

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re *Inter Partes* Review of:

Inventor: Hayashi et al.

U.S. Patent No. 9,207,701

Issued: Dec. 8, 2015

For: Supply Voltage Generating
Circuit

IPR2023-01230

**DECLARATION OF DR. BRUCE JACOB
IN SUPPORT OF PATENT OWNER'S RESPONSE
RE: *INTER PARTES* REVIEW OF U.S. PATENT NO. 9,207,701**

TABLE OF CONTENTS

I. Overview	3
A. Introduction	3
B. Legal Principles	4
C. Qualifications	6
II. Level of Skill in the Art	11
III. Technical Background	12
A. Basics of Boosting on-Chip Voltage Levels	12
B. Putting Capacitors and Booster Stages in Parallel and in Series	18
IV. Overview of U.S. Patent 9,207,701 ('Hayashi')	26
V. Response to Claims Analysis in '701 Petition	32
A. Overview of the Prior Art Used for Grounds 1 and 2	32
B. The Proposed Combination of Zhang and Javanifard	36
C. There Is no Motivation to Combine Zhang With Javanifard, as the Two Solve Opposing Problems, Meaning That the Combination Does no More Than Javanifard by Itself	43
D. Ground 2: Dr. Stiegler Misrepresents Javanifard's Figure-2 Embodiment and Proposes a Nonsensical Combination of Javanifard With Itself ..	58
E. Claims 1–18 Are Therefore Valid	77

I. OVERVIEW

A. Introduction

(1) I have been asked by counsel for plaintiff Longitude Licensing Limited (“Longitude”) to provide my opinions concerning the petition for *Inter Partes* Review of U.S. Patent No. 9,207,701 (“the ’701 Patent”).

(2) I am being compensated at my standard hourly rate plus any out-of-pocket expenses. My compensation does not depend in any way on the particular testimony or opinions that I express in this disclosure, or on the outcome of this litigation.

(3) For this document, I have been asked to prepare a disclosure of my opinions regarding the ’701 patent and associated documents.

(4) I reserve the right to respond to any disclosed opinions, declarations, or arguments made by Defendant or its experts during claim construction. I have prepared this disclosure having reviewed the patent-in-suit and the IPR filed regarding the patent-in-suit. I have also reviewed the material specifically cited within this Declaration.

B. Legal Principles

(5) In rendering my opinion herein and considering the subject matter of the patent, I am relying on legal principles related to validity that counsel for Patent Owner has explained to me. These principles are summarized below.

(6) I understand that a patent may be rendered “obvious” based on an alleged prior art reference or a combination of such references plus what a person of ordinary skill in the art would understand based on his or her knowledge and the references. I understand that a patent cannot be properly granted for subject matter that would have been obvious to a person of ordinary skill in the art at the time of the alleged invention, and that a patent claim directed to such obvious subject matter is invalid under 35 U.S.C. § 103. It is also my current understanding that in assessing the obviousness of claimed subject matter one should evaluate obviousness over the prior art from the perspective of one of ordinary skill in the art at the time the invention was made (and not from the perspective of either a layman or a genius in that art).

(7) I understand that obviousness is a question of law premised on underlying facts including: (a) the scope and content of the prior art, (b) the differences between the prior art and the claims, and (c) the level of ordinary skill in the pertinent art. I understand that the question is not whether the differences

themselves would have been obvious, but whether the claimed invention as a whole would have been obvious.

(8) In arriving at an obviousness determination, I understand that one must sufficiently explain and support the conclusions that the prior art references disclose all the elements recited in the Challenged Claims, and that a POSITA not only could have, but would have been motivated to combine all the prior art references in the way the patent claims, and would have had an expectation of success in doing so. That is, even if all the claim elements are found across a number of references, an obviousness determination must consider whether a person of ordinary skill in the art would have the motivation to combine those references. I understand that this combinability determination, as supported by an articulated motivation to combine, requires a plausible rationale for why those prior art references would have worked together. I understand that absent some articulated rationale, a “common sense” finding is no different than the conclusory statement “would have been obvious.” I understand that it is also important that knowledge of a problem and motivation to solve it are entirely different from motivation to combine particular references.

(9) I understand that even if all the claim elements are found across a number of combinable references with sufficient motivation to combine those references, a

POSITA would still need a “reasonable expectation of success” in making the claimed invention as a whole, in order for the invention to be found obvious. I understand that in order to possess the requisite “reasonable expectation of success” in combining all the references to make the claimed invention as a whole, the POSITA must be motivated to do more than merely vary all parameters with no indication of critical parameters or try all of a number of possible choices with no direction until successful. I understand that the prior art must provide more than a mere “general guidance” that seems to be a “promising field of experimentation.”

(10) I understand that the obviousness inquiry must exclude hindsight and avoid reading into the prior art the patent's teachings. I understand that imbuing a POSITA with knowledge of the invention when no prior art reference or references of record convey or suggest that knowledge, is such example of improper hindsight bias. I understand that invoking design choice to add many distinct components to a system evidences impermissible hindsight, not obviousness.

C. Qualifications

(11) In 1988, I earned the *Ars Baccalaureate* degree in mathematics from Harvard University in Cambridge, Massachusetts. In 1995 and 1997, respectively, I earned M.S. and Ph.D. degrees, both in Computer Science & Engineering, from the

University of Michigan in Ann Arbor, Michigan. In 2020, I was named a Fellow of the IEEE (Institute of Electrical and Electronics Engineers) “for contributions to computer memory design and analysis.”

(12) From 1997 to 2003, I was an Assistant Professor in the Department of Electrical & Computer Engineering at the University of Maryland at College Park, Maryland. I was promoted to Associate Professor and was granted tenure in 2003. In 2007, I was promoted to Director of Computer Engineering, a position I held until my sabbatical beginning Fall 2010. In the Spring of 2011, I was promoted to full Professor. In 2022, I joined the United States Naval Academy in Annapolis Maryland, where I am a Professor of Cyber Science and study the security of memories and computer hardware systems.

(13) For the past twenty-five years I have designed high-performance semiconductor memory systems and computer-systems architectures for both industry and government. For example, I helped Micron design their Hybrid Memory Cube architecture, redesigned Cray's memory controller for their Black Widow memory system, helped Northrop Grumman design a memory-system interconnect for their experimental ultra-low-power datacenter, designed a high-performance memory system for the 1024-core Teraflux chip funded by the European Commission, and have collaborated with the Departments of Energy

and Defense on the design of their next-generation supercomputers. My group also fabricated semiconductor memory chips and a CPU chip as part of an early DOD research program on die stacking using through-silicon vias.

(14) I received the National Science Foundation CAREER Award in 1999 for my work in semiconductor memory devices and memory systems. At Maryland, I was decorated seven times as a Clark School “Rainmaker” (renamed “Research Leader” in 2006) for my research program, being one of the top 100 researchers and fundraisers in the University.

(15) I have published two books on computer memory systems: Memory Systems: Cache, DRAM, Disk (Morgan Kaufmann) 2007; and The Memory System (You Can't Avoid It; You Can't Ignore It; You Can't Fake It) (Morgan & Claypool) 2009. My MS thesis at Michigan was on the design of multi-level cache hierarchies; my Ph.D. thesis at Michigan was on the topic of virtual memory system architectures. I have written over eighty refereed articles on memory systems, cache systems, computer design, embedded systems, operating systems design, astrophysics, and algorithmic composition. These can be found listed in my *curriculum vitae*. As of Summer 2023, my body of research has been cited over 8000 times, with over 80 cites per work on average. I hold three patents in the area of circuit design and another two in the area of memory-systems design.

(16) I founded the annual *International Symposium on Memory Systems* and am frequently invited to give keynote speeches and high-level briefings on the topic of memory systems around the world. I have been the keynote speaker at meetings including EMS, MCHPC, MODSIM, SAMOS, and Computing Frontiers. I was asked to speak before a National Academies panel on the topic of memory systems and to participate in a roundtable discussion with the Department of Energy (DOE) to create a roadmap for computer science and applied mathematics research at DOE in the coming decades.

(17) Topics my research group has researched extensively include *DRAM and main memory technologies* (e.g., controllers, devices, bus organizations, system architectures, and memory-system power), *non-volatile memories* (e.g., flash memories, numerous resistive RAM variants, and controller & system designs for both main memory and storage), and *embedded systems*. The last area includes investigation of topics in design verification for heterogeneous systems (e.g., the detailed, hardware-verified simulation of heterogeneous systems, for the design and verification of embedded systems that control various hardware devices in real time), highly integrated, heterogeneous systems-on-a-chip (e.g., design and implementation of a monolithic *micro electro mechanical systems* [MEMS]-based gas sensor component, integrated on silicon with a digital microcontroller), and the hardware/software co-design of embedded microcontrollers and real-time operating systems.

(18) In summary, I have extensive technical experience and expertise relating to aspects of designing and implementing computer memories and memory systems, including the fabrication of semiconductor devices and chips, all of which are the topics of concern in this matter.

(19) My *curriculum vitae* is attached as Exhibit 1.

(20) The following is a brief overview of my opinions.

II. LEVEL OF SKILL IN THE ART


(21) In my opinion, a person of ordinary skill in the relevant art of the asserted patents, around the time of their filing, would have had at least a bachelor's degree or equivalent in electrical or computer engineering, with circuit design and memory-systems coursework, and two years of experience in the design of memory systems and associated components. More education can supplement practical experience and vice versa.

III. TECHNICAL BACKGROUND

(22) The '701 patent is concerned with charge-pump design, including the use of a dynamically configurable number of charge-pump stages to achieve higher voltage outputs. Some of the prior art used in the '701 Petition is concerned with wiring multiple charge-pump stages in either parallel or series. The following discusses some charge-pump fundamentals.

A. Basics of Boosting on-Chip Voltage Levels


(23) In many scenarios, a chip is given an input voltage source at one voltage level, and one would like to produce an output voltage level that is higher than the input level provided. A common way to boost a voltage level on an integrated circuit is to use a *charge pump*. Charge pumps come in several different varieties. The original design used switches, as shown below, in an example four-stage pump.



(24) As the clock oscillates, the even/odd switches alternate between open and closed positions, and current is only allowed to flow forward, where the charge

builds up on each successive capacitor. When the switches flip, the reversing of the polarity causes an increase of the voltage across each capacitor, thereby increasing the voltage that is output by the pump, proportionally to the number of stages. The main principle of voltage boosting is that the amount of charge (typically designated “q” in equations) is conserved in a capacitor, and it is conserved even when one of the capacitor's terminals is switched to a different potential.

(25) This behavior is most easily seen in the operation of a simple one-stage boost circuit, shown below.



(26) In phase 1, at the top of the figure above, switch S1 is closed, and the input (for instance an oscillating clock) is in its “low” phase, putting GND at the bottom


plate of capacitor C. The capacitor is charged to VDD by the power supply through switch S1. Switch S2 is open, preventing the capacitor from discharging to the output.

(27) In phase 2, shown at the bottom of the figure above, switches S1 and S2 have changed orientation, and the clock input is in its “high” phase, during which it is equal to VDD. Now, the top plate of capacitor C is connected to the output, and the bottom plate of capacitor C is connected to VDD (the clock’s “high” value). Due to charge conservation, the electrons have gone nowhere (yet) and are all still stored within the capacitor. Thus, there must remain a potential of VDD across the plates of the capacitor. At this point, the clock is high, and so the potential on the bottom plate is VDD. If the difference between top plate and bottom plate must be VDD, which it must, then the top plate must then assume a potential of VDD on top of VDD, or $2 \times \text{VDD}$, relative to GND.


(28) At this point, if only $2 \times \text{VDD}$ is needed, no further boosting stages are necessary. If the high voltage is only needed during a single cycle, the operation is over; the periodic input signal need not continue. In other words, even though the examples herein show four-stage pumps with continuously oscillating clocks as timing signals, not all charge pumps need to be that complex.

(29) As mentioned, charge pumps come in numerous forms and organizations.

The Dickson charge pump is one of the better known designs, and it replaces the switches of the original design with diodes. An example four-stage Dickson pump is shown below.




(30) As the Dickson pump illustrates, diodes can replace the switches in the previous design, as they perform the same function by allowing current to flow easily only in one direction. In integrated circuits (i.e., on computer chips), transistors are used to create both the diodes and the capacitors of the charge pump, wired as follows:



(31) MOSFETs with their gates tied to their source terminal are *de facto* diodes; MOSFETs with their source and drain (and/or body) tied together are *de facto* capacitors. This type of charge-pump design is very common in CMOS circuit designs for precisely this reason.


(32) Given that MOSFETs can act simultaneously as diodes and capacitors, one could provide both functions, capacitor and diode, in a single MOSFET, producing yet another circuit design that is topologically equivalent to the Dickson charge pump:



(33) Here, each MOSFET provides both the capacitor and the diode.

(34) For simplicity the previous figures use as their timing signals the clock and clock-bar signals (the clock signal and its inverse), but as suggested earlier, not all applications require a continuous timing signal, as not all applications require a continuously running charge pump. Some applications require only a high voltage for specific actions at specific brief periods, and these applications can simply use

one or more pulses to control a pump instead of a continuously running clock. For example, if an application requires a higher voltage than $2 \times VDD$, but for only a short period of time, it would be possible to extend the earlier single-pulse example to a multi-stage pump. This amounts to the figure below.




(35) As described earlier, a single stage boosts the input voltage by a theoretical amount of VDD (the voltage swing of the pulse signal) on top of its input whenever it sees a pulse. Thus, if one sends a timing pulse down a multi-stage pump, provided that each stage's timing pulse does not overlap with that of the previous stage, this provides a one-time voltage boost on the output similar to a single cycle output of the continuously operational pumps above. Inserting delays in the signal chain enforces this requirement.

B. Putting Capacitors and Booster Stages in Parallel and in Series

(36) It can be desirable to wire and operate charge pumps in parallel, just as it is desirable to wire charge-pump stages in series as in the examples above. The following discussion illustrates.

(37) To begin, the figure below shows the two stages of voltage-boost operation in a single-stage boost circuit (e.g., a fundamental stage of a multi-stage charge pump).



(38) The top illustration represents what we can call the *precharge* phase of a boosting circuit; the bottom illustration represents what we can call the *boost* phase. The question: what is the potential X when the boost capacitor has been charged

to VDD, and the input/output switches S_i and S_o flip, replacing GND by VDD at the bottom plate of the boost capacitor? We know that, given a capacitor with a voltage across it, the following is true:

$$q = CV$$

(39) This simply tells us that the amount of charge q that is held in the capacitor is equal to the capacitance C times the voltage potential V across the capacitor.

(40) In the *precharge* phase (the top illustration in the figure above), the following holds:

$$q_{\text{precharge}} = C_B * (VDD - GND) = C_B * VDD$$

(41) In the *boost* phase (the bottom illustration in the figure above), the following holds:

$$q_{\text{boost}} = C_B * (X - VDD)$$

(42) Due to the principle of charge conservation, it is the same amount of charge, the same q , in both scenarios, i.e., $q_{\text{precharge}} = q_{\text{boost}}$. Therefore the following holds:

$$C_B * VDD = C_B * (X - VDD)$$


$$X = 2 * VDD$$

(43) This is the same result as we found earlier: that, due to charge conservation, the voltage potential of the top plate of the capacitor is boosted to a higher level. The derivation here simply takes a different path.

1. Wiring Capacitors (or Charge-Pump Stages) in Parallel

(44) Wiring capacitors or charge-pump stages in parallel maintains the same voltage-boosting level but increases the pump's current output. Putting two capacitors in parallel, thereby changing the boost capacitor's effective size, will make no difference to the voltage level produced, because two capacitors in parallel are equivalent to a single capacitor with the sum of the capacitances.


(45) The following illustrates. First, note that the following are electrically equivalent:



(46) In both scenarios, we have switches between the C_B boost capacitors that selectively couple the boost capacitors to GND and VDD. There is no significant

difference between two parallel wires and switches to ground and one wire and switch to ground.

(47) Assuming that the switches operate in tandem, this becomes the following, which mirrors the previous example, the only difference being the size of the boost capacitor:



(48) As before, the top illustration represents the *precharge* phase of a boosting circuit; the bottom illustration represents the *boost* phase. The S_o *output* switch chooses between VDD during the precharge phase and the output of the circuit during the boost phase; the S_i *input* switches choose between GND during the precharge phase and VDD during the boost phase. When wiring capacitors in

parallel, the effective total capacitance is the sum of the capacitances. Therefore, in this case, the effective capacitance is $2C_B$.

(49) As before we have the following:

$$q_{\text{precharge}} = 2C_B * VDD$$

$$q_{\text{boost}} = 2C_B * (X - VDD)$$

(50) Due to charge conservation, $q_{\text{precharge}} = q_{\text{boost}}$, and therefore we have the following:

$$2C_B * VDD = 2C_B * (X - VDD)$$

$$X = 2 * VDD$$


(51) This is identical to the previous example in that the equations yield the same $2 * VDD$ boost level for the output voltage X . The double-capacitance $2C_B$ does not affect the output voltage level. Instead, what is affected is the current, as shown in the following relation governing the behavior of capacitors:

$$I = C dv/dt$$

(52) Doubling the capacitance of a charge pump, either by doubling the size of the capacitor or by wiring two pumps in parallel, doubles the current I , because the current output of a capacitor is proportional to the capacitance C .

2. Wiring Capacitors (or Charge-Pump Stages) in Series

(53) The way to increase voltage output when combining two or more boost capacitors, or two or more boost stages, is to put those capacitors or stages in series, as shown below:



(54) As before, the top illustration represents the *precharge* phase of a boosting circuit; the bottom illustration represents the *boost* phase. The S_i/S_o input/output switches behave as before: the S_o *output* switch chooses between VDD during the

precharge phase and the output of the circuit during the boost phase; the S_i *input* switches choose between GND during the precharge phase and the output of the previous stage (or VDD, the input of the circuit) during the boost phase. The interior S_p *precharge* switches take advantage of the fact that the next stage has an S_i input switch; the S_p switches choose between connecting to VDD during the precharge phase and acting as an open circuit during the boost phase.

(55) In this example, each capacitor is charged to VDD during the precharge phase. When the switches are thrown, and the bottom potential of the first capacitor is swapped from GND to VDD, the first capacitor sees its top plate boosted to $2 * VDD$, as detailed previously. This voltage becomes the new potential seen at the bottom plate of the next capacitor, which, due to charge conservation, must retain a voltage potential of VDD across its top and bottom plates (unless discharged, its top plate must remain at a potential of VDD above its bottom plate), and therefore it sees its top plate boosted to $3 * VDD$. This is fed to the next stage's bottom plate, which sees its top plate boosted to $4 * VDD$.

(56) In general, this is what one means when one says that the final output of a charge pump is $(n+1)*VDD$ for an n-stage charge pump. As opposed to ideal designs like the ones above, real implementations, such as the Dickson charge pumps illustrated earlier, use diodes to avoid backflush. In this case, the boosted

voltage is reduced by the voltage drops through each diode, so one does not reach the actual theoretical maximum boost voltage. Nonetheless, the output voltage is boosted in proportion to the number of stages in series.

IV. OVERVIEW OF U.S. PATENT 9,207,701 ('HAYASHI')

(57) The '701 patent ('Hayashi') is concerned with the trade-off between over-designing a charge pump and under-designing the charge pump. As Hayashi describes, early in the design of chips, the final voltage specifications are not necessarily known. A designer has two bad choices to make: to over-design the charge pump, in which case it will work irrespective of the final operating voltage but will be larger and more expensive than necessary, or to under-design the charge pump, in which case it could fail to work, if the final operating voltage is too much for the circuit to handle:

If specifications relating to a DRAM are changed or the specifications have not been fully decided, there are cases where the supply voltage supplied from the outside in accordance with a request from the user differs. A supply voltage supply circuit for driving a word line is one circuit for which different supply voltages are required of externally applied power. Designing such a supply voltage supply circuit involves preparing a plurality of types of transistors having oxide-film thicknesses for which circuit operation is optimum, and designing the circuit based upon the plurality of types of transistors. *If a booster circuit is designed assuming a case where the power supplied will differ depending upon the specifications, the problem of oxide-film withstand voltage and a current-consumption problem ascribable to gate leakage must be taken into account. If, in order to avoid these problems, a transistor for thick-film capacitance is used in the circuit in accordance with the higher voltage among voltages that have been set*

depending upon specifications, there is the danger that this will result in a booster circuit of larger area. [Hayashi, 1:23–41, emphasis added]


The booster of the supply voltage generating circuit has a boosting pump capacitor element that utilizes the oxide-film capacitance of a transistor. *There are many cases where the boosting capacitor element requires a very large capacitance. This necessitates a large area and has a major impact upon the area of the overall chip.* Since the capacitance of the oxide film is inversely proportional to the thickness of the oxide film, usually use of a thin-film capacitor element (thin-film transistor) is desired. However, *in a case where the supply voltage supplied from the outside differs depending upon the specifications, as mentioned above, a thick-film transistor usually is used as the boosting capacitor element in order that the supply voltage supplied will not exceed the breakdown (or withstand) voltage of a thin-film transistor.* Normally the value of a voltage for driving a word line will exhibit almost no change even if the supply voltage supplied from the outside changes.

Accordingly, regardless of the fact that the lower the external supply voltage, the greater the number of pump capacitor elements required, *a thick film having little capacitance must be used even in a case where voltage is low because of the breakdown voltage of the oxide film.* Thus, with the conventional circuitry, the design must be such that transistor breakdown voltage will be in conformity with the high side of the external supply voltage. Consequently, part of the booster capacitance in the booster circuit, e.g., the first capacitor in Patent Document 1 [prior art], must be made a thick-film capacitor element and, hence, the degree to which chip area can be reduced is limited. [Hayashi, 1:64–2:24, emphasis added]

(58) Both scenarios are undesirable, and so Hayashi proposes a charge-pump design that uses thick oxides and extra die area only in the areas where it is absolutely needed, thereby achieving a trade-off between over- and under-designing the pump. The pump is designed to be configurable in the number of stages that it uses, and it uses thin-film oxides to reduce the required die area. If the incoming voltage is too high, a step-down mechanism reduces the voltage level, and the charge-pump responds by using more stages to recover the loss in voltage.

Figure 1 is shown below:

FIG.1



The '701 patent, Fig 1

(59) The figure shows the input voltage provided to the charge pump (BOOSTER 20) being one of two possibilities: either the VDDP1 or a stepped-down voltage VDLP. The booster circuit is reconfigurable, so that the input voltage S0 is either boosted by one booster circuit (e.g., Booster Circuit 21b, as would be the case in the configuration shown in the figure above), or it is boosted by the two booster circuits 21a and 21b wired in series, as would be the case if switch SW2 were thrown to the other position.

