In Memoriam: Michael J. B. Duff

January 17, 1933 - December 29, 2021

We are deeply saddened to report the passing of Michael Duff, a distinguished researcher who was a long-time leader in the international pattern recognition community. Michael was instrumental in the founding of the British Machine Vision Association and was also very involved in the IAPR from its early days. From the early 1980s to the early 2000s, he had roles on the Governing Board and ExCo (as Secretary and as President from 1990-92). He later led the Constitution & Bylaws and Advisory Committees and served as Editor of the IAPR Newsletter. *Michael was elected an IAPR Fellow in 1994 for "contributions to architectures for parallel processing, and outstanding leadership."*

On behalf of the entire IAPR community, the ExCo offers its deepest condolences to Michael's wife Susan, the other members of his family, and his many colleagues and friends.

~ Daniel Lopresti, IAPR President

Michael J. B. Duff

by Virgionio Cantoni

At the end of this troubled 2021, the news came that Michael J. B. Duff has left us. Michael played a very important role in the development of multiprocessor architectures for image processing. He inspired us, transmitting his passion for research, and he boosted our engagement and skills.

Michael started his career as a physicist in the 1950s searching for particles on bubble chamber images. He quickly realized it was necessary to speed up processing at the array level, as inputs were large but structured. Therefore, during the 1950s, cellular automata (quoted even in the early work by John von Neumann) were emulated using the general-purpose computers that were available at the time. Michael began building his own computer that matched the 2D image data structure, and he started his brilliant career in computer architectures.

Cellular Logic Operations (CLOs) are performed digitally to transform a data array P(I,J) into a new data array P'(I,J). The value of each element in the new array is determined by its value in the original array and the original values of its nearest neighbors constituting the "cell"; whence the term "cellular logic". During thirty years there was a vast community of researchers, both academic and industrial, that committed themselves to the design and implementation of innovative parallel architectures that could be efficiently used for image processing. Many of them remained on paper, but just to quote two outstanding examples following the CLOs approach, we can mention the so-called SIMD architectures such as CLIP4 (Duff, 1978) and the pipelined (MISD) such as Cytocomputer (Sternberg, 1981). From the 60s, for over 40 years, computer scientists have suggested, designed, built and sometimes even marketed, new computer architectures for image processing. This was the genesis of modern image systems, following these lines it was hoped that real-time processing could be achieved.

Michael, heading the UCL Image Processing Group with Terry Fountain and other co-workers (I had the opportunity to strongly interact also with Tony Reeves and Kim Matthews) developed a series of eight increasingly complex systems (CLIP0 to CLIP7), ranging from arrays of 25 to 9216 processors. In some details, the basic characteristics of the Processor Elements (PEs) were: the broadcasting of single bit data to the 4/8 neighbors and the gating of the data in input following the 4/8 connectivity; the propagation that means recursion CLOs operations (synchronous/asynchronous) and detection of a stable condition on the array (by the so called 'OR-Sum-Tree', having through this the connected component as 'atomic' data; the PE was a single bit processor operating in SIMD.

All these subjects were exciting, at the cutting edge of image processing research of the time. Beside the activity on CLOs on flat array, it is worth mentioning Michael's strong activity on the analysis of how our topic field was evolving, the evaluation of the potentialities of new research lines, and the different aspects of matching algorithms to architectures.

In this connection, a meaningful example is the one on the collection of multiresolution, or "pyramid" techniques, for rapidly extracting global structures (features, regions, patterns) from an image. One of

Remembering Michael J. B. Duff

the most successful paradigms of pyramidal architectures was the planning strategy: processing images at low resolution, with a subset of data, and subsequently refining the resolution at the required level of detail. In 1986, we organized a NATO Advanced Research Workshop on Pyramidal Systems for Image Processing with the participation of seven groups engaged in the design of pyramidal architectures (fine grained [bin quad, four quad] and coarse grained, SIMD/MIMD/MISD, using existing or ad hoc chip or custom made]; seven groups were engaged in pyramidal algorithms for image analysis; six groups on the implementation of pyramidal algorithms on different architectures (array, hypercube and prism) and analysis of expected performance.

Michael entitled his chapter: “Pyramid. Expected Performances”. I consider this contribution very exciting. Besides an evaluation, he suggested how to change the flat array architecture and how to perform efficiently pyramidal processing on arrays. This has been a very successful suggestion because it has been what we later pursued to propose the ‘logical pyramid’! In fact, in our final hardware, we did not change the CLOs paradigm, but we found the way to avoid extra connections in the array, by including the bypass of the PEs and log2 N (N is the square side length) controllers (one for each plane of the quad pyramid), the instruction being distributed in row parallelism. At the maximum resolution it was a flat array; when this resolution was not required, all the low resolutions of the quad pyramid could work in parallel in Multi-SIMD mode.

Broadly speaking, hardware evolution is, in fact, a constant adaptation of technology to demand-driven processes along time and it may be considered as a steadily changing evolution.

Nowadays, a substantial thread in hardware development passes through Graphical Processing Units (GPU) architectures, also driven by the Deep Learning application paradigm. Typical GPU design schemas are based on arrays (exploiting also multiresolution) of cores using shared memory for communication, whereas software applications make a mixed use of CPU (general computing, coarse grain) and GPU (data parallel computation, fine grain) On this purpose Michael once stated: "Many hands make light work is a well-known saying, but then so is too many cooks spoil the broth". Of course, technology has evolved enormously, but the primitives of that time are curiously not that far away.

Michael’s contributions are relevant also in service to the British and International Pattern Recognition communities. In 1967 he founded a discussion group on Pattern Recognition, which developed in 1976 into the British Pattern Recognition Association and in the mid-1980s, now the British Machine Vision Association. Throughout all these years, Michael has been an outstanding member of IAPR. Fellow of IAPR since its institution, he has served as president from 1990 to 1992, secretary for four years, chairman of various IAPR committees and Editor of the *IAPR Newsletter*.

His research led to publications of high scientific value, written in a brilliant style, which, at the same time, were concrete and precise.

I had the opportunity to meet him at several international conferences and workshops. His open-minded view was combined with a profound intellectual honesty. To say his own, calmness and 'humor' were his strength and even when the ideas were different and the objection oversized, he ended with his 'really?' with a legendary distinction.

Unfortunately, he left us 'really', without a question mark, and we will miss him forever 'really!', with an exclamation point.