

# Habits for Happy Success: Reflections from My Anti-Bio

**Sarita Adve**

**University of Illinois at Urbana-Champaign**

**sadve@illinois.edu**

Recent research is supported in part by DARPA, NSF, and the Applications Driving Architecture (ADA) Research center (JUMP center co-sponsored by SRC & DARPA)



You are going to be successful

The only question is

what shape will your success take?

C. Sidney Burrus, Rice University

# My Bio

- Grew up in India
- 1983-87: B.Tech., Indian Institute of Technology, Bombay
- 1987-88: One year at Johns Hopkins
- 1988-93: PhD., University of Wisconsin-Madison
- 1993-99: Assistant Prof, Rice University
- 1999-05: Associate Prof, University of Illinois
- 2005 - : Professor, Illinois

Lots of papers, students, awards, funding, ...

# My Bio

- Grew up in India
- 1983-87: B.Tech., Indian Institute of Technology, Bombay **Met future husband**
- 1987-88: One year at Johns Hopkins **Long-distance relationship**
- 1988-93: PhD., University of Wisconsin-Madison **Married**
- 1993-99: Assistant Prof, Rice University **Husband was research scientist (2BP)**
- 1999-05: Associate Prof, University of Illinois **Husband assistant prof, had Kid 1**
- 2005 - : Professor, Illinois **Husband got tenure, Kid 2**

Lots of papers, students, awards, funding, ...

# Anti-Bio

- Grew up in India
- 1983-87: B.Tech., Indian Institute of Technology, Bombay *Met future husband*
- 1987-88: One year at Johns Hopkins *Long-distance relationship*
- 1988-93: PhD., University of Wisconsin-Madison *Married*
  - Denied admission
  - First paper got in, but had a hard time convincing industry about approach
  - Key journal paper rejected several times
- 1993-99: Assistant Prof, Rice University *Husband was research scientist (2BP)*
  - Brother diagnosed with cancer in '94
  - In first three years, had only one paper from my group  
(one led by senior colleague, another tutorial on my thesis work)
- 1999-05: Associate Prof, University of Illinois *Husband assistant prof, had Kid 1*
  - Almost gave up, too hard with newborn, asst prof husband, new job, ...
  - Java memory model: almost gave up
  - GRACE project: first cross-layer (arch/OS/network/app) energy/quality/performance adaptation for multimedia apps,  
but couldn't get a paper published in architecture conferences
- 2005 - : Professor, Illinois *Husband got tenure, Kid*
  - Initial PnT meeting declined to promote
  - DeNovo cache coherence: many rejections
  - ...

Don't give up

Believe in yourself

# Don't Give Up: Believe in Yourself

- Grew up in India
- 1983-87: B.Tech., Indian Institute of Technology, Bombay *Met future husband*
- 1987-88: One year at Johns Hopkins *Long-distance relationship*
- 1988-93: PhD., University of Wisconsin-Madison *Married*
  - Denied admission
  - First paper got in, but had a hard time convincing industry about approach
  - Key journal paper rejected several times
- 1993-99: Assistant Prof, Rice University *Husband was research scientist (2BP)*
  - Brother diagnosed with cancer in '94
  - In first three years, had only one paper from my group  
(one led by senior colleague, another tutorial on my thesis work)
- 1999-05: Associate Prof, University of Illinois *Husband assistant prof, had Kid 1*
  - Almost gave up, too hard with newborn, asst prof husband, new job, ...
  - Java memory model: almost gave up
  - GRACE project: first cross-layer (arch/OS/network/app) energy/quality/performance adaptation for multimedia apps,  
but couldn't get a paper published in architecture conferences
- 2005 - : Professor, Illinois *Husband got tenure, Kid*
  - Initial PnT meeting declined to promote
  - DeNovo cache coherence: many rejections
  - ...

