HPC & Big Data Convergence

The Cutting Edge & Outlook

Rashid Mehmood
King Abdulaziz University

@ IXPUG 2018, KAUST
13 March 2018

Data analytics and computing ecosystem compared

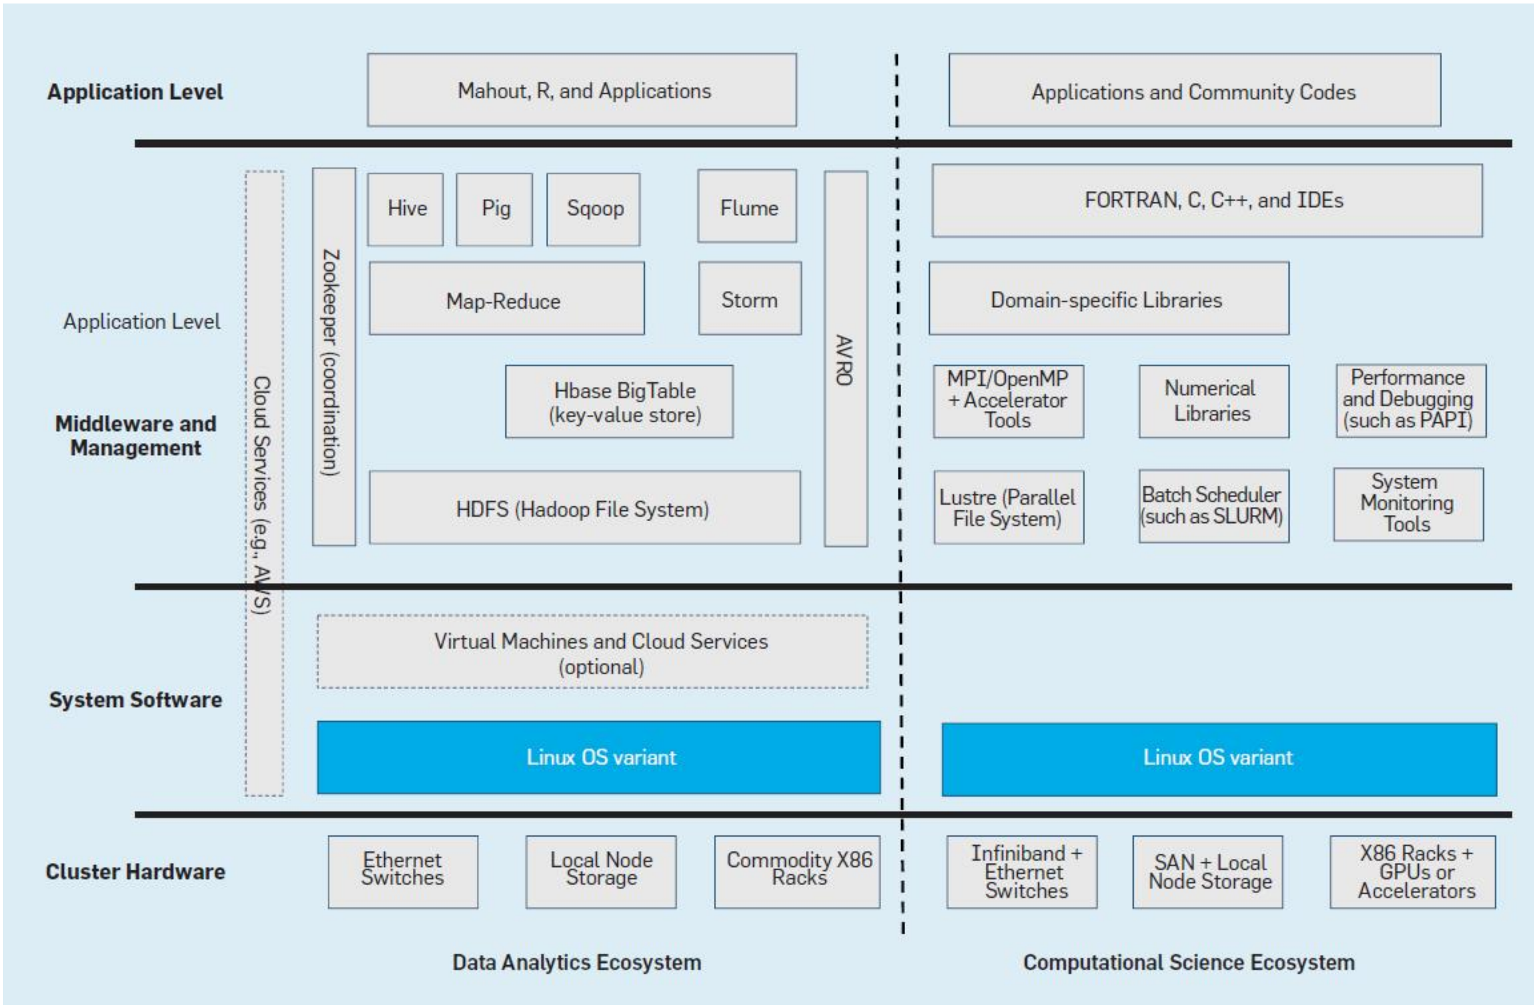

HPC and Big Data

- HPC technologies are needed by Big Data to deal with the ever increasing Vs of data in order to forecast and extract insights from existing and new domains, faster, and with greater accuracy.
- Increasingly more data is being produced by scientific experiments from areas such as bioscience, physics, and climate, and therefore, HPC needs to adopt data-driven paradigms.
- Moreover, there are synergies between them with unimaginable potential for developing new computing paradigms, solving long-standing grand challenges, and making new explorations and discoveries.

High Performance Data Analysis

Needs HPC resources
- High complexity (algorithms)
- High time-criticality
- High variability
- (On premise or in cloud)

Simulation & analytics
- Search, pattern discovery
- Iterative methods
- Established HPC users + new commercial users

Data of all kinds
- The 4 V's: volume, variety, velocity, value
- Structured, unstructured
- Partitionable, non-partitionable
- Regular, irregular patterns

IDC Analyze the Future

Global High Performance Data Analytics Market, by Region, 2014-2024 (in BN USD)