(60) Hayashi describes the figure as follows:

FIG. 1 is a block diagram illustrating the configuration of a supply voltage generating circuit according to an exemplary embodiment of the present invention. As shown in FIG. 1, the supply voltage generating circuit includes a power-supply step-down unit 10, a booster 20 and a switch element SW1.

The booster 20 generates a boosted voltage by the charge pumping of capacitor elements and outputs the boosted voltage generated. *In a case where the voltage of an external power supply exceeds the breakdown voltage of the capacitor elements, the power-supply step-down unit 10 steps down the external supply voltage to a voltage within the range of breakdown voltage of the capacitor elements and applies the stepped-down voltage to the power supply of the booster 20. [...]*

Preferably, the booster 20 is so adapted that the number of booster stages (e.g., one stage or two stages) constituted by the multiple booster circuits 21 a, 21 b is changed over by a switch element SW2 in accordance with the power-supply voltage applied to the booster 20. [...]

In accordance with the supply voltage generating circuit having the above-described configuration, the voltage of the external power supply is stepped down by the power-supply step-down unit 10 in a case where the external supply voltage is so high as to exceed the breakdown voltage of the capacitor elements (thin-film transistors). This is based upon the condition that the external supply voltage will differ depending upon the specifications. *It is possible for thin-film capacitor elements having a large capacitance to be used in the booster 20 and to be used conjointly also in a case where the supply voltage is low. Accordingly, layout area in the chip can be reduced by a wide margin. [Hayashi, 3:63–4:47, emphasis added]*

(61) As Hayashi relates, the trade-off allows one to design the charge pump so that a minimal die size for the pump is required, by using thin-film capacitors. Simultaneously, the potential for damage due to under-design is eliminated by the dynamic ability to reduce the voltage level when needed and recover the loss in voltage level by using more charge-pump stages in series. As described earlier, the voltage-level output of a charge pump is proportional to the number of stages in the pump, so wiring more stages in series achieves a higher output voltage level.

V. RESPONSE TO CLAIMS ANALYSIS IN '701 PETITION

(62) In this section, I respond to the technical-expert analysis presented in the report titled “Declaration of Dr. Havey [sic] J. Stiegler under 37 C.F.R. §1.68 in Support of Petition for *Inter Partes* Review” (“Stiegler Report”). The Stiegler Report and '701 Petition present effectively the same analysis.

A. Overview of the Prior Art Used for Grounds 1 and 2

(63) The following sections describe the prior-art documents used in Dr. Stiegler's analysis for Grounds 1 and 2.


1. United States Patent 7,427,888 ('Zhang')

(64) United States Patent 7,427,888 ('Zhang') describes a means to speed up the charge-pump operation within a chip, by supplying an optional external high-voltage input that is conditionally applied to the charge pump (“high-voltage” meaning higher than the operating voltage of the rest of the chip). Providing a higher-voltage input to the charge pump allows the charge pump to reach its target voltage output more quickly, as it allows the charge pump to start at a higher-voltage point that is closer to the desired final output.

(65) This mechanism can be used to improve the time that it takes to test chips during the manufacturing process. For instance, chip designs are tested while still in

the wafer form, before packaging the known good dies into chips, and then those packaged chips are tested again after being packaged, before shipping the good chips to the final destination. Supplying a higher voltage to the chip would speed up the testing process, thereby saving money, at the cost of an extra pin and the control circuitry required to enable to use of the higher voltage.

(66) Zhang's system is shown in the figure below:



(67) The charge pump accepts two different voltage-level inputs over two different wires: V_{CC} and V_{PP} shown on the far left. V_{CC} is the chip's operating

voltage, and V_{PP} is the optional voltage input that, if provided, allows the charge pump to reach its target output level more quickly. The switch circuit 200 at the front of the charge pump contains a detector and selector — it detects when the high-voltage V_{PP} input is present and, if that voltage is detected, it is selected and provided to the charge pump. If V_{PP} is not detected, then V_{CC} is provided to the charge pump.

(68) Zhang uses a reference voltage and a voltage-level detector to shut down the pump's operation when the desired output level is reached. This happens in the Voltage Level Sensing Block 400, which turns the oscillator 300 on and off. While the pump is on, the voltage-sensor circuit 400 monitors the output of the charge pump, disabling its operation once the charge pump's output reaches a desired reference level V_{ref} . Zhang discloses that the Voltage Level Sensing Block 400 is prior art, but it is uncommon. That type of circuit is not present in many charge-pump designs including, for example, those in the '701 patent or Javanifard described below. It is necessary in Zhang, because that is how Zhang controls its output, which can vary significantly based upon the voltage level of its input, which as described can vary between the normal operating voltage and something significantly higher.

2. United States Patent 5,767,735 ('Javanifard')

(69) United States Patent 5,767,735 ('Javanifard') describes a configurable charge-pump design in which the stages can be rearranged electrically to produce different numbers of charge-pump stages in series and in parallel. By setting switches appropriately, the circuit topology is changed to reconfigure the same circuit elements (charge-pump stages) into different circuit topologies. One of Javanifard's embodiments is shown in the figure below:




Figure 1

Javanifard, Fig 1

(70) By setting switches 130 and 131 appropriately, one can put charge pumps 110 and 120 into either a *series* configuration (switch 130 on, switch 131 off) or a *parallel* configuration (switch 130 off, switch 131 on). As described above in the *Background* section, putting charge-pump stages in *series* increases the *voltage level* of the final output, and putting charge-pump stages in *parallel* increases the amount of *current* that the charge pump generates as its output. The voltage level of the charge pump's output is proportional to the number of stages in series; the current output of the charge pump is proportional to the number of stages in parallel.

B. The Proposed Combination of Zhang and Javanifard

(71) Dr. Stiegler proposes a combination of Zhang and Javanifard as part of his Grounds-I analysis. The following sections outline a few salient points regarding the proposed combination.


1. It Produces a Single Output Voltage Level

(72) Dr. Stiegler describes his proposed combination of Zhang with Javanifard as follows:

[A] variable stage charge pump system, which can both “accommodate different output levels for a given charge pump power supply input level,” “accommodate a given output level for different charge pump power supply input levels,” and provide flexibility for

design needs (e.g., power saving mode, and current requirements), area reduction, and performance improvement, as taught by Javanifard.
[Stiegler, par. 90]

(73) Dr. Stiegler illustrates the proposed combination as follows:



(74) ... and as follows:




Figure 1

EX1006, Javanifard, FIG. 1, annotated

Stiegler, par. 110

(75) While Dr. Stiegler describes the combination as one that can “accommodate different output levels for a given charge pump power supply input level” [Stiegler, par. 90], it is clear from the descriptions and illustrations that the proposed combined charge pump only ever generates one voltage per configuration, i.e., only

one voltage for a given configuration and input voltage. Tellingly, there is no verbal description or illustration of any control mechanism that would enable the charge pump to produce different multiple levels, with dynamic control that could achieve a particular desired or specified level out of a set of values. Additionally, nowhere does the Stiegler Report describe how different voltages would or could be used. Finally, the configuration table on page 65 (shown below) is only set in the OFF position or series-coupled ON position, which indicates that the pump is configured statically and outputs a single voltage from that configuration.

TABLE 1			
Switch 130	Switch 131	Charge Pump 110	Charge Pump 120
OFF	OFF	ON	OFF
OFF	ON	ON	parallel-coupled to charge pump 110
ON	OFF	ON	series-coupled to charge pump 110
ON	ON	N/A	N/A

a second number of booster stages, the second number being less than first number

a first number of booster stages

EX1006, Javanifard, Table 1, annotated

Stiegler, p. 65

(76) Thus, it is clear that the proposed combination of Zhang with Javanifard involves providing only a single output voltage level.

2. The “Disadvantages” That Dr. Stiegler Enumerates To Motivate the Combination Do Not Apply to Zhang

(77) Dr. Stiegler provides the following “disadvantages” as motivation to combine Zhang with Javanifard:

Javanifard explicitly provides motivations to apply its teaching of the variable stage charge pump to replace the fixed stage charge pump in Zhang’s charge pump system. Javanifard addresses various **disadvantages** (e.g., **area usage, performance including pump efficiency, lack of flexibility**) of using the fixed stage charge pump, including, e.g., “us[ing] a number of charge pump circuits each having a fixed number of stages ...each charge pump circuit requires a dedicated space in the integrated circuit,” “the fixed stage method typically sacrifices performance at one power supply voltage to accommodate another power supply voltage.” [...]


A POSITA facing the need to improve its charge pump system would have looked to Javanifard and been **motivated to apply Javanifard’s teachings of the variable stage charge pump in the multi-power-source charge pump system of Zhang** using fixed stage charge pump, because Javanifard’s teachings provide the obvious, beneficial, and predictable results of a power supply generation circuit that can both “accommodate different output levels for a given charge pump power supply input level,” “accommodate a given output level for different charge pump power supply input levels,” with improved flexibility for design needs (e.g., power saving mode, and current requirements), area reduction, and performance improvement, as taught by Javanifard. [...]

Zhang’s teachings including its selector circuit are not limited for

only “a test or general operation mode.” See, EX1005, Zhang, 5:23-51 (claiming a charge pump circuit including a selector circuit without any operation mode requirement). As such, a POSIA [sic] would have been motivated to apply Javanifard's teachings in Zhang to provide a multi-power-source charge pump to support a device having multiple operation modes, including a device having “a test or general operation mode” and devices having other operation modes (e.g., a memory device having different read, write, programming modes). [Stiegler, pars. 92–94, emphasis added]

(78) Nowhere does Dr. Stiegler describe how these disadvantages apply to Zhang, and it is clear from Zhang itself that Zhang does not fall prey to any of the listed disadvantages (***area usage, performance including pump efficiency, lack of flexibility***). Zhang can be made smaller than other charge pumps, due to the additional voltage input (see element 200 in the figure below), and it is a single-pump design, not a multiple-pump design; the extra voltage input allows Zhang to achieve higher voltages more quickly than other charge pumps (i.e., Zhang is extremely efficient); and Zhang's sensor-controlled shut-off (Voltage Level Sensing Block 400 in the figure below) is at least as flexible¹ as Javanifard's reconfigurable stages in its ability to target different output levels.

¹ Zhang's sensor-driven approach is arguably *more* flexible than Javanifard, because it allows one to reach any desired output level by adjusting the Vref voltage level accordingly (Zhang causes the charge pump to stop when the desired level is reached), whereas Javanifard can only reach output levels that correspond to the number of charge-pump stages in series. Zhang can thus target in-between voltage levels that Javanifard cannot.



(79) In sum, there is no reason to believe that a person of skill would have been inclined to see Javanifard's "disadvantages" as being applicable to Zhang.

(80) An additional "disadvantage" that is described in Javanifard but not in the Stiegler Report is the use of voltage-divider networks. I note that Zhang is silent on the use of voltage-divider networks, and so there is no reason to expect a person of skill to consider this as being applicable to Zhang.

C. There Is no Motivation to Combine Zhang With Javanifard, as the Two Solve Opposing Problems, Meaning That the Combination Does no More Than Javanifard by Itself

(81) As described above, the two references used in the analysis for Grounds 1 and 2 solve the following problems:

Zhang takes *many* voltages as input and produces *one* voltage as output

Javanifard takes *one* voltage as input and produces *many* voltages as output

(82) These are very different problems: one (Zhang) produces the same voltage level as output, no matter the input voltage level, by comparing the boosted output to a voltage reference and adjusting the pump's operation accordingly; the other (Javanifard) is the opposite: it produces multiple different voltage levels as output, given one input voltage level, by allowing its charge-pump stages to be reorganized into different configurations.

(83) The proposed combination of Zhang and Javanifard raises several significant issues. It is not clear (a) that their combination produces anything beyond Javanifard by itself, (b) that Dr. Stiegler understands the purpose of Zhang's switch 200, because the switch is used incorrectly in Dr. Stiegler's proposed combination of Zhang with Javanifard, and (c) that there is anything, anywhere, describing a need for a system that combines these two opposite-facing solutions.

(84) The following sections discuss these issues, and others, in more detail.


1. Combining Zhang With Javanifard Produces a Control Problem That Dr. Stiegler Does Not Address

(85) As described above, an important component of Zhang's charge-pump invention is the use of a voltage sensor to determine whether the boosted output voltage has reached the desired level or not. At the start of voltage boosting, the DET signal is set to turn the charge pump ***on***. Once the desired output level has been reached, the voltage sensor detects this, and it clears the DET signal, thereby turning the charge pump ***off***.

(86) Zhang describes this process as follows:

... if the output voltage V_{out} is smaller than the reference voltage V_{ref} , the level sensing signal DET is enabled. If the output voltage V_{out} is larger than the reference voltage V_{ref} , on the other hand, the level sensing signal DET is disabled. [Zhang, 3:7–12]

(87) The mechanism, Voltage Level Sensing Block 400, is shown in the figure below at the bottom of the drawing:



(88) The Stiegler Report does not address the control aspects of Zhang at all, which is problematic because there are only two main points of Zhang's invention (described earlier in the overview on Zhang), and control is one of them. In particular, Javanifard enables the production of many different output voltage levels, and to the extent the Petition's combination involves the production of many different output levels, the control of such a system would be non-trivial.

(89) The Stiegler Report is silent on this. Indeed, Dr. Stiegler's analysis seems to be little more than the following: "here are two references; I will now pick and choose disclosures from each of these references that individually satisfy the claim

limitations.” Such simplistic analysis could possibly work when considering references whose combination produces nothing more than the sum of the collected parts, but this pair of references is not of that type. Zhang and Javanifard do not obviously and straightforwardly complement each other; instead, they do opposite things in a non-complementary way. It is the responsibility of the one presenting the analysis to explain how any incompatibilities would have been predictably overcome by a person of ordinary skill in the art, and Dr. Stiegler makes no attempt to do this, despite the obvious and significant incompatibilities that the pair present. Charge-pump design is non-trivial, as an analog design living in a digital implementation environment, which presents numerous challenges ... this is well known. For instance:

[C]harge pump design, being purely analog in nature and involving high voltage, needs meticulous design techniques, intensive semiconductor device analysis, careful design layout planning, and accurate parasitic extraction process to produce excellent results in real implementation on silicon.

— F. Pan, T. Samaddar. Charge Pump Circuit Design. McGraw-Hill, 2006.

(90) It is Dr. Stiegler’s responsibility to assure the reader that the proposed design will, in fact, work. However, Dr. Stiegler’s design has obvious flaws that I have pointed out, and Dr. Stiegler offers no explanations or assurances.

(91) In sum, Zhang and Javanifard have significant incompatibilities with regard to the control of the proposed combined charge pump. Dr. Stiegler makes no attempt to address the incompatibilities and effectively sweeps the entire issue under the rug. To the extent Dr. Stiegler's proposed combination involves generating multiple output voltages, it is my opinion that Dr. Stiegler has not established a reasonable expectation of success. [e.g., "Charge pump, being a pure analog design, carries an inherent risk of unanticipated effects, which when overlooked can significantly reduce circuit performance or cripple the circuit operation." Pan & Samaddar 2006, *Preface*]

2. Switch 200 From Zhang Does Not Do What the Stiegler Report Says It Does, and It Is Used Incorrectly in the Proposed Combination

(92) Regarding the issue above, one could potentially say that Dr. Stiegler is simply not interested in the control provided by Zhang and only wants the front-end switch ("selecting circuit block 200") provided by Zhang. Dr. Stiegler does not explain how or why the proposed combination would use only a portion of Zhang, so I can only guess here. However, the problem is that Dr. Stiegler's use of Zhang's switch 200 is inappropriate: Zhang's "selecting circuit block 200" does not do what Dr. Stiegler uses it for in the proposed combination.

(93) As described above in the overview on Zhang, the purpose of switch 200 is to choose between a low voltage and a high voltage so that, for instance when performing wafer tests before packaging or chip tests before shipping, one can use an externally provided higher input voltage for boosting and thereby speed up the testing of the part. If one starts at a higher voltage, then the time it takes for the charge pump to reach a higher voltage level is reduced, as compared to normal operation. Thus, if one provides the extra external high voltage, one can speed up the operation ... this would not be used normally, for instance in battery-driven systems in which there is one voltage supply: the battery. Providing a second voltage source would be simple in a testing environment, however. Zhang calls the low voltage the “power voltage” or “supply voltage,” i.e., the normal operating voltage of the chip, and the higher test voltage the “external high voltage.”

(94) Dr. Stiegler seems to be aware of this distinction [Stiegler, par. 91] and yet mischaracterizes Zhang when trying to explain the motivation to combine Zhang and Javanifard:

A POSITA facing the need to improve its charge pump system would have looked to Javanifard and been motivated to apply Javanifard's teachings of the variable stage charge pump in **the multi-power-source charge pump system of Zhang** [Stiegler, par. 93, emphasis added]

(95) Zhang is very clear that the two input voltages are not **power** sources. Again, Zhang calls the low voltage the “power voltage” or “supply voltage,” indicating that the low voltage is intended to be used as the normal operating voltage of the chip; Zhang calls the higher test voltage the “external high voltage,” thereby indicating that it is some special voltage not meant to be used as a generic supply. Zhang makes clear that the high voltage level is not intended for chip operation but only to be used for rapid testing of the charge pump’s output (e.g., to test erase and write operations in a nonvolatile memory chip). For instance, Zhang explains, “The voltage VPP refers to a high voltage provided externally during a test operation mode. The high voltage VPP may be applied externally in the test operation mode as well as during the general operation mode.” [Zhang, 3:24–27] It would have been clear to a person of ordinary skill that the higher voltage is intended for the operation of the charge pump and not the operation of the general digital circuits within the chip, because operating the entire chip at higher voltages would irreparably damage all of the circuitry expecting the normal power voltage. Because the Stiegler Report represents the charge-pump voltage as equivalent to a supply voltage in the bolded excerpt above, Dr. Stiegler’s mischaracterization is extremely misleading.

(96) Dr. Stiegler evidently recognizes that the description in the excerpt above (Stiegler, par. 93) is indeed a mischaracterization of Zhang and attempts in the paragraph following the mischaracterization to remedy or justify it:

Furthermore, while Zhang describes “a charge pump circuit that operates responsive to a test or general operation mode,” Zhang’s teachings including its selector circuit are not limited for only “a test or general operation mode.” See, EX1005, Zhang, 5:23-51 (claiming a charge pump circuit including a selector circuit without any operation mode requirement)². As such, a POSIA [sic] would have been motivated to apply Javanifard’s teachings in Zhang to provide a multi-power-source charge pump to support a device having multiple operation modes, including a device having “a test or general operation mode” and devices having other operation modes (e.g., a memory device having different read, write, programming modes). See, e.g., EX1006, Javanifard, FIG. 2 (illustrating variable stage charge pumps used for a multi-power-source memory device with different modes). [Stiegler, par. 94]

(97) Dr. Stiegler’s remediation fails because the issue is not one of a test mode vs. an operating mode; the issue is what voltage levels are being selected between.

Zhang makes it clear that the two voltage levels are (1) the chip’s operating voltage

² Note: just because the mode is not explicitly required by the claim has no bearing on whether a person of skill in the art would have used the Vpp voltage correctly. It could have been used for test mode or, for example, in specialized systems requiring higher performance through rapid production of boosted voltages. Just because the claim doesn’t mention a test mode is no reason to expect that the high voltage Vpp should be interpreted as a normal operating supply, and it is instructive that nowhere does Zhang suggest it should be used thus.

(“power voltage” or “supply voltage”) and (2) the externally provided boost voltage that will *only* be used by the charge pump, and nothing else in the chip³.

Dr. Stiegler, however, uses Zhang's invention to choose between different operating voltages (“a multi-power-source” system [Stiegler, pars. 79, 90, 93–96, 163, ...]), which is not Zhang's invention. Furthermore, and Dr. Stiegler provides no explanation as to *why* Zhang's voltage selection circuit would be modified to select between different operating voltages. Dr. Stiegler provides no motivation at all for this change.

(98) Dr. Stiegler does not describe how Zhang would have been modified by a person of ordinary skill to accommodate this new purpose and behavior. Like with the previous issue, Dr. Stiegler makes no attempt to address the inconsistency and effectively sweeps the entire issue under the rug.

³ For instance, Zhang says, “In the specification power supply voltage VCC refers to the power voltage operating in general operation mode. The voltage VPP refers to a high voltage provided externally during a test operation mode. The high voltage VPP may be applied externally in the test operation mode as well as during the general operation mode.” [Zhang, 3:22–27] Even when saying that the high voltage can be provided during a general operation mode, Zhang still characterizes the higher voltage VPP as being distinct and different from the operating voltage VCC.

3. One Cannot Use Zhang's Selecting Circuit Block 200 Without Also Providing a Control Mechanism To Determine When To Stop Boosting (Like Zhang's Voltage Level Sensing Block 400)

(99) As described above in the overview on Zhang, the purpose of Zhang's voltage-level sensor is to stop the charge pump's operation when the output level reaches the desired voltage, and this type of control is required because the input voltage is a variable. By its nature, a charge pump produces as output a specific *multiple* of the input voltage, not a specific absolute *voltage level*. Thus, if the input voltage changes, so does the output voltage level that the charge pump naturally generates.

(100) Dr. Stiegler's proposed combination uses only Zhang's input switch 200 and seems to ignore the control aspect (see previous section). This is problematic, because the output generated by the proposed combination will therefore be dependent on *both* the voltage input that is chosen *and* the configuration of the charge pump's stages that is chosen, which means that the output voltage level could be one of many different values. Because the circuits that use these boosted voltages are dependent on very specific voltage levels (e.g., the nonvolatile memory cells in flash-memory chips), this would be problematic. Either the control must be accounted for, or there must be a mechanism similar to Zhang's Voltage Level Sensing Block 400 that stops the voltage boosting once a desired level is reached.

Dr. Stiegler does not describe this either. Dr. Stiegler does not provide any

description of control, except to state that pumps would be turned on and/or off according to the values in his Table 1. While this configures the pumps, it does not *control* them, meaning that it does not demonstrate how to achieve any particular desired voltage level as output. Note that the whole point of Zhang's voltage level sensing block 400, which Dr. Stiegler seems to omit, is to control the output level.

(101) Dr. Stiegler does not describe how the variable voltage-level input is accommodated such that a specific desired output voltage is generated by the proposed Zhang+Javanifard charge pump. Like with the previous issues, Dr. Stiegler makes no attempt to address the technical challenge and effectively sweeps the entire issue under the rug.

