these needs

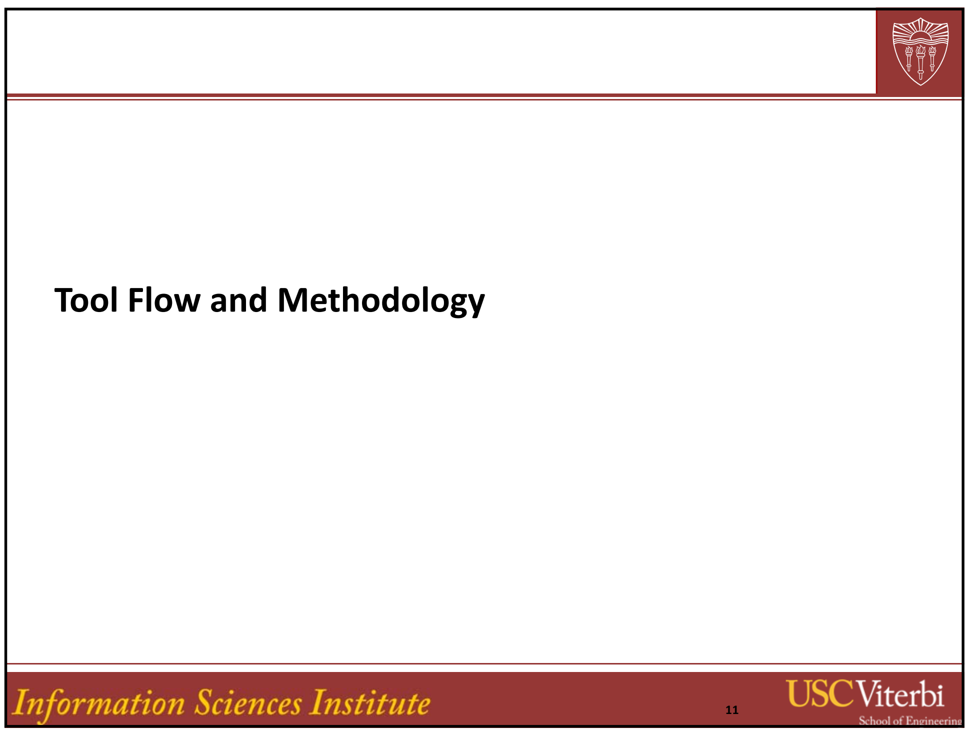
Tool Flow and Methodology
Information Sciences Institute
11
USC Viterbi
School of Engineering

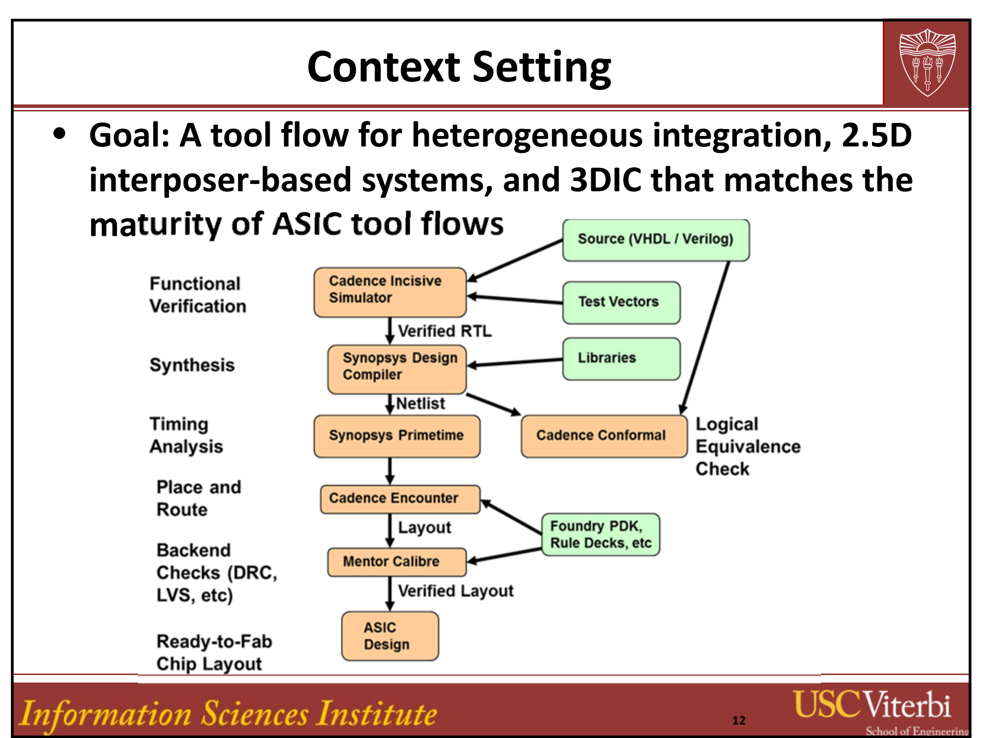
Context Setting
• Goal: A tool flow for heterogeneous integration, 2.5D interposer-based systems, and 3DIC that matches the maturity of ASIC tool flows
Source (VHDL / Verilog)
Functional Verification
Cadence Incisive Simulator
Test Vectors
Verified RTL
Synthesis
Synopsys Design Compiler
Libraries
Netlist
Timing Analysis
Synopsys Primetime
Cadence Conformal
Logical Equivalence Check
Place and Route
Cadence Encounter
Layout
Foundry PDK, Rule Decks, etc
Backend Checks (DRC, LVS, etc)
Mentor Calibre
Verified Layout
Ready-to-Fab Chip Layout
ASIC Design
Information Sciences Institute
12
USC Viterbi
School of Engineering