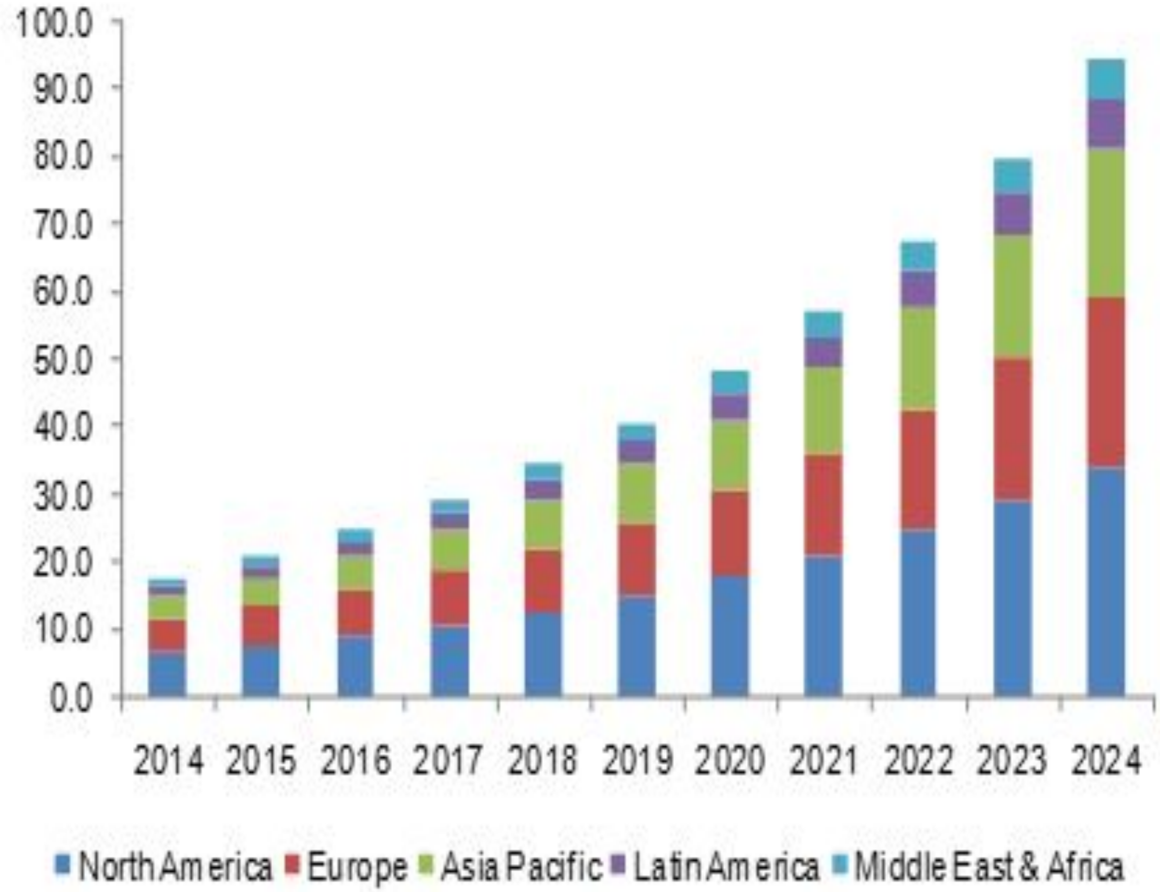

Attribute	Details
Base Year	2016
Historic Analysis	2014 & 2015
Forecast	2017 to 2024
Complete Free Customization*	Equivalent to 50 analyst hours

Big Data Analytics Workloads

Of all the workloads on your systems, what percentage of your total system utilization is devoted to Big Data analytics (MapReduce/Hadoop, graph analytics, semantic analysis, etc.)?			
	Percentage of System in Government	Percentage of System in Academia	Percentage of System in Industry
Currently used for Big Data analytics - Percent	18.2%	8.4%	32.6%
Expected 6 to 18 months from now - Percent	19.8%	15.6%	33.4%
N = 128			

Source: IDC, 2015

Technical Challenges for a Converged System

- Higher Energy Efficiency
  - Is the most crucial challenge
  - Sunway TaihuLight: 15.37 MW power consumption for 93.014 Pflops/s (Rmax), 125.435 Pflops/s (Rpeak)
  - Tianhe-2: 17.81 MW for 33.86 Pflops/s (54.9 PFLOPS peak)
  - Power and cooling technologies
  - circuit technologies
  - software
  - We need exascale performance while keeping the power consumption to the existing levels
  - Power aware algorithms and software would be another trend

Communications and networking

- Thousands of processes and billions of threads
- Fine-grained interaction between the processes is needed
- Advances in communication technologies have been slower compared to advances in processors
- High bandwidth and low latency technologies are needed
- Algorithms and software are required that localise and minimise communication
  - Use additional computing rather than communication

Challenges (memory and data locality)

- Cost of moving data is higher than a floating point operation
  - Energy cost of add operation is around 0.9nj, L1 to register is 1nJ, while moving data off-chip to register is around 100nJ
  - 28 to 40% of the total energy consumption is spent moving data
  - 19 to 36% of it is wasted in stalled cycles
- Bandwidth
  - the Xeon bandwidth is only up to $60GBs^{-1}$
  - IBM's Power8 bandwidth of up to $192GBs^{-1}$
  - NVIDIA's latest GPUs have over $280GBs^{-1}$
- Memory capacity, latency and bandwidth are critical for exascale computing
- New technologies such as stacked memory may help
- But memory per core is decreasing
- Design of algorithms requiring lower memory would be the trend
- Data locality would be among the most important goals
- Needs to be intrinsic in the programming models

