Radio Astronomy at the Fringe
ASP Conference Series, Vol. 300, 2003
J.A. Zensus, M.H. Cohen, & E. Ros, eds.

Early Work in Imaging

J. A. Högbom
Stockholm Observatory, AlbaNova University Center, SCFAB, SE-106 91 Stockholm, Sweden

Abstract. In this paper I will focus on two stories in which I have been personally involved: Earth rotation synthesis and the CLEAN algorithm.

1. Introduction

It's an excellent idea to arrange a workshop at this occasion; for many of us an opportunity to meet with friends from that ancient generation to which I myself belong. Marshall H. Cohen first suggested that I speak about "CLEAN and Aperture Synthesis". Such a title gives the impression of a thoughtful and systematic presentation of the development of these ideas throughout history. Fortunately he changed his mind. Many people, including Ken Kellermann (Kellermann & Moran 2001), have done the proper research and traced the development of radio image formation—from the famous cliff interferometer paper by McCready, Pawsey & Payne-Scott (1947) over the Mill's Cross, the Christiansen solar arrays, Ryle's aperture synthesis, the VLA and so on, all the way to the sophisticated image making procedures used today. I cannot improve on these fine papers. But here, in the history session, and with a more flexible title, I may be excused to do some micro-history and focus on two stories in which I have been personally involved: Earth rotation synthesis and the CLEAN algorithm.

2. Earth Rotation Synthesis

I came to Cambridge in 1955 to do thesis work and stayed there until 1959. It was in a time when one began to realize that Radio Astronomy had the potential of becoming much more than just a poor cousin of traditional optical astronomy. Angular resolution was perhaps not doomed to remain vastly inferior to that of optical astronomy just because nobody could build a steerable paraboloid several miles in diameter.

It was a turbulent period at Cambridge. The 2C survey had just been published and come under fire. I understood that there were some very unpleasant people down in Sydney, but not much more. Tony Hewish was my supervisor and I used simple two-antenna interferometers to observe the scattering of radio waves by the solar corona and to position radio bursts on the solar disc. Much of the heated arguments about source statistics passed me by.

However, I attended Ratcliffe's now famous lecture series on Fourier analysis and like many others I became excited by the beauty of it all and, in particular, by its application to aperture synthesis.

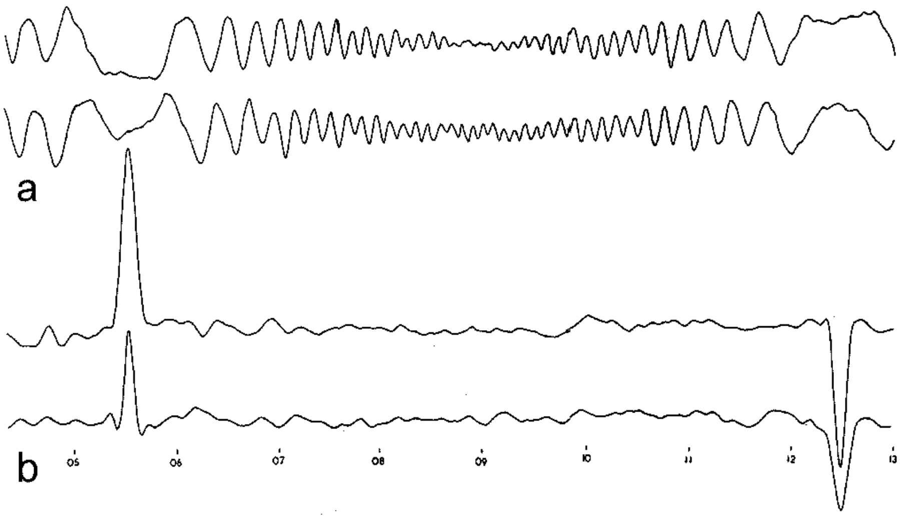

Figure 1. a) The simple small antenna N-S interferometer: Cos- and Sin- recordings showing the Crab Nebula and M 87. b) Synthesised N-S interferometer: Cos- and Sin- recordings.

The NS and EW simple interferometers that I used for the positioning of solar bursts were kept running during the night to get amplitude and phase calibrations from the Crab Nebula and M 87. Virtually all telescopes at Cambridge were meridian passage instruments. All interferometer records coming off the pen recorders in the receiver hut looked much the same - a regular sine wave under some envelope determined by the antenna dimensions. My NS interferometer on the other hand produced something very different, a non regular periodicity with an obvious central fringe (Fig. 1a). I understood of course how that could be, but I was puzzled by some properties of the pattern. First, in contrast to my EW interferometer pattern, there was an obvious central fringe, and the variable fringe period seemed to make it possible to separate the two sources even if they had been close enough to overlap. Convolving the record with the theoretical source response - which is what you should do when looking for a particular pattern - would clearly give a strong dominant peak just at the meridian passage of the source and little else. The same operation on the EW interferometer record would just produce a similar uninteresting sine wave with no obvious central peak. Why this apparently superior amount of information available from the NS interferometer?