4. Combining Zhang With Javanifard Provides No Advantages Beyond Javanifard by Itself

(102) As described above, the two references used in the analysis for Grounds 1 and 2 solve the following problems:

Zhang takes *many* voltages as input and produces *one* voltage as output

Javanifard takes *one* voltage as input and produces *many* voltages as output

(103) Their combination should produce something more than either one by itself. Creating a charge pump that takes *one* voltage as input and produces *one* voltage as

output has already been done. The Stiegler Report suggests that the purpose of the combination is to produce a charge pump that takes *many* voltages as input and produces *many* voltages as output:

A POSITA would have been motivated to apply Javanifard's teaching of a variable stage charge pump in Zhang's multi-power-source charge pump system to produce the obvious, beneficial, and predictable results of a variable stage charge pump system, which can both "accommodate different output levels for a given charge pump power supply input level," "accommodate a given output level for different charge pump power supply input levels," and provide flexibility for design needs (e.g., power saving mode, and current requirements), area reduction, and performance improvement, as taught by Javanifard. [Stiegler, par. 90]

(104) The problem is that Javanifard already does exactly this. Javanifard provides the following objects of the invention:

One object is to provide a charge pump that can accommodate ***different output levels*** for a given charge pump power supply input level.

Another object is to provide a charge pump that can accommodate a given output level for ***different charge pump power supply input levels***. [Javanifard, 2:19–24, emphasis added]

(105) Javanifard already accomplishes all of the stated goals of "a variable stage charge pump system, which can both 'accommodate different output levels for a

given charge pump power supply input level,' 'accommodate a given output level for different charge pump power supply input levels,' and provide flexibility for design needs (e.g., power saving mode, and current requirements), area reduction, and performance improvement," as described by Dr. Stiegler [par. 90]. Zhang adds nothing more that, in the proposed combination, would address these design goals..

(106) The only thing that Dr. Stiegler uses from Zhang is the input switch 200, but Dr. Stiegler never explains why that switch is needed — e.g., what does Javanifard lack that Zhang's input switch 200 provides? This is a critical question that must be addressed. Otherwise, there is no motivation to combine the two references, as it is clear that Javanifard by itself already accomplishes all of the design goals and benefits that Dr. Stiegler alleges his proposed combination of Zhang and Javanifard would produce.

5. Nobody Asked for the Combination of Zhang and Javanifard, as Zhang Already Does It by Itself

(107) As described above in the immediately preceding section, the prior art and the two references used in the analysis for Grounds 1 and 2 all address the different combinations of

takes *one/many* voltage/s as input
and
produces *one/many* voltage/s as output

(108) There are two variables (one/many) and two scenarios in which they can be used (input/output), which creates four combinations. All four are covered by either the prior art or the two references used in the analysis for Grounds 1 and 2. In other words, all of the following problems are solved:

Prior designs take *one* voltage as input and produce *one* voltage as output

Zhang takes *many* voltages as input and produces *one* voltage as output

Javanifard takes *one* voltage as input and produces *many* voltages as output

Javanifard takes *many* voltages as input and produces *many* voltages as output

(109) The bases, as they say, are covered. In the domain of charge pumps, nobody anywhere has defined anything outside of this set of combinations or described a need for it. In particular, nowhere has anyone described a system in which the invention of Zhang and the invention of Javanifard are both needed together — it is found neither in the Stiegler Report, nor in Zhang, nor in Javanifard, nor in any other referenced art.

(110) Moreover, if the goal is the *spirit* of what Dr. Stiegler proposes: “a ~~variable~~ ~~stage~~ charge pump system, which can both ‘accommodate different output levels for a given charge pump power supply input level,’ ‘accommodate a given output

level for different charge pump power supply input levels,” [Stiegler, par. 90, text strike-through added], then Zhang already provides that. A person of skill would have easily recognized that, if he varied the V_{ref} voltage, Zhang's Voltage Level Sensing Block 400 would trigger at different voltage levels and thus would provide different output voltages. The problem would thereby be solved. This could be a determination that is either static or dynamic⁴. If a substantially higher voltage level is required, then one would simply add charge-pump stages, a common technique well within the capabilities of a person of ordinary skill.

6. Conclusion

(111) My conclusion is that the only way to arrive at the proposed combination of Zhang and Javanifard is through improper hindsight, using the '701 claims as a roadmap. Dr. Stiegler's combination takes from Zhang only the input switch 200. The input switch seems to be included only to satisfy specific claim limitations from the '701 patent, as the switch's purpose and use within the combination is not explained anywhere by Dr. Stiegler. Indeed, as I discussed, Javanifard by itself can accomplish all of the purported benefits described by Dr. Stiegler without needing

⁴ Zhang does not explicitly state that V_{ref} cannot be dynamically determined according to a desired output voltage level. Dr. Stiegler opines that Zhang, in some instances, does not explicitly *require* that the input voltages are for test and general operation, and therefore Zhang encompasses other scenarios. By this logic, the fact that Zhang does not explicitly *require* that its V_{ref} stay at a constant voltage would likewise encompass situations where V_{ref} is dynamically determined.

Zhang at all — Javanifard however cannot satisfy all of the '701 claim limitations by itself, which seems to be the main reason for including Zhang. Such reasoning would be well outside the scope of a person of ordinary skill seeking a technical result and is instead the purview of a lawyer seeking a legal result.

D. Ground 2: Dr. Stiegler Misrepresents Javanifard's Figure-2 Embodiment and Proposes a Nonsensical Combination of Javanifard With Itself

(112) In his Grounds-2 analysis, Dr. Stiegler misrepresents the behavior of Javanifard's main working embodiment, misrepresents the purpose of Javanifard's design, and misrepresents circuit design in general. The result is a conclusory combination that would have made no sense to a person of ordinary skill.

(113) Javanifard provides an embodiment with two voltage supplies that can be coupled to a charge pump; this is shown in Figure 2, reproduced below.




Figure 2

Javanifard, Fig 2

(114) The figure shows two voltage sources: V_{PP} (top left) and V_{CC} (bottom left). There is a switch that conditionally connects power supply V_{PP} to *Low Current 12V Charge Pump* 210, and there is a switch that chooses between the two power supplies V_{PP} and V_{CC} and connects one of them to *High Current 5V/12V Charge Pump* 220.

1. The Stiegler Report Misrepresents Javanifard's Figure 2

(115) Dr. Stiegler describes the figure above as follows:

205. Embodiment 2 (described with example of FIG. 2 below and Tables 2 and 3) describes ***selecting from multiple power sources (e.g., VCC and VPP), and providing the selected power supply source to the variable stage charge pumps (e.g., variable stage charge pumps 210 and 220)***. EX1006, Javanifard, 6:61-64.

Each of the variable stage charge pumps is configured based on the input power supply voltage (e.g., VCC or VPP) to provide required output voltage levels. [Stiegler, par. 205, emphasis added]

(116) Dr. Stiegler asserts, incorrectly, that both power supplies can be provided to both of the charge pumps. This is incorrect, as the figure clearly shows that only one of the charge pumps can take VCC as an input, and Javanifard explains very clearly the scenario in which it is allowed to happen.

(117) One can see in the figure above that either VPP or VCC can be used as an input to the *High Current 5V/12V Charge Pump* 220 (lower center). Javanifard makes clear that the two power supplies are not interchangeable: VCC is the chip's "primary power source," meant to power all of the general functions of the chip, and VPP is a "supplemental voltage" used specifically for the charge pump. Javanifard says that the chip's primary power source, voltage input VCC, is to be

used as a voltage input to the charge pump 220 only when V_{PP} is unavailable or non-existent, e.g., if V_{PP} is at 0V:

During a read operation V_{pp} might be zero. Therefore, if high current charge pump 220 is needed it will receive its power from V_{cc} instead of V_{pp} during the read mode. [Javanifard, 8:31–34]

(118) Dr. Stiegler tries to show through annotation that both voltages go to both charge pumps:




Figure 2

EX1006, Javanifard, FIG. 2, annotated

Stiegler, p. 102

(119) Despite Dr. Stiegler's figure annotations suggesting that both voltage inputs V_{PP} and V_{CC} can connect to both charge pumps 210 and 220 (as well as and verbal insistence quoted above), charge pump 210 can only be connected to the voltage supply V_{PP} . Therefore, Dr. Stiegler's conclusion that "Each of the variable stage

charge pumps is configured based on the input power supply voltage (e.g., VCC or VPP) to provide required output voltage levels” [Stiegler, par. 205] is simply wrong.

2. The Stiegler Report Misrepresents A Well-Understood Concept in Circuit Design

(120) In the section on “Reasons to Combine Embodiments of Javanifard,” Dr. Stiegler misrepresents what is meant in circuit design by the notion of a circuit working “for [or with] different power supply levels.”

(121) What the phrase normally means is that the given circuit is designed to function correctly over a *range* of different voltage levels as input. This is what is described in Javanifard — in many places, including the following:

In one embodiment, VCC is approximately 5 volts. Vpp, however, might be 3.3, 5, or 12 volts. [Javanifard, 6:51–52]

To maximize the utility of a flash memory system, however, the variable stage charge pump must be able to generate approximately 5 volts, 9 volts, and 12 volts from either a 3.3 volt or 5 volt power supply. [Javanifard, 6:56–60]

FIG. 2 illustrates the power supply circuitry for a flash memory including variable stage charge pumps for accommodating a 5 volt, 9 volt, and 12 volt voltage level from either a 3.3 volt or 5 volt power supply. [Javanifard, 6:61–64]

In addition, reasonable power supply tolerances typically require that the circuitry function properly when the input voltage is within a given percentage of the nominal value (10%, for example). This means that the output voltage from the charge pumps must depend on V_{pp} input voltages of $3.3 \text{ volts} \pm 10\%$ or $5 \text{ volts} \pm 10\%$. Assuming a 10% lower V_{pp} than nominal, this means that the flash power supply circuitry must be able to provide the proper voltage and current levels to the flash arrays from a V_{pp} of approximately 3.0 volts or 4.5 volts. [Javanifard, 6:67–7:9]

Another disadvantage is that the integrated circuit might not be able to be interchangeably used with both a 3.3 volt power supply and a 5 volt power supply without additional external circuitry. [Javanifard, 1:53–56 — noting that a circuit “might not” support a range of different voltages, which indicates that many circuits would]

(122) Additionally, Tables 2 and 3 both show that the inputs for the charge pumps 210 and 220 can range from 3.3V to 5V in the Figure-2 embodiment:

TABLE 2

High Current Variable Stage Charge Pump (220)

V_{PP}	Mode		
	program	verify	erase
5	18/3	18/3	18/3
3.3	18/3	18/3	9/6

Javanifard, col 8

TABLE 3


Low Current Variable Stage Charge Pump (210)

V_{PP}	Mode		
	program	verify	erase
5	4/3	4/3	4/3
3.3	2/6	2/6	2/6

Javanifard, col 8

(123) Javanifard makes clear that the system does not dynamically choose between multiple different input voltages (e.g., 3.3V vs 5V) but instead can accept either of those as its voltage-supply input.

(124) Javanifard has a switch in place as a redundancy backup in case the required voltage source is not present in a given implementation. Figure 2 is presented again below, annotated to show the power-source switch for charge pump 220, and Javanifard's description of the embodiment is presented below it:



During a read operation V_{pp} might be zero. Therefore if high current charge pump 220 is needed it will receive its power from V_{CC} instead of V_{pp} during the read mode. If V_{CC} is less than 4.0 volts (as determined by the 3/5 volt detector 232), then high current charge pump 220 is needed to provide 5 volts for the wordlines (i.e., gates of the flash memory cells) on line 262. In such a case, charge pump 220 is in the 18/3 configuration in this embodiment. Otherwise, V_{CC} is presumed to be sufficient and line 262 is switched to provide the wordline voltage from V_{CC} instead of from charge pump 220 (which can be turned off because it is no longer required). [Javanifard, 8:31–44, emphasis added]


(125) Despite Javanifard's clear explanations, Dr. Stiegler misrepresents the behavior as one in which the circuits are designed to choose dynamically between different power-supply levels 3.3V and 5V through a selector (e.g., a switch), claiming that Javanifard “explicitly provides motivations to apply its teaching of selection from multiple power sources ... by providing the selected power supply to the variable stage charge pump”:

207. Javanifard explicitly provides motivations to apply its teaching of selection from multiple power sources of Embodiment 2 to the variable stage charge pump of Embodiments 1 or 3, by providing the selected power supply to the variable stage charge pump. For example, one of Javanifard's objects is to “provide a charge pump that can accommodate a given output level for **different charge pump power** supply input levels.” [Stiegler, par. 207, emphasis original]

(126) Javanifard makes clear that the circuits are designed to accept a range of input voltage levels, but Dr. Stiegler interprets this to mean something that Javanifard neither disclosed nor intended: the dynamic choosing between multiple positive voltage levels, e.g., 3.3V vs. 5V.

3. The Stiegler Report Combines Javanifard With Itself, but at Cross-Purposes to Itself

(127) Dr. Stiegler proposes a combination of two different embodiments in Javanifard, as illustrated below:



(128) This proposed combination takes only a single switch from the embodiment in Figure 2 and attaches it, not to one of the existing charge pumps of Figure 2 (variable stage charge pumps 210 and 220) whose configuration and operation are

well described in Javanifard, but rather to a completely different charge pump that is not part of the Figure-2 embodiment.

(129) A more natural combination would have been simply to suggest, as Javanifard itself does⁵, that either or both of charge pumps 210/220 in the Figure-2 embodiment would be implemented as what is shown in Javanifard's Figure 1 or Figure 3. However, doing so would have produced a combination that fails to meet Hayashi's claim limitations. Instead, in the Stiegler Report, Dr. Stiegler invents a new configuration in which Javanifard's description of the charge pumps and their relationships to the voltage supplies can be ignored. In doing so, Dr. Stiegler uses Javanifard's components in ways that are at cross purposes, as described by Javanifard itself.

(130) As described above, Dr. Stiegler proposes a combination in which the switch from Figure 2 is used to choose dynamically between two positive voltage supplies and provide that voltage to the back-end charge pump. Dr. Stiegler claims that Javanifard itself provides the motivation to combine the pieces in this way:

⁵ For instance, charge pumps 210 and 220 in Figure 2 are called "variable stage" charge pumps (see, for example, Tables 2 and 3), and Javanifard describes the charge pump in Figure 1 as a "variable stage" charge pump — e.g., "FIG. 1 illustrates a variable stage charge pump." [Javanifard, 2:34], and "FIG. 1 illustrates one embodiment of a variable stage charge pump." [Javanifard, 2:42–43]

207. Javanifard explicitly provides motivations to apply its teaching of selection from multiple power sources of Embodiment 2 to the variable stage charge pump of Embodiments 1 or 3, by providing the selected power supply to the variable stage charge pump. [...] It was well-known that often different charge pump power supply input levels are provided by different power supply source lines. See, e.g., EX1006, Javanifard, 6:45-46 and EX1005, Zhang, Abstract, FIG. 4. [Stiegler, par. 207, emphasis added]

(131) As described previously, Javanifard provides no such motivations. Javanifard is very clear about the configuration of the charge pumps and their relation to voltage supplies, in that they can accept a range of values, and the charge pump 220 in particular only uses the backup voltage V_{CC} in read operations during which the V_{PP} supply is not present.

(132) Dr. Stiegler uses the components in Javanifard in ways that are contradictory to Javanifard's teachings, and therefore the proposed combination is not one that would have been obvious to a person of ordinary skill.

4. Combining Javanifard With Itself Provides No Advantages Beyond Javanifard's Figure-2 Embodiment Alone

(133) As described in the previous section regarding Ground 1, Javanifard solves the following problems:

Javanifard takes *one* voltage as input and produces *many* voltages as output

Javanifard takes *many* voltages as input and produces *many* voltages as output

(134) The Stiegler Report suggests that the purpose of the combination of Javanifard with itself is to produce a charge pump that takes *any* voltage as input and produces *any* given voltage as output:

[A] POSITA would have been motivated to apply Javanifard teachings of an input power supply selector for selecting from input nodes supplied with different power supply voltages (as taught in Embodiment 2) in Javanifard's variable stage charge pump system (e.g., as taught by variable stage charge pump 100 of Embodiment 1 or the variable stage charge pump with n-charge [sic, probably "n-stage"] pumps as taught in Embodiment 3). Such combinations have the benefit of providing a charge pump that can "**accommodate a given output level for different charge pump power supply input levels**", e.g., for a device supplied with multiple power sources using two source lines (e.g., a primary power source and a supplemental power source).

[Stiegler, par. 216, emphasis original]

(135) The problem is that Javanifard already does exactly this. Javanifard provides the following objects of the invention:

One object is to provide a charge pump that can accommodate ***different output levels*** for a given charge pump power supply input

level.

Another object is to provide a charge pump that can accommodate a given output level for ***different charge pump power supply input levels***. [Javanifard, 2:19–24, emphasis added]

(136) Javanifard's Figure-2 embodiment already accomplishes all of the stated goals of "providing a charge pump that can 'accommodate a given output level for different charge pump power supply input levels', e.g., for a device supplied with multiple power sources using two source lines (e.g., a primary power source and a supplemental power source)," as described by Dr. Stiegler [par. 216]. Dr. Stiegler's proposed combination adds nothing more that would address these design goals.

(137) The only thing that Dr. Stiegler adds is the switch from Figure 2, configured in a way that Javanifard neither illustrates nor describes, but Dr. Stiegler never explains why that switch is needed or why a person of ordinary skill would be motivated to do that. Javanifard's Figure 2 already has everything one would ever need for a system. What does Figure 2 lack that is improved by eliminating all of its functionality and then replacing it with a single charge pump circuit? This is a critical question that must be addressed. Otherwise, there is no motivation for a person of ordinary skill to produce the proposed combination, as it is clear that Javanifard's Figure 2 by itself already accomplishes all of the design goals and benefits that Dr. Stiegler alleges his proposed combination would produce.

(138) More to the point, Dr. Stiegler's proposal combines a heavily modified Figure-2 embodiment with a variable-stage charge pump, and it is clear that the modification is quite significant, in that it completely changes the principle under which the Figure-2 embodiment operates. For example, as is well known in the industry, the design and control of charge pumps is non-trivial [see, for example, Pan & Samaddar 2006, *Preface*], and Javanifard confirms this. Javanifard says the following about the functional requirements and control aspects of the Figure-2 embodiment:

Although 5 volts is a popular standard for V_{pp} , 3.3 volts is gaining in popularity. The variable stage charge pumps might be incorporated into the flash memory device in order to simplify external circuitry. *To maximize the utility of a flash memory system, however, the variable stage charge pump must be able to generate approximately 5 volts, 9 volts, and 12 volts from either a 3.3 volt or 5 volt power supply.*

FIG. 2 illustrates the power supply circuitry for a flash memory including variable stage charge pumps for accommodating a 5 volt, 9 volt, and 12 volt voltage level from either a 3.3 volt or 5 volt power supply. *The variable stage charge pumps are designed to exceed the nominal required voltages because its output is routed to a voltage regulator.* In addition, reasonable power supply tolerances typically require that the circuitry function properly when the input voltage is within a given percentage of the nominal value (10%, for example). This means that *the output voltage from the charge pumps must depend on V_{pp} input voltages of 3.3 volts $\pm 10\%$ or 5 volts $\pm 10\%$.* Assuming a 10% lower V_{pp} than nominal, this means that the flash power supply circuitry must be able to

provide the proper voltage and current levels to the flash arrays from a V_{pp} of approximately 3.0 volts or 4.5 volts. [Javanifard, 6:53–7:10, emphasis added]

*The stage control for the variable stage charge pumps 210 and 220 is determined by the operational mode (i.e., erase, read, or program) and by the V_{pp} level. Level detectors 230, 231, and 232 determine V_{pp} and VCC levels. 5/12 volt V_{pp} detector 230 is used to determine whether V_{pp} is at 5 volts or 12 volts. Detector 230 indicates a V_{pp} of approximately 12 volts. Variable stage charge pump 220 is a high current variable stage charge pump used during the read, program, and erase operations. Variable stage charge pump 210 is a low current variable stage charge pump used in addition to charge pump 220 only during program and erase operations. The output of the V_{pp} detectors is provided to circuitry 290 to select the appropriate erase and programming algorithms for the various V_{pp} levels. These algorithms may vary depending upon V_{pp} and VCC. Circuitry 290 controls the appropriate switches (e.g., switch 274) in accordance with the erase and programming algorithms. *These switches are used, for example, for controlling power to charge pumps 210 and 220 and for selecting the appropriate source of power to be supplied to the flash memory array on lines 260 and 264.* [Javanifard, 7:21–42, emphasis added]*

For the read mode, low current variable stage charge pump 220 is not used and can be turned off. *During a read operation V_{pp} might be zero. Therefore if high current charge pump 220 is needed it will receive its power from VCC instead of V_{pp} during the read mode. If VCC is less than 4.0 volts (as determined by the 3/5 volt detector 232), then high current charge pump 220 is needed to provide 5 volts for the wordlines (i.e., gates of the flash memory cells) on line 262. In such a case, charge pump 220 is in the 18/3 configuration in this*

embodiment. Otherwise, VCC is presumed to be sufficient and line 262 is switched to provide the wordline voltage from VCC instead of from charge pump 220 (which can be turned off because it is no longer required). [Javanifard, 8:30–44, emphasis added]

(139) There are numerous aspects of control here that Dr. Stiegler eliminated when he chose to eliminate *all* of the control circuitry from the Figure-2 embodiment, retaining only one voltage switch. These include voltage-level tolerances, recognition of and response to operational mode, detection of and response to different voltage input levels, detection and selection of appropriate output voltage levels, etc.

(140) Dr. Stiegler's choice to eliminate all the mechanisms and control circuitry designed to perform the functions described above significantly changes the entire principle of operation of the Figure-2 embodiment. Dr. Stiegler does not explain anywhere why such a significant modification to the Figure-2 embodiment would have been obvious to a person of ordinary skill; he does not explain why it would have been necessary; and he does not explain how the system would be capable of functioning in the way that Javanifard describes (e.g., "the variable stage charge pump must be able to generate approximately 5 volts, 9 volts, and 12 volts from either a 3.3 volt or 5 volt power supply"), after the removal of all the requisite control circuitry.

