

<p>Institution: University of Cambridge</p>
<p>Unit of Assessment: UoA10</p>
<p>Title of case study: Design and optimization of communication networks</p>
<p>1. Summary of the impact (indicative maximum 100 words) This case study provides an account of work on a mathematical framework for the design and optimization of communication networks, and some examples of the framework's influence upon the development of the network congestion control schemes that underlie modern communication networks, notably the Internet. The impact on protocol development and on network architectures has been significant; in particular on the development of congestion control algorithms and multipath routing algorithms that are stable and fair. Several of the insights on large scale system behaviour have been transferred to help understand cascading failures in other large scale systems, including transport infrastructures.</p>
<p>2. Underpinning research (indicative maximum 500 words) In the 1980s, Professor Kelly (Professor of Mathematics of Systems at the Department of Pure Mathematics and Mathematical Statistics, University of Cambridge 1990 to present) and his collaborators made important contributions to the design of routing algorithms for circuit-switched networks such as telephony networks. In the early 1990s Professor Kelly was involved with the considerable international research effort between telecommunication companies and academics looking for ways to design what were then called multiservice networks, using developments of the models which had been so successful for telephony. It was widely thought in the telecommunications world that packet-switching (as used in the then miniscule Internet) could not scale to support systems comparable in size to the telephony network. This was a view not shared in the computer communication world, but the two worlds spoke different languages. From 1997-2005 Professor Kelly and his co-authors published a series of papers [1-6], based on their research in that period in the University of Cambridge. These papers provided a mathematical framework showing that the important distinction between circuit- and packet-switched networks was not whether bits are arranged in streams or files, but between different control architectures. The framework made clear, in a precise language comprehensible to both worlds, that the different control architectures corresponded to either a primal or a dual approach to the solution of a certain optimization problem. Algorithms based on the two approaches were shown to be stable about a system optimum characterized by a proportional fairness criterion. Stability was established by showing that, with an appropriate formulation of the overall optimization problem, the network's implicit objective function provided a Lyapunov function for the dynamical system defined by the control algorithm. Both classes of algorithm were generalized to include routing control, and provided natural implementations of proportionally fair pricing. The framework allowed natural treatments of stochastic effects and of time delays, with simple and scalable conditions for fairness and stability that are especially important for wide-area networks where links vary in capacity by many orders of magnitude, as do the propagation delays that are a consequence of geographical diversity and the finite speed of light. The framework made clear that there was no <i>in principle</i> reason why a (primal) architecture based on end-system control (such as the Internet) could not scale to arbitrary size. Research on communication networks is a subtle dialogue conducted over decades and involving many disciplines. Mathematical models provide the language that allows the disciplines to speak to each other. The key importance of end-system control was apparent to the pioneering engineers of the Internet. Indeed it was Jacobson's reference to earlier work of mathematical scientists (Aldous, Hajek and Kelly) on random access communication in his 1988 invention of TCP's congestion control algorithm that first prompted Professor Kelly to think about the possibility that the developing algorithms of the Internet could be addressed systematically as distributed mechanisms to solve large-scale optimization problems.</p>
<p>3. References to the research (indicative maximum of six references) [1] "Charging and rate control for elastic traffic", F.P. Kelly, European Transactions on</p>

Impact case study (REF3b)

Telecommunications, volume 8 (1997) pages 33-37.

[2]* "Rate control in communication networks: shadow prices, proportional fairness and stability", Frank Kelly, Aman Maulloo and David Tan, Journal of the Operational Research Society 49 (1998) 237-252, DOI: 10.1057/palgrave.jors.2600523.

[3] "Resource pricing and the evolution of congestion control", R.J. Gibbens and F.P. Kelly, Automatica 35 (1999) 1969-1985, DOI: 10.1016/S0005-1098(99)00135-1.

[4] "Models for a self-managed Internet", F.P. Kelly, Philosophical Transactions of the Royal Society A358 (2000) 2335-2348, DOI: 10.1098/rsta.2000.0651.

[5] "Fairness and stability of end-to-end congestion control", Frank Kelly, European Journal of Control 9 (2003) 159-176, DOI: 10.3166/ejc.9.159-176.

[6] "Stability of end-to-end algorithms for joint routing and rate control", Frank Kelly and Thomas Voice, Computer Communication Review 35:2 (2005) 5-12.

*References which best reflect the quality of the underpinning research.

4. Details of the impact (indicative maximum 750 words)

Within the Internet control is end-to-end rather than hidden in the network, and the algorithms concerned had been outstandingly successful, as the Internet evolved from a small-scale research network to an interconnection of tens of millions of endpoints and links. However, by the turn of the millennium there were several developments that posed challenges. The huge capacity of the links being deployed together with the desire to carry delay-sensitive traffic was causing an evolution towards a network with much smaller queueing delays, and it was unclear if end-to-end congestion control would remain feasible in such a network. And our demands for reliability and robustness were increasing as the Internet came to be a critical part of the infrastructure of a modern society and economy. The telephony network provided reliability and robustness through the simultaneous use of parallel paths, but would such a development compromise the stability of the Internet, much as "route-flap" had in the Internet's early years? The research of Professor Kelly and his collaborators provided a mathematical framework that allowed these challenges to be addressed in systematic and principled manner. The framework allowed congestion control and routing algorithms to be interpreted as a distributed mechanism solving a global optimization problem. The form of the optimization problem made explicit the equilibrium resource allocation policy of the algorithms, which could often be restated in terms of a fairness criterion. And the dynamics of the models allowed the machinery of control theory to be used to study stability, and to develop algorithms that scale to arbitrary capacities.