Current State / Risks

- **Open Foundry type of approach to 2.5D and 3D has been slow to develop**
  - Most successes have been product oriented: Micron HMC, Intel Stratix10, AMD, GlobalFoundries / OpenSilicon collaboration, etc
  - Even companies like Tezzaron/Novati that are touting such a model are not having significant success
- **The tools are more PCB-oriented than ASIC-design**
  - For open foundry 2.5D / 3D to thrive, tools need to have a more holistic system view
- **For CHIPS to succeed longterm, Open Foundry model involving substrate suppliers and tools will have to mature significantly**

Current CAD Tool Support

- **All the major vendors (Cadence, Mentor, Synopsys) claim tool support for 3DIC and 2.5D interposer-based systems**
  - https://www.cadence.com/content/cadence-www/global/en_US/home/solutions/3dic-design-solutions.html
  - https://www.mentor.com/solutions/3d-ic-design/
  - https://www.synopsys.com/Solutions/EndSolutions/3d-ic-solutions/Pages/default.aspx
- **But the approach is more PCB-oriented than ASIC-design**
  - For example, much of Cadence's tool support builds off PCB design tools rather than ASIC tools
  - Pros: support for physical verification, extraction, simulation and testing
  - Cons: missing timing analysis, power network analysis, signal integrity, etc.

2.5D/ 3D Current CAD Tool Challenges

- **Tools are missing a holistic system-view approach for all verification**
  - The support for physical verification, extraction, simulation and testing is really more focused on verifying connectivity than complete functionality
- **Can "fake" the tools into thinking a 2.5D or 3D system is an ASIC but that's an error-prone approach that would introduce inconsistencies between design and verification**
- **There really is a need for better tool development, and the hope is that a chiplet-based ecosystem will drive that**

Heterogeneous Integration Tool Flow

- **In general, the challenges with heterogeneous integration tool flows are more about the PDK than the tools themselves**
  - An adequate PDK consistent with CAD tools solves the problem
- **Northrop Grumman (NG) PDK and design flow established under the DAHI program works well for target technologies**
  - Initial versions had some minor issues like inconsistencies between availability of Spectre verus Spice models of various structures
    - **Easily fixable and corrected in later versions**
- **Serves as a solid template for other PDK development targeting other heterogeneous material**

Conclusions

- **Tool companies do not appear to have essential partnerships for 2.5D / 3D akin to their chip foundry relationships**
  - Evaluation of Cadence 2.5D tools depends on a substrate-specific PDK, and such a PDK is not openly available
  - CHIPS must invest in maturing an Open Foundry model for 2.5D / 3D or such technology will always be product-oriented with a narrow market
- **For further heterogeneous integration, a generator framework could aid in PDK development**
  - Rather than repeating the NG PDK development for every different heterogeneous material, build a generator tool that can output the PDK based on a few input parameters that characterize the heterogeneous material