5. Conclusion

(141) My conclusion is that the only way to arrive at the proposed combination of Javanifard with itself is through improper hindsight, using the '701 claims as a roadmap. Dr. Stiegler's combination takes from Javanifard's Figure-2 embodiment only the input switch and applies it in a way that is not described in Javanifard and is instead described only in the '701 patent claims. The input switch seems to be included only to satisfy specific claim limitations from the '701 patent, as the switch's purpose and use within the combination is not explained anywhere by Dr. Stiegler. Indeed, Javanifard's Figure-2 embodiment can accomplish all of the purported benefits described by Dr. Stiegler as presented — Javanifard however cannot satisfy all of the '701 claim limitations by itself, which seems to be the main reason for including the switch in the way that Dr. Stiegler does. Such reasoning would be well outside the scope of a person of ordinary skill seeking a technical result and is instead the purview of a lawyer seeking a legal result.

E. Claims 1–18 Are Therefore Valid

(142) Because the Stiegler Report's Ground 1 analysis for Claims 1–18 of the '701 patent requires various combinations of Zhang and Javanifard, and as described above there is significant motivation *not* to combine Zhang with Javanifard, the Stiegler Report's Ground 1 obviousness analysis fails.

(143) Because the Stiegler Report's Ground 2 analysis for Claims 1–18 of the '701 patent requires a combination of Javanifard with itself that is in contradiction with Javanifard's teachings, the Stiegler Report's Ground 2 obviousness analysis fails.

Executed on Tuesday, May 7, 2024 in Annapolis, Maryland



Dr. Bruce Jacob

FIELDS Hardware-Based Security (*Side-Channel Analysis, Robust Computer Architectures & Operating Systems*)
Memory Systems (*Caches, DRAM, Flash & Other Non-Volatile Technologies — from Devices to Systems*)
Technology Innovation (*Start-Up Companies, Entrepreneurship, High Tech, & Electric Guitars*)

In the middle of COVID lock-downs, Bruce Jacob retired from the University of Maryland, after spending the previous 25 years there as a professor. In 2022, he joined the U.S. Naval Academy's Department of Cyber Science, where he studies the intersection between hardware and security, including the following:

- the security and protection of memory devices and memory systems
- side-channel exploits and their mitigation
- the development of robust computer architectures and operating systems, a pursuit that combines his knowledge of memory systems and operating systems with methodologies derived from the side-channel experiments

In addition, he continues to do research in the design of novel memory architectures and systems that exploit them.

In 2020, Jacob was named a Fellow of the IEEE "for contributions to computer memory design and analysis." He has designed computer-system architectures and computer memory-system architectures for both industry and government, national and international, focusing on highly efficient designs at the High-Performance Computing level, as well as at the high-performance embedded-systems level. For instance, he helped Micron design their new *Hybrid Memory Cube* DRAM architecture, he redesigned Cray's memory controller for their *Black Widow* memory system, he helped Northrop Grumman design a system interconnect for their experimental ultra-low-power datacenter, he designed a high-performance memory system for the 1024-core *Teraflux* chip funded by the European Commission, he has collaborated with researchers at the Department of Energy on the design of their next-generation supercomputers, and with collaborators at the University of Maryland, he currently works on the design of monolithic computer systems that integrate high-capacity restive RAMs directly into CMOS logic chips for high-performance, low power, single-chip systems.

Recognized internationally as a leading expert on computer memory systems, Jacob founded the annual *International Symposium on Memory Systems* and is regularly invited to give keynote speeches and high-level briefings on the topic of memory systems around the world. He has been the keynote speaker at meetings including EMS, MCHPC, SAMOS, Multicore World, and Computing Frontiers. In 2007, he was one of fifteen members of academia and industry invited to brief a National Academies panel on the then-current state of computing, resulting in the highly-cited NRC exascale report *The Future of Computing Performance — Game Over or Next Level?* In 2012, he was one of twenty members of academia and industry asked to brief the Secretary of Energy, Steven Chu, on the state of computing relative to the Department of Energy's plans to achieve exascale-class computer performance. In 2016, he was one of a dozen members of academia and industry invited to participate with DOE in creating a roadmap for computer science and applied mathematics research at DOE in the coming decades. In 2020, he was one of four members of academia and two dozen members of industry and national labs invited to work with Semiconductor Research Corp., Semiconductor Industry Assoc., and DOE to develop a roadmap for memory devices and systems in the coming decade.

In addition to his academic credentials, Jacob has significant industry experience. Between college and graduate school, he worked in the start-up industry in Boston for two different telecommunications companies, serving as a software engineer at Boston Technology and then as the chief engineer and system architect at Priority Call Management. Both of these start-up companies were successful — in particular, Priority Call Management, for which Jacob designed and developed the product's system-level architecture, its distributed middleware code, and its object-oriented applications framework, was purchased for \$162M in the late 1990s.

Jacob has has written two textbooks on computer memory systems and over one hundred articles ranging over memory systems, computer design, embedded systems, operating system design, astrophysics, and algorithmic composition. He holds two patents in memory-systems design and three patents in the circuit design of electric guitars. In 2009, the electric guitars were featured on Washington DC local television and radio stations; in articles appearing in *The Washington Post*, *Los Angeles Times*, and *The Chronicle of Higher Education*; and on National Public Radio (Jacob and his graduate student were interviewed by Robert Siegel on *All Things Considered* on July 10, 2009).

I. EDUCATION AND EMPLOYMENT

A. EDUCATION

- 1997 *Ph.D.* **Computer Science & Engineering** **University of Michigan–Ann Arbor**
Dissertation: *Software-Oriented Memory-Management Design*
Results published in *HPCA 1997*; *IEEE Computer*, June 1998; *IEEE Micro*, July/August 1998
- 1995 *M.S.* **Computer Science & Engineering** **University of Michigan–Ann Arbor**
Thesis: *Optimization of Storage Hierarchies*
Results published in *IEEE Transactions on Computers*, October 1996
- 1988 *A.B.* **Mathematics (cum laude)** **Harvard University**
Senior Project: *A Radio Study of the Galactic Center*
Results published in *Interstellar Matter*, Gordon & Breach NY, 1987
Starting PK on 1987 Ivy League Championship football team (yes, I have a ring)

B. PROFESSIONAL EXPERIENCE

- 2022–present **Professor**, Department of Cyber Science, United States Naval Academy, Annapolis MD
- 2021–2022 **Professor/Chair**, Department of Electrical & Computer Engineering, Florida Polytechnic University
- 2020 Named **Fellow**, *Institute of Electrical and Electronics Engineers (IEEE)*
The IEEE Board of Directors, at its November 2020 meeting, elevated Jacob to IEEE Fellow, effective as of 1 January 2021, with the following citation:
for contributions to computer memory design and analysis
- 2015 **Founder**, International Symposium on Memory Systems (*published by ACM*)
First meeting held October 2015; every year sees over 100 attendees, ~30% international, who come for a 3-day technical meeting that has now become the world's preeminent forum for research in memory systems
- 2014 **Visiting Professor**, Foundation for Research & Technology – Hellas (FORTH), Heraklion Greece
*Two-month research collaboration on EC-sponsored datacenter project **EuroServer**, designing nonvolatile main memories for improved system-level power dissipation*
- 2014 **Visiting Professor**, University of Cyprus, Nicosia Cyprus
*One-month research collaboration on EC-sponsored supercomputer project **Teraflux**, studying dataflow issues within single-chip multiprocessors*
- 2012–2013 **Visiting Professor**, University of Siena, Siena Italy
*Fourteen-month research collaboration on EC-sponsored supercomputer project **Teraflux**, designing a memory system for their 1024-core CPU chip*
- 2011–2021 **Professor (Full)**, Department of Electrical & Computer Engineering, University of Maryland, College Park MD
*Helped design Micron's **Hybrid Memory Cube** (revolutionary DRAM architecture announced in 2012); designed a system interconnect for Northrop Grumman's experimental ultra-low-power datacenter*
- 2007–2010 **Director of Computer Engineering**, Department of Electrical & Computer Engineering, University of Maryland, College Park MD
Updated the curriculum of the undergraduate degree program and increased the program's visibility and enrollment (applications were up 25% the year after changes were made)
- 2006–2010 **Founder**, Coil Guitars, College Park MD
*Patented three circuit-design technologies and served as public relations emissary for the University of Maryland; featured in **The Washington Post**, **Los Angeles Times**, and on **National Public Radio***

- 2006 Named **Keystone Professor**, University of Maryland's Clark School of Engineering
Honorific for teaching excellence; helped to redesign the Clark School's introductory engineering class to become more interdisciplinary and to use more relevant projects
- 2003–2011 **Associate Professor**, Department of Electrical & Computer Engineering, University of Maryland, College Park MD
Redesigned Cray's memory controller, significantly improving their system's sustained memory bandwidth; wrote and published Memory Systems: Cache, DRAM, Disk (1000-page reference work)
- 1997–2003 **Assistant Professor**, Department of Electrical & Computer Engineering, University of Maryland, College Park MD
*Received National Science Foundation **CAREER Award** for work on DRAM devices and architectures*
- 1992–1997 **Graduate Student Research Assistant**, University of Michigan, Ann Arbor MI
Designed the memory-management architecture for the PUMA processor
- 1991–1992 **System Architect/Chief Engineer**, Priority Call Management, Wilmington MA
Employee #2; designed the company's real-time distributed telecommunications system; developed the first system prototype used to attract seed funding; designed, developed, and installed into the Dallas Convention Center the company's first real product
- 1989–1991 **Software Engineer**, Boston Technology, Cambridge MA
Developed real-time telecommunications applications
- 1988–1989 **Mathematics Teacher**, Thayer Academy, Braintree MA

II. RESEARCH, SCHOLARLY, AND CREATIVE ACTIVITIES

A. CONTRACTS AND GRANTS

Office of Naval Research

Investigating "Complex" Instruction-Set Design and Multicore Processors for Improved Computer Security; Summer 2023; \$46,000. PI.

Department of Defense

E-VERIFY: ReRAM Architecture Exploration; 8/2022–2/2024; \$420,000. Co-PI, with Yeung (PI) and Peckerar.

Department of Defense

E-VERIFY: Advanced Memory Modeling, Simulation, and Architectural Studies — Task 8; 8/2022–12/2025; \$400,000. PI.

Department of Defense

E-VERIFY: High Bandwidth Memory System Development; 8/2021–8/2022; \$420,000. Co-PI, with Yeung (PI).

Department of Defense

E-VERIFY: Advanced Memory Modeling, Simulation, and Architectural Studies — Task 8; 8/2021–8/2022; \$210,000. PI.

Department of Defense

E-VERIFY: High Bandwidth Memory System Development; 5/2020–5/2021; \$270,000. Co-PI, with Yeung (PI).

Department of Defense

E-VERIFY: High Bandwidth Memory Model/Simulation Development; 5/2020–5/2021; \$175,000. PI.

Semiconductor Research Corporation

Monolithic Integration of CPU and Main Memory Systems; 1/2020–12/2022; \$270,000. Co-PI, with Yeung (PI).

Department of Defense

Memory Modeling and Simulation; 5/2019–5/2020; \$300,000. PI.

Department of Defense

Emerging Memory Technologies; 5/2019–5/2020; \$200,000. Co-PI, with Yeung (PI).

Department of Defense

Memory Modeling: Infrastructure Extension and RRAM Plug-In; 5/2018–5/2019; \$175,000. PI.

Department of Defense

Using RRAM to Build CPUs with On-Chip Memory Systems; 5/2018–5/2019; \$275,000. Co-PI, with Yeung (PI).

Department of Defense

High Bandwidth Memory Modeling and Simulation; 2/2018–2/2020; \$500,000. PI.

National Science Foundation

3DSIM: A Unified Framework for 3D CPU Co-Simulation; 9/2016–9/2019; \$500,000. Co-PI, with Srivastava (PI).

Northrop Grumman Corporation

Memory-System Design for Cold Logic Program; 9/2015–12/2015; \$625,000. PI.

Department of Energy Office of Science

Data Movement Dominates II: Advanced Memory Technology to Address the Real Exascale Power Problem (Extension); 10/2013–9/2015; \$5,100,000. Co-PI, with investigators at Sandia National Laboratories (Lead Institution), Lawrence Berkeley National Laboratory, Micron Technology, and Columbia University.

Sandia National Laboratories

E-VERIFY: High Performance Memory Systems Design and Optimization; 2/2012–2/2013; \$150,000. PI.

Defense Advanced Research Projects Agency

X-caliber Ubiquitous High Performance Computing (UHPC) Proposal; 10/2010–9/2014; \$25,800,000. Academic Subcontractor, on team including Sandia National Laboratories (Lead Institution), FTL Systems, Micron Technologies, LexisNexis Special Services, Indiana Integrated Circuits LLC, Georgia Institute of Technology, Louisiana State University, North Carolina State University, Stanford University, University of Illinois, University of Notre Dame, and University of Southern California.

Department of Energy Office of Science

Data Movement Dominates: Advanced Memory Technology to Address the Real Exascale Power Problem; 10/2010–9/2013; \$5,100,000. Co-PI, with investigators at Sandia National Laboratories (Lead Institution), Lawrence Berkeley National Laboratory, Micron Technology, and Columbia University.

Intel Corporation

Novel Applications of Persistent Memory; 10/2010–9/2013; \$440,000. PI.

Sandia National Laboratories

Non-Volatile Memory Systems; 10/2010–9/2013; \$200,000. PI.

Maryland Industrial Partnerships

Novel Signal-Processing Prototypes II; 1/2010–12/2010; \$135,000. PI.

Intel Corporation

Memory Hierarchy Architecture Research; 11/2009 [Foundation gift]; \$120,000. PI.

Cypress Semiconductor Corporation

Memory-System Modeling; 11/2009–7/2010; \$200,000. PI.

Maryland Industrial Partnerships

Novel Signal-Processing Prototypes; 1/2009–12/2009; \$135,000. PI.

Department of Defense

Next-Generation DRAM-System Architectures and Optimizations; 11/2008–10/2013; \$500,000. PI.

Note: this contract covered the collaborative design and development of Micron's Hybrid Memory Cube DRAM architecture ("HMC"), announced as a commercial product in 2012.

Sandia National Laboratories

High Performance Memory-Systems Design and Optimization; 6/2008–12/2010; \$255,000. PI.

National Institute of Standards and Technology

Integration Methodologies and Full System Verification for Heterogeneous Systems on Chip; 1/2007–1/2009; \$150,000. PI.

Office of the Secretary of Defense

Focused Logistics Transformation and Sense & Respond Support; 9/2005–12/2007; \$120,000 (my portion). Associated Researcher, on team including investigators from the School of Public Policy (Lead Institution); the College of Computer, Mathematical, and Natural Sciences; and the College of Engineering.

National Science Foundation, MRI Award

MRI: Development of Energy-Efficient Embedded Systems for Wireless Sensor Networks; 8/2005–7/2008; \$400,000. Co-PI, with Ephremides (PI), Abshire, Barua, Petrov, Qu, Ulukus, and Vishkin. Equipment grant.

Cray Computer Corporation

High-Bandwidth Controller and System Design; 5/2005–11/2005; \$60,000. PI.

IBM Corporation, SUR Award

DRAM Device Configuration and Performance Scalability; 2/2005; \$25,000 in equipment. PI.

National Institute of Standards and Technology

System-on-Chip Integration of Microcontroller with Gas Sensor VC; 1/2005–1/2007; \$140,000. PI.

Cray Computer Corporation

DDR3 Memory-System Configuration Sensitivity; 4/2004–9/2004; \$50,000. PI.

National Institute of Standards and Technology

Embedded Sensor System-on-Chip Microcontroller Core and System Architecture Design; 8/2003–5/2004; \$30,000. PI.

National Science Foundation, ITR Award

ITR: Parallel Random-Access Model (PRAM)-On-Chip; 9/2003–9/2008; \$750,000. Co-PI, with Vishkin (PI), Barua, Franklin, and Qu.

Department of Defense — Maryland Procurement Office

Joint Program for Advanced Electronic Materials; 6/2003–5/2006; \$2,400,000. Associated Researcher, with Lee (PI), Goldsman, and Yang.

Air Force Office of Scientific Research, MURI Award

Basic Studies of the Effect of Microwave Pulses on Electronics; 5/2001–4/2006; \$3,300,000. Co-PI, with Granatstein (PI), Anlage, Antonsen, Baker, Carmel, Goldsman, Iliadis, Melngailis, O'Shea, Ott, Ramahi, and Rodgers.

National Science Foundation

Hardware-Software Co-Design of Real-Time Operating Systems and Embedded Microprocessors; 3/2001–3/2005; \$1,400,000. PI, with Stewart (subcontractor).

National Science Foundation, CAREER Award

Architecture Issues in DRAM Devices and Systems; 1/2000–1/2004; \$230,000. PI.

Compaq Computer Corporation

Evaluation of EV6 with Contemporary DRAM Architectures; 8/1999; \$40,000 in equipment. PI.

IBM Corporation

The Study of Design Parameters for DDR DRAM and DDRII DRAM; 11/1999–1/2003; joint study agreement with access to IBM proprietary mainframe memory traces. PI.

National Science Foundation

Hardware-Software Co-Design of an Experimental Real-Time Operating System and a Microcontroller Architecture; 9/1998–9/2000; \$280,000. Co-PI, with Stewart (PI).

UNIVERSITY OF MARYLAND INTERNAL

UMD Minta Martin Aeronautical Research Fund

Microarchitecture Support for Real-Time Embedded Systems; 5/1998–10/1999; \$35,000.

UMD General Research Board Summer Research

A Better Simulation Environment for Computer Architecture; 6/1998–9/1998; \$6,000.

Note: In the following sections on publications, lead author is first; student advisees are in **bold**; number of citations is given only for works whose citation count is notable (e.g., double digits). Citation counts come from Google Scholar and include self-citations.

Total number of citations > 8200; **mean citations** per work ~80; **i1000-index** of citations = 2 (works with 1000+ citations); **i100-index** of citations = 24 (works with 100+ citations); **i10-index** of citations = 67 (works with 10+ citations)

B. BOOKS

B. Jacob. The Design of Modern Computers from Transistor to Kernel — Building Osprey and the RiSC-17. Forthcoming.

B. Jacob, with contributions by **S. Srinivasan** and **D. T. Wang**. The Memory System (You Can't Avoid It; You Can't Ignore It; You Can't Fake It). ISBN 978-1598295870. Morgan & Claypool Publishers: San Rafael CA, 2009.

Citations: 97

B. Jacob, S. Ng, and **D. Wang**, with contributions by **S. Rodriguez**. Memory Systems: Cache, DRAM, Disk. ISBN 978-0123797513. Morgan Kaufmann: San Francisco CA, 2007.

Citations: 1,180

➡ ~500,000 words, 950 pages. An authored work, not an edited work.

C. REFEREED WORK (JOURNAL ARTICLES AND CONFERENCE PAPERS)

B. Jacob. “Thoughts on merging the file system and virtual memory system — Design decisions and their ramifications in developing the Osprey kernel.” In *Proc. ACM/IEEE International Symposium on Memory Systems (MEMSYS 2023)*, pp. 173–179. Washington DC, October 2023.

M. Peckerar, P.-C. Huang, R. A. Jamil, B. Jacob, and D. Yeung. “Critical issues in advanced ReRAM development.” In *Proc. ACM/IEEE International Symposium on Memory Systems (MEMSYS 2023)*, p. 285. Washington DC, October 2023.

L. Kang, Y. Xue, W. Jia, X. Wang, J. Kim, C. Youn, M. Kang, H. Lim, B. Jacob, J. Huang. “IceClave: A trusted execution environment for in-storage computing.” In *Proc. 54th International Symposium on Microarchitecture (MICRO 2021)*, pp. 199–211. Athens Greece, October 2021.

Citations: 18

L. Kang and B. Jacob. “Zoned FTL: Achieve resource isolation via hardware virtualization.” In *Proc. ACM/IEEE International Symposium on Memory Systems (MEMSYS 2021)*, pp. 1–8. Washington DC, October 2021.

C. Walden, D. Singh, **M. Jagasivamani**, **S. Li**, **L. Kang**, M. Asnaashari, S. Dubois, B. Jacob, and D. Yeung. “Monolithically integrating non-volatile main memory over the last-level cache.” *ACM Transactions on Architecture and Code Optimization*, vol. 18, no. 4, article 48. July 2021.

M. Jagasivamani, C. Walden, D. Singh, **L. Kang**, M. Asnaashari, S. Dubois, B. Jacob, D. Yeung. “Tileable monolithic ReRAM memory design,” In *Proc. IEEE Symposium in Low-Power and High-Speed Chips (COOL CHIPS 2020)*, pp. 1–3. Kokubunji Japan, April 2020.

S. Li, Z. Yang, **D. Reddy**, A. Srivastava, B. Jacob. “DRAMsim3: A cycle-accurate, thermal-capable DRAM simulator.” *IEEE Computer Architecture Letters*, vol. 19, no. 2, pp. 106–109. July–December 2020. Published on-line February 2020.

Citations: 118

M. Jagasivamani, C. Walden, D. Singh, **L. Kang**, **S. Li**, M. Asnaashari, S. Dubois, D. Yeung, and B. Jacob. “Analyzing the monolithic integration of a ReRAM-based main memory into a CPU’s die.” *IEEE Micro*, vol. 39, no. 6, pp. 64–72. November/December 2019. (Special Issue on Monolithic 3D Architectures)

S. Li, R.S. Verdejo, P. Radojković, and B. Jacob. “Rethinking cycle accurate DRAM simulation.” In *Proc. ACM/IEEE International Symposium on Memory Systems (MEMSYS 2019)*, pp. 184–194. Washington DC, October 2019.

Citations: 15

M. Jagasivamani, C. Walden, D. Singh, **L. Kang**, **S. Li**, M. Asnaashari, S. Dubois, D. Yeung, and B. Jacob. “Design for ReRAM-based main-memory architectures.” In *Proc. ACM/IEEE International Symposium on Memory Systems (MEMSYS 2019)*, pp. 342–350. Washington DC, October 2019.

Citations: 11

S. Li and B. Jacob. “Statistical DRAM modeling.” In *Proc. ACM/IEEE International Symposium on Memory Systems (MEMSYS 2019)*, pp. 521–530. Washington DC, October 2019.

Citations: 10

M. Radulovic, R. Verdejo, P. Carpenter, P. Radojković, B. Jacob, and E. Ayguadé. “PROFET: Modeling system performance and energy without simulating the CPU.” *Proceedings of the ACM on Measurement and Analysis of Computing Systems (POMACS)*, vol. 3, no. 2, pp. 34:1–33. June 2019.