I did the celestial sphere geometry and soon realized that convolving the pattern with the theoretical response was much the same as synthesising the output from an EW (curved) line antenna traced out by the rotation of the Earth. I gave a preliminary write-up of my ideas to Ryle in July 1958 and somewhat sheepishly asked him to “read it in his bath”. Apparently he did, because the next time we met he said that he had had an “interesting bath”. As far as I can remember, he made no further comment. I also showed it to Tony Hewish who helped me by running the convolution—now synthesis—calculations on the Cambridge EDSAC II computer (Fig. 1b).

I thought at the time that Earth rotation synthesis was a fun idea but of limited use. The trick, I imagined, required very small antennas with wide beams: the source must stay in the beam long enough for a significant synthesis. Why couldn't I see that fully steerable antennas would make all the difference between a fun idea and something really useful? Well, I just didn't. Fully steerable antennas did not exist at Cambridge! That shows the difference between us ordinary guys and people like Martin Ryle.

My thesis work on the Sun produced 78 typewritten pages and looked rather thin. Previous radio astronomy theses in the library were quite solid and none contained less than 100 pages. I searched for something more to say about the Sun—but failed. In the end I crossed my fingers and added two chapters on a completely different subject: aperture synthesis theory and the Earth rotation experiment. That widened my thesis to a more comforting 104 pages (Högbom 1959). That's why the Earth rotation story ever got written down!

I did of course not 'invent' Earth rotation synthesis. I 'rediscovered the wheel' (or at least some part of the wheel!). Martin Ryle had it all thought out long before—including the one mile telescope design. But he kept it to himself, worrying that others with plenty of money—read the Americans—would exploit his ideas before he had a chance to do much himself. I was in the Netherlands working on the Benelux Cross project when in 1962 he published a short Nature article containing a simple drawing of the future One Mile Telescope (Ryle 1962). I saw the article and felt a bit stupid not to have thought about such a beautiful application. From there on I began propagating for a super-synthesis array, possibly as a complement to a shortened cross for extended objects. As time went by, the cross faded out of the project altogether to make place for the Westerbork Synthesis Telescope completed in 1970 (Högbom & Brouw 1974).

3. The CLEAN Algorithm

The second story I'd like to tell is the story about CLEAN and, in fact, that story started right here in Green Bank in 1968. I had time on the 3-antenna interferometer for a program of mapping the structure and polarisation of some 60 sources selected mainly from the 3CR catalogue. The plan was to make short "snapshot mode" observations at six hour angles and with a few different antenna spacing arrangements. That should give me 50-60 complex visibilities distributed over the (u, v)-plane for each source. At the time I had no firm idea how to reduce the data, but naïvely imagined that it shouldn't be too much of a problem to determine source structures from such measurements.

Back in Leiden I realized that there was indeed a problem! Straight interpolation between the measurements is just mathematically wrong as everybody knows. Model fitting was messy and unreliable. I tried an iteration scheme, flipping between the visibility and map domains. Starting with the direct transform of the individual measurements—i.e., what is now known as the 'dirty map', I shaved off negative values in the map domain as being impossible, returned to the visibility domain, restored the visibilities to their measured values and so on. It half- worked in some simple cases but never very well.

During these trials I found myself looking at 'dirty maps' of many sources including some calibration point sources. It was then a small step to ask: if I subtract a full theoretical point source pattern, a suitably scaled and positioned 'dirty beam', from the map then there should be nothing left - unless of course there is something else out there! Often there was and I went on subtracting. Returning to the map only the nice part - the central lobe - of each subtracted pattern was a temptation I couldn't resist and it actually seemed to work! From there it rolled on. When things got messy it often helped to do things cautiously, a little at a time, i.e., using a reduced 'loop gain'. So CLEAN had a very simple-minded beginning but in the end turned out to be more useful than I had ever expected. Dave Rogstad was in Groningen at the time and we had some good discussions. He used it with his Caltech interferometer measurements of M 101 and published the first CLEANed map (Rogstad & Shostack 1971).

It interesting to note that the paper describing the CLEAN algorithm was later presented just where the story started: in Charlottesville, at the 1972 NRAO symposium on 'The Collection and Analysis of Astrophysical Data' (Högbom 1974). Now 30 years later, it is good to be back here again and to look upon this whole story as part of history!

References

Högbom, J. A. 1959, Ph.D. Thesis, Cambridge University, UK

Högbom, J. A. 1974, A&AS, 15, 417

Högbom, J. A., & Brouw, W. N. 1974, A&A, 33, 289

Kellermann, K. I., & Moran, J. M. 2001, ARA&A, 39, 457

McCready, L. L., Pawsey, J. L., & Payne-Scott, R. 1947, Proc. Roy. Soc. London, Series A, 190, 357

Rogstad, D. H., Shostack, G. S. 1971, A&A, 13, 99

Ryle, M. 1962. Nature, 194, 517