

# Yonghong Yan

Department of Computer Science and Engineering  
University of South Carolina, Columbia, SC 29208

yanyh@cse.sc.edu  
<http://cse.sc.edu/~yanyh>

## CURRENT POSITION AND MEMBERSHIP

---

Assistant Professor, University of South Carolina *08/2017 - present*  
Member of OpenMP Architectural Review Board (ARB) *07/2013 - present*

## AREA OF INTERESTS

---

Parallel and High Performance Computing, Parallel Programming Languages and Compilers,  
Computer Systems and Architectures, and Distributed Computing Systems

## RESEARCH AND EDUCATION BACKGROUND

---

**Oakland University**, Rochester MI *08/2014 - 08/2017*  
**Assistant Professor**, Department of Computer Science and Engineering

**University of Houston**, Houston TX *06/2011 - 08/2014*  
**Research Assistant Professor**, Computer Science Department

**Rice University**, Houston TX *07/2008 - 06/2011*  
**Postdoctoral Research Associate**, Computer Science Department, Sponsor: Vivek Sarkar

**University of Houston**, Houston TX *01/2002 - 05/2007*  
**Ph.D. of Computer Science**, GPA 3.7/4.0, Advisor: Barbara M. Chapman  
Dissertation: Scheduling Scientific Workflow Applications in Computational Grids

**University of Shanghai for Science and Technology**, Shanghai China *09/1998 - 02/2001*  
**Master of Computer Engineering**  
Thesis: A Common Object Request Broker Architecture(CORBA) Resource Management  
Framework for Distributed Computing

**Shandong University of Science and Technology**, Shandong China *09/1994 - 07/1998*  
**Bachelor of Electrical and Mechanical Engineering**, GPA 3.7/4.0  
Minor in Computer Science, GPA: 3.7/4.0  
Thesis: Embedded Computer System for Automatic Bottle Filling

## CURRENT PROJECTS

---

NSF:CAREER Programming the Existing and Emerging Memory Systems for Extreme-scale Parallel Performance, Lead PI, NSF CISE CCF/SHF-1652732, **Awarded**, \$600,00.00, 02/01/2017-01/31/2022

NSF SHF:Small:Collaborative: Application-aware Energy Modeling and Power Management for Parallel and High Performance Computing, Lead PI, NSF CISE SHF-1551182, **Current**, \$250,000.00 (with Virginia Tech), 09/01/2014-08/31/2081

NSF SHF:Medium:Collaborative: Compute on Data Path: Combating Data Movement in High-Performance Computing, Co-PI, NSF CISE SHF-1409946, **Current**, \$130,000.00 (with Texas Tech and Northwestern U), 08/01/2014-07/31/2018

High Performance Implementation using GPU and FPGA of Deformable Image Registration for Cancer Treatment, Lead PI, Beaumont Cancer Institute of William Beaumont Health System, **Current**, 08/2015 current