Challenges (Fault Tolerance)

- HPC software are developed with the assumptions of low fault occurrence probability
- Typically proportional to the number of processes and interactions
- Billions of fine-grained processes in a computation will increase the likelihood of fault occurrence
- Big data systems are typically more fault-tolerant partly due to machine virtualisation, and failover technologies
  - But not to the desired level
- This will push towards requiring more loosely-coupled, intelligent, somewhat failure aware algorithms and software design
- Resilience needs to be intrinsic in programming paradigms

Challenges (4Vs of Big Data)

- Volume, Velocity, Variety and Veracity
- Large volume of data typically mean more memory, more computing, and more interaction
- Fast data requires faster I/O, memories, computational elements, workflows, applications and strategies for data management
- Big data may have many more files than the current file systems are able to deal with
- Restructuring and management of current scientific workflows is required to meet the current and future developments in HPC
- In situ data analysis would also needs to be integrated where efficient
- Variety and veracity require intelligent self-aware methods

Challenges (programming models)

- MPI (Message Passing Interface) is relatively tightly coupled and a typical choice among the HPC community
- BSP (Bulk Synchronous Parallel) and Map-Reduce are typically used by Big Data users
- Big data environments are more expressive and productive but offer lower performance compared to HPC
- New programming models are needed to express the fine-grained parallelism among millions of cores and billions of threads.
  - Plus heterogeneity of systems
  - Resilience
  - Data locality

Software and Hardware Gap

- The hardware has already changed at a much faster rate than the software
- Exascale level parallelism would create even more heterogeneity and change
- The effort in developing software is huge for a given architecture
- maintaining and adapting an existing software for new architectures is daunting
- Even more challenging is to adapt existing codes to fine-grained (cores) heterogeneous system environments
- Reformulation of scientific problems, algorithms and workflows are needed to move to exascale computing
  - E.g. compute rather than fetching data if possible

Additional Technical Challenges

- Correctness of algorithms and software
- Efficiency of the scientific process to set up experiments and use HPC/BD resources
- Usability and Impact is more important than benchmarks
- Application scientists are understandably more focussed on results and accuracy
  - and less on energy, system and workflow / process efficiency
  - Coordination to consider the bigger picture would improve efficiency, productivity and time to innovation

Current Convergence Efforts

Convergence Approach	Specific Solutions
HPC oriented MapReduce solution	Mellanox UDA, RDMA-Hadoop, DataMPI, Hadoop-IPoIB, HMOR
Hadoop on-demand on traditional HPC resources	myHadoop
HPC application's interface to HDFS	LibHDFS
Parallelization of many task applications with different workflow systems	MPI, ad-hoc Hadoop, CloudBlast, Spark, HTCondor
Overlapping of map, shuffle and merge phases	dataMPI
Map-reduce like framework for in situ data analysis	Virtualized Analytics Shipping (VAS), Spark on demand
Solutions to deal with massive amount of data in data intensive applications	iRODS, MapReduce-MPI, Pilot-MapReduce, SRM etc
Hybrid design to reduce I/O bottleneck in HDFS	(Triple-H)