# Don't Give Up: Believe in Yourself

- Grew up in India
- 1983-87: B.Tech., Indian Institute of Technology, Bombay *Met future husband*
- 1987-88: One year at Johns Hopkins *Long-distance relationship*
- 1988-93: PhD., University of Wisconsin-Madison *Married*
  - Denied admission
  - First paper got in, but had a hard time convincing industry about approach
  - Key journal paper rejected several times

## How to learn this habit?

- 1999-05: Associate Prof, University of Illinois *Husband assistant prof, had Kid 1*
  - Almost gave up, too hard with newborn, asst prof husband, new job, ...
  - Java memory model: almost gave up
  - GRACE project: first cross-layer (arch/OS/network/app) energy/quality/performance adaptation for multimedia apps, but couldn't get a paper published in architecture conferences
- 2005 - : Professor, Illinois *Husband got tenure, Kid*
  - Initial PnT meeting declined to promote
  - DeNovo cache coherence: many rejections
  - ...



Get comfortable being uncomfortable

Work just outside your comfort zone

# Working Outside Comfort Zone

Happens slowly

PhD:

**Data-race-free memory models:** Software-centric approach to perceived hardware problem

A third of my (architecture) PhD thesis was on proofs, no simulations/experiments

Tenure:

**RSIM simulator:** First simulator to model out-of-order processors in a multiprocessor

First quantifications of using speculation for strong memory models

Memory level parallelism

Full professor:

**GRACE:** Global Resource Adaptation through CoopEratioN

Working with signal processing, network, OS, multimedia faculty

Cross-layer, performance/power/quality co-design before they were buzzwords

**Java memory model:** Convince programming languages community of DRF approach

Post-promotions:

**DeNovo coherence:** Initially co-designed with a language with a correctness-motivated type system

**SW-driven hardware resiliency:** Software engineering for hardware errors

**ILLIXR: Illinois Extended Reality** testbed: First fully open-source AR/VR/MR/XR system

**SIGARCH chair:** Changing culture



# Working Outside Comfort Zone

Happens slowly

PhD:

**Data-race-free memory models:** Software-centric approach to perceived hardware problem

A third of my (architecture) PhD thesis was on proofs, no simulations/experiments

Tenure:

**RSIM simulator:** First simulator to model out-of-order processors in a multiprocessor

First quantifications of using speculation for strong memory models

## But isn't this too hard?

**CRATER:** Global Resource Adaptation through Cooperation

Working with signal processing, network, OS, multimedia faculty

Cross-layer, performance/power/quality co-design before they were buzzwords

**Java memory model:** Convince programming languages community of DRF approach

Post-promotions:

**DeNovo coherence:** Initially co-designed with a language with a correctness-motivated type system

**SW-driven hardware resiliency:** Software engineering for hardware errors

**ILLIXR:** Illinois Extended Reality testbed: First fully open-source AR/VR/MR/XR system

**SIGARCH chair:** Changing culture



Be passionate

Believe in what you do, go all in

# Be Passionate, Go All In

- What is the problem?  
Understand, read
- Question fundamentals  
Don't be afraid, take risks
- Focus on impact  
Impact = Change minds. Takes time.

An example...

# Passionate About Memory Models

*[With Mark Hill]*

- 1988 to 1989: What is a memory consistency model?
  - Simplest model: sequential consistency (SC) [Lamport79]
    - Memory operations execute one at a time in program order
    - Simple, but inefficient


– Implementation/performance-centric view

- Order in which memory operations execute
- Different vendors w/ different models (orderings)

– Alpha, Sun, x86, Itanium, IBM, AMD, HP, Cray, ...

- Complex, many ambiguities, ...

– A new memory model virtually everyday



- 1988 to 1989: What is a memory consistency model?

Memory model = What value can a read return?

Initially X=Y=Flag=0

<i>Thread 1</i>	<i>Thread 2</i>
X = 26	if (Flag == 1) {
Y = 90	... = Y ← 90
...	... = X ← 26
Flag = 1	... }

*Note: In the original image, the values 90 and 26 in the Thread 2 code are crossed out with a red 'X', and a red '0' is written to the right of the second line.*

HW/SW Interface: affects performance, programmability, portability

- 1990-93: Software-centric view: Data-race-free (DRF) model
  - Sequential consistency for data-race-free programs [Adve, Hill ISCA90]
  - Distinguish **data** vs. **synchronization (race)**
    - Data can be optimized  $\Rightarrow$   $\uparrow$  performance for DRF programs