## PUBLICATIONS

---

- HOMP: Automated Distribution of Parallel Loops in Highly Parallel Accelerator-Based Systems  
**Yonghong Yan**, Jiawen Liu, and Kirk W. Cameron  
accepted to The IEEE International Parallel & Distributed Processing Symposium (IPDPS) 2017
- A Cross-Layer Solution in Scientific Workflow System for Tackling Data Movement Challenge  
Dong Dai, Robert Ross, Dounia Khaldi, **Yonghong Yan**, Matthieu Dorier, Neda Tavakoli, Yong Chen  
IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC) 2016
- Comparison of Spark Resource Managers and Distributed File Systems  
Soulmaz Salehian, **Yonghong Yan**  
First IEEE International Workshop on Big Data and Cloud Applications (BDCloudApp 2016), October, 2016
- A Proposal to OpenMP for Addressing the CPU Oversubscription Challenge  
**Yonghong Yan**, Jeff R. Hammond, Chunhua Liao, Alexandre E. Eichenberger  
International Workshop on OpenMP (IWOMP) 2016, October 2016
- Extending OpenMP to Support Multiple Accelerators  
Pei-Hung Lin, **Yonghong Yan**, Chunhua Liao, and Daniel J. Quinlan  
submitted to The International Journal of High Performance Computing Applications, 2016
- An Adaptive and Lightweight Synchronization Mechanism for Asynchronous Tasks on Manycore Architectures  
**Yonghong Yan**, Sanjay Chatterjee, Daniel Orozco, Elkin Garcia, Zoran Budimlic, Robert Pavel, Jun Shirako, Guang R. Gao, and Vivek Sarkar  
accepted to the Journal of Concurrency and Computation: Practice and Experience, 2016
- Compiler Transformation of Nested Loops for GPGPUs  
Xiaonan Tian, Rengan Xu, **Yonghong Yan**, Sunita Chandrasekaran, and Barbara Chapman  
Journal of Concurrency and Computation: Practice and Experience, August 2015
- Supporting Multiple Accelerators in High-Level Programming Models  
**Yonghong Yan**, Pei-Hung Lin, Chunhua Liao, Bronis R. de Supinski, and Daniel J. Quinlan  
2015 International Workshop on Programming Models and Applications for Multicores and Manycores (PMAM 2015) in conjunction with PPOPP, February 2015
- NAS Parallel Benchmarks for GPGPUs using a Directive-based Programming Model  
Rengan Xu, Xiaonan Tian, Sunita Chandrasekaran, **Yonghong Yan** and Barbara Chapman  
27th International Workshop on Languages and Compilers for Parallel Computing (LCPC2014)
- Predicting Cache Contention for Multithread Applications at Compile Time  
Munara Tolubaeva, **Yonghong Yan** and Barbara Chapman  
16th Workshop on Advances in Parallel and Distributed Computational Models to be held in conjunction with IPDPS 2014, May 2014
- Reduction Operations in Parallel Loops for GPGPUs  
Rengan Xu, Xiaonan Tian, **Yonghong Yan**, Sunita Chandrasekaran, and Barbara Chapman  
2014 International Workshop on Programming Models and Applications for Multicores and Manycores (PMAM 2014) in conjunction with PPOPP, February, 2014
- Compiling a High-level Directive-Based Programming Model for GPGPUs  
Xiaonan Tian, Rengan Xu, **Yonghong Yan**, Zhifeng Yun, Sunita Chandrasekaran, and Barbara Chapman  
26th International Workshop on Languages and Compilers for Parallel Computing (LCPC2013)
- Compile Time Modeling of Off-Chip Memory Bandwidth for Parallel Loops  
Munara Tolubaeva, **Yonghong Yan**, and Barbara Chapman  
26th International Workshop on Languages and Compilers for Parallel Computing (LCPC2013)
- A Prototype Implementation of OpenMP Task Dependency Support