S. Li, **D. Reddy**, and B. Jacob. “A performance & power comparison of modern high-speed DRAM architectures.” In *Proc. IEEE International Symposium on Memory Systems (MEMSYS 2018)*, pp. 331–343. Washington DC, October 2018.

Citations: 36

R.S. Verdejo, K. Asifuzzaman, M. Radulovic, P. Radojković, E. Ayguade, and B. Jacob. “Main memory latency simulation: The missing link.” In *Proc. IEEE International Symposium on Memory Systems (MEMSYS 2018)*, pp. 107–116. Washington DC, October 2018.

Citations: 18

M. Jagasivamani, C. Walden, D. Singh, **L. Kang**, **S. Li**, M. Asnaashari, S. Dubois, B. Jacob, and D. Yeung. “Memory-system design challenges in realizing monolithic computers.” In *Proc. IEEE International Symposium on Memory Systems (MEMSYS 2018)*, pp. 98–104. Washington DC, October 2018.

Citations: 10

S. Li, **P.-C. Huang**, and B. Jacob. “Exascale interconnect topology characterization and parameter exploration.” In *Proc. 20th IEEE International Conference on High Performance Computing and Communications (HPCC 2018)*. Exeter UK, June 2018.

D. M. Mathew, E. F. Zulian, M. Jung, K. Kraft, C. Weis, B. Jacob, N. Wehn. “Using run-time reverse-engineering to optimize DRAM refresh.” In *Proc. International Symposium on Memory Systems (MEMSYS 2017)*, pp. 115–124. Washington DC, October 2017.

S. Li, **P.-C. Huang**, **D. Banks**, **M. DePalma**, **A. Elshaarany**, S. Hemmert, A. Rodrigues, **E. Ruppel**, **Y. Wang**, J. Ang, and B. Jacob. “Low latency, high bisection-bandwidth networks for exascale memory systems.” In *Proc. 2016 International Symposium on Memory Systems (MEMSYS 2016)*, pp. 62–73. Washington DC, October 2016.

P. Tschirhart, **J. Stevens**, Z. Chishti, and B. Jacob. “The case for associative DRAM caches.” In *Proc. 2016 International Symposium on Memory Systems (MEMSYS 2016)*, pp. 211–219. Washington DC, October 2016.

I. Bhati, **M.-T. Chang**, Z. Chishti, S.-L. Lu, and B. Jacob. “DRAM refresh mechanisms, penalties, and trade-offs.” *IEEE Transactions on Computers*, vol. 65, no. 1, pp. 108–121. January 2016.

Citations: 145

P. Tschirhart, **J. Stevens**, Z. Chishti, S.-L. Lu, and B. Jacob. “Bringing modern hierarchical memory systems into focus: A study of architecture and workload factors on system performance.” In *Proc. 2015 International Symposium on Memory Systems (MEMSYS 2015)*, pp. 179–190. Washington DC, October 2015.

J. Stevens, **P. Tschirhart**, and B. Jacob. “The semantic gap between software and the memory system.” In *Proc. 2015 International Symposium on Memory Systems (MEMSYS 2015)*, pp. 43–46. Washington DC, October 2015.

I. Bhati, Z. Chishti, S.-L. Lu, and B. Jacob. “Flexible Auto-Refresh: Enabling scalable and energy-efficient DRAM refresh reductions.” In *Proc. 42nd International Symposium on Computer Architecture (ISCA 2015)*, pp. 235–246. Portland OR, June 2015.

Citations: 125

B. Jacob. “The 2 petaflop, 3 petabyte, 9 TB/s, 90 kW cabinet: A system architecture for Exascale and Big Data.” *IEEE Computer Architecture Letters*, vol. 15, no. 2, pp. 125–128. July–December 2016. Published on-line June 2015.

B. Jacob. “The case for VLIW-CMP as a building block for Exascale.” *IEEE Computer Architecture Letters*, vol. 15, no. 1, pp. 54–57. January–June 2016. Published on-line April 2015.

I. Bhati, Z. Chishti, and B. Jacob. “Coordinated Refresh: Energy-efficient techniques for DRAM refresh scheduling.” In *Proc. 2013 International Symposium on Low Power Electronics and Design (ISLPED 2013)*, pp. 205–210. Beijing China, September 2013.

Citations: 31

J. Stevens, **P. Tschirhart**, **M.-T. Chang**, **I. Bhati**, **P. Enns**, J. Greensky, Z. Chishti, S.-L. Lu, and B. Jacob. “An integrated simulation infrastructure for the entire memory hierarchy: Cache, DRAM, nonvolatile memory, and disk.” *Intel Technology Journal*, vol. 17, no. 1, pp. 184–200. April 2013.

Citations: 28

B. Jacob. “Forward ... A Foreword.” Foreword to Special Issue on Memory Systems, *Intel Technology Journal*, vol. 17, no. 1. April 2013.

M.-T. Chang, **P. Rosenfeld**, S.-L. Lu, and B. Jacob. “Technology comparison for Large Last-Level Caches (L3Cs): Low-leakage SRAM, low write-energy STT-RAM, and refresh-optimized eDRAM.” In *Proc. 19th International Symposium on High Performance Computer Architecture (HPCA 2013)*, pp. 143–154. Shenzhen China, February 2013.

Citations: 288

E. Cooper-Balis, **P. Rosenfeld**, and B. Jacob. “Buffer On Board memory systems.” In *Proc. 39th International Symposium on Computer Architecture (ISCA 2012)*, pp. 392–403. Portland OR, June 2012.

Citations: 68

K. Bergman, G. Hendry, P. Hargrove, J. Shalf, B. Jacob, K. S. Hemmert, A. Rodrigues, and D. Resnick. “Let there be light! The future of memory systems is photonics and 3D stacking.” In *Proc. ACM SIGPLAN Workshop on Memory Systems Performance and Correctness (MSPC 2011)*, pp. 43–48. San Jose CA, June 2011.

Citations: 14

P. Rosenfeld, **E. Cooper-Balis**, and B. Jacob. “DRAMsim2: A cycle-accurate memory-system simulator.” *Computer Architecture Letters*, vol. 10, no. 1, pp. 16–19. January 2011.

Citations: 1,104

A. Rodrigues, K. Hemmert, B. Barrett, C. Kersey, R. Oldfield, M. Weston, R. Risen, J. Cook, **P. Rosenfeld**, **E. Cooper-Balis**, and B. Jacob. “The Structural Simulation Toolkit.” *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 4, pp. 37–42. March 2011.

Citations: 373

E. Cooper-Balis and B. Jacob. “Fine-grained activation for power reduction in DRAM.” *IEEE Micro*, vol. 30, no. 3, pp. 34–47. May/June 2010.

Citations: 84

C. Dirik and B. Jacob. “The performance of PC solid-state disks (SSDs) as a function of bandwidth, concurrency, device architecture, and system organization.” In *Proc. 36th International Symposium on Computer Architecture (ISCA 2009)*, pp. 279–289. Austin TX, June 2009.

Citations: 294

S. Srinivasan, L. Zhao, **B. Ganesh**, B. Jacob, M. Espig, and R. Iyer. “CMP memory modeling: How much does accuracy matter?” In *Proc. Fifth Annual Workshop on Modeling, Benchmarking and Simulation (MoBS 2009)*, pp. 24–33. Austin TX, June 2009.

Citations: 27

A. Jaleel, R. S. Cohn, C.-K. Luk, and B. Jacob. “CMP\$im: A Pin-based on-the-fly multi-core cache simulator.” In *Proc. Fourth Annual Workshop on Modeling, Benchmarking and Simulation (MoBS 2008)*, pp. 28–36. Beijing China, June 2008.

Citations: 243

A. Varma, E. Debes, I. Kozintsev, P. Klein, and B. Jacob. “Accurate and fast system-level power modeling: An XScale-based case study.” *ACM Transactions on Embedded Computing Systems*, vol. 7, no. 3, pp. 25:1–20. April 2008.

Citations: 42

➡ Identical to September 2007 article of same title; reason for reprinting unknown; reason for reordering author list also unknown.

A. Varma, B. Jacob, E. Debes, I. Kozintsev, and P. Klein. “Accurate and fast system-level power modeling: An XScale-based case study.” *ACM Transactions on Embedded Computing Systems*, vol. 6, no. 4, pp. 26:1–20. September 2007.

Citations: see April 2008 printing of article, above

B. Ganesh, **A. Jaleel**, **D. Wang**, and B. Jacob. “Fully-Buffered DIMM memory architectures: Understanding mechanisms, overheads and scaling.” In *Proc. 13th International Symposium on High Performance Computer Architecture (HPCA 2007)*, pp. 109–120. Phoenix AZ, February 2007.

Citations: 108

A. Varma, Y. Afridi, A. Akturk, P. Klein, A. Hefner, and B. Jacob. “Modeling heterogeneous SoCs with SystemC: A digital/MEMS case study.” In *Proc. International Conference on Compilers, Architectures, and Synthesis for Embedded Systems (CASES 2006)*, pp. 54–64. Seoul Korea, October 2006.

S. V. Rodriguez and B. Jacob. “Energy/power breakdown of pipelined nanometer caches (90nm/65nm/45nm/32nm).” In *Proc. International Symposium on Low Power Electronics and Design (ISLPED 2006)*, pp. 25–30. Tegernsee Germany, October 2006.

Citations: 83

A. Jaleel and B. Jacob. “In-line interrupt handling and lock-up free translation lookaside buffers (TLBs).” *IEEE Transactions on Computers*, vol. 55, no. 5, pp. 559–574. May 2006.

Citations: 12

A. Jaleel, M. Mattina, and B. Jacob. “Last-level cache (LLC) performance of data-mining workloads on a CMP — A case study of parallel bioinformatics workloads.” In *Proc. 12th IEEE International Symposium on High Performance Computer Architecture (HPCA 2006)*, pp. 88–98. Austin TX, February 2006.

Citations: 179

H. Wang, **S. V. Rodriguez**, **C. Dirik**, and B. Jacob. “Electromagnetic interference and digital circuits: An initial study of clock networks.” *Electromagnetics*, vol. 26, no. 1, pp. 73–86. January 2006. (Special Issue on RF Effects on Digital Systems)

Citations: 23

D. Wang, **B. Ganesh**, **N. Tuaycharoen**, **K. Baynes**, **A. Jaleel**, and B. Jacob. “DRAMsim: A memory-system simulator.” *SIGARCH Computer Architecture News*, vol. 33, no. 4, pp. 100–107. September 2005.

Citations: 412

J. Teller, C. Silio, and B. Jacob. “Performance characteristics of MAUI: an intelligent memory system architecture.” In *Proc. Workshop on Memory System Performance (MSP 2005)*, pp. 44–53. Chicago IL, June 2005.

K. Albayraktaroglu, **A. Jaleel**, X. Wu, M. Franklin, B Jacob, C. Tseng, and D. Yeung. “BioBench: A benchmark suite of bioinformatics applications.” In *Proc. 2005 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS 2005)*, pp. 2–9. Austin TX, March 2005.

Citations: 226

A. Jaleel and B. Jacob. “Using virtual load/store queues (VLSQs) to reduce the negative effects of reordered memory instructions.” In *Proc. 11th IEEE International Symposium on High Performance Computer Architecture (HPCA 2005)*, pp. 191–200. San Francisco CA, February 2005.

Citations: 14

➡ Aamer Jaleel won the *Best Paper Presentation* award for his talk presenting this paper (*International Symposium on High Performance Computer Architecture*, February 2005).

A. Varma, E. Debes, I. Kozintsev, and B. Jacob. “Instruction-level power dissipation in the Intel XScale embedded microprocessor.” In *Proc. SPIE’s 17th Annual Symposium on Electronic Imaging Science & Technology*, vol. 5683, pp. 1–8. San Jose CA, January 2005.

Citations: 42

H. Wang, S. V. Rodriguez, C. Dirik, A. Gole, V. Chan, and B. Jacob. “TERPS: The Embedded Reliable Processing System.” A finalist in the University LSI Design Contest. *Asia and South Pacific Design Automation Conference 2005 (ASP-DAC 2005)*. Shanghai China, January 2005.

M. Afridi, A. Hefner, D. Berning, C. Ellenwood, **A. Varma**, B. Jacob, S. Semancik. “MEMS-based embedded sensor virtual components for system-on-a-chip (SoC).” *Solid-State Electronics*, vol. 48, no. 10/11, pp. 1777–1781. October/November 2004.

Citations: 65

H. Wang, C. Dirik, S. V. Rodriguez, A. V. Gole and B. Jacob. “Radio frequency effects on the clock networks of digital circuits.” In *Proc. IEEE International Symposium on Electromagnetic Compatibility (EMC 2004)*. Santa Clara CA, August 2004.

Citations: 17

B. Iyer, S. Srinivasan, and B. Jacob. “Extended Split-Issue: Enabling flexibility in the hardware implementation of NUAL VLIW DSPs.” In *Proc. 31st Annual ACM/IEEE International Symposium on Computer Architecture (ISCA 2004)*, pp. 364–375. München Germany, June 2004.

Citations: 17

M. Afridi, A. Hefner, D. Berning, C. Ellenwood, **A. Varma**, B. Jacob, S. Semancik. “MEMS-based embedded sensor virtual components for SoC.” In *Proc. International Semiconductor Device Research Symposium (ISDRS 2003)*. Washington DC, December 2003.

K. Baynes, C. Collins, E. Fiterman, B. Ganesh, C. Smit, T. Zhang, and B. Jacob. “The performance and energy consumption of embedded real-time operating systems.” *IEEE Transactions on Computers*, vol. 52, no. 11, pp. 1454–1469. November 2003.

Citations: 105

A. Varma, B. Ganesh, M. Sen, S. R. Choudhary, L. Srinivasan, and B. Jacob. “A control-theoretic approach to dynamic voltage scheduling.” In *Proc. International Conference on Compilers, Architectures, and Synthesis for Embedded Systems (CASES 2003)*, pp. 255–266. San Jose CA, October 2003.

Citations: 145

P. Kohout, B. Ganesh, and B. Jacob. “Hardware support for real-time operating systems.” In *Proc. First IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS 2003)*, pp. 45–51. Newport Beach CA, October 2003.

Citations: 180

B. Jacob. "A case for studying DRAM issues at the system level." *IEEE Micro*, vol. 23, no. 4, pp. 44–56. July/August 2003.

Citations: 30

V. Cuppu, B. Jacob, B. Davis, and T. Mudge. "High performance DRAMs in workstation environments." *IEEE Transactions on Computers*, vol. 50, no. 11, pp. 1133–1153. November 2001. (Special Issue on High-Performance Memory Systems)

Citations: 83

K. Baynes, C. Collins, E. Fiterman, B. Ganesh, P. Kohout, C. Smit, T. Zhang, and B. Jacob. "The performance and energy consumption of three embedded real-time operating systems." In *Proc. International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2001)*, pp. 203–210. Atlanta GA, November 2001.

Citations: 54

S. Srinivasan, V. Cuppu, and B. Jacob. "Transparent data-memory organizations for digital signal processors." In *Proc. International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2001)*, pp. 44–48. Atlanta GA, November 2001.

A. Jaleel and B. Jacob. "In-line interrupt handling for software-managed TLBs." In *Proc. 19th IEEE International Conference on Computer Design (ICCD 2001)*, pp 62–67. Austin TX, September 2001.

Citations: 28

V. Cuppu and B. Jacob. "Concurrency, latency, or system overhead: Which has the largest impact on uniprocessor DRAM-system performance?" In *Proc. 28th Annual ACM/IEEE International Symposium on Computer Architecture (ISCA 2001)*, pp. 62–71. Göteborg, Sweden, June 2001.

Citations: 116

B. Jacob and T. Mudge. "Uniprocessor virtual memory without TLBs." *IEEE Transactions on Computers*, vol. 50, no. 5, pp. 482–499. May 2001.

Citations: 50

A. Jaleel and B. Jacob. "Improving the precise interrupt mechanism of software-managed TLB miss handlers." In *High Performance Computing, Lecture Notes in Computer Science*, Vol. 2228, pp. 282–293. B. Monien, V. Prasanna, and S. Vajapeyam, Editors. Springer Publishing: Berlin, Germany, 2001.

B. Davis, T. Mudge, V. Cuppu, and B. Jacob. "DDR2 and low-latency variants." In *Proc. Workshop on Solving the Memory Wall Problem*, pp. 15–29. Vancouver, British Columbia, June 2000.

Citations: 36

B. Davis, T. Mudge, and B. Jacob. "The new DRAM interfaces: SDRAM, RDRAM and variants." In *High Performance Computing, Lecture Notes in Computer Science*, Vol. 1940, pp. 26–31. M. Valero, K. Joe, M. Kitsuregawa, and H. Tanaka, Editors. Springer Publishing: Tokyo, Japan, 2000.

Citations: 20

V. Cuppu, B. Jacob, B. Davis, and T. Mudge. "A performance comparison of contemporary DRAM architectures." In *Proc. 26th Annual ACM/IEEE International Symposium on Computer Architecture (ISCA 1999)*, pp. 222–233. Atlanta GA, May 1999.

Citations: 313

B. Jacob. "Hardware/software architectures for real-time caching." In *Proc. International Conference on Compiler and Architecture Support for Embedded Systems (CASES 1999)*, pp. 135–138. Washington DC, October 1999.

D. Stewart and B. Jacob. "Hardware/software co-design of I/O interfacing hardware and real-time device drivers for embedded systems." In *Proc. International Conference on Compiler and Architecture Support for Embedded Systems (CASES 1999)*, pp. 115–119. Washington DC, October 1999.

B. Jacob. "Cache design for embedded real-time systems." In *Proc. Embedded Systems Conference–Summer*. Danvers MA, June 1999. CD-ROM proceedings.

Citations: 31

B. Jacob. "Software-managed caches: Architectural support for real-time embedded systems." In *Proc. International Conference on Compiler and Architecture Support for Embedded Systems (CASES 1998)*. Washington DC, December 1998.

B. Jacob and T. Mudge. "A look at several memory-management units, TLB-refill mechanisms, and page table organizations." In *Proc. Eighth ACM/IEEE International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 1998)*, pp. 295–306. San Jose CA, October 1998.

Citations: 162

B. Jacob and T. Mudge. "Virtual memory in contemporary microprocessors." *IEEE Micro*, vol. 18, no. 4, pp. 60–75. July/August 1998.

Citations: 155

B. Jacob and T. Mudge. "Virtual memory: Issues of implementation." *IEEE Computer*, vol. 31, no. 6, pp. 33–43. June 1998.

Citations: 122

➡ Listed as the editor's must-read pick in the "Literature Watch" column of *Microprocessor Report*, vol. 12, no. 12, p. 24. September 14, 1998.

B. Jacob and T. Mudge. "Software-managed address translation." In *Proc. Third IEEE International Symposium on High Performance Computer Architecture (HPCA 1997)*, pp. 156–167. San Antonio TX, February 1997.

Citations: 97

B. Jacob, P. Chen, S. Silverman, and T. Mudge. "A comment on 'An analytical model for designing memory hierarchies.'" *IEEE Transactions on Computers*, vol. 46, no. 10, p. 1151. October 1997.

B. Jacob. "Algorithmic composition as a model of creativity." *Organised Sound*, Cambridge University Press, vol. 1, no. 3, pp. 157–165. December 1996.

Citations: 146

B. Jacob, P. Chen, S. Silverman, and T. Mudge. "An analytical model for designing storage hierarchies." *IEEE Transactions on Computers*, vol. 45, no. 10, pp. 1180–1194. October 1996.

Citations: 119

B. Jacob. "Service Discovery: Access to local resources in a nomadic environment." *OOPSLA 1996 Workshop on Object Replication and Mobile Computing*. San Jose CA, October 1996. **Note:** The original extended abstract submitted to the conference was entitled "Support for nomadism in a global environment," which is the title given on some websites.

B. Jacob and T. Mudge. "The trading function in action." In *Proc. 1996 ACM SIGOPS European Workshop*, pp. 241–247. Connemara, Ireland, September 1996.

B. Jacob. "Composing with genetic algorithms." In *Proc. 1995 International Computer Music Conference (ICMC 1995)*, pp. 452–455. Banff Alberta, September 1995.

Citations: 204

B. Jacob. "The use of distributed objects and dynamic interfaces in a wide-area transaction environment." *SIGCOMM 1995 Workshop on Middleware*. Cambridge MA, August 1995.

D. CHAPTERS IN BOOKS

- B. Jacob. "Reward: How to Foster a Technology-Innovation Culture within a Large Organization (What You Can Learn from Startup Companies)." In The Handbook of Technology Management, pp. 1:964–1:977. H. Bidgoli, Editor. John Wiley & Sons: Hoboken NJ, 2009.
- B. Jacob. "Virtual Memory." In Computer Sciences. R. Flynn, Editor. Macmillan Reference USA: Farmington Hills MI, 2002.
- B. Jacob. "Virtual Memory Systems and TLB Structures." In The Computer Engineering Handbook, pp. 5:55–5:65. V. Oklobdzija, Editor. CRC Press: Boca Raton FL, 2002.

Citations: 14

- ➡ In February, 2003, The Computer Engineering Handbook was honored as an Outstanding Academic Title for 2002 by *Choice Magazine*, a publication of the American Library Association.

- B. Jacob, J. Szczepanski, and P. Ho. "SiO J=1–0 Emission Near the Galactic Center: Tidal Stripping and Shocks?" In Interstellar Matter. Gordon and Breach: New York NY, 1987.

E. TECHNICAL REPORTS & WHITE PAPERS

- Y. Xue, **L. Kang**, W. Jia, X. Wang, J. Kim, C. Youn, M. J. Kang, H. J. Lim, B. Jacob, J. Huang. "Building a trusted execution environment for in-storage computing." *arXiv preprint arXiv:2205.06361*. May 2022.
- P. Rosenfeld**, **E. Cooper-Balis**, T. Farrell, D. Resnick, and B. Jacob. "Peering over the memory wall: Design space and performance analysis of the Hybrid Memory Cube." Tech. Rep. UMD-SCA-2012-10-01, University of Maryland Systems & Computer Architecture Group. October 2012.

Citations: 21

- A. Rodrigues, K. Bergman, D. Bunde, **E. Cooper-Balis**, K. Ferreira, K.S. Hemmert, B. Barrett, C. Versaggi, R. Hendry, B. Jacob, H. Kim, V. Leung, M. Levenhagen, M. Rasquinha, R. Riesen, **P. Rosenfeld**, M. Varela, S. Yalamanchili. "Improvements to the Structural Simulation Toolkit." Tech. Rep. SAND2012-2185C, Sandia National Lab. March 2012.

