PROFESSOR JACEK GONDZIO

Managing the unmanageable

Through an ambitious project, **Professor Jacek Gondzio** and his collaborators are developing the crucial tools which will allow us to manage an almost exponentially increasing morass of data

How did your interest in the theory and implementation of large-scale optimisation methods come about?

While studying electronics, I was exposed to a number of interesting applications of optimisation across areas such as control theory or the design of very large-scale integrated microprocessors. The solutions to such problems often require a combination of appropriate mathematical modelling with the use of well-crafted optimisation algorithms; it is a discipline I find both challenging and fascinating.

Can you outline the main aims and objectives of your current research engagement?

We must rise to the challenge created by the excessive amount of data which is currently created and stored. It requires that we propose new methods for its analysis and the retrieval of useful information from it. We 'Google' information several times a day to find news, check the weather at a destination, find the result of an interesting sport event, etc., and we enjoy the fact that there is ingenious mathematics underlying Google's search engine, which almost instantaneously provides us with high-quality responses to our queries.

Our research team has been designing new algorithms which can learn from vast digital data resources. In particular, we are focused on classification and ranking problems, and new methods which attempt to condense information, ie. reduce its dimension.

In your programme dealing with accessibility, how are you approaching the issue of increasing quality of service while regulating access to digital resources?

We are developing new models of queueing systems. We assume that there are different classes of customer arriving in the queueing system and that they impose positive externalities on each other; the arrival of a new customer in the queue is beneficial for customers of a different type. For instance, in an employment portal the arrival of a new jobseeker benefits the employers, and the arrival of a new employer increases the number of job offers, which benefits the employees. My colleague Dr Burak Büke is an expert in queueing systems and is leading developments in this area.

Are you and your collaborators disseminating this important work throughout the community?

Dr Peter Richtárik has been successful in attracting a UK Engineering and Physical Sciences Research Council (EPSRC)/Defence Science and Technology Laboratory (DSTL) grant, and he and our project partner SAS have been involved in another new successful EPSRC bid. Dr Jakub Marecek – a postdoctoral research associate working on the grant – has been given a permanent position at IBM Research Dublin, while Martin Takac – a PhD student working with Richtárik on the grant – has won several national and international prizes for his posters, talks and papers, both as runner-up in the Best Student Paper Award and a finalist in the INFORMS Computing Society and the 16th IMA Leslie Fox Prize in Numerical Analysis.

Richtárik has organised two very successful workshops, the first being Advances in Large-Scale Optimisation, which took place in Edinburgh in May 2012, and also Optimisation and Big Data, which happened in May 2013. Both gathered leading optimisation researchers from around the world. Richtárik has been invited to become a long-term visitor for Theoretical Foundations of Big Data Analysis, a semester-long programme running at the newly funded Simons Institute for the Theory of Computing at the University of California, Berkeley.

I have been invited to give plenary talks at two major optimisation conferences in 2013 – the 11th EUROPT Workshop on Advances in Continuous Optimisation, which takes place in Florence, Italy, June 2013, and the XIII International Conference on Stochastic Programming, in Bergamo, July 2013.

You are now approaching the end of the second year of the project, have you made any noteworthy advances?

Büke has been working on a novel queueing model to analyse the performance of web portals popular in British society, while I have been working on matrix-free optimisation methods and their applications in solving very large problems which are intractable by other approaches. Richtárik has made major contributions to the theory of coordinate descent methods and has helped to establish them as one of the leading algorithms in big data optimisation. He has also written two packages of efficient code which are available online.

The digital deluge

Mathematics for Vast Digital Resources will provide meaningful analysis of ever-growing digital data resources using the latest queuing theory and optimisation techniques

HUMANITY IS WITNESSING a global and accelerating digitisation of government, economy and society. Enjoying unprecedented levels of electronic connectivity, consumers now routinely shop online for everything from books to jobs to houses, searching for information and building social networks, while governments seek to provide faster, more efficient, cheaper and more comprehensive services to citizens. Bricks-and-mortar businesses have become digital stores, using websites, targeted advertisements and delivery services to eliminate traditional geographical and logistical barriers to customer access.

Whilst a connected online world has many benefits, the main drawback is that humanity must find a way to produce, store and manipulate huge quantities of data. This unprecedented and urgent challenge requires new science to organise and interpret data from sources like the internet, commercial databases, scientific experiments and government records.

OPTIMAL STRATEGY

Central to these efforts will be techniques from the mathematical discipline of optimisation – a key consideration in any decision-making process and a concept whose real-life applications are diverse and manifold. Examples might include choosing the optimum route to minimise distance or travel time to a destination, selecting a portfolio to maximise the return on an investment, or designing the most aerodynamic shape of an aircraft or car.

Hence, whilst the web presents huge opportunities for boosting economic growth and improving quality of life, current and future generations of researchers will have to rise to the considerable challenge of creating the tools to best capitalise upon these opportunities, using optimisation techniques to manage an almost overwhelming torrent of data.

ENTIRELY NEW PROBLEMS

At the University of Edinburgh's School of Mathematics, three researchers from the Edinburgh Research Group in Optimization (ERGO) – Professor Jacek Gondzio and colleagues Drs Burak Büke and Peter Richtárik – coordinate Mathematics for Vast Digital Resources, a project investigating and developing methodologies for addressing these new challenges and, at the same time, maximising opportunities arising from the increasing vastness in both the size and accessibility of digital resources. With the implicit acknowledgement that current approaches to optimisation are unfit for dealing with emerging applications, the researchers aim to generate novel mathematical insights into underlying problems in the digital economy, providing industry and society with tools to address them and meet public expectations for the next decade and beyond.