- Priyanka Ghosh, **Yonghong Yan**, Deepak Eachempati and Barbara Chapman  
International Workshop on OpenMP (IWOMP) 2013, September 2013
- Early Experiences With The OpenMP Accelerator Model  
Chunhua Liao, **Yonghong Yan**, Bronis R. de Supinski, Daniel J. Quinlan and Barbara Chapman  
International Workshop on OpenMP (IWOMP) 2013, September 2013
- Integrating Asynchronous Task Parallelism with MPI  
Sanjay Chatterjee, Zoran Budimlić, Vincent Cavé, Milind Chabbi, Max Grossman, Saĝnak Taşırlar,  
**Yonghong Yan**, Vivek Sarkar; IPDPS 2013
- Performance and Power Characteristics of Matrix Multiplication Algorithms  
on Multicore and Shared Memory Machines  
**Yonghong Yan**, Jeremy Kemp, Xiaonan Tian, Abid Muslim Malik, and Barbara Chapman  
Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (ScalA) at SC12, Nov 2012
- Support for Dependency Driven Execution among OpenMP Tasks  
Priyanka Ghosh, **Yonghong Yan**, Barbara Chapman  
Workshop on Data-Flow Execution Models for Extreme Scale Computing (DFM 2012) at PACT, 2012/09
- Compile-Time Detection of False Sharing via Loop Cost Modeling  
Munara Tolubaeva, **Yonghong Yan**, Barbara Chapman  
Workshop on High-Level Parallel Programming Models and Supportive Environments at IPDPS, 2012/05
- Integrating MPI with Asynchronous Task Parallelism using Habanero-C  
**Yonghong Yan**, Sanjay Chatterjee, Zoran Budimlic, and Vivek Sarkar, EuroMPI 2011
- Hardware and Software Tradeoffs for Task Synchronization on Manycore Architectures  
**Yonghong Yan**, Sanjay Chatterjee, Daniel Orozco, Elkin Garcia, Jun Shirako, Zoran Budimlic,  
Robert Pavel, Guang R. Gao, and Vivek Sarkar, Euro-Par 2011
- Parallel Object-Oriented Scientific Computing with Habanero-Java  
Zoran Budimlic, Vincent Cave, Jun Shirako, **Yonghong Yan**, Jisheng Zhao,  
Vivek Sarker, Michael Glinsky, James Gunning  
9th Workshop on Parallel/High-performance Object-Oriented Computing at SPLASH 2010
- The Habanero Multicore Software Research Project  
Rajkishore Barik, Zoran Budimlic, Vincent Cave, Sanjay Chatterjee, Yi Guo, David Peixotto,  
Raghavan Raman, Jun Shirako, Sagnak Tasirlar, **Yonghong Yan**, Yisheng Zhao, and Vivek Sarkar,  
24th ACM SIGPLAN Conference on OOPSLA, 2009
- JCUDA: a Programmer-Friendly Interface for Accelerating Java Programs with CUDA  
**Yonghong Yan**, Max Grossman, and Vivek Sarkar; Euro-Par 2009, August 2009
- Hierarchical Place Trees: A Portable Abstraction for Task Parallelism and Data Movement  
**Yonghong Yan**, Jisheng Zhao, Yi Guo, and Vivek Sarkar  
The 22nd Workshop on Languages and Compilers for Parallel Computing (LCPC), October 2009
- A Feature-Rich Workflow Description Language that Supports Resource Co-allocations  
**Yonghong Yan** and Barbara Chapman  
High Performance Computing and Grids in Action, edited by Lucio Grandinetti,  
Volume 16 of Advances in Parallel Computing, March 2008, IOS Press Amsterdam
- Scientific Workflow Scheduling in Computational Grids - Planning, Reservation and Data/Network-Awareness  
**Yonghong Yan** and Barbara Chapman  
The 8th IEEE/ACM International Conference on Grid Computing, September 2007
- Air Quality Prediction in a Production Quality Grid Environment  
Barbara Chapman, Priya Raghunath, Babu Sundaram, and **Yonghong Yan**  
Engineering the Grid: status and perspective, edited by Jack Dongarra,

Hans Zima, Adolfo Hoisie, Laurence Yang and Beniamino Di Martino, January 2006

Campus Grids Meet Applications: Modeling, Metascheduling and Integration

**Yonghong Yan** and Barbara Chapman; Journal of Grid Computing, Volume 4(2), pp159-175, June 2006

Air Quality Forecasting on Campus Grid

**Yonghong Yan**, Barbara Chapman, and Babu Sundaram

Global Grid Forum - Workshop on Grid Applications: from Early Adopters to Mainstream Users, June 2005

Grid Environment With Web-Based Portal Access For Air Quality Modeling

Barbara Chapman, Hari Donepudi, Yupeng Li, Priya Raghunath, **Yonghong Yan**, Babu Sundaram and Jiwen He

Parallel and Distributed Scientific and Engineering Computing: - Practice and

Experience Advances in Computation: Theory And Practice, NOVA Publisher 2004

An OGSi-Compliant Portal for Campus Grids

Barbara Chapman, Hari Donepudi, Yupeng Li, Priya Raghunath, Babu Sundaram

**Yonghong Yan**, and Jiwen He

International Conference on Concurrent Engineering: Research and Applications, July 2003