Initially  $X=Y=Flag=0$

<i>Thread 1</i>	<i>Thread 2</i>
$X = 26$	if ( $Flag == 1$ ) {
$Y = 90$	... = Y
...	... = X
$Flag = 1$	...
	}

**Worked on many HW-SW contracts, dug as deep as possible**

**Most of this unpublished  
But necessary for passion!**

Ack: Jim Goodman, Bart Miller, Rob Netzer, Kouros Gharachorloo




*[With Vijay Pai and Partha Ranganathan]*

- **1993-99: Performance benefits of relaxed models**
  - New out-of-order processors emerging, new speculation techniques
  - No tools to understand performance implications
  - RSIM: Built first publicly available multiprocessor simulator with out-of-order processors [Pai et al. ASPLOS'96, ISCA'97, ...]
- More confidence in DRF!
  - Called out compiler and PL community
  - Proceedings of IEEE paper caught attention of Bill Pugh

*[with Bill Pugh, Jeremy Manson, Doug Lea, Hans Boehm, et al.]*

- 2000-05: Java memory model [Manson, Pugh, Adve POPL'05]
  - DRF model BUT racy programs need semantics

⇒ No out-of-thin-air values



Can  $r1=r2=42$ ?

Problem: Incredibly hard to formalize a spec that prohibits this result without prohibiting common optimizations

Java memory model = DRF + big mess

*[With Hans Boehm et al.]*

- 2005-08: C++ memory model [Boehm, Adve PLDI'08]
  - DRF model BUT need high performance; mismatched hardware
  - Baseline DRF (DRF0) requires synchronization/atomics to be SC
  - Hardware vendors, software developers complained, but no option
  - Compromise: Relaxed atomics (only for experts)
    - ⇒ DRF + big mess

Good news: After 20 years, convergence at last!

But: How to debug racy programs, how to avoid out of thin air values, no semantics for relaxed atomics, ...

CACM'10: Memory Models: A Case for Rethinking Parallel Languages and Hardware

*[With Matt Sinclair and John Alsup]*

## C++17 "specification" for relaxed atomics

- Races that don't order other accesses
  - Implementations should ensure no "out-of-thin-air" values are computed that circularly depend on their own computation
- "C++ (relaxed) atomics were the **worst idea ever**. I just spent days (and days) trying to get something to work. ... My example **only has 2 addresses and 4 accesses, it shouldn't be this hard**. Can you help?"

- Email from employee at major research lab

Again, software use driven solution in ISCA'17, still unfolding

[With Vikram Adve, Byn Choi, Rakesh Komuravelli, Matt Sinclair, Hyojin Sung]

- 2008-14: Software-centric view for coherence: DeNovo protocol
    - More performance-, energy-, and complexity-efficient than MESI
      - Began with DPJ's disciplined parallelism
      - Identified fundamental, minimal coherence mechanisms
      - Loosened s/w constraints, but still minimal, efficient hardware
- Ack: Marc Snir, UPCRC
- Meanwhile: the end of Dennard and Moore's laws
    - Architecture enters golden age
    - *Déjà vu for coherence and consistency*
  - Next phase with John Alsop, Matt Sinclair, Weon Tak Na (+ Huzaifa Muhammad, Wes Darwin, Sam Grayson) with *Spandex coherence interface, still unfolding*

# And Now... ILLIXR

**[with Muhammad Huzaifa and large ILLIXR team]**

- Golden age of computer architecture: domain-specific systems
  - But what domain
  - What systems?
  - Cross-layer co-design, but no end-to-end workloads!
  
- Last 3+ years
  - Illinois Extended Reality Testbed
  - Illixr.org

**Passion is easy!**



Build relationships that make you fly  
And pay it forward

# It's All About People

Collaborators: faculty, industry colleagues, students, ...

Mentors

Family

**Pay it Forward**





# My Habits for Happy Success

- Don't give up
- Don't be afraid to work outside your comfort zone
- Be passionate
- Build relationships and pay it forward

You are going to be successful

The only question is

what shape will your success take?

C. Sidney Burrus, Rice University