Citations: 27

- A. Jaleel**, R. S. Cohn, C-K. Luk, and B. Jacob. "CMP\$im: A binary instrumentation approach to modeling memory behavior of workloads on CMPs." Tech. Rep. UMD-SCA-2006-01, University of Maryland Systems & Computer Architecture Group. June 2006.

Citations: 23

- B. Jacob. *Synchronous DRAM Architectures, Organizations, and Alternative Technologies*. Technical Report, University of Maryland. December 2002.

Citations: 41

- B. Jacob. *The RiSC-16 Instruction-Set Architecture*. ENEE 446: Digital Computer Design, University of Maryland. Fall 2000.

Citations: 10

- ➡ This is a description of a teaching architecture that I developed at the University of Maryland and which has been used as a teaching aid in numerous other universities around the world and realized in silicon by my own research group and several others as well.

- B. Jacob and S. Bhattacharyya. "Real-time memory management: Compile-time techniques and run-time mechanisms that enable the use of caches in real-time systems." Tech. Rep. UMIACS-TR-2000-60, University of Maryland Institute for Advanced Computer Studies (UMIACS). September 2000.

- C. Collins**, **E. Fiterman**, **T. Zhang**, and B. Jacob. "SimBed: Accurate microarchitecture-level simulation of embedded real-time operating systems." Tech. Rep. UMD-SCA-2000-1, University of Maryland Systems & Computer Architecture Group. April 2000.

V. Cuppu and B. Jacob. “Organizational design trade-offs at the DRAM, memory bus, and memory controller level: Initial results.” Tech. Rep. UMD-SCA-1999-2, University of Maryland Systems & Computer Architecture Group. November 1999.

Citations: 16

E. Berkovich, J. Nuzman, M. Franklin, B. Jacob, and U. Vishkin. “XMT-M: A scalable decentralized processor.” Tech. Rep. UMIACS-TR-99-55. University of Maryland Institute for Advanced Computer Studies (UMIACS). October 1999.

B. Jacob. “Segmented addressing solves the virtual cache synonym problem.” Tech. Rep. UMD-SCA-1997-1, University of Maryland Systems & Computer Architecture Group. December 1997.

B. Jacob. *Software-Oriented Memory-Management Design*. PhD thesis, The University of Michigan. July 1997.

B. Jacob. “Notes on calculating computer performance.” Tech. Rep. CSE-TR-231-95, University of Michigan. March 1995.

Citations: 13

B. Jacob. *Optimization of Mass Storage Hierarchies*. M.S. thesis, Tech. Rep. CSE-TR-228-95, The University of Michigan. May 1994.

F. EDITORIAL WORK

B. Jacob and S. Bhattacharyya, Editors. Embedded Systems: The Memory Resource. Special (and inaugural) issue of *ACM Transactions on Embedded Computing Systems*, vol. 1, no. 1, September 2002.

G. KEYNOTES, SPECIAL SEMINARS, AND HIGH-LEVEL BRIEFINGS

“Whither External Memory?” **Panelist Talk**, *SRC/SIA/DOE Decadal Plan for Semiconductors: New Trajectories for Memory and Storage Technologies*. Fall 2020.

One of roughly thirty members of academia, government labs, and industry invited to participate with SRC in creating a roadmap for memory devices and systems in the coming decade.

“DRAMsimL — Current Issues in Its Development.” **Invited Talk**, *DOE Brainstorming Workshop on Machine Learning for Modeling and Simulation*. Fall 2020.

“Machine Learning is the Holy Grail of Memory-System Modeling & Simulation.” **Panelist Talk**, *Modeling & Simulation of Systems and Applications (ModSim)*. Fall 2020.

“All Tomorrow’s Memories ... for Multicore.” **Keynote**, *Multicore World 2020*. Wellington NZ, Winter 2019/20.

“Faster and Accurater — The Future of Memory-System Modeling and Simulation.” **Invited Talk**, *Modeling & Simulation of Systems and Applications (ModSim)*, Seattle WA. Fall 2019.

“All Tomorrow’s Memories.” **Invited Talk**, *University of Texas at Austin Electrical & Computer Engineering Department*. Austin TX, Winter 2018/19.

“Tomorrow’s [High-Performance] Memory Systems, Rev. 2018.” **Keynote**, *Supercomputing 2018 Memory-Centric HPC Workshop (MCHPC 2018)*. Dallas TX, Fall 2018.

“Modeling Tomorrow’s Memory Systems.” **Invited Talk**, *Modeling & Simulation of Systems and Applications (ModSim)*. Seattle WA, Fall 2018.

“All Tomorrow’s Memory Systems.” **Panelist Talk**, *ICCD Special Panel — Non-Volatile Memory: will it be memory, storage, or neither of them? (ICCD 2018)*. Orlando FL, Fall 2018.

“All Tomorrow’s Memories.” **Panelist Talk**, *18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2018)*. Washington DC, Spring 2018.

“Tomorrow’s Memory Systems and What They Enable.” **Invited Talk**, *Chesapeake Large-Scale Analytics Conference*. Annapolis MD, Fall 2017.

CLSAC is an annual limited-attendance DOE conference on analytics in high-performance computing. Attendance is by invitation only, and all talks are invited (e.g., no talks are contributed; all are solicited by the program committee).

“What needs to happen to restore balance and performance efficiency as we approach exascale?” **Panelist Talk**, *International Supercomputing Conference*. Frankfurt Germany, Fall 2017.

“Modeling Tomorrow’s Memory Systems.” **Invited Talk**, *Modeling & Simulation of Systems and Applications (ModSim)*. Seattle WA, Fall 2017.

“Tomorrow’s Memory Systems, Rev. 2017.” **Chaired Professorship Talk**, *Texas A&M University Computer Science & Engineering Department*. College Station TX, Spring 2017.

“Tomorrow’s Memory Systems, Rev. 2017.” **Keynote**, *DATE Workshop on Emerging Memory Solutions*. Dresden Germany, Spring 2017.

“A Couple Decades of (accidental) Hardware/Software Co-Optimization.” Computer Engineering **Eminent Scholar Seminar Series**, *Texas A&M University Electrical & Computer Engineering Department*. College Station TX, Fall 2016.

“Highly Efficient Systems from Embedded to Exascale — All Begin with Memory.” **Keynote**, *FutureWorld.tech*. London Great Britain, Fall 2016. Conference postponed to 2017.

“Execution Models and the Memory System.” **Invited Talk**, *Computing 2025 and Beyond: DOE round table held at Argonne National Laboratory*. Lemont IL, Summer 2016.

One of roughly a dozen members of academia and industry invited to participate with DOE in creating a roadmap for computer science and applied mathematics research at DOE in the coming decades.

“Exascale Begins at the Memory System.” **Invited Talk**, *Salishan Conference on High-Speed Computing*. Glenden Beach OR, Spring 2016.

Salishan is an annual limited-attendance DOE conference on high-performance computing. Attendance is by invitation only, and all talks are invited (e.g., no talks are contributed; all are solicited by the program committee).

“Exascale Begins at the Memory System.” **Keynote**, *DATE Workshop on Emerging Memory Solutions*. Dresden Germany, Spring 2016.

“Tomorrow’s High-Bandwidth, High-Capacity, Low-Power Memory Systems.” **Invited Talk**, *High-End Computing Inter-Agency Workgroup (HEC IWG), Networking and Information Technology Research and Development (NITRD)*. Washington DC, Fall 2015.

Briefing to an inter-agency workgroup interested in high-performance computing and comprised of division directors at DARPA, DOE, NASA, and NSA.

“Tomorrow’s High-Bandwidth, High-Capacity, Low-Power Memory Systems.” **Keynote**, *Modeling & Simulation of Systems and Applications (ModSim)*. Seattle WA, Fall 2015.

“Next-Generation Memory Systems.” **Keynote**, *International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS)*. Samos Island Greece, Summer 2014.

“High-Bandwidth, High-Capacity, Low-Power Memory Systems.” **Keynote**, *ACM International Conference on Computing Frontiers*. Cagliari Italy, Spring 2014.

“Data Access, Data Movement, Data Integrity.” **Briefing to Secretary Steven Chu**. Spring 2012.

One of roughly twenty members of academia, industry, and the national laboratories who were asked to brief US Energy Secretary Steven Chu on the status of supercomputing today and what is required to achieve exascale computing performance in the near future.

“Big Memories.” **Keynote**, *Session on Large Memory Systems & Challenges at International Supercomputing Conference*. Hamburg Germany, Spring 2012.

“Memory Systems: Then, Now, and To Come.” **Keynote**, *Session on Memory Systems at International Supercomputing Conference*. Hamburg Germany, Spring 2010.

“Wealth, Innovation, Design, & Entrepreneurship.” **Invited Talk**, *Advisory Board to the University of Maryland Department of Electrical & Computer Engineering*. College Park MD, Fall 2009.

“Memory Systems.” *National Academies CSTB Briefing: Trends in Computing Performance*. Mountain View CA, Fall 2007.

One of roughly fifteen people across all areas of computing (including sub-disciplines of Storage, Power, Programming, and Applications) invited to brief a National Academies committee on the state of affairs in our various sub-disciplines. Briefers were asked to inform the committee on the technical issues and present/future trends that can hinder or aid the continued growth of computer performance. The resulting NRC report, *The Future of Computing Performance — Game Over or Next Level?*, can be found online at <http://www.nap.edu/catalog/12980/the-future-of-computing-performance-game-over-or-next-level>.

H. OTHER INVITED TALKS AND PRESENTATIONS

“Careers in (Computer) Engineering.” *University of Maryland ECE TA Training and Development Program Workshop*. College Park MD, Spring 2019.

“All Tomorrow’s Memory Systems (2018 version).” *DOE ASCR Workshop on Extreme Heterogeneity*. Gaithersburg MD. Winter/Spring 2018.

“Tomorrow’s Memory Systems and Monolithic Processors.” *National Energy Research Scientific Computing Center*. Berkeley CA, Fall 2017.

“Tomorrow’s Memory Systems, Rev. 2017.” *Barcelona Supercomputing Center*. Barcelona Spain, Spring 2017.

“A Couple Decades of (accidental) Hardware/Software Co-Optimization.” *University of Michigan Department of Electrical Engineering & Computer Science*. Ann Arbor MI, Winter 2016.

“A Bad Talk, A Good Talk, A Bad Talk, A Good Talk, Some Thoughts (or: How Not To Suck).” *University of Maryland ECEGSA Seminar Series*. College Park MD, Fall 2016.

“Some of Our Work on Hybrid Memory Cube.” Invited Talk, *Workshop on Hybrid Memory Cube*. Fort Meade MD, Fall 2015.

“Tomorrow’s High-Bandwidth, High-Capacity, Low-Power Memory Systems.” *Department of Energy*. Washington DC, Fall 2015.

“Wealth, Innovation, Design, & Economic Growth, and How it Begins with Entrepreneurship.” *University of Cyprus*. Nicosia Cyprus, Spring 2014.

“(Exascale) Data Movement.” *Foundation for Research and Technology (FORTH)*. Heraklion (Crete) Greece, Winter 2013.

“(Exascale) Data Movement.” *University of Cyprus*. Nicosia Cyprus, Winter 2013.

“Research and Development in Electrical & Computer Engineering.” Invited talk, *Career Day at Baltimore Polytechnic Institute*. Baltimore MD, Spring 2010.

“Trends in Memory Systems.” *Sun-DARPA UNIC Architecture Workshop*. Menlo Park CA, Winter 2010.

“Innovation & Design ... and Modern Engineering Entrepreneurship.” *University of Maryland Clark School Board of Visitors*. College Park MD, Fall 2009.

“High-Tech Design as Modern Engineering Entrepreneurship.” *University of Maryland ECEGSA Seminar Series*. College Park MD, Fall 2009.

“Embedded Systems, Memory Systems, and Embedded Memory Systems.” *ACACES 2009 Fifth International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems*. Terrassa Spain, Summer 2009.

“R&D in Electrical & Computer Engineering.” *National Student Leadership Conference*. College Park MD, Summer 2009.

“The Memory System and You: A Love/Hate Relationship.” Invited Talk, *SOS-13 Workshop on Distributed Supercomputing*. Hilton Head SC, Spring 2009.

“R&D in Electrical & Computer Engineering.” *National Student Leadership Conference*. College Park MD, Summer 2008.

“Engineering Entrepreneurship: Towards a High-Tech Cottage Industry.” *University of Maryland Inventis Lecture Series on Professional Concepts in Engineering*. College Park MD, Spring 2008.

“Design as Modern Engineering Entrepreneurship.” Invited lecture. *UNIV 100, The Student and the University*. College Park MD, Fall 2007.

“Embedded and Non-Classical Systems.” *Dialogue with the Dean lecture series, University of Maryland College of Engineering*. College Park MD, Spring 2006.

“A Holistic Approach to DRAM (and Systems).” *North Carolina State University Electrical & Computer Engineering Department*. Raleigh NC, Spring 2006.

“A Holistic Approach to DRAM-System Design.” *IBM T.J. Watson Research Center*. Yorktown Heights NY, Winter 2004.

“A Bad Talk, A Good Talk, A Bad Talk, A Good Talk, Some Thoughts (or: How Not To Suck).” First presented at *University of Maryland ECEGSA Seminar Series*. College Park MD, Spring 2004.

“Research Issues in Embedded Systems.” *Dialogue with the Dean lecture series, University of Maryland College of Engineering*. College Park MD, Fall 2003.

“Embedded Systems at Maryland.” *University of Maryland College of Engineering, Department of Electrical Engineering, Department of Computer Science, Institute for Advanced Computer Studies, and Institute for Systems Research Research Review Day*. College Park MD, Spring 2003.

“How To Do Research (Well).” First presented at *University of Maryland ECEGSA Seminar Series*. College Park MD, Spring 2003.

“Computers and Memory Systems.” *Dialogue with the Dean lecture series, University of Maryland College of Engineering*. College Park MD, Fall 2001.

“Career Paths in Computer Engineering, or So What Are You Going To Do With Your Life?” *University of Maryland IEEE Student-Professional Awareness Conference (S-PAC)*. College Park MD, Spring 2001.

“How Not to Configure Your DRAM System.” *Compaq Systems Research Group*. Marlborough MA, Fall 2000.

“High-Performance DRAM Systems.” *Compaq Alpha Group*. Shrewsbury MA, Fall 2000.

“Contemporary DRAM Architectures and Beyond.” *Compaq Systems Research Center*. Palo Alto CA, Fall 1999.

“Cache Design for Embedded Real-Time Systems.” *Embedded Systems Conference*. Danvers MA, Spring 1999.

“DRAM Memory Systems.” *IBM T.J. Watson Research Center*. Yorktown Heights NY, Spring 1999.

“The Top 10 Lessons I Learned as a Grad Student.” First presented at *University of Maryland Electrical & Computer Engineering Department*. College Park MD, Fall 1998.

“Software-Managed Memory Systems.” *University of Wisconsin Electrical & Computer Engineering Department*. Madison WI, Spring 1998.

“The IA-64 Architecture.” *University of Maryland IEEE Chapter*. College Park MD, Fall 1997.

I. CONFERENCE-PAPER PRESENTATIONS

“Thoughts on Merging the File System and Virtual Memory System.” Presented at *The 9th ACM/IEEE International Symposium on Memory Systems (MEMSYS 2023)*. Washington DC, October 2023.

“Exascale Interconnect Topology Characterization and Parameter Exploration.” Presented at *The 20th IEEE International Conference on High Performance Computing and Communications (HPCC 2018)*. Exeter United Kingdom, June 2018.

“Energy/Power Breakdown of Pipelined Nanometer Caches (90nm/65nm/45nm/32nm).” Presented at *The International Symposium on Low Power Electronics and Design (ISLPED 2006)*. Tegernsee Germany, October 2006.

- “Concurrency, Latency, or System Overhead: Which Has the Largest Impact on Uniprocessor DRAM-System Performance?” Presented at *The 28th International Symposium on Computer Architecture (ISCA 2001)*. Göteborg Sweden, June 2001.
- “Hardware/Software Architectures for Real-Time Caching.” Presented at *The Second International Workshop on Compiler and Architecture Support for Embedded Systems (CASES 1999)*. Washington DC, December 1999.
- “A Performance Comparison of Contemporary DRAM Architectures.” Presented at *The 26th Annual ACM/IEEE International Symposium on Computer Architecture (ISCA 1999)*. Atlanta GA, May 1999.
- “Software-Managed Caches: Architectural Support for Real-Time Embedded Systems.” Presented at *The First International Workshop on Compiler and Architecture Support for Embedded Systems (CASES 1998)*. Washington DC, December 1998.
- “A Look at Several Memory Management Units, TLB-Refill Mechanisms, and Page Table Organizations.” Presented at *The Eighth ACM/IEEE International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 1998)*. San Jose CA, October 1998.
- “Software-Managed Address Translation.” Presented at *The Third IEEE International Symposium on High Performance Computer Architecture (HPCA 1997)*. San Antonio TX, February 1997.
- “Method Invocation in the Face of Anarchy.” Presented at *The 1996 OOPSLA Workshop on Object Replication and Mobile Computing (ORMC 1996)*. San Jose CA, October 1996.
- “Composing with Genetic Algorithms.” Presented at *The 1995 International Computer Music Conference (ICMC 1995)*. Banff Alberta, September 1995.

J. PATENTS ISSUED

Integrating a Resistive Memory System into a Multicore CPU Die to Achieve Massive Memory Parallelism. US Patent 11,126,550; issued September 21, 2021; filed September 4, 2018. Donald Yeung, Bruce L. Jacob, Mehdi Asnaashari, and Sylvain Dubois, Crossbar Inc. and University of Maryland at College Park.

Programmable Switch for Configuring Circuit Topologies. US Patent 8,445,770, issued May 21, 2013; filed September 21, 2008; claiming priority to provisional patent filed June 14, 2008. Bruce Jacob.

Citations: 20

Volume-Adjustment Circuit for Equilibrating Pickup Settings. US Patent 8,324,495, issued December 4, 2012; filed February 13, 2009. Bruce Jacob.

Citations: 11

Electronic Guitar Harness Component Connector. US Patent 7,915,506, issued March 29, 2011; filed August 10, 2010. Bruce Jacob.

System and Method for Performing Multi-Rank Command Scheduling in DDR SDRAM Memory Systems. US Patent 7,543,102, issued June 2, 2009; filed April 17, 2006; claiming priority to provisional patent filed April 18, 2005. Bruce Jacob and Dave Wang, University of Maryland.

Citations: 89

K. COMPUTATIONAL ARTIFACTS

DRAMsimL, developed 2019–present (ongoing)

A machine-learning-based memory-system simulation framework that can be used in place of traditional cycle-accurate simulation, for early experimental results. The framework is described in the following publications:

- S. Li, R.S. Verdejo, P. Radojković, and B. Jacob. “Rethinking cycle accurate DRAM simulation.” In Proc. ACM/IEEE International Symposium on Memory Systems (MEMSYS 2019), pp. 184–194. Washington DC, October 2019.

S. Li and B. Jacob. “Statistical DRAM modeling.” In Proc. ACM/IEEE International Symposium on Memory Systems (MEMSYS 2019), pp. 521–530. Washington DC, October 2019.

The framework relies upon the results of cycle-accurate simulators, which are used to create the statistical models.

NVsim, developed 2009–present (ongoing)

A framework for simulating non-volatile memories of all types (including NAND flash, NOR flash, PCM, and SST-RAM). It has been developed by my graduate students and is described in the following publication:

J. Stevens, P. Tschirhart, M.-T. Chang, I. Bhati, P. Enns, J. Greensky, Z. Chishti, S.-L. Lu, and B. Jacob. “An integrated simulation infrastructure for the entire memory hierarchy: Cache, DRAM, nonvolatile memory, and disk.” *Intel Technology Journal*, vol. 17, no. 1. April 2013.

The framework accurately models configurations of non-volatile memories in SSD architectures as well as in standalone configurations such as hybrid main memories.

DRAMsim, developed 1999–present (ongoing)

A memory-system simulation framework that interacts with popular CPU simulators such as MARSSx86, SST, GEMS, MASE, and SimAlpha. It has been developed by my graduate students and represents many student-decades of development effort. The framework is described in the following publications:

D. Wang, B. Ganesh, N. Tuaycharoen, K. Baynes, A. Jaleel, and B. Jacob. “DRAMsim: A memory-system simulator.” *SIGARCH Computer Architecture News*, vol. 33, no. 4, pp. 100–107. September 2005.

P. Rosenfeld, E. Cooper-Balis, and B. Jacob. “DRAMsim2: A cycle-accurate memory-system simulator.” *Computer Architecture Letters*, vol. 10, no. 1, pp. 16–19. January 2011.

S. Li, Z. Yang, D. Reddy, A. Srivastava, B. Jacob. “DRAMsim3: A cycle-accurate, thermal-capable DRAM simulator.” *IEEE Computer Architecture Letters*, published on-line February 2020.

The software has been released to the community under the GNU Public License and can be found at github.com/dramninja/UMD/DRAMSim2.

BioBench/BioParallel, developed 2004–2005

A collection of benchmarks, each a different application area within the larger domain of bioinformatics and data mining. BioBench, a suite of workloads collected by students and faculty at the University of Maryland, was released in 2005. The suite expanded in 2006 with the addition of parallel workloads collected by researchers at Intel (distinguished from the original *BioBench* suite as the *BioParallel* suite). Described in the following publications:

K. Albayraktaroglu, A. Jaleel, X. Wu, M. Franklin, B. Jacob, C. Tseng, and D. Yeung. “BioBench: A benchmark suite of bioinformatics applications.” In *Proc. ISPASS 2005*, pp. 2–9. Austin TX, March 2005.

A. Jaleel, M. Mattina, and B. Jacob. “Last-level cache (LLC) performance of data-mining workloads on a CMP — A case study of parallel bioinformatics workloads.” In *Proc. HPCA 2006*, pp. 88–98. Austin TX, February 2006.

The benchmark suites have been released to the community and can be found at www.ece.umd.edu/biobench.

Variations, developed 1994–1996

A music composition system that uses genetic algorithms to generate and organize complex tonal and microtonal music for orchestral organizations. The software is described in the following publications:

B. Jacob. “Composing with genetic algorithms.” In *Proc. 1995 International Computer Music Conference*, pp. 452–455. Banff Alberta, September 1995.

B. Jacob. “Algorithmic composition as a model of creativity.” *Organised Sound*, vol. 1, no. 3, pp. 157–165. December 1996.

High-fidelity audio examples of the system’s output (in MP3 format) as well as the software itself can be found at www.ece.umd.edu/~blj/algorithmic_composition.