Design Patterns Based Converged Future

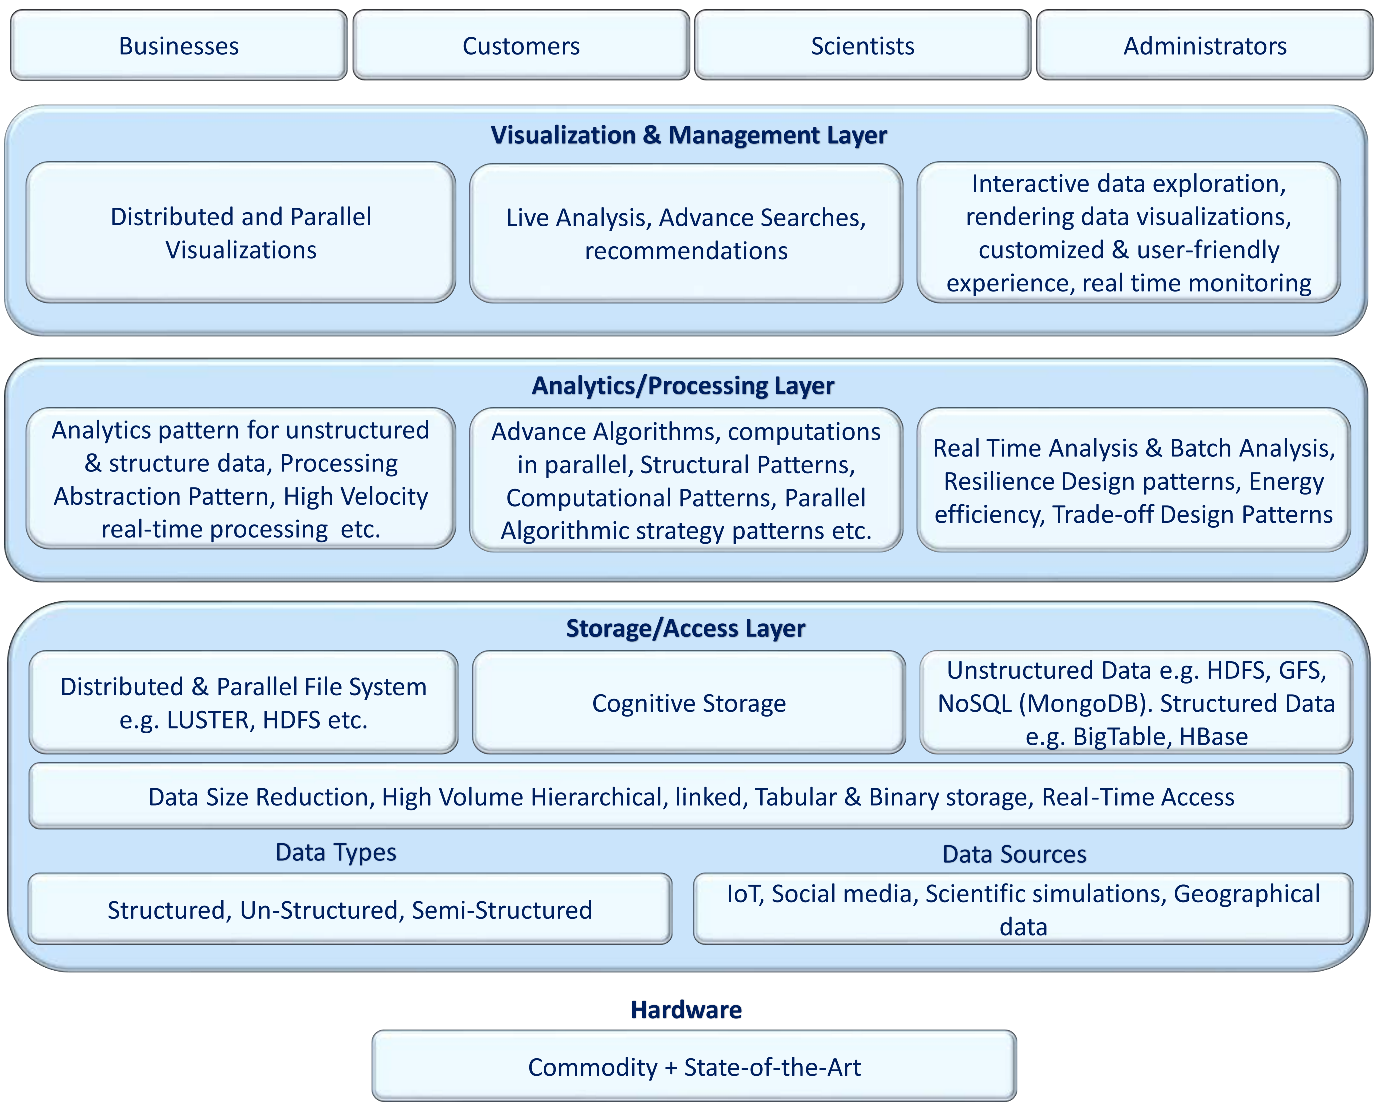

Acknowledgement

- Sardar Usman
- Furqan Alam

References

[1] Eric D. Isaacs. (2010, Nov.) Huffpost Chicago. [Online]. http://www.huffingtonpost.com/eric-d-isaacs/why-america-must-win-the_b_785652.html