### Technical Reports

Support for Specification and Scheduling of Workflow Applications on the Grid

**Yonghong Yan** and Barbara Chapman

Department of Computer Science, University of Houston, August 2006

Workflow Support in a GRACCE Metascheduling Architecture

**Yonghong Yan** and Barbara Chapman

Department of Computer Science, University of Houston, October 2005

Comparative Study of Distributed Resource Management Systems - SGE, LSF, PBS Pro, and LoadLeveler

**Yonghong Yan** and Barbara Chapman

Department of Computer Science, University of Houston, May 2005

Adaptive Execution Framework for Air Quality Modeling Application

Priya Raghunath, **Yonghong Yan**, Barbara Chapman and Jiwen He

Department of Computer Science, University of Houston, July 2003

Introduction to Tivoli Management Environment Framework

**Yonghong Yan**, Peinan Shao, Shaoqing Yan

System Software Group, East China Institute of Computing Technology, September 2000

### INVITED TALKS AND TUTORIALS

---

Parallel Programming Models for Heterogeneous and Manycore Architectures

(SICM)2 Parallel Computing Workshop, Software Institute for Chemistry, Materials, and Biomolecular Simulation, New York, NY, March 2014

REX: REthinking the "X" in the "MPI+X" for Exascale Nodes

Productive Programming Models for Exascale, PNNL

Portland, OR, August 14-15, 2012

Tutorial on OpenMP Parallel Programming

DoE CScADS Summer Workshops, July of 2011 and 2012

Introduction to OpenACC Programming on GPGPUs

DoE CScADS Summer Workshops, July 2012

Air Quality Forecasting on UH Campus Grid

Global Grid Forum 14, Chicago, IL, June 2005

Air Quality Prediction on the Grid

Sun Microsystems HPC Consortium, SC04, Pittsburgh PA, October 2004

## PROFESSIONAL ACTIVITIES

---

**Program Committee** of International Workshop on Programming Models and Applications for Multi-cores and Manycores (PMAM) 2015, 2016, 2017

**Program Committee** of International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), 2015, 2016, 2017

**Program Committee** of The 22nd IEEE International Conference on Parallel and Distributed Systems (ICPADS) 2016

**Program Committee** of SC15

**Program Committee** of Fourth International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing (WOLFHPC of SC14, SC15 and SC16)

**NSF Panel** and proposal reviewers, 2014

**Program Chair** of Programming Models, Languages and Compilers Workshop, for Manycore and Heterogeneous Architectures (PLC) at IPDPS 2013, 2014

**Program Committee** of ICPP 2012, PGAS 2013, IEEE CloudCom 2013

**Publicity and Session Chair** of PGAS 2011 Conference, Galveston, TX

**Reviewer** for several journals and conferences. Recent review activities include ACM International Conference on Computing Frontiers (CF), ACM International Conference on Supercomputing (ICS), Journal of Scientific Programming and Journal of Parallel Computing.

**OpenMP subcommittee member**, OpenMP ARB standard organizations

**ACM SIGHPC and IEEE member**

## STUDENT ADVISING

---

**Advised and graduated** one Ph.D. student and two master students

**Currently advising** three Ph.D. students and two master students

## TEACHING EXPERIENCE

---

### **Department of Computer Science and Engineering, Oakland University**

**Lecturer:** Data Structures with Object Oriented Programming II *Fall 2014, Fall 2015 and Spring 2017*

**Lecturer:** Concurrent and Multicore Programming *Spring 2015, Summer 2015, Spring 2016*

**Lecturer:** Graduate Computer Architecture *Fall 2016*

### **Computer Science Department, University of Houston**

**Co-lecturer:** Programming Languages and Compilers *Fall 2011, Fall 2012*

**Co-lecturer:** Distributed and Grid Computing *Fall 2004, Spring 2005, Fall 2005*

**Lab instructor:** Computer Architecture *Fall 2003, Spring 2004, Spring 2005, Spring 2007*

**Teaching assistant (both undergraduate and graduate levels):** *2002 - 2007*

Compiler and Programming Languages, High Performance Computing, Operating Systems, Computer Networks, Database, Data Structures, and Computer Architecture.