Telecom Services Architecture, developed 1991–1996

A distributed system of active network objects whose service methods are invoked transparently via a service broker. The software is described in the following publications:

B. Jacob. “The use of distributed objects and dynamic interfaces in a wide-area transaction environment.” *SIGCOMM 1995 Workshop on Middleware*. Cambridge MA, August 1995.

B. Jacob and T. Mudge. “The trading function in action.” In *Proc. 1996 ACM SIGOPS European Workshop*, pp. 241–247. Connemara Ireland, September 1996.

This architecture was created for the telecommunications startup Priority Call Management. The software forms the core technology in Priority Call’s product line, and the architecture was still in use as late as 1999.

L. HONORS AND AWARDS

Named Fellow, Institute of Electrical and Electronics Engineers (2020, effective 2021)

The IEEE Board of Directors, at its November 2020 meeting, elevated to IEEE Fellow, effective 1 January 2021, with the following citation:

for contributions to computer memory design and analysis

Invited to serve on ACM/IEEE CS George Michael Memorial HPC Fellowship Committee (2018–2020)

Served on the international Fellowship Committee for three years. The Fellowship “honors exceptional PhD students throughout the world whose research focus is on high-performance computing applications, networking, storage, or large-scale data analysis using the most powerful computers that are currently available.

The award committee is selected by the two societies and includes past winners as well as leaders in the field.”

University of Maryland “Research Leader” (2017)

Recognized by the University of Maryland’s Division of Research as “having made significant contributions to the University of Maryland College Park research program.”

University of Maryland “Research Leader” (2016)

Recognized by the University of Maryland’s Division of Research as “having made significant contributions to the University of Maryland College Park research program.”

University of Maryland “Research Leader” (2012)

Recognized by the University of Maryland’s Division of Research as “having made significant contributions to the University of Maryland College Park research program.”

University of Maryland “Research Leader” (2010)

Recognized by the University of Maryland’s Division of Research as “having made significant contributions to the University of Maryland College Park research program.”

Philip Merrill Presidential Scholar Faculty Mentor (2009)

Recognized as the faculty mentor who has made the most impact on the academic achievement of one of the University of Maryland’s Philip Merrill Presidential Scholars.

Clark School of Engineering Keystone Professor (2006)

One of six professors named to Keystone: The Clark School Academy of Distinguished Professors.

University of Maryland “Research Leader” (2006)

Recognized by the Office of Research and Graduate Studies for being one of the top 100 researchers in the University of Maryland, as measured by research funding. Was previously called the University’s “Rainmaker” award.

University of Maryland “Rainmaker” (2005)

Recognized by the Office of Research and Graduate Studies for being one of the top 100 researchers in the University of Maryland, as measured by research funding.

Best Paper Presentation (HPCA 2005)

My student, Aamer Jaleel, won the Award for Best Paper Presentation for our paper in the *11th Int’l Symposium on High Performance Computer Architecture (HPCA 2005)*, February 2005.

Award for Teaching Excellence (2004)

Recognized by the University of Maryland (campus-wide) for excellence in teaching at the undergraduate level.

University of Maryland “Rainmaker” (2001)

Recognized by the Office of Research and Graduate Studies for being one of the top 100 researchers in the University of Maryland, as measured by research funding.

National Science Foundation CAREER Award (1999)

Project Title: “Architecture Issues in DRAM Devices and Systems.”

George Corcoran Memorial Award (1998–9 Academic Year)

“In recognition of teaching and educational leadership at the University of Maryland, College Park campus, effective contribution at the national level, and creative and other scholarly activities related to Electrical and Computer Engineering education.”

M. PROFESSIONAL MEMBERSHIPS

1993– Member, ACM
 1994– Member, IEEE
 1993–1998 Member, USENIX
 1999– USENIX Campus Liaison at University of Maryland

III. TEACHING AND ADVISING**A. COURSE INSTRUCTION**

Teaching evaluation scores are often in the top half dozen scores in the department (out of roughly 70 faculty); they have been the highest-ranking score in the department on numerous occasions.

Academic Year	Course Number	Course Title	No. of Students	Teaching Eval. Total	Average for Similar
1997/1998	ENEE 350H	Computer Organization (Honors)	12	3.81	3.36
	ENEE 759M	Advanced Topics in Microarchitecture	15	3.71	3.37
1998/1999	ENEE 350H	Computer Organization (Honors)	15	4.92 (of 5)	4.20 (of 5)
	ENEE 698b	Computer Engineering Seminar	17		
	ENEE 647	Design of Distributed Systems	50	3.58	3.25
1999/2000	ENEE 350	Computer Organization	45	3.73	3.33
	ENEE 759M	Advanced Topics in Microarchitecture	18	3.87	3.28
2000/2001	ENEE 446	Design of Digital Computers (undergrad)	27	3.74	3.25
	ENEE 350	Computer Organization	20	3.62	3.34
2001/2002	ENEE 646	Design of Digital Computers (grad)	70	3.50	3.14
	ENEE 759M	Advanced Topics in Microarchitecture	45	3.60	3.20
2002/2003	ENEE 646	Design of Digital Computers (grad)	66	3.56	3.24
	ENES 100	Introduction to Engineering Design	36	3.17	2.99
	ENEE 759H	High-Speed Memory Systems	17	3.98	3.23

Academic Year	Course Number	Course Title	No. of Students	Teaching Eval. Total	Average for Similar
2003/2004		On sabbatical			
2004/2005	ENEE 302H	Digital Electronics (Honors)	24	3.57	3.09
	ENEE 759H	High-Speed Memory Systems	12		
	HONR 218V	Digital Sound and Fury on Mac OSX: Desktop Publishing, Music, Photography, and Video (University Honors)	24		
2005/2006	ENEE 302H	Digital Electronics (Honors)	11	3.52	3.29
	ENEE 446	Design of Digital Computers (undergrad)	23	3.40	3.34
2006/2007	ENEE 646	Design of Digital Computers (grad)	27	3.70	3.30
	ENES 100	Introduction to Engineering Design	36		
	ENEE 359A	Digital VLSI Circuits	15	3.50	3.16
	ENES 100	Introduction to Engineering Design	36		
2007/2008	ENEE 759H	High-Speed Memory Systems	12	3.43	3.11
	ENES 100	Introduction to Engineering Design	40	3.50	3.24
	ENEE 359A	Digital VLSI Circuits	12	3.82	3.24
	ENEE 159B	Start-Up 101: Electric Guitar Design	12	3.65	3.39
	ENES 100	Introduction to Engineering Design	40	3.40	3.32
2008/2009	ENEE 646	Design of Digital Computers (grad)	21	3.63	3.31
	ENES 100	Introduction to Engineering Design	41	3.57	3.10
	ENEE 159B	Start-Up 101: Electric Guitar Design	12	3.71	3.29
	ENEE 459J	Start-Up 101: Consumer Electronics	12	3.70	3.34
	ENES 100	Introduction to Engineering Design	40	3.67	3.25
2009/2010	ENEE 350H	Computer Organization (Honors)	25	3.79	3.24
	ENES 100	Introduction to Engineering Design	40	3.55	3.02
	ENEE 359R	Reverse Engineering (Supervisor)	11	3.90	3.24
	ENEE 159B	Start-Up 101: Electric Guitar Design	12	3.71	3.22
	ENES 100	Introduction to Engineering Design	40	3.68	3.21
2010/2011		On sabbatical			
2011/2012	ENEE 646	Design of Digital Computers (grad)	25	3.47	3.29
	ENEE 446	Design of Digital Computers (undergrad)	30		
2012/2013	Visiting at University of Siena (Italy), working on EC-funded TERAFLUX project				
	PHD-SEM	Advanced Memory Systems (25-hour grad seminar at University of Siena)	7	n/a	n/a

Academic Year	Course Number	Course Title	No. of Students	Teaching Eval. Total	Average for Similar
2013/2014	Visiting at University of Siena, FORTH (Crete, Greece), and University of Cyprus, working on EU-funded TERAFLUX and EuroServer projects				
	CS-590.26	High-Speed Memory Systems (30-hour grad seminar at University of Crete)	15	n/a	n/a
2014/2015	ENEE 646	Design of Digital Computers (grad)	42	3.63	3.27
	ENEE 245	Digital Circuits & Systems Lab	45	3.11	3.18
	ENEE 245	Digital Circuits & Systems Lab	36	3.59	3.10
2015/2016	ENEE 447	Operating Systems	34	3.33	3.16
2016/2017	ENEE 245	Digital Circuits & Systems Lab	112	3.37	3.18
2017/2018	ENEE 646	Design of Digital Computers (grad)	16	3.38	3.30
	ENEE 408L	Capstone Design: Electric Guitars	12	3.75	3.24
	ENEE 446	Design of Digital Computers (undergrad)	46	3.62	3.18
2018/2019	ENEE 408L	Capstone Design: Electric Guitars	13	3.26	3.23
	ENEE 447	Operating Systems	60	3.74	3.24
2019/2020	On sabbatical				
2020/2021	ENEE 646	Design of Digital Computers (grad)	16		
	ENEE 447	Operating Systems			
2021/2022	CompSci 2271	Introduction to Computation and Programming	96		
	CompEng 4930	The Design of Modern Computer Systems			
2022/2023	SY110	Introduction to Cyber Security	60		
	SY204	Cyber Systems Programming & OS Fundamentals	20		

B. UNDERGRADUATE INDEPENDENT STUDY

- ENEE499 *On-Board Low-Power Audio Electronics*, with Rob Brisentine, Noel Cervino, and Matt Wienke. Development of low-power signal-processing electronics to mix synthetic and real-time audio, all run on 9V batteries.
Spring 2010
- ENEE499 *A Comparison of Preamps: Transistor-based versus OpAmp-based*, with Ilia Kobrinsky. A comparison of the fidelity and power requirements of two different audio preamp designs.
Spring 2010
- ENEE 499L *High-Resolution Low-Power A/D Conversion on an Electric Guitar*, with Spencer Black. Development of analog-to-USB and analog-to-firewire DSP units to go on-board an electric guitar.
Summer 2008
- ENEE 499 *Analog Circuit Design for Electric Guitars*, with Justin Ahmanson and Timothy Babich. Design and development of a novel preamp and variable EQ circuit for electric guitar, to be used on-board the guitar and powered by 9V battery.
Fall 2007
- ENEE 499 *Applications Engineering*, with Garrett Lang, Han Dao, and Jason Borrero. Design and development of web-based application software. One step in the development of a new course to be called *Start-Up 101: Applications Engineering*.
Fall 2007
- ENEE 499 *Verilog Modeling of Real-Time Embedded Systems*, with Vincent Chan, Lei Zong, and David Xia. Development of a Verilog ARM core with support for precise interrupts and programmable I/O controllers.
Fall 2001
- ENEE499 *Development of an Embedded-System Simulation Environment*, with Eric Fiterman. Design and development of an embedded-system simulation environment that emulates the embedded CPU and peripheral IO devices as well as a real-time operating system.
Spring 2000
- ENEE499 *A Pipelined Microprocessor on an FPGA*, with Andrew Ripple. Implementation of a high-performance pipelined processor on programmable hardware, including timing specifications and performance measurements.
Fall 2000
- ENEE499 *Full Computer-System Emulation*, with Aamer Jaleel, Jiwanjot Tulsi, and Yun Hua Wang. Design and development of a simulation environment that emulates both a full workstation (including CPU and peripheral IO devices) and a sophisticated operating system (including support for virtual memory, a file system, and multitasking).
Spring 1999
- ENEE 499 *Implementing the Ridiculously Simple Computer*, with Julian Requejo. Design and development of a small digital computer from IC components on circuit bread-boards. The computer will be able to run simple applications stored in memory.
Spring 1999

C. EDUCATIONAL DEVELOPMENT AT THE UNIVERSITY OF MARYLAND

New courses introduced into the regular curriculum

Developed and introduced *ENEE 759M: Advanced Topics in Microarchitecture*

Developed and introduced *ENEE 759H: High-Speed Memory Systems*

Developed and introduced *ENEE 447: Operating Systems*

Developed and introduced *ENEE 408L: Electric Guitar Design Capstone*

Developed and introduced *ENEE 359A: Digital VLSI Circuits*

Developed and introduced *ENEE 159B: Electric Guitar Design (Start-Up 101)*

Course development — ENEE 408L: Electric Guitar Capstone Design

At the Department's request, developed a Capstone Design course on electric guitars. The course uses a popular topic (electric guitars) to motivate students to learn about and do research in the fundamentals of electrical & computer engineering, and express their understanding through their design project.

Modern electric guitars contain both EE and CE components; the fundamentals are EE-based, but an increasing number of current designs incorporate CE-based concepts such as active digital control circuits, digital signal

processors, and many embedded digital effects units. The course builds upon previous disciplines and development skills students have learned, such as circuit design & test and working with analog and digital circuit components, as well as the interaction between magnetic fields and electric currents. It also introduces skills such as PCB design, fabrication, and assembly, for those students who have not already learned them previously.

The students first start with an assigned experiment: using the guitars, amplifiers, and guitar parts provided to them, they develop and validate an electrical-engineering model for how the guitar's pickup works. Many models exist, but they are often simplistic and overlook critical elements, such as the fact that the bridge must be grounded; the strings are an integral part of the circuit; and the pickup's inductors are not actually necessary: one could (and the students do) put alligator clips on the ends of the strings and send the signal to an amplifier, produced only by the string's vibration over a fixed magnet. The second project is for the students to spend a class period experimenting with different "tone" circuits that have appeared in guitars, to see how the various filtering functions interact with the pickups' inductors and their resonance properties. After coming to a better understanding of the components involved and how they interact, the students then propose their design and spend the rest of the semester implementing and characterizing it.

Course development — ENEE 447: Operating Systems

Asked by department chair to develop and teach an Operating Systems course for the Computer Engineering degree program. Developed a treatment of Operating Systems specifically targeted to Computer Engineers, presenting a deep look at the hardware and how the OS interacts with it. In this course, students build a rudimentary but fully complete operating system from scratch on actual hardware, because, in most Engineering courses, the best way for one to learn how something works is to build it for oneself. The material is covered much in the way that CE and EE courses cover their material: as a sequence of problems, each followed by a range of solutions, proofs of their correctness (in this case, primarily intuitive rather than formal or mathematical proofs), and the advantages and disadvantages of each solution. Topics given special emphasis include multicore processors and nonvolatile memories, both of which call for a shake-up in operating systems design.

In particular, multicore is here to stay, and therefore scheduling, coherence, control, and basic operating-system design all need to be reevaluated. Non-volatile main memory will be here tomorrow, which means that separate virtual memory systems and file systems will become a thing of the past. Issues such as security, reliability, and power/energy efficiency are extremely important today but were less important (in some cases not at all important) when today's operating systems were designed; these warrant reevaluation of existing design choices. Thus, it made sense to address these issues specifically in the course.

The resulting course can serve as a template for how Operating Systems is taught to Computer Engineering students everywhere, as Computer Engineering students have needs and learning styles that are very distinct and different from those of Computer Science students.

Formal course revision & improvement

As Director of Computer Engineering, restructured and reoriented the program's curriculum to address its age and concerns by students, faculty, and graduates that the curriculum needed better structure, cohesiveness, and topical focus. In 2007/8 brought together faculty, current students, and recent graduates to work on the program over numerous focus groups and meetings. The new proposed curriculum strengthened the program's focus on hardware, circuits, and electronics and was informally adopted by the Computer Engineering group via vote in the Fall of 2008. The computer group is working on a formal transition and implementation plan.

Member of a Dean's committee appointed to review and revise *ENES 100 — Introduction to Engineering Design*, charged with the task of making the class more attractive to freshmen and more relevant in general. The committee's recommendations resulted in, among other things, the formation by the Dean of a cadre of faculty called *Keystone: The Clark School Academy of Distinguished Professors*, the members chosen by the Dean for their proven teaching skills and ability to reach students, and then charged with realizing the committee's vision. The course is now taught by Keystone professors and has undergone a significant overhaul for the better.

New courses offered on a trial basis

Developed and introduced several pilot courses that are aimed at educating students on the technical aspects of building real-world artifacts. Due to the confluence of several economic factors, now is an ideal time for students and recent graduates to start up medium- and high-tech companies on their own. However, most entrepreneurship classes, workshops, and seminars focus on the business aspects of startups (e.g., funding, attracting venture capitalists, writing business plans, etc.), which ignores the technical/design aspects of entrepreneurship. This group of classes is an attempt to bring into the classroom real-world design & development skills — skills that, unlike manufacturing and construction, are *not* being out-sourced.

Start-Up 101: Electric Guitar Design (ENEE 159B). Design and development of circuits in electric guitars, such as analog switches, volume, and EQ; basic skills and knowledge such as wiring, soldering, electromagnetism, the physics of sound; printed circuit-board design, construction, and assembly. Offered Spring 2008. Note: at the Department's request, this became *ENEE 408L: Capstone Design*.

Start-Up 101: Consumer Electronics (ENEE 459J). Design and development of hardware/software systems, focusing on issues in real-time operating systems, firmware development, circuit-board design and construction, and correctness of design. Offered Spring 2009.

Start-Up 101: Applications Engineering (as ENEE 499). Design and development of large, industrial-strength software applications, 10–100x the size of those typically developed in the classroom. Offered Fall 2007.

Informal course revision & improvement — ENEE 350

Revised ENEE 350 (computer organization) to become a project-oriented course, which is very different from the way it had been taught previously, both in content and in instructional delivery. This approach was successful in both the honors-level course and the regular version of ENEE 350. The approach was introduced to the regular section in the Fall of 1999, and now the projects are used by other faculty in the ECE department and CS department as well.

Informal course revision & improvement — ENEE 446/646

ENEE 446 (computer architecture) focuses on the implementation of microprocessors and systems; the class projects were revised to reflect how industrial-strength designs work and how commercial hardware development is done. To begin with, students implement a microprocessor (the “RiSC-16,” designed by me) in the Verilog hardware description language; Verilog is a language used in industry that represents hardware in its VLSI implementation, as opposed to a C-language emulation as is used in 350 and was previously used in 446. To this Verilog implementation students add the necessary hardware constructs to allow an operating system to run on the microprocessor: namely, facilities for precise interrupts and virtual memory. Lastly, the students write the rudiments of an operating system, in assembly code, to run on their Verilog microprocessor.

The projects are used by other faculty in the ECE department. Even more interesting is the response from outside the university. For instance, the RiSC-16 architecture has been used in universities all around the world, such as the University of Illinois at Urbana-Champaign, Université Libre de Bruxelles (Belgium), and Universidad Pontificia Comillas (Spain); the various design documents posted on the class website show up in on-line databases as cited works; numerous developers have contacted me over the years to help them build VLSI implementations of the processor design (google “RiSC-16 maryland” to see examples); and the implementation of precise interrupts, an extremely important mechanism that had previously been neglected in *all* computer-systems textbooks — despite the fact that it is nearly impossible to build a working microprocessor without precise interrupts — began to show up in undergraduate texts several years after the RiSC-16 design documents were first posted.

D. ADVISING AT THE UNIVERSITY OF MARYLAND

Graduated Ph.D. Students

Jennie Hill, Ph.D. 2023. *Evaluation of Selected Side-Channel Analysis Methods for Ransomware Classification and Detection*. (Unites States Naval Academy)

Luyi Kang, Ph.D. 2022. *Tackling Performance and Security Issues for Cloud Storage Systems*. (Intel)

Meenatchi Jagasivamani, Ph.D. 2020. *Resistive RAM Based Main-Memory Systems: Understanding the Opportunities, Limitations, and Tradeoffs*. (Northrop Grumman)

Shang Li, Ph.D. 2019. *Scalable and Accurate Memory System Simulation*. (Cadence Design Systems)

Paul Tschirhart, Ph.D. 2015. *Bringing Modern Multi-Level Main Memory Systems Into Focus*. (Northrop Grumman)

Ishwar Bhati, Ph.D. 2014. *Scalable and Energy-Efficient DRAM Refresh Techniques*. (Oracle)

Paul Rosenfeld, Ph.D. 2014. *Performance Exploration of the Hybrid Memory Cube*. (Micron)

Mu-Tien Chang, Ph.D. 2013. *Technology Choices for Large Last-Level Caches*. (Samsung)

Elliott Cooper-Balis, Ph.D. 2012. *Buffer-on-Board Memory Systems*. (Micron)

Cagdas Dirik, Ph.D. 2009. *Performance and Analysis of NAND Flash Memory Solid-State Disks*. (SanDisk)

Sadagopan Srinivasan, Ph.D. 2007. *Prefetching vs. the Memory System: Optimizations for Multi-core Server Platforms*. (Intel)

Brinda Ganesh, Ph.D. 2007. *Understanding and Optimizing High-Speed Serial Memory-System Protocols*. (Intel)

Ankush Varma, Ph.D. 2007. *High-Speed Performance, Power, and Thermal Co-Simulation for SoC Design*. (Intel)

Nuengwong (Ohm) Tuaycharoen, Ph.D. 2006. *Disk Design-Space Exploration in Terms of System-Level Performance, Power, and Energy Consumption*. (Dhurakijpundit University, Thailand)

Samuel Rodriguez, Ph.D. 2006. *myCACTI: A New Cache-Design Tool for Pipelined Nanometer Caches*. (AMD)

Aamer Jaleel, Ph.D. 2005. *The Effects of Out-of-Order Execution on the Memory System*. (Intel)

David Tawei Wang, Ph.D. 2005. *Modern DRAM Memory Systems: Performance Analysis and a High Performance, Power-Constrained DRAM-Scheduling Algorithm*. (MetaRAM, a Silicon Valley start-up)

Brian Davis, Ph.D. 2000. *Modern DRAM Architectures*. Davis was a doctoral student at the University of Michigan; I was his thesis co-chair with Trevor Mudge. (Michigan Technological University)

Graduated M.S. Students

(only MS with theses listed)

Dhiraj Reddy Nallapa Yoge, M.S. 2018. *Performance Study of Various Modern DRAM Architectures*.

