SKI Report 97:14 CSSP R

Fuel Assembly Assessment from CVD Image Analysis: A Feasibility Study

Clark S. Lindsay and Thomas Lindblad

May 1997

ISSN 1104-1374
ISRN SKI-R--97/14--SE

SKI
STATENS KÄRNKRAFTINSPEKTION
Swedish Nuclear Power Inspectorate

Atomic Energy Control Board
Commission de contrôle de l'énergie atomique

SKI Report 97:14 CSSP Report 110

Fuel Assembly Assessment from CVD Image Analysis: A Feasibility Study

Clark S. Lindsay and Thomas Lindblad

Department of Physics, Royal Institute of Technology,
Stockholm, Sweden

May 1997

SKI Project Number 96184

Prepared for the Swedish and Canadian Safeguards Support Programs to IAEA under IAEA Task JNT A00704.

NORSTEDTS TRYCKERI AB
Stockholm 1997

Abstract

SKI commissioned a feasibility study of automatic assessment of fuel assemblies from images obtained with the digital Cerenkov viewing device (CVD) currently in development. The goal is to assist the IAEA inspectors in evaluating the fuel since they typically have only a few seconds to inspect an assembly. We report results here in two main areas: (1) investigation of basic image processing and recognition techniques needed to enhance the images and find the assembly in the image; (2) study of the properties of the distributions of light from the assemblies to determine whether they provide unique signatures for different burn-up and cooling times for real fuel or indicate presence of non-fuel.

Sammanfattning

Denna förstudie har genomförts på uppdrag av Statens Kärnkraftsinspektion, SKI, och avser automatisk utvärdering av bränslet. Studien har genomförts med hjälp av bilder tagna med den digitala Cerenkov kamera (CVD), som just utvecklas. Syftet med detta arbete har varit att bistå IAEA inspektörerna vid bedömning/utvärdering av bränsle eftersom de ofta endast har några sekunder till förfogande vid inspektion av lagrat bränsle. De resultat som här redovisas är i huvudsak att hänföra till två områden, nämligen: (1) undersökning av grundläggande metoder för bildbehandling och bild-igenkänning, inklusive förbättringar av bilderna för snabbare och säkrare identifieringar. (2) studier av de egenskaper som karaktäriserar den speciella ljus fördelningen, speciellt med avsikt att bestämma huruvida dessa utgör unika "signaturer" för verkligt bränsle (av olika typ, fabrikat, användning, nedkylningstider, etc) och/eller för att indikera närvaron av andra element än bestrålade bränslestavar.

Contents

1. Introduction

A SKI-AECL team recently developed the concept of a digital Cerenkov viewing device (CVD), a UV-sensitive CCD camera, to observe fuel assemblies in storage (Attas 1996). To test the camera the team recorded a number of images of fuel assemblies at the storage ponds at CLAB and Ringhals in Sweden (Trepte 1996, Chen 1996). SKI commissioned us to examine these images to determine if useful information could be extracted and provided to an inspector by an image analysis system connected to the CVD. Ideally this system would provide information in *real time* or near real time since the inspectors only spend 10-20 seconds examining a particular assembly. The information might help the inspector enhance the images (e.g. indicate that the gain should be adjusted), help align the camera over the centre of the assembly (i.e. by maximising the collimation effect), and perhaps provide a detailed assessment of the fuel assembly properties (e.g. calculate the probability that it is fuel or non-fuel.). This study only examined what kind of information is feasible to obtain. Development of a working system will require extensive interaction and field testing with the inspectors to determine what assistance with the images they really want and need. So, as a feasibility study of limited duration (5 months), we concentrated on general techniques and proof of principles.

Two main areas were explored: (1) image processing and recognition methods; (2) properties of the Cerenkov light distribution across the assembly face. The first area involves the processing of the images to find the assembly in the image, identify its type, and obtain the intensity distribution. Basic image processing techniques include, for example, intensity equalisation since the brightness can vary considerably from a hot and bright assembly to a cool and dim one. Most of the image recognition techniques work best with binary images so the grey scale image must be converted to a black and white image. The assembly must be found in the image and its type identified. This requires a pattern recognition algorithm to scan the image. Filtering methods help to reduce the background noise and clutter to make the recognition task easier.

Once the assembly has been found in the image, the distribution of light over the assembly face provides the only clues to the properties of the fuel (or, perhaps, non-fuel) rods. Previous Monte Carlo studies indicated that the light distributions could reveal even a single dummy rod within an assembly (Kulka 1995). Real world problems make this difficult. Handles and cover plates, for example, obscure many of the rods. Therefore we had to study the properties of the light distributions to determine if they indeed provide information unique to the types of assemblies (e.g. fuel with different cooling times and burn-up or faked non-fuel). We concentrated on the gross features of the distributions to see if they consistently correlated with the type of assembly, the burn-up and cooling time, and fuel and non-fuel type.

