ARTICLE IN PRESS

Measurement xxx (2013) xxx–xxx

Contents lists available at SciVerse ScienceDirect

Measurement

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

The convergence of measurement science and computer science: A scientific conversation

A.C.W. Finkelstein [a,*], L. Finkelstein [b]

[a] *Department of Computer Science, University College London, UK*
[b] *Measurement and Instrumentation Centre, City University London, UK*

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Measurement science
Computer science
Sensors
Software
Design

ABSTRACT

This short paper is based on the first named author's final scientific conversations with Ludwik Finkelstein, it forms part of what was an on-going discussion whose aim was to bring together our work and perhaps to contribute to the further development of systems engineering. The paper argues for the creation of a new design discipline founded upon the convergence of measurement science and computer science. The convergence of these disciplines proceeds from both an improved appreciation of the fundamental relationships between the disciplines and from the technological developments that have brought measurement and computing ever closer together. The paper points to the underlying relationships between measurement science and computer science and suggests that these relationships can form a principled basis for tackling the design challenges that technological convergence presents. The paper concludes by setting out an initial agenda for the new discipline.

1. Foundations

In earlier work [1–3] we have presented an account of measurement science that is based on the treatment of measurement as an information process and of instruments as information machines. This view is founded on the concepts and principles of the sciences of information and systems. Instruments receive a power or material flow from an object under measurement at the input, assign to it a symbol and carry out operations on the symbol. Information is carried by the magnitude, or attributes of the time variation, of a physical variable termed the signal. The principles underlying the analysis of signals and of signal processes are based on information theory. Systems engineering forms a basis for design and for the construction of instrument systems from component sub-systems.

This account directly suggests the relationship with computer science, which is the study of a related class of information machines and is equally founded on the concepts and principles of the sciences of information and systems. Computational systems store, transform and communicate information. These processes are realised, generally in software, by algorithms that are executed, ideally efficiently, by the computational system. Systems engineering forms a basis for design and for the construction of software systems from component sub-systems.

Despite the close relationships between the foundations of the two disciplines they have largely developed independently and the potential for syntheses between the disciplines has been ignored. We argue below that the commonalities between measurement science and computer science can form a basis for a new converged design discipline.

2. Design heuristics

The design of complex systems requires strategies to manage complexity so as to limit the number of potential solutions that require working through and evaluation.

* Corresponding author.
E-mail address: a.finkelstein@cs.ucl.ac.uk (A.C.W. Finkelstein).

0263-2241/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.measurement.2013.04.047

Please cite this article in press as: A.C.W. Finkelstein, L. Finkelstein, The convergence of measurement science and computer science: A scientific conversation, Measurement (2013), http://dx.doi.org/10.1016/j.measurement.2013.04.047

The means for achieving this are design heuristics. These heuristics provide a broad characterisation of the limitations of particular solution technologies, which permit a designer to rapidly prune the design space. The problem arises, of course, when, as the result of technological advances, the heuristics no longer hold true.

We believe that changes in computing and measurement technologies have fundamentally changed the game (our early analysis was set out in [4]), invalidating existing design heuristics and requiring a fundamentally new design approach. We set out initially our view of the game-changing transformations in each area.

2.1. Changes in computing technologies

Computation as a commodity resource. The price of computing has steadily reduced as both the cost of hardware – driven by Moore's Law, multicore architectures, competition and manufacturing efficiencies – and the cost of software – driven by improved development practices and tools – have decreased. Virtual machine technology has provided a separation between computational capabilities and hardware. We have now moved to a tipping point at which compute is a commodity resource, purchasable or, in the case of the cloud, rentable, very cheaply in whatever quantity is required.

Cheap large-scale data storage. Data storage has followed a similarly upward trend of increased capacity and reducing price. Vast amounts of storage can be made available locally or remotely – from the cloud. Cheap SSD data storage means that a robust and stable storage medium is available that can have a variable form-factor and a performance that is close to, and in some respects better than, hard disk drive.

Universal network connectivity. The availability of networking through Ethernet, 802.11 wifi, Bluetooth, 3G wireless and similar is ubiquitous and built into virtually all computing devices. A standardised networking software stack is available for all commonly used platforms and provides a simple interface to communication capabilities.

Framework of standards to support interoperation. Interoperation, meaning the ability for components of a, generally distributed, computing system to cooperate to provide a complex composite service, has conventionally been difficult to achieve. Increasingly however there are open standards for interfaces and data formats that allow systems to interoperate with relative ease. Interoperation is now largely a design rather than a technical challenge.

