Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.

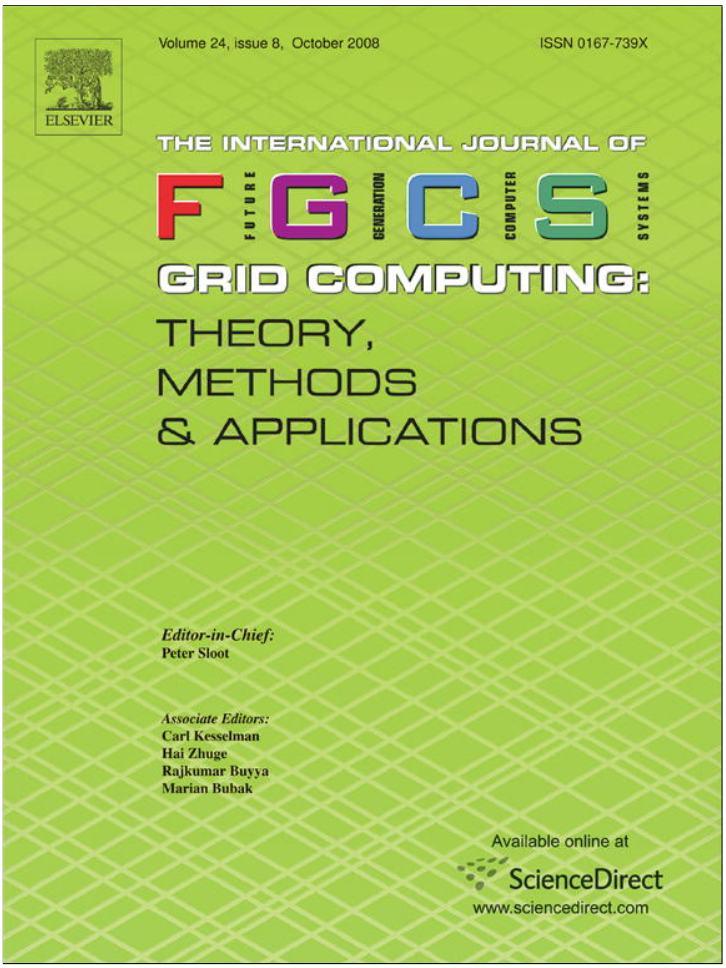

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright

Author's personal copy

Future Generation Computer Systems 24 (2008) 833

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Editorial

Special section: Scalable information systems

Xiaohua Jia [a], Francis C.M. Lau [b,*]

[a] Department of Computer Science, City University of Hong Kong, Hong Kong
[b] Department of Computer Science, The University of Hong Kong, Hong Kong

ARTICLE INFO

Article history:
Received 6 May 2008
Accepted 2 June 2008
Available online 7 June 2008

To give a description of a "scalable system" is trivial for many of us who work in the area of information/computer systems, but the fact is that "scalability" has been and remains as an illusive quality of many hardware and software systems because of the lack of a rigorous definition other than for very specific domains. Scalability in the past as well as today has invariably been taken as an indicator of good design, but experienced designers would all agree that a scalable design is not necessarily the best design for a particular set of operating conditions. A design with a vaguely defined requirement for scalability could easily turn into an overkill when the reality deviates from the conceived. Scalability is as good a quality as it is practically useful. Scalability for many years has been used to benchmark parallel systems and applications, and because of that it is almost synonymous with "performance". There are, however, other dimensions in system design where scalability may make practical sense and can serve as a tangible design goal. A speech-to-text system, for instance, may be labeled a scalable system if it can cope comfortably with a wide range of spoken languages.

This special section is derived from the collection of papers presented at the First International Conference on Scalable Information Systems (Infoscale 2006), held on 29 May–1 June, 2006 in Hong Kong. Being the inaugural event of a series of many conferences to come, Infoscale 2006 provided an open forum for researchers and practitioners to exchange their views on the meaning and implications of scalability and to come to grips with how scalability may translate into tangible benefits in real systems. The experiences shared during the conference will serve as good references for future efforts in designing scalable information systems. The conference featured 46 submitted or invited papers, of which the best ones were invited for re-submission to this special section. These extended versions of the original papers went through a new round of rigorous reviews, which resulted in the four papers being featured in this special section on scalable information systems. Interestingly, the four papers address four different kinds of popular computer systems, respectively: peer-to-peer systems, grid systems, wireless sensor/ad-hoc networks and the ubiquitous computing environment. The first paper, by Batko et al., presents several distributed scalable techniques for similarity search in peer-to-peer networks. The second paper, by Liu et al., presents an approximation method for mapping service level agreements in grid systems. The last two papers address the issue of security and protection. The paper by Li et al. proposes a device grouping approach and its associated key management scheme for wireless ad-hoc/sensor networks. The final paper of the special section, by Wang et al., describes an access control model for the protection of devices and services in an ubiquitous computing environment.

We congratulate the authors and thank all the reviewers who have helped to make this publication possible.

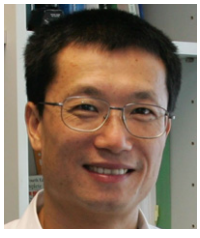

Xiaohua Jia received his BSc (1984) and MEng (1987) from the University of Science and Technology of China, and DSc (1991) in Information Science from Univ. of Tokyo, Japan. He is currently a Professor in Dept of Computer Science at City University of Hong Kong and Cheung Kong Professor with the School of Computing, Wuhan University, China. Prof. Jia is an editor of IEEE Trans. on Parallel and Distributed Systems, The Journal of Supercomputing, Journal of World Wide Web, Journal of Combinatorial Optimization, etc. He was a general Chair, PC Chair, PC member and OC member of many international conferences.

Francis C.M. Lau received his PhD (1986) in computer science from the University of Waterloo, Canada. He is currently a professor of computer science at The University of Hong Kong where he has served since 1987. He is a visiting professor of the Institute of Theoretical Computer Science of Tsinghua University, China.

* Corresponding author.
E-mail address: fcmlau@cs.hku.hk (F.C.M. Lau).

0167-739X/$ – see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2008.06.002