## HONORS AND AWARDS

---

**Graduate Assistant Tuition Fellowship** 2002-2004  
by Department of Computer Science, University of Houston

**Exxon Mobil Scholarship** 2003  
by Exxon Mobil Corporation

**Intel Qualification Award** 2001

**Intel Excellent Achievement for PD2 FCBGA Startup** 2001  
by Intel China

**Outstanding Graduating Student** 2001

<b>Scholarship for Outstanding Students</b>	1998-1999, 1999-2000
<b>Excellent Student Organization Leader</b> by University of Shanghai for Science and Technology	1999-2000, 2000-2001
<b>Scholarship for Outstanding Students</b>	7 times, 1994 - 1998
<b>Three-Good Students</b>	3 times, 1994 - 1998
<b>Excellent Student Organization Leader</b> by Shandong University of Science and Technology	1996-1997

## SERVICE AND LEADERSHIP

---

University of Shanghai for Science and Technology <b>Vice-Chairman of the Graduate Student Council</b>	2000-2001
Shandong University of Science and Technology <b>Chairman of the Student Association of Research</b> <b>Vice-Chairman of the Department Student Council</b>	1997-1998 1996-1997

## RESEARCH PROJECTS AND WORKING EXPERIENCE

---

**Computer Science Department, University of Houston** 06/2011 - present  
*Research Assistant Professor, HPCTools Group (<http://www.cs.uh.edu/~hpctools>)*

- OpenUH compiler research and development (<http://www.cs.uh.edu/~openuh>) at HPCTools Research Group
- PI of Image Processing Cloud (IPC) project funded through NSF. The UH's role in this project is to create a high-level domain specific language for image processing applications using MapReduce framework and on the cloud computing environments.
- UH Technical Lead of eXascale PRogramming Environment and System Software (XPRESS) funded by DoE. The UH's role in this project is to integrate OpenMP and MPI and runtime environments with the HPX ParallelX execution model implementations created by project partners from LSU and Indiana University.
- Project Lead of OpenACC Research Compiler project funded by Nvidia. The goal of this project is to create an open-source research compiler of OpenACC programming models for accelerators.
- HOMP (the ROSE OpenMP compiler support): the goal of this project is to support the latest OpenMP standard and to innovate on new language extension and implementation techniques in ROSE source-to-source compiler framework. This collaborative work between UH and LLNL will leverage the current ROSE OpenMP support and the OpenMP runtime in the OpenUH compiler.
- Led the Extreme OpenMP project funded by NSF. In this project, we have extended OpenMP with several features important for using OpenMP for HPC applications, including subteams, locations, data-driven tasks, and dynamic compiler optimizations. Some of those extensions are now under discussion to be included in OpenMP standard. The project completed August, 2012.

**Computer Science Department, Rice University** 07/2008 - 06/2011  
*Postdoctoral Research Associate with Prof. Vivek Sarkar, Habanero Group (<http://habanero.rice.edu>)*

- Led the design and implementation of Habanero-C (HC) compiler and runtime system. The HC language extends C with constructs for task parallel programming, dynamic phase-based synchronization and locality-aware data/task co-allocation. The HC compiler extends and uses the ROSE compiler framework for source-to-source translation. The HC runtime employs a work-stealing runtime. The HC system is being extended to support programming and scheduling tasks on heterogeneous and distributed computing systems, including GPU and FPGA and clusters of hybrid CPU/GPU/FPGA machines. This work is part of the NSF-funded Center for Domain Specific Computing (CDSC) project.
- Led the research and development of the runtime for manycore architectures in a collaborative NSF-funded project between Rice Habanero Group and the University of Delaware CAPSL group led by Prof. Guang R. Gao. We are using the IBM Cyclops64 manycore processors in this research.

- Research and development in the DARPA-funded Platform Aware Compilation Environment (PACE) project and NSF-funded Customizable Domain-Specific Computing (CDSC) project, including a unified runtime system for heterogenous architectures and compiler optimizations related to ROSE.
- High Productivity GPU Programming: Design and implementation of the JCUDA Java language extensions, library and compiler for programming GPU in Java. The JCUDA constructs for Java allows programmer directly write GPU calls from Java, and the JCUDA compiler is able to generate CPU-GPU bridge code and calls for data transfers. This work was published in the Euro-Par 09 conference proceedings.