Due to the limited time available we took several shortcuts to simplify the task. For example, we concentrated on just three types of assemblies although there are several dozens of assembly types in use. We only had images of a few types and assumed that the

methods would generalise to all assembly types. For programming we used the Matlab interactive environment since it provides quick prototyping capabilities and includes an extensive image processing toolbox. A working system would need to be very much faster to provide real time processing. We briefly discuss possible ways to speed up the image processing with C code and implementation in image processing accelerator cards.

In section 2 we discuss the available data and the software used in the analysis. Section 3 discusses the image processing and recognition techniques. Section 4 gives the results of the analysis of the assembly light distributions. Section 5 discusses some implementation schemes for a real world system. Tables give descriptions of the data.

2. Data and Software

The SKI-AECL team made three image taking runs: two at the Ringhals reactor site and one at the CLAB storage facility. The data was placed on a single CD-ROM. Table 1 describes the data for each run. To avoid an excessive number of variables and to maximise the statistics, the images used in the analysis were restricted to a *standard* camera set-up that required: telephoto lens, 2μs readout, Visilog 3 file format (Visilog was the frame grabber software used to make the images) and no binning of the CCD pixels. Most images were with the Ultra UV filter but a few were taken with the Lite filter (Trepte 1995). Table 2 and 3 list the assemblies used in the analysis and the measurements for each (note that for some assemblies there were multiple images, particularly for different exposure times, and so average values were used.)

The images were all 528x560 pixels. The Visilog 3 files used 12 bit values for the intensities. For most of the functions used here, 8 bit values were expected. So the values were re-scaled to 8 bits. However, if left shift of 4 bits were used rigidly, then in many of the images the assembly was too dark to see. Therefore, the scaling was left as an option when analysing the images to optimise the quality of the assembly view.

Figure 1 shows examples of the types of assemblies seen in these images. The Pressurised Water Reactor (PWR) assemblies examined hold 15x15 rods (13x13 visible due to a cover plate around side). Unfortunately, the helium non-fuel PWR assembly was a 17x17 type and so was not used in the analysis since there were no similar fuel assemblies to which to compare it. The Boiling Water Reactor (BWR) assemblies hold 8x8 matrices of rods or four sub-bundles with 4x4 rods each (SVEA64).

The BWR assemblies are labelled here, as shown in figure 1, according to the handle and bracket configuration. The SVEA64 type has a X cross. SVEA64-TP has a cross plus a top plate that covers the outer rows. The 8X8 assemblies have only a single handle running diagonally. Only a few examples of the SVEA64 assemblies were available so the analysis concentrated on the SVEA64-TP and 8X8 assemblies. Unfortunately, as for the PWR helium assembly, the helium non-fuel assembly was an uncommon type in the images, a SVEA64 type, and so was not used in the analysis.

The analysis was carried out with Matlab ver. 4.2 and the Image Processing Toolkit ver. 1.0. We ran primarily on a 200 MHz MMX Pentium. Appendix A describes the functions written in Matlab to carry out the tasks not provided by the standard Matlab functions. Matlab processes run fairly slowly compared to C routines, especially for the scripted functions that must be interpreted rather than compiled. However, the interactive nature of Matlab provides a powerful tool for prototyping and experimenting with analysis methods. For an actual working image analysis system, the code would need to be transferred to C or to a parallelized code to run on special accelerators (discussed further in Section 5.)

3. Image Processing and Recognition

An automatic fuel assessment system would initially need to find the assembly in the image. Although this task is trivial for humans, it is actually quite difficult to program a computer to do it. However, this sort of task is now common in the field of image recognition and a number of tools and techniques have been developed. Usually the image must be simplified and filtered before presenting it to the recognition algorithms. We divide this pre-recognition processing into two parts: pixel processing (often in the literature referred to as pre-processing) and morphology operations. Pixel processing involves operations that work only on the pixel intensity values and ignore any higher level (i.e. multi-pixel) structure or shapes in the image. Pixel operations include equalisation and thresholding. Filtering to simplify the images was tested primarily with morphology operations. Morphology or shape, operations refer to those algorithms that work on connected (i.e. adjacent pixels of the same value) pixel groupings.

After the image is prepared by the pixel operations and filters, the resulting image must be scanned by the recognition algorithm to determine where the assembly is located. The assembly type must be determined so that the light distribution across it can be compared to what is expected for that kind of assembly. This comparison may indicate whether it consists of fuel or non-fuel rods.

3.1 Pixel Processing

The images of the assemblies vary greatly in their intensity. This can arise from the actual brightness variations among assemblies with fuels of different burn-up and cooling times. It can also arise from different exposure times. (The team was testing the camera and exposure time was one of the parameters they varied.) Figure 2 shows two images and their histograms of grey scale values. In one case the distribution is peaked near low values while in the other case the distribution is wider and peaks at higher values. A single threshold to create a binary (B/W) image would blank out most of the assembly in the first case if set too high, while adding excessive background noise if set too low.

Histogram equalisation provides a way of making the histograms more similar. Histogram equalisation is a common image enhancement technique (Gonzalez 1992) and is available with Matlab. Here the initial grey levels