Mature software architectures to support distributed computing. The design of robust and scalable distributed computing systems has largely been a dark art. There are now a set of well understood software architectures, some built into programming frameworks, that support the construction of distributed software systems and can be used as a constructive basis for assuring the properties of the resultant systems.

2.2. Changes in measurement technologies

Novel sensor technologies. There is a wide range of novel sensor technologies, deploying for example, chemical, biological, optical and semiconductor components. Many of these sensors are capable of being used in environments where sensing has hitherto been difficult.

Cheap, potentially disposable, sensors. Low cost electronics, optical components and new materials have also given rise to a range of cheap sensors that can be deployed in a redundant manner and where reliability is no longer an issue because the sensor can be disposed of.

On board communications capability. These cheap, potentially disposable, sensors can be equipped with a low-cost integrated communications capability, connecting either to a standard 'commercial' network or an ad hoc peer-based network, generally wirelessly.

Flexible display. Until relatively recently the display of measurement data has been the most expensive part of an instrument system. Specialised components were required both for output and for input to control the display. Now cheap computer-based display with input by way of standard devices such as the mouse or touch screen are universally available. These permit a sophisticated visualisation of instrument data as well as offering alternative display modalities.

2.3. Implications

None of these changes should, by itself, be very surprising to the experienced systems engineer but the cumulative effects of the changes are worthy of attention. Existing design practice, notably in the area of instrumentation, is based on a set of invalid assumptions: that data processing is expensive; that accuracy and data filtering is important because you can only store and process a limited amount of data within a limited window of time; that communication requires a built-for-purpose network infrastructure and specialised protocols; that the construction of distributed computing to analyse data is difficult and cannot be made sufficiently robust for measurement and control applications; that sensors and displays are costly so that careful strategies must be devised in order to optimise their use. Removing these assumptions permits us to design new classes of systems and also demands a different kind of design discipline.

3. Vision

The new classes of system that we envisage are based on ubiquitous sensing, measurement and computing. The basic idea is that environments will be saturated with cheap, disposable, sensors that can communicate with each other and with a ubiquitous network infrastructure. They will draw on a pervasive computing 'fabric' that will permit 'transparent' ultra-large scale data storage and effectively limitless compute power. Display on mobile devices and tablets will be at any point with rich interactivity and the ability to drill down into data. All the physical elements of the environment will be trackable through tagging and geolocatable.

The paradigm examples for ubiquitous sensing, measurement and computing are in so-called 'smart cities' [5]. Data from across a large range of embedded sensors for instance in cars, roads, public transport, mobile devices carried by pedestrians, street furniture and so on, can be collected, fused and analysed to optimise routing traffic control and energy use. Packaging that can sense environmental conditions and inform the consumer of the state of the food it protects but can also be tracked through the logistics chain to reduce waste and improve sustainability. These examples are not in fact particularly far-fetched.

4. Challenges

Ubiquitous sensing, measurement and computing systems present particular design challenges that lie at the intersection of measurement and computer science. The challenges arise in significant part from the scale and openness of the systems. These are ultra-large scale involving potentially many tens of thousands of sensors and hundreds of thousands of computing devices spread across large distributed networks. The systems have no centralised control and heterogeneous sensors and computing devices may join and leave. The computing fabric is unreliable and it cannot be assumed that nodes – either sensors or computing devices – are trustworthy (for example a vehicle may falsely report its position and speed). Many of the sensors and mobile computing devices will use batteries or more radically will be power scavenging. Ensuring that the systems will be 'globally' energy efficient is critical.

5. Agenda

These challenges and the realisation of the kinds of system described in the vision above require new design strategies that bridge measurement science and computer science.

Machine learning: It is now often cheaper to acquire and store data than it is to make a real-time (or possibly even a design-time) decision as to whether to keep the data or not. This is particularly the case when this decision must necessarily be taken locally but the data may have non-local relevance. Once having accumulated the data it is amenable to machine learning, which is the use of algorithms that can learn properties of interest from observations [6]. Machine learning techniques can both predict, having learned properties from training data or discover unknown properties from data. Recent developments in computer science have introduced new classes of algorithms and developed a basis for understanding the properties of these algorithms. The real breakthrough since our early observations on this topic [7] has been the development of widely available software that essentially packages machine learning techniques, which have hitherto required substantial expertise, and makes them readily available to systems developers. It is clear that machine learning allows a balance to be struck between measurement and post hoc analysis and prediction. Making this balance requires a new theoretical framework which brings together measurement and computer science.

Adaptive system design: In a large-scale system of the kind we envisage the environment can change dynamically and we would want the system to adapt to these changes while maintaining the overall system properties [8]. In other words we would want the sensor system to detect unanticipated changes in the environment and for the software system semi-autonomously to deploy new behaviours in response to these changes. The design of such adaptive systems lies precisely at the meeting point of measurement science, computer science and control theory.