**Computer Science Department, University of Houston**

05/2007 - 07/2008

*Information Technology System Administrator and Supervisor*

- At the time of Ph.D. graduation from UH, our CS department had an urgent need of a senior person to overall reorganize the department's IT infrastructure. There are two goals of this reorganization: first, to apply advanced computing technology, such as video conference and streaming, to provide state-of-the-art support for teaching and interdisciplinary research collaboration; second, to make more efficient use of university-wide lab computers and supercomputers for scientific and engineering applications using volunteer and grid computing technology. I was entrusted with this mission based upon my prior research in grid computing and my extensive practical experiences in managing research clusters and supercomputers.
- Virtual classroom: A video conference room has been set up using AccessGrid and related grid middleware. A remote person is able to interact with the people in the room not only via video/audio, but also to share a presentation, web browser, whiteboard, or even desktop. It can be used for remote teaching, seminar and defense presentations, and research collaborations.
- Virtual campus supercomputer (VCSC): A volunteer computing infrastructure has been set up using BOINC software. Campus lab computers, clusters and supercomputers can donate their computation power to our VCSC for computation-intensive applications that have been deployed. By utilizing high network bandwidth across campus, our VCSC provides much more computation power and supports a wider range of applications than world-wide volunteer computing that relies on low-bandwidth Internet.

**Sun Microsystems Laboratory, Menlo Park, CA**

06/2006 - 09/2006

*Research Intern in Embedded Solaris Project*

- Research and development of Solaris solutions for embedded computing devices, such as dynamic module load, link and scheduling for reduced power and memory footprint.
- Implementation of Solaris NFS boot on PowerPC architecture, including the whole TCP/IP protocol stack, BOOTP, RARP, TFTP and NFS protocol, the program interface to Open Firmware, and the Solaris kernel boot loader.
- Implementation of Solaris kernel runtime linker (krtld) for PowerPC Application Binary Interface (ABI), including krtld self relocation and linking, module loading, relocation and linking to kernel address space on demand, and cache flushing.
- Project URL: <http://www.opensolaris.org/os/project/ppc-dev/>.

**HPCTools Group, Computer Science Department, University of Houston**

01/2002 - 05/2007

*Ph.D. Research: Scheduling Scientific Workflow Applications in Computational Grids*

- The design and implementation of a workflow system for executing scientific workflow applications in grid environments, which includes:
  1. A workflow description language that supports for the descriptions of workflow structures and scheduling required details; implemented using XML-Schema.
  2. Workflow scheduling algorithms that apply such advanced scheduling techniques as resource advanced reservation, execution planning, performance prediction and resource co-allocations.
  3. A grid simulation environment for evaluating workflow scheduling algorithms and studying performance issues of applications.
- The work was supported by the Sun Microsystems Center of Excellence in Geosciences and in High Performance Application Development.

*File Pre-fetching for I/O Intensive Jobs in Cluster Schedulers*

- Enhancement of SGE scheduler with NFS file pre-fetching for parallel I/O intensive jobs.
- Linux kernel modification to support userland memory mapping.
- Application profiling for predicting file access patterns.

*Representative Course Project - MicroMIPS Processor Design*

- Design of an 8-bit MIPS processor in digital level, including ALU, cache, data and control logic, etc.
- Design of a 16-bit MIPS using Verilog; synthesis and verification using Synopsys EDA tools and Xilinx FPGA.

**Intel Technology (China) Ltd**, Shanghai China

01/2001 - 02/2002

*Project Lead and System Engineer*

- Development, deployment and support of applications for chip manufacturing, including business requirement analysis, application module design and development, product testing and deployment, training to customers, etc.
- Implementation and deployment of failover clusters with fail-safe RAID systems and Oracle servers.
- Oracle database administration and Oracle security administration.