[2] Wikipedia. Supercomputer. [Online]. https://en.wikipedia.org/wiki/Supercomputer

[3] Top500. http://www.top500.org/. [Online]. http://www.top500.org/

[4] BBC. (2015, Apr.) US nuclear fears block Intel China supercomputer update. [Online]. http://www.bbc.com/news/technology-32247532

[5] Giles M. B. and Reguly I., "Trends in high-performance computing for engineering calculations. ," Phil. Trans. R. Soc. A, vol. 372, no. 2022, 2014, http://dx.doi.org/10.1098/rsta.2013.0319.

[6] The White House, Office of the Press Secretary. Executive Order -- Creating a National Strategic Computing Initiative, 29 July 2015. [Online]. https://www.whitehouse.gov/the-press-office/2015/07/29/executive-order-creating-national-strategic-computing-initiative

[7] Robert F. Service, "Obama orders effort to build first exascale computer," Science, AAAS, July 2015, http://www.sciencemag.org/news/2015/07/obama-orders-effort-build-first-exascale-computer

[8] Lawrence Livermore National Laboratory Dona Crawford. (2016, Jan.) The Impact of the U.S. Supercomputing Initiative Will Be Global. [Online]. http://www.top500.org/blog/the-impact-of-the-us-supercomputing-initiative-will-be-global/

[9] Daniel A. Reed and Jack Dongarra, "Exascale Computing and Big Data," Communications of the ACM, vol. 58, no. 7, pp. 56--68, July 2015, http://doi.acm.org/10.1145/2699414

References

10. Niroshinie Fernando, Seng W. Loke, Wenny Rahayu, Mobile cloud computing: A survey, Future Generation Computer Systems, Vol. 29, Issue 1, pp 84–106, 2013.
11. Hoang T. Dinh, Chonho Lee, Dusit Niyato, Ping Wang, A survey of mobile cloud computing: architecture, applications, and approaches, Vol 13 Issue 18, 2013.
12. Scott Jarr, Fast Data and the New Enterprise Data Architecture. First Edition. October 2014. O'Reilly Media, Inc.
13. Bob Marcus, "Data Processing in Cyber-Physical Systems", January 2016
14. Barry Bolding, 5 Predictions for Supercomputing in 2016
15. https://www.hpcwire.com/2015/11/18/hpc-roi-invest-a-dollar-to-make-500-plus-reports-idc/
16. http://www.enterprisetech.com/2016/11/16/idc-ai-hpda-driving-hpc-high-growth-markets/?eid=328369061&bid=1593803
17. Saudi Arabia Invests US$70 Billion in Economic Cities Project, Cisco.
18. Technology Holds the Key to Success for Saudi Arabia's Vision 2030, Says IDC, 21 May 2016.
19. Big data essential to cancer 'moonshot', CIO, 11 May 2016. http://www.cio.com/article/3068571/government/big-data-essential-to-cancer-moonshot.html
20. Vice President Biden Says Better Data, Computing Make Cancer Beatable, CIO, 19 September 2016. http://blogs.wsj.com/cio/2016/09/19/vice-president-biden-says-better-data-computing-make-cancer-beatable/
21. http://qz.com/811199/apple-aapl-is-scaling-back-its-autonomous-car-ambitions-and-focusing-on-creating-self-driving-software/
22. "Is Hadoop the New HPC," 2016. [Online]. Available: http://www.admin-magazine.com/HPC/Articles/Is-Hadoop-the-New-HPC. [Accessed: 17-March-2018].
23. References list to be updated