Model based measurement: Measuring systems can only sense active, or directly observable, quantities, those that characterise the flow of energy and matter. Quantities, which are passive, concerned with storage, transformation or transmission of energy or matter, can only be observed by a process of inferential measurement. Inferential or model based measurement entails interrogating the system under observation by exciting it and estimating the quantities under observation as parameters of a model of the system. This problem of system identification is a specifically measurement problem. The efficient construction of such models in a large distributed systems setting is however a problem which computer science must address.

Analysis and estimation of errors: The theory of errors and uncertainty in measurement remains a core and distinctive problem of measurement science. There are specific problems associated with the analysis and estimation of errors and uncertainty in digital systems and processing. In open systems with heterogeneous and unreliable computing nodes and sensors this problem comes clearly to the fore because the characteristics of the nodes and sensors may not be known in advance of their deployment nor can it be assumed that the data can be trusted. Thus the analysis of errors and uncertainty from measurement science is brought into conjunction with the problems of security and trust from computer science.

Semantic metadata: It is now possible to store very large quantities of data from across multiple data sources, of which many may be sensors, and to use this data in different ways and in different interpretive contexts. To ensure that this can be done while preserving the integrity of the information it is essential to associate with the data appropriate semantic metadata that supports the accurate interpretation of its meaning [9]. This has been the subject of considerable work in computer science under the broad, and somewhat misleading, heading of 'semantic web technologies'. Measurement science directly addresses the principles of assigning symbols to objects and events of the real world in such a way as to describe them. Representing these assignments alongside information that makes the provenance and characteristics of measurement information explicit constitutes a further important challenge at the meeting point of computer science and measurement science.

6. Conclusion

The virtual, the domain of computer science, and the physical, the domain of measurement science, are touching in the new generation of large scale 'cyber-physical systems'. A new design discipline is required if the full

Please cite this article in press as: A.C.W. Finkelstein, L. Finkelstein, The convergence of measurement science and computer science: A scientific conversation, Measurement (2013), http://dx.doi.org/10.1016/j.measurement.2013.04.047

potential of these systems is to be realised. There is a significant agenda of demanding scientific problems whose solution requires interdisciplinary collaboration.

Acknowledgements

This paper was assembled by Anthony Finkelstein from shared notes of his last scientific conversations with Ludwik Finkelstein who died in August 2011. Ludwik Finkelstein would undoubtedly have wished to acknowledge the intellectual contributions of his colleagues and his students. Anthony Finkelstein acknowledges his own profound intellectual and personal debt to Ludwik Finkelstein.

References

[1] L. Finkelstein, State and advances of general principles of measurement and instrumentation science, Measurement 3 (1) (1985) 2–6.

[2] L. Finkelstein, Measurement and instrumentation science an analytical review, Measurement 14 (1) (1994) 3–14.

[3] L. Finkelstein, Measurement science state and trends, in: B. Zajc, A.Trost (Eds.), Proceedings of the Sixteenth International Electrotechnical and Computer Science Conference, Vol. A of ISSN 1581-4572, 2007, pp. I–IV.

[4] L. Finkelstein, Measurement and instrumentation technology – reflections on the present and future, Measurement + Control 35 (2002) 110–113.

[5] D. Gann, M. Dodgson, D. Bhardwaj, Physical–digital integration in city infrastructure, IBM Journal of Research and Development 55 (1.2) (2010) 8.1–8.10.

[6] D. Barber, Bayesian Reasoning and Machine Learning, Cambridge University Press, 2012.

[7] L. Finkelstein, D. Hofmann, Intelligent measurement – a view of the state of the art and current trends, Measurement 5 (4) (1987) 151–153.

[8] B. Cheng, R.d. Lemos, H. Giese, P. Inverardi, J. Magee, J. Andersson, B. Becker, N. Bencomo, Y. Brun, B. Cukic, G.D.M. Serugendo, S. Dustdar, A. Finkelstein, C. Gacek, K. Geihs, V. Grassi, G. Karsai, H. Kienle, J. Kramer, M. Litoiu, S. Malek, R. Mirandola, H. Muller, S. Park, M. Shaw, M. Tichy, M. Tivoli, D. Weyns, J. Whittle, Software engineering for self-adaptive systems: a research roadmap, Software Engineering for Self-Adaptive Systems, vol. 5529, Springer, 2009, pp. 1–26.

[9] N. Shadbolt, W. Hall, T. Berners-Lee, The semantic web revisited, IEEE Intelligent Systems 21 (3) (2006) 96–101.