The era of computational abundance – and its risks

The advent of the internet of things and the explosion of computational power it portends will transform banking information systems. David Rowe warns, though, that many institutions will be constrained by a reluctance to surrender decades-old assumptions

As a user of computing resources for almost 50 years, I have seen changes that were unimaginable in prospect and are still dramatic in retrospect. Nevertheless, throughout my career I have always felt I could rely on one iron-clad proposition, namely, ‘Computers are never fast enough’. It seemed clear to me that the human imagination would always run well ahead of the ability of computers to keep pace. In effect, we were operating in a world of computational scarcity, and always would be.[1]

At some level, I feel sure this Law of Computational Scarcity will continue to apply. Nevertheless, within my limited focus on the domain of information systems for economics, finance, accounting and risk management, I have become convinced that I must surrender the conviction of a lifetime. In this domain, I believe we are entering an Era of Computational Abundance (see pages 12–16).

David Rowe is senior strategist for risk and regulation at Misys in London.
Email: david.rowe@misys.com

A recent straw in the wind was Google’s announcement that it is abandoning the use of Captchas – distorted text that is used online to distinguish humans from web robots. It seems the robots can decipher even the most difficult Captchas 99.8% of the time. If we contemplate the challenge of controlling fleets of driverless cars, where the image to be deciphered is rapidly changing, the massive computing power required becomes evident.

While growth in the internet of things – devices that are connected to the web, such as driverless cars – will likely be constrained by limited computing capacity, the resources needed to make this a reality at all will swamp traditional computing domains such as financial risk management. A relevant recent example of this phenomenon is the revolution in international voice communications. Fifteen years ago, international phone calls were notably expensive; today, they are practically free on Skype, Apple FaceTime, Vonage and other alternatives. How did this happen? In effect, we built a huge global communication capacity geared to support the new market of streaming video. In this context, voice became almost a rounding error. It could be a virtual giveaway relative to the massive expansion in global communication capacity.

As global computing capacity expands exponentially, driven by things talking to each other rather than to people, traditional business computing will start to take on a role similar to voice communications over the past 15 years. It will be swimming in an ocean of computational abundance. This raises important technological and sociological difficulties.

Virtually all existing business computing systems have been built around an architecture rooted in the mentality of computational scarcity. Minimising storage requirements and CPU cycles has always been a key goal for system architects and developers. This has a serious downside. A crucial technique for achieving computational parsimony is to embed information in formal structures – the schemas that underlie relational database designs are pervasive examples of this. Such schemas economise on storage and CPU requirements, but effectively mask important metadata needed to consolidate the underlying information across databases with multiple, inconsistent schemas.

An architecture based on computational abundance will look very different and can only be implemented incrementally. During the transition, these two architectures will have to operate in tandem. One way this can work is to build enterprise data stores on the basis of the new realities. This would employ traditional extract, transform and load tools to access data from source systems but store the results in transparent, self-describing documents rather than in a complex and inflexible relational data warehouse.

At least as big a problem will be the sociological one. Hardly any IT professional under 50 years of age has stored data in any fashion other than in relational databases. This deeply ingrained bias will be hard to overcome. In the 1960s, it was accepted wisdom that, ‘No-one is ever fired for buying IBM’. Today it is accepted wisdom that, ‘No-one is ever fired for buying Oracle database licences’. The effectiveness of relational databases within narrow applications for which they are designed reinforces this bias.

Top management needs to impress the importance of enterprise-wide data transparency on its IT staff. Those that do not will risk being dangerously slow in reaping the benefits of computational abundance. **R**

---

[1] *Many will remember computer pioneer Grace Murray Hopper, who loved to scold her students about the egregious carelessness in wasting a microsecond (one millionth of a second) of computer time. She illustrated the difference by comparing a nanosecond to a one-foot-long piece of wire, with a microsecond as a 1,000-foot-long piece of wire, these being how far an electronic signal could travel in a nanosecond and a microsecond respectively.*