References

- Eric D. Isaacs. (2010, Nov.) Huffpost Chicago. [Online]. http://www.huffingtonpost.com/eric-d-isaacs/why-america-must-win-the_b_785652.html
- Wikipedia. Supercomputer. [Online]. https://en.wikipedia.org/wiki/Supercomputer
- Top500. http://www.top500.org/. [Online]. http://www.top500.org/
- BBC. (2015, Apr.) US nuclear fears block Intel China supercomputer update. [Online]. http://www.bbc.com/news/technology-32247532
- Giles M. B. and Reguly I., "Trends in high-performance computing for engineering calculations. ," Phil. Trans. R. Soc. A, vol. 372, no. 2022, 2014, http://dx.doi.org/10.1098/rsta.2013.0319.
- The White House, Office of the Press Secretary. Executive Order -- Creating a National Strategic Computing Initiative, 29 July 2015. [Online]. https://www.whitehouse.gov/the-press-office/2015/07/29/executive-order-creating-national-strategic-computing-initiative
- Robert F. Service, "Obama orders effort to build first exascale computer," Science, AAAS, July 2015, http://www.sciencemag.org/news/2015/07/obama-orders-effort-build-first-exascale-computer
- Lawrence Livermore National Laboratory Dona Crawford. (2016, Jan.) The Impact of the U.S. Supercomputing Initiative Will Be Global. [Online]. http://www.top500.org/blog/the-impact-of-the-us-supercomputing-initiative-will-be-global/
- Daniel A. Reed and Jack Dongarra, "Exascale Computing and Big Data," Communications of the ACM, vol. 58, no. 7, pp. 56--68, July 2015, http://doi.acm.org/10.1145/2699414

References

- Niroshinie Fernando, Seng W. Loke, Wenny Rahayu, Mobile cloud computing: A survey, Future Generation Computer Systems, Vol. 29, Issue 1, pp 84–106, 2013.
- Hoang T. Dinh, Chonho Lee, Dusit Niyato, Ping Wang, A survey of mobile cloud computing: architecture, applications, and approaches, Vol 13 Issue 18, 2013.
- Scott Jarr, Fast Data and the New Enterprise Data Architecture. First Edition. October 2014. O'Reilly Media, Inc.
- Bob Marcus, "Data Processing in Cyber-Physical Systems", January 2016
- Barry Bolding, 5 Predictions for Supercomputing in 2016
- https://www.hpcwire.com/2015/11/18/hpc-roi-invest-a-dollar-to-make-500-plus-reports-idc/
- http://www.enterprisetech.com/2016/11/16/idc-ai-hpda-driving-hpc-high-growth-markets/?eid=328369061&bid=1593803
- Saudi Arabia Invests US$70 Billion in Economic Cities Project, Cisco.
- Technology Holds the Key to Success for Saudi Arabia's Vision 2030, Says IDC, 21 May 2016.
- Big data essential to cancer 'moonshot', CIO, 11 May 2016. http://www.cio.com/article/3068571/government/big-data-essential-to-cancer-moonshot.html
- Vice President Biden Says Better Data, Computing Make Cancer Beatable, CIO, 19 September 2016. http://blogs.wsj.com/cio/2016/09/19/vice-president-biden-says-better-data-computing-make-cancer-beatable/
- http://qz.com/811199/apple-aapl-is-scaling-back-its-autonomous-car-ambitions-and-focusing-on-creating-self-driving-software/
- References list to be updated

