

Parallel Programming on Cloud Platforms

— Hand Coding or Automatic Translation

Dr. Yi Pan

Regents' Professor of Computer Science
Interim Associate Dean and Chair of Biology
Georgia State University
Atlanta, Georgia, USA
<http://www.cs.gsu.edu/pan>

Abstract:

Cloud computing has gradually evolved into an infrastructural tool for many scientific and business applications with intensive data or computing requirements. One of the challenges in cloud computing now is how to run software efficiently on cloud platforms since lots of classic sequential codes are not ready to be executed in parallel in cloud environments, resulting in long execution time and low efficiency. It is also costly and labor intensive to redesign and convert current sequential codes into cloud codes running on cloud programming models such as MapReduce or Spark. Thus, automatic translation from sequential codes to cloud codes is one of the directions that could resolve the problem of slow code migration from traditional computing platforms to cloud infrastructures. In this talk, I will present several automatic translators (M2M, J2M and J2S) for cloud programming models MapReduce and Spark. I will provide details of the design of our translators and their performance results based on many experiments. Performance comparisons with hand coded cloud programs will also be studied. Our experimental results indicate that the translators not only can precisely translate the sequential codes such as MATLAB and Java codes into cloud codes, but also can achieve almost linear speedup in performance if the data sizes are big enough. In addition, their limitations and shortcomings will be identified and future directions in this area will also be provided.

Biography:

Yi Pan is a Regents' Professor of Computer Science, an Interim Associate Dean and Chair of Biology at Georgia State University, USA. He is also a visiting Changjiang Chair Professor at Central South University, China. Dr. Pan received his B.Eng. and M.Eng. degrees in computer engineering from Tsinghua University, China, in 1982 and 1984, respectively, and his Ph.D. degree in computer science from the University of Pittsburgh, USA, in 1991. His profile has been featured as a distinguished alumnus in both Tsinghua Alumni Newsletter and University of Pittsburgh CS Alumni Newsletter. Dr. Pan's research interests include parallel and cloud computing, wireless networks, and bioinformatics. Dr. Pan has published more than 180 journal papers with over 70 papers published in various IEEE journals. In addition, he has published over 150 papers in refereed conferences. He has also co-authored/co-edited 40 books.

His work has been cited more than 8000 times. Dr. Pan has served as an editor-in-chief or editorial board member for 15 journals including 7 IEEE Transactions. He is the recipient of many awards including IEEE Transactions Best Paper Award, several conference Best Paper Awards, IBM Faculty Award, JSPS Senior Invitation Fellowship, IEEE BIBE Outstanding Achievement Award, NSF Research Opportunity Award, and AFOSR Summer Faculty Research Fellowship. He has organized many international conferences and delivered keynote speeches at over 50 international conferences around the world.