"Recently, we have observed an exceptional increase in digital data collection – the amount of data we produce and store in just one hour today is comparable to the total our species produced throughout its whole existence until 2003," Gondzio reveals. "The need to process this staggering amount of data creates entirely new problems for mathematics and computer science."

INVALUABLE PARTNERSHIPS

Analysing this ocean of information demands the creation of new optimisation tools capable of working with huge volumes of data and retrieving useful information, and which ideally access data no more than a few times and provide answers quickly enough to be used online. Collaborating with industrial partners and digital economy hubs to ensure impact, Gondzio's team is analysing the mathematical properties of these problems and designing novel techniques to exploit their structure and implement them in efficient algorithms.

The group comprises broad expertise across interior point methods, first-order gradient-based methods and queueing systems, so is well-equipped to respond to challenges relating to the size and accessibility of vast digital resources. A number of external collaborators from industry include Yahoo!, SAS, Orange/FT and Stochastic Solutions – modern companies whose businesses are focused on the gathering and processing of information. As such, they are invaluable sources of real-life problems for the researchers to tackle: Yahoo! built its business around their web search engine; SAS brings numerous software tools to the table designed for sophisticated statistical analyses; Orange/FT has access to large amounts of customer data and has an interest in its intelligent processing; and Stochastic Solutions is seeking novel approaches to data mining. All of the parties involved have an interest in effectively managing large portfolios of data in order to improve their customers' experience.

DESCENDING ON SOLUTIONS

Richtárik is an expert in first-order methods and has been leading the development of coordinate descent methods, the logic of which is based on the idea that a function of many variables may be minimised by a series of steps along each coordinate independently. While seemingly simplistic,

MATHEMATICS FOR VAST DIGITAL RESOURCES

OBJECTIVES

- To construct and analyse mathematical models of vast digital resources using techniques at the interface of modern queueing theory and optimisation, which are capable of dealing with the unprecedented scale of modern digital resources and will upgrade the access management to these new resources
- To gain new mathematical insights into the underlying problems in the digital economy and provide industry and society with new tools to address these problems appropriately to meet the public's expectations over the next decade

TEAM MEMBERS

Professor Jacek Gondzio, School of Mathematics, University of Edinburgh, UK

Dr Burak Büke, School of Mathematics, University of Edinburgh, UK

Dr Peter Richtárik, School of Mathematics, University of Edinburgh, UK

PARTNERS

Yahoo!

SAS

Orange/FT

Stochastic Solutions

FUNDING

Engineering and Physical Sciences Research Council (EPSRC) grant EP/I017127/1

CONTACT

Professor Jacek Gondzio
Principal Investigator

School of Mathematics
University of Edinburgh
James Clerk Maxwell Building
King's Buildings
Edinburgh
EH9 3JZ, UK

T +44 131 650 8574
F +44 131 650 6553
E j.gondzio@ed.ac.uk

PROFESSOR JACEK GONDZIO works on the theory and implementation of algorithms for optimisation. He is best known for his contributions in the area of interior point methods. Gondzio applies these methods to solve various very large-scale optimisation problems arising in finance, telecommunications, signal processing and machine learning.

Whilst a connected online world has many benefits, the main drawback to this new paradigm is that humanity must find a way to produce, store and manipulate huge quantities of data

this approach can solve problems with a high degree of separability between coordinates. Consequently, coordinate descent methods are well-suited to dealing with problems which exhibit a high degree of data randomness, translating into limited connectivity between coordinates, and constitute an attractive approach for certain classes of easy machine learning problems. As they are relatively simple to implement – their basic steps involve a small subset of coordinates, possibly even a single coordinate, which guarantees that they do not need any sophisticated linear algebraic operations – they are leading algorithms in big data optimisation. This means that the extra memory needed to implement the optimisation method is negligible, making this an attractive approach for application to certain huge-scale optimisation problems.

OOPS

The project is also relying on the earlier-developed Object-Oriented Parallel interior point Solver (OOPS) which was used to solve the largest quadratic problem ever solved with a direct method. "OOPS is a modern implementation of the primal-dual interior point method," Gondzio explains. "It is able to exploit any block structure of the problem, so is able to solve problems which originate from modelling spatial or inter-temporal relations, such as uncertainty."

Optimisation problems of this kind commonly appear in financial planning and other sectors of the economy, as well as telecommunications and power systems, such as electricity transmission. OOPS accesses the problem's data and performs its computations in parallel, and has already demonstrated almost perfect speed-ups when run on high performance computers.

DEFIANCE

Naturally, new algorithms demand novel ways of proving convergence and analysing their worst-case complexity, and the development of some of these new techniques is fraught with difficulty. "Of course we have encountered difficulties," Gondzio admits. "Some of them may be overcome by exploring alternative algorithm analysis techniques, others defy current knowledge."

The size of the optimisation problems which need to be tackled has been growing incredibly quickly over the last few years, and the coming decade will only see their size increase further and faster. Some of the 'easier', more straightforward examples of these large problems – especially those which display a high degree of randomness and yield a high level of separability in data – can be solved, but others stubbornly continue to rebuff existing optimisation methods. "New techniques have to be developed to deliver solutions to these problems, so optimisation remains a vital area of computational mathematics," Gondzio concludes. "Ultimately, the quality of optimisation algorithms is determined by their ability to solve real-life problems, and their reliability, robustness and practical speed in delivering solutions."

ERGO

Centred around the Operational Research and Optimization research group in the School of Mathematics at the University of Edinburgh, the Edinburgh Research Group in Optimization (ERGO) is a loose association of researchers at the University's Management School, Institute for Energy Systems and the Agriculture/Institute of Ecology and Resource Management department, together with the Optimization group at the University of Dundee, and the Numerical Analysis groups at the Oxford University Computing Laboratory and the Rutherford Appleton Laboratory.