References

- Giffinger, Rudolf; Christian Fertner, Hans Kramar, Robert Kalasek, Nataša Pichler-Milanovic, Evert Meijers (2007). "Smart cities – Ranking of European medium-sized cities". http://www.smart-cities.eu/. Vienna: Centre of Regional Science.
- Rashid Mehmood and M. Nekovee, Vehicular Ad hoc and Grid Networks: Discussion, Design and Evaluation, In Proc of the 14th World Congress on Intelligent Transport Systems, October 2007
- SMART 2020: Enabling the low carbon economy in the information age. A report by The Climate Group on behalf of the Global eSustainability Initiative (GeSI). 2008.
- Nicholas Stern. Key Elements of a Global Deal on Climate Change, London School of Economics and Political Science. 2008. http://www.lse.ac.uk/collections/climateNetwork/publications/KeyElementsOfAGlobalDeal_30Apr08.pdf
- Nicholas Stern. Executive Summary, Stern Review on the Economics of Climate Change, HM Treasury. 2006.
- http://abhi-carmaniacs.blogspot.co.uk/2012/02/vehicular-ad-hoc-network.html
- http://mubbisherahmed.wordpress.com/2011/11/29/the-future-of-intelligent-transport-systems-its/
- R. Mehmood, *Disk-based Techniques for Efficient Solution of Large Markov Chains*, School of Computer Science, University of Birmingham, UK, October 2004
- R. Mehmood, J A. Lu, *Computational Markovian analysis of large system*, In Special issue on Intelligent Management Systems in Operations, Journal of Manufacturing Technology Management, Vol. 22, Issue 6, pp.804 – 817, 2011, DOI: 10.1108/17410381111149657
- Rashid Mehmood, Jie A. Lu. Computational Markovian analysis of large system. In Special issue on Intelligent Management Systems in Operations. Journal of Manufacturing Technology Management, Vol. 22, Issue 6, pp.804 – 817, 2011. DOI: 10.1108/17410381111149657

References

- N. Komninos, “Intelligent cities: Variable geometries of spatial intelligence,” Intell. Build. Int., vol. 3, no. 3, pp. 172–188, 2011
- Rashid Mehmood, Furqan Alam, Nasser N. Albogami, Iyad Katib, Aiiad Albeshri and Saleh M. Altowaijri, UTiLearn: A Personalised Ubiquitous Teaching and Learning System for Smart Societies, IEEE Access, March 2017
- Zubaida AlAzawi, Omar Alani, Mohmmad B. Abdljabar, Saleh Altowaijri, and Rashid Mehmood, A Smart Disaster Management System for Future Cities, In Proceedings of the ACM Workshop on Wireless and Mobile Technologies for Smart Cities (WiMobCity 2014), in conjunction with the 15th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2014), Philadelphia, USA, August 11-14, pp 1-10, 2014.
- Z. Alazawi, M. Abdljabar, S. Altowaijri, A. M. Vegni and R. Mehmood, ICDMS: An Intelligent Cloud based Disaster Management System for Vehicular Networks, Communication Technologies for Vehicles, Lecture Notes in Computer Science, Vol. 7266/2012, April 2011, DOI: 10.1007/978-3-642-29667-3_4
- Zubaida Alazawi, Mohmmad Abdljabar, Saleh Altowaijri and Rashid Mehmood. Invited Paper: Intelligent Disaster Response System based on Cloud-Enabled Vehicular Networks. 11th International Conference on Intelligent Transportation Systems (ITS) Telecommunications, Saint-Petersburg, Russia, August 2011. DOI: 10.1109/ITST.2011.6060083
- Rashid Mehmood and Jon Crowcroft, Parallel Iterative Solution Method for Large Sparse Linear Equation Systems, Technical Report UCAM-CL-TR-650, Computer Laboratory, University of Cambridge, October 2005, http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-650.html
- http://cities.media.mit.edu/
- Wikipedia. Smart City. [Online]. https://en.wikipedia.org/wiki/SmartCity
- A. Caragliu, C. Del Bo, and P. Nijkamp, “Smart Cities in Europe,” 3rd Cent. Eur. Conf. Reg. Sci., pp. 45–59, 2009.

Be Part of our Journey towards Realizing Saudi Vision 2030…

Thank you…

Rashid Mehmood
RMehmood@kau.edu.sa