The impact on protocol development and on network architectures has been significant. Detailed evidence is provided by use of the framework and acknowledgements of the research papers in section 3 in standardization documents for protocols (e.g. [7]), technical reports in companies, and patent applications (a Google search for "patent" within the more than four thousand citations of [2] gives over eighty results including over twenty patents or patent applications). An indicative summary is given in [8], and we quote from this paper:

"After a decade of work by many researchers, there is now a substantial set of theory, algorithms, applications and even commercialization based on the NUM model of networks."

"An analogy can be drawn with Shannon's seminal work in 1948. By turning the focus from the design of finite-block length codes to the regime of infinite-block length codes, thus enabling the Law of Large Numbers to take effect, Shannon's work provided architectural principles (e.g. source-channel-separation) and established fundamental limits (e.g. channel capacity) in his mathematical theory of communication. "A similar development has happened in networking over the past decade. By turning the focus from coupled queueing dynamics to an optimization problem based on deterministic fluid models, thus enabling the views of decomposition and distributed control, Kelly's work [1,2] leads to our current study on architectural principles in networks and a mathematical foundation for network protocol design."

The comparison with the work of Shannon is clearly an exaggeration, but the quotation accurately reflects the fundamental change in conceptual framework from a microscopic queueing view to a macroscopic optimization view.

A former executive at multinational telecommunications equipment company Lucent (and then Lucent-Alcatel) states that Kelly's research within the period "established the global stability and

fairness of TCP/IP, the basic rate control mechanism of the Internet. It is fair to say that Prof. Kelly's work established the framework in which almost all subsequent work concerning these issues was done. The nature of the impact may be characterized thus: - it rendered unnecessary other mechanisms, which might have had a deleterious effect on the Internet, - it rendered unnecessary other research programs, - it gave confidence to service providers and users that the Internet would remain stable despite rapid growth as long as certain conditions were met, and this was essential for continued investment in the Internet infrastructure". [13]

A senior Microsoft researcher writes "MultiPath TCP is being standardised by the IETF, and Kelly's work is one of the fundamental pieces of work used in its architectural design, particularly with regard to the control loop. This has been implemented in the Linux kernel, in FreeBSD and recently saw its first large scale deployment on September 18th 2013 when it was part of Apple's iOS7 release - hence it will be used on 10s of millions of devices. For cloud computing, Data Center TCP (DCTCP) developed by Stanford and Microsoft has been implemented in Linux, prototyped by Microsoft, and also uses insight from Kelly and Voice [6] to design the feedback controller." [14]

In 2008 Kelly was awarded the John von Neumann Theory Prize of INFORMS for "his profound contributions to the mathematical theory of stochastic networks, and for applications of these theories to the understanding, performance evaluation, and design of telecommunications networks"[9]. For other awards citing the impact of this research see [10]. In 2012 Professor Kelly was elected a Foreign Associate of the National Academy of Engineering "for contributions to the theory and optimization of communication networks"[11].

Several of the insights on large scale system behaviour developed in the study of communication networks transfer to help understand cascading failures in other large scale systems, such as transport infrastructures [12].

5. Sources to corroborate the impact (indicative maximum of 10 references)

[7] Evidence of impact on protocol development and standardization of protocols:

Multipath TCP Resources: <http://nrg.cs.ucl.ac.uk/mptcp/>

Internet Research Task Force RFC 6077: <http://tools.ietf.org/html/rfc6077>

[8] Evidence of impact on architectural principles and network protocol design:

"Stochastic network utility maximisation - a tribute to Kelly's paper published in this journal a decade ago", Yung Yi and Mung Chiang, European Transactions on Telecommunications. 2008; 19:421-442

[9] Evidence of impact of mathematical framework on the networking community:

<http://www.informs.org/Recognize-Excellence/Award-Recipients/Frank-P.-Kelly>

[10] Evidence of impact on computer and communication systems, and other domains:

<http://sigmetrics.org/achievementaward-2009.shtml>

<http://www.theorsociety.com/Pages/Awards/Beale.aspx#2011>

http://www.tue.nl/fileadmin/content/universiteit/Academische_plechtigheden/Dies_Natalis/2011/Frank_Kelly.pdf

[11] Evidence of engineering impact: <http://www.nae.edu/56166.aspx/>

[12] Evidence of insight transferring to other large-scale systems:

"A national infrastructure for the 21st century", Council for Science and Technology, 2009. (Page 34, cascade effects):

<http://webarchive.nationalarchives.gov.uk/+/http://www.cst.gov.uk/reports/files/national-infrastructure-report.pdf>

[13] Statement from a former executive at Lucent (and then Lucent-Alcatel) on the impact of Kelly's research in shaping the thinking of the technical community and in the engineering of the Internet.

[14] Statement from senior researcher at Microsoft.