Jeffrey Scott Smith, M.S. 2006. *Distributed Two-Dimensional Fourier Transforms on DSPs with Applications for Phase Retrieval*. (NASA)

Rami Nasr, M.S. 2005. *FBsim and the Fully Buffered DIMM Memory System Architecture*. (Lutron)

Amol Gole, M.S. 2003. *TERPS: The Embedded Reliable Processing System*. (U.S. Patent Office)

Bharath Iyer, M.S. 2003. *Extended Split-Issue Mechanism in VLIW DSPs to Support SMT and Hardware-ISA Decoupling*. (AMD)

Nuengwong Tuaycharoen, M.S. 2003. *RTOS-Based Dynamic Voltage Scaling*. (continued as U. Maryland Ph.D. student)

Lei Zong, M.S. 2003. *Nanoprocessors: Configurable Hardware Accelerators for Embedded Systems*. (Army Research Lab)

Brinda Ganesh, M.S. 2002. *Architectural Support for Embedded Operating Systems*. (continued as U. Maryland Ph.D. student)

Aamer Jaleel, M.S. 2002. *In-Line Interrupt Handling and Lockup-Free TLBs*. (continued as U. Maryland Ph.D. student)

Paul Kohout, M.S. 2002. *Hardware Support for Real-Time Operating Systems*. (EVI Technology)

Tiebing Zhang, M.S. 2001. *RTOS Performance and Energy Consumption Analysis Based on an Embedded System Testbed*. (Aeptec Microsystems)

Christopher Collins, M.S. 2000. *Emulation of Embedded Microcontrollers and Real-Time Operating Systems*. (Intel)

Undergraduate Research

Emily Ruppel (2015–2016). *Modeling High-Performance Memory Systems*.

Spencer Black (as part of ENEE 499L, Summer 2008). *High-Resolution Low-Power A/D Conversion on an Electric Guitar*.

Paul Rosenfeld (Summer 2007). *Modeling High-Performance DRAM Systems*.

James Shen (Summer 2007). *Embedded Wireless Networks*.

Brian Davis (as part of the MERIT/ICE program, Summer 2002). *Direct-DRAM-to-DRAM-Access to Support High-Level Data Movement*.

Kevin Ghozati (Summer 2002). *Nanoprocessor Modeling for Real-Time Embedded Systems*.

Sandy Klemm (as part of the MERIT/ICE program, Summer 2002). *Applying Control Theory to Dynamic Voltage Scaling*.

Jeremy Monaldo (Spring/Summer 2002). *Nanoprocessor Modeling for Real-Time Embedded Systems*.

Vincent Chan (as part of ENEE 499, Fall 2001). *Verilog Modeling of Real-Time Embedded Systems*.

David Xia (as part of ENEE 499, Fall 2001). *Verilog Modeling of Real-Time Embedded Systems*.

Lei Zong (as part of ENEE 499, Fall 2001). *Verilog Modeling of Real-Time Embedded Systems*.

Daniel O'Brien (as part of the MERIT/ICE program, Summer 2001). *High-Performance Benchmarks for Real-Time Embedded Systems*.

Lei Zong (as part of the MERIT/ICE program, Summer 2001). *High-Performance Benchmarks for Real-Time Embedded Systems*.

Kathleen Baynes (as part of the MERIT/ICE program, Summer 2000). *Power Modeling of Embedded Systems*.

Christine Smit (as part of the MERIT/ICE program, Summer 2000). *Power Modeling of Embedded Systems*.

Aamer Jaleel (1999–2000). *A Performance Study of Several Virtual Memory Mechanisms*.

Eric Fiterman (Summer 1999). *Microarchitecture-Level Emulation of Real-Time Embedded Systems*.

Ken Powers (Summer 1999). *A Study of DRAM Architectures and Organizations*.

Jiwanjot Tulsi (as part of ENEE 499, Spring 1999). *Design of a Provably Correct Operating System*.

Yun Hua Wang (as part of ENEE 499, Spring 1999). *Design of a Provably Correct Operating System*.

Aamer Jaleel (as part of ENEE 499, Spring 1999). *Design of a Provably Correct Operating System*.

Benjamin Garrett (Fall 1998). *A Speculative Architecture for Handling Precise Interrupts*.

Benjamin Garrett (as part of ENEE 759M, Spring 1998). *Miss Handling in Software-Managed TLBs*.

Joe Nuzman (as part of ENEE 759M, Spring 1998). *A Simulation of Explicit Multi-Threading*.

High-School Student Research

Rachel Fallon (research intern from River Hill High School, Summer 2016). *Network Design for Supercomputer Memory Systems*.

Kevin Hu (high-school research internship for the Science, Mathematics, and Computer Science program at Poolsville High School, Summer 2016). *Processor Design for High-Performance Memory Systems*.

Demetri Wolfe (high-school research internship for the School Without Walls Senior High School in Washington, DC, Spring 2016). *Building an Electric Guitar out of Cardboard*.

Andrea Ng (high-school research internship for the Magnet Program at Montgomery Blair High School, Summer 2009). *Hardware Validation of Simulators*.

Colin Schmidt (high-school research internship for the Magnet Program at Montgomery Blair High School, Summer 2008). *Validating a DRAM-System Simulator*.

Ben Li (high-school research internship for the Magnet Program at Montgomery Blair High School, Summer 2006). *Performance Modeling of Advanced DRAM Systems*.

E. ADVISING (NON-RESEARCH)

Undergraduate Cyber Mentoring — United States Naval Academy. The Cyber Science department has a formal advising/mentoring program in which students are assigned to faculty. Currently advise roughly ten Cyber Science undergraduates.

Inventis Mentoring — University of Maryland. The Inventis mentoring program spans the Engineering College; the mentors are selected by the program organizers. Advised roughly a dozen students in this program.

Undergraduate ECE Mentoring — University of Maryland. The ECE department has a formal advising/mentoring program in which students are assigned to faculty; students cannot register for classes without meeting with their faculty advisor. Advised roughly fifteen ECE undergraduates at a time.

F. EDUCATION-RELATED HONORS AND AWARDS

Asked to serve as the Moderator of the 9th Annual ECEGSA Academic Careers Panel, an all-day event held November 4, 2016. The discussion topic was “Careers in Academia — Is Academia for Me? and How Do I Get An Academic Job?” Featured Guests included Professor Nicole McFarlane from the University of Tennessee, and Professor Aswin Sankaranarayanan from Carnegie Mellon.

Invited to speak at the *National Student Leadership Conference*, a week-long seminar for high-school students who have shown academic talent. Summer 2008, and again Summer 2009.

Named *Clark School of Engineering Keystone Professor*. One of six professors named to Keystone: The Clark School Academy of Distinguished Professors, 2006.

Received the *Award for Teaching Excellence*, University of Maryland (campus-wide), in recognition of excellence in teaching at the undergraduate level, May 2004.

In February, 2003, *The Computer Engineering Handbook* (a title to which I contributed the following chapter) was honored as an Outstanding Academic Title for 2002 by *Choice Magazine*, a publication of the American Library Association.

B. Jacob. “Virtual Memory Systems and TLB Structures.” In *The Computer Engineering Handbook*, pp. 5:55–5:65. V. Oklobdzija, Editor. CRC Press: Boca Raton FL, 2002.

Selections for this honor are made based upon “their excellence in scholarship and presentation, the significance of their contribution to the field, and their value as important — often the first — treatment of their subject.”

General Co-Chair for 5th Workshop on Computer Architecture Education, January 1999.

Received the *George Corcoran Memorial Award*, 1998/9 academic year, “[i]n recognition of teaching and educational leadership at the University of Maryland, College Park campus, effective contribution at the national level, and creative and other scholarly activities related to Electrical and Computer Engineering education.”

IV. SERVICE

A. PROFESSIONAL SERVICE

Award/Honors Committees

Member, *ACM/IEEE CS George Michael Memorial HPC Fellowship Committee*, 2018–2020

Editorial Service

Associate Editor, *IEEE Computer Architecture Letters (CAL)*, 2007–2011

Associate Editor, *ACM Transactions on Embedded Computing Systems (TECS)*, 2006–2009

Co-Editor, Special Inaugural Issue of *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 1, no. 1.
December 2002

Proposal-Review Panels

Panelist, Department of Energy, July 2018

Panelist, Department of Energy, April–May 2018

Panelist, National Science Foundation, February 2017

Panelist, National Science Foundation, March 2008

Panelist, National Science Foundation, October 2003

Panelist, National Science Foundation, November 2003

Panelist, National Science Foundation, February 2002

Organizing Committees

General Chair, *9th International Symposium on Memory Systems (MEMSYS 2023)*, October 2023

General & Program Chair, *8th International Symposium on Memory Systems (MEMSYS 2022)*, October 2022
(virtual)

General & Program Chair, *7th International Symposium on Memory Systems (MEMSYS 2021)*, October 2021
(virtual)

General & Program Chair, *6th International Symposium on Memory Systems (MEMSYS 2020)*, October 2020
(virtual)

General & Program Chair, *5th International Symposium on Memory Systems (MEMSYS 2019)*, October 2019

General & Program Chair, *4th International Symposium on Memory Systems (MEMSYS 2018)*, October 2018

Organizing Committee Member, *DOE Workshop on Modeling & Simulation of Systems and Applications (ModSim 2018)*, August 2018

Moderator/Panel Chair, *2018 Symposia on VLSI Technology and Circuits*, June 2018

General & Program Chair, *3rd International Symposium on Memory Systems (MEMSYS 2017)*, October 2017

General & Program Chair, *2nd International Symposium on Memory Systems (MEMSYS 2016)*, October 2016

General & Program Chair, *1st International Symposium on Memory Systems (MEMSYS 2015)*, October 2015

Panel Discussion Coordinator: *Memory Systems — What's the Real Problem? What's the Next Solution?* Held at the
Workshop on Memory System Performance and Correctness (MSPC 2008), March 2008

Session Chair, *13th International Symposium on High-Performance Computer Architecture (HPCA 2007)*,
February 2007

Session Chair, *International Symposium on Low Power Electronics and Design (ISLPED 2006)*, October 2006

Workshops Chair, *8th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2005)*, September 2005

Local Arrangements Chair, *7th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2004)*, November 2004

Coordination Vice-Chair, *5th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2002)*, November 2002

Publicity Chair, *2nd Workshop on Compiler and Architecture Support for Embedded Systems (CASES 1999)*, October 1999

General Co-Chair, *5th Workshop on Computer Architecture Education (WCAE 1999)*, January 1999

Local Arrangements Chair, *3rd IEEE Symposium on High-Assurance Systems Engineering (HASE 1998)*, November 1998

Program Committees

Program Committee Member, *The 48th International Symposium on Computer Architecture (ISCA 2021)*, May 2021

Program Committee Member, *The 27th International Symposium on High-Performance Computer Architecture (HPCA-27)*, February 2021

Program Committee Member, *The 45th International Symposium on Computer Architecture (ISCA 2018)*, June 2018

Program Committee Member, *Supercomputing — Architectures & Networks (SC17)*, November 2017

Program Committee Member, *14th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2011)*, October 2011

Program Committee Member, *43rd International Symposium on Microarchitecture (MICRO 2010)*, December 2010

Program Committee Member, *13th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2010)*, October 2010

Program Committee Member, *41st International Symposium on Microarchitecture (MICRO 2008)*, November 2008

Program Committee Member, *11th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2008)*, November 2008

Program Committee Member, *International Conference on Computer Design (ICCD 2008)*, October 2008

Program Committee Member, *Workshop on Memory System Performance and Correctness (MSPC 2008)*, March 2008

Program Committee Member, *13th International Symposium on High-Performance Computer Architecture (HPCA 2007)*, February 2007

Program Committee Member, *39th International Symposium on Microarchitecture (MICRO 2006)*, December 2006

Program Committee Member, *9th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2006)*, October 2006

Program Committee Member, *International Conference on Computer Design (ICCD 2006)*, October 2006

Program Committee Member, *8th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2005)*, September 2005

Program Committee Member, *Memory Systems Performance Workshop (MSP 2005)*, June 2005

Program Committee Member, *7th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2004)*, November 2004

Program Committee Member, *6th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2003)*, November 2003

Program Committee Member, *4th International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES 2001)*, November 2001

Program Committee Member, *Workshop on Memory Performance Issues*, June 2001

Program Committee Member, *4th IEEE Symposium on High-Assurance Systems Engineering (HASE 1999)*, November 1999

General Reviewing

Reviewer for papers in the following journals:

IEEE Transactions on Computers

IEEE Transactions on Parallel and Distributed Systems

IEEE Transactions on Very Large Scale Integration Systems

IEEE Computer

ACM Transactions on Embedded Computing Systems

ACM Transactions on Architecture and Code Optimization

ACM Computing Surveys

IBM Journal of Research and Development

The Computer Journal

Reviewer for papers in the following annual meetings:

ACM/IEEE International Symposium on Computer Architecture (ISCA)

ACM/IEEE International Symposium on Microarchitecture (MICRO)

ACM International Conference on Supercomputing (ICS)

IEEE International Symposium on High-Performance Computer Architecture (HPCA)

IEEE International Symposium on High-Assurance Systems Engineering (HASE)

IEEE International Conference on Computer Design (ICCD)

Int'l Conf. on Compilers, Architectures, and Synthesis for Embedded Systems (CASES)

Workshop on Computer Architecture Education (WCAE)

B. UNIVERSITY SERVICE AT THE NAVAL ACADEMY

2023–present Chair, Cyber Science Faculty-Search Committee

C. UNIVERSITY SERVICE AT MARYLAND

Departmental

Winter 2020 Speaker at the ECE Department's TA Teaching and Development (TATD) event

Fall 2018 Member, ECE APT Committee

Spring 2017 Keynote speaker at the ECE Department's TA Teaching and Development Fellowship event

Spring 2017 Member, ECE Faculty Search Sub-Committee for Computer Engineering

Fall 2016 Member, ECE Faculty Search Committee for Microelectronics

Fall 2016 Served as the Moderator of the 9th Annual ECEGSA Academic Careers Panel, an all-day event held November 4, 2016

Fall 2016 Invited by ECEGSA to speak to the ECE Department's graduate students on how to give a good research presentation

2015–2016	Developed from scratch and taught <i>ENEE 447: Operating Systems</i> at Dept. Chair's request, to improve the Computer Engineering program
2015–2017	Member, ECE Faculty Search Committee
2014–2015	Brought in to teach and overhaul the laboratories of <i>ENEE 245: Digital Circuits & Systems Lab</i> at Dept. Chair's request, to improve the Computer Engineering program
2013–2014	Chair, ECE Promotion and Tenure Subcommittee
2013–2016	Member, ECE Promotion and Tenure Committee
2007–2010	Director, Computer Engineering Program — first faculty to hold the position; charged by Dept. Chair to overhaul the curriculum (in particular, improve the curriculum's consistency across topics), increase the program's external visibility, and increase enrollment; all tasks completed successfully before sabbatical in 2010
2006–2011	Member, ECE PhD Qualifying Exam Committee
2006–2007	Prepared accreditation materials for Middle States Commission on Higher Education
2005–2007	Member, ECE Faculty Search Committee
2005–2006	Prepared materials for ABET accreditation process
Spring 2005	ECE search committee for <i>Chairperson of the ECE Department</i>
2004–2007	Member, ECE Promotion and Tenure Committee
Spring 2004	ECE search committee for <i>Director of Computing Facilities</i> position
2002–2004	Member, ECE Department Council
Spring 2001	ECE search committee for <i>Chairperson of the ECE Department</i>
2000–2002	Member, ECE Facilities and Services Committee
1999–2000	Member, ECE Undergraduate Affairs Committee
1998–1999	Chair, ECE Undergraduate Affairs Committee
1998–2000	Organized the weekly <i>Research in Computer Engineering Colloquium</i>
1998–2000	Member, ECE General Academic Affairs Committee
1998–2000	Member, ECE Curriculum Revision Committee
1998–2000	Computer Engineering representative for ABET transition
Spring 1998	ECE search committee for Personnel and <i>Research Coordinator</i> position
Spring 1998	ECE search committee for <i>Systems Analyst</i> position
1997–1998	Organized a weekly Microarchitecture Seminar Series

College

2005–2008	Member, Engineering College Promotion and Tenure Committee
2004–2005	Member, Dean's committee formed to review and revise the college-wide <i>ENES 100 — Introduction to Engineering Design</i>
1999–2001	Member, ECE College of Engineering Council

University

Fall 2018	Invited to present Tech+Music seminar to Technica/Hackathon
Winter 2017	Guest speaker at <i>Foundations of Entrepreneurship & Innovation (HEIP143)</i> course

Fall 2017 Invited to present Innovation & Design seminar to Mtech Startup Hours/Shell Talk

2017–2019 Faculty consultant for the EIP/NIST Technology Market Assessment (T-Map) Project

2006–2007 ECE representative/member of CS Department Curriculum Review Committee

1999–2003 Faculty advisor for university chapter of Eta Kappa Nu (EE Honor Society)

➡ The Maryland chapter received the Certificate of Merit from the national Eta Kappa Nu Association for Outstanding Chapter Activities, 1999–2000 academic year. The two other chapters receiving the same award were from Carnegie Mellon and MIT.

V. RESEARCH, TESTIMONY, AND DEPOSITION AS EXPERT WITNESS

Dr. Jacob has written numerous expert reports, has testified in court, has been deposed, has performed prior-art searches, has done claims construction and patent-validity analysis. He has worked on roughly fifty cases in the past twenty-five years and has worked for most semiconductor memory manufacturers, many systems manufacturers, and several general semiconductor manufacturers. He has served on the side of the plaintiff/complainant/petitioner as well as on the side of the defendant/respondent.

He is a leading expert in memory systems and has served as a testifying expert in that and related fields.

The following are his appearances at trial and deposition (out of respect for client confidentiality and protective orders, only publicly available information is provided here):

U.S. District Court for the Central District of California, Case No. 8:23-cv-00039-JWH-DFM, July 2023–present

Provided technical opinions to Plaintiff on claim-term constructions regarding memory arrays and devices, and their I/O subsystems both volatile and non-volatile, planar and 3D.

Deposition of January 3, 2024

U.S. District Court for the Western District of Texas, Case No. 6:20-cv-00636-ADA, June 2023–February 2024

Provided technical opinions to Plaintiff, and generated performance analyses, on the relationship between transistor delay and memory performance in DRAM products.

Trial testimony of February 13, 2024

Deposition of November 8, 2023

U.S. Patent and Trademark Office Inter Partes Review No. IPR2021-01064, March 2022

Petitioner submitted a public declaration of mine from a previous case, to institute the IPR, and I agreed to a deposition to defend the analysis in the declaration.

Deposition of March 31, 2022

U.S. District Court for the Western District of Texas, Case No. 6:20-cv-501-ADA, July 2020–February 2021

Provided technical opinions to Plaintiff on claim-term constructions, and performed technical analysis regarding the management of flash-based memory arrays and devices, the design of flash translation layers (FTLs), and the internals of solid-state drives (SSDs).

Deposition of February 8, 2021

U.S. Patent and Trademark Office Inter Partes Review No. IPR2020-00533, July 2019–November 2020

Provided technical opinions to Petitioner, performed prior-art search, helped with legal strategy, and performed detailed validity analysis regarding processor-cache hierarchies and memory-coherence protocols within them.

Deposition of November 2, 2020

U.S. Patent and Trademark Office Inter Partes Review No. IPR2019-00104, June 2018–November 2019

Provided technical opinions to Petitioner, performed prior-art search, helped with legal strategy, and performed detailed validity analysis regarding voltage boosting circuits.

Deposition of November 21, 2019

U.S. Patent and Trademark Office Inter Partes Review No. IPR2019-00102, June 2018–November 2019

Provided technical opinions to Petitioner, performed prior-art search, helped with legal strategy, and performed detailed validity analysis regarding memory-array sense amplifier design.

Deposition of November 21, 2019

U.S. Patent and Trademark Office Inter Partes Review No. IPR2018-01794, June 2018–November 2019

Provided technical opinions to Petitioner, performed prior-art search, helped with legal strategy, and performed detailed validity analysis regarding leakage techniques for memory arrays.

Deposition of November 20, 2019

U.S. Patent and Trademark Office Inter Partes Review No. IPR2018-01784, June 2018–November 2019

Provided technical opinions to Petitioner, performed prior-art search, helped with legal strategy, and performed detailed validity analysis regarding memory-array sense amplifier design.

Deposition of November 20, 2019

U.S. International Trade Commission Investigation No. 337-TA-1097, October 2017–April 2019

Provided technical opinions to Complainants, helped with legal strategy, and performed detailed schematics- and firmware-based infringement analysis on SSD and stacked memory architectures.

Deposition of December 17 and 18, 2018

U.S. Patent and Trademark Office Inter Partes Review No. IPR2017-00381, September 2016–January 2018

Provided technical opinions to Petitioner, performed prior-art search, helped with legal strategy, and performed detailed validity analysis regarding memory-system and memory-device signaling interfaces.

Deposition of January 23, 2018

Deposition of September 11, 2017

U.S. Patent and Trademark Office Inter Partes Review No. IPR2015-01068, January 2015–April 2016

Provided technical opinions to Petitioner, helped with legal strategy, and performed detailed, schematics-based non-infringement and validity analysis regarding memory-system and memory-device signaling interfaces.

Deposition of January 7, 2016

U.S. International Trade Commission Investigation No. 337-TA-941, January 2015–August 2015

Provided technical opinions to Respondents, helped with legal strategy, and performed detailed, schematics-based non-infringement and validity analysis regarding memory-system and memory-device signaling interfaces.

Trial testimony of August 20, 2015

Deposition of July 2, 2015

U.S. International Trade Commission Investigation No. 337-TA-909, April 2014–October 2014

Provided technical opinions to Complainants, helped with legal strategy, and performed detailed, schematics-based infringement analysis on NOR and NROM flash-device architectures.

Trial testimony of October 21, 22, and 27, 2014

Deposition of September 5, 2014

U.S. International Trade Commission Investigation No. 337-TA-803, November 2011–September 2012

Provided technical opinions to Respondents, helped with legal strategy, and performed research on prior art regarding memory-system and memory-device architectures.

Deposition of June 7, 2012

Markman-Hearing testimony of April 17, 2012

Deposition of April 9, 2012

U.S. International Trade Commission Investigation No. 337-TA-753, March 2011–October 2011

Provided technical opinions to Respondents, helped with legal strategy, and performed research on prior art regarding high-speed memory-system architectures.

Trial testimony of October 14, 2011

Deposition of August 23 and 24, 2011

Rambus vs. Nanya USA, May 2007–March 2008

Provided technical opinions to Nanya USA regarding DDR2 SDRAM interfaces and alternative technologies.

Deposition of October 26, 2007

U.S. Federal Trade Commission vs. Rambus, Docket No. 9302, July 2002–July 2003

Provided technical opinions to the FTC regarding SDRAM interfaces, Rambus intellectual property, infringement of Rambus patents, and alternative technologies.

Trial testimony of June 16 and 17, 2003

Deposition of March 19, 2003