**East China Institute of Computing Technology**, Shanghai China

12/1999 - 03/2001

*Research Intern in System Software Group*

- Design and implementation of ECIsys distributed system management tools using IONA Orbix CORBA implementation; and the integration of ECIsys with Tivoli.
- Deployment of IBM Tivoli computing resource management framework, including identity management, software deployment, security management, network management, etc.
- Software development and integration of applications for Tivoli using Tivoli CORBA SDK.

**University of Shanghai for Science and Technology**

09/1998 - 02/2001

*Master Research: A Common Object Request Broker Architecture(CORBA) Resource Management Framework for Distributed Computing*

- Design and implementation of an agent-based resource management system (ECIsys) using CORBA standard and C++/Java.
- Implementation of a cross-platform ECIsys interface for Windows, Linux and Solaris.
- Integration of ECIsys with IBM Tivoli Framework.

*Representative Course Project - An Integrated Board for Electrics and Electronics Experiments*

- Design of a PCB board using Protel that supports course experiments in college electrics and electronics courses.

**Shandong University of Science and Technology**

09/1994-07/1998

*Bachelor Research: Embedded Computer System for Automatic Bottle Filling*

- Design of an Intel 8086 controller board with RAM, IO, sensor interface and firmware.
- Integration of sensors for detecting bottle position, movement, and liquid flow rate.
- Design of a sequential and parallel interface card for Intel 8086 microprocessor.

## TECHNICAL SKILLS

---

**General Programming:** C/C++, Java, Perl, Fortran and x86/PowerPC assembly languages

**Compilers:** ROSE, Open64, LLVM and EDG

**Parallel Programming:** MPI, OpenMP, OpenACC, Cilk, Hadoop, and pthread

**Distributed and grid computing:** Globus, CORBA, Java RMI, and XML web services

**Kernel and firmware development:** Solaris kernel, IEEE 1275 Open Firmware

**System administration:** Solaris, Linux, Oracle DBA and Security, MySQL, IBM Tivoli, SGE/LSF, Postfix, DNS, DHCP, iptable firewall, router/switch and Apache



## REFERENCES FOR YONGHONG YAN

---

Barbara Chapman

*Professor, Computer Science Department, Stony Brook University*

IACS Building L158, Stony Brook, NY 11794-5250

Email: chapman@cs.stonybrook.edu

Phone: 631-632-2351

Home Page: <http://www.cs.uh.edu/~chapman>

Vivek Sarkar

*E.D. Butcher Professor of Computer Science, Rice University*

*Joint appointment, Department of Electrical and Computer Engineering*

*Associate Director, Center for Domain-Specific Computing*

*Adjunct Faculty, CSE department, Indian Institute of Technology, Kanpur*

*Visiting Professor, Waseda University, Japan*

*ACM Fellow*

MS-132 6100 Main St., Houston, TX 77005

Email: vsarkar@rice.edu

Phone: 713-348-5304; Fax: 713-348-5014

Home Page: <http://www.cs.rice.edu/~vsarkar>

Guang R. Gao

*Endowed Distinguished Professor, Department of Electrical and Computer Engineering, University of Delaware*

*Founder and Director of the Computer Architectures and Parallel Systems Laboratory (CAPSL)*

*IEEE Fellow, ACM Fellow*

140 Evans Hall, Department of Electrical and Computer Engineering, University of Delaware, Newark, DE 19716

Email: ggao@capsl.udel.edu

Tel: 302-831-8218; Fax: 302-831-4316

Home Page: <http://www.capsl.udel.edu/~ggao>

Bronis R de Supinski

*Chief Technology Officer (CTO) for Livermore Computing (LC) at Lawrence Livermore National Laboratory (LLNL)*

*Chair of OpenMP Language Committee*

*Adjunct Associate Professor in the Department of Computer Science and Engineering at Texas A&M University*

Lawrence Livermore National Laboratory, Box 808, L-554, Livermore, CA 94551-0808

Email: bronis@llnl.gov

Tel: 925-422-1062

Kirk W. Cameron

*Professor*

Department of Computer Science

Virginia Tech

114 McBryde Hall (0106), 225 Stanger Street, Blacksburg, VA 24061

Email: cameron@vt.edu

Phone: 540-231-4238

Home Page: <http://people.cs.vt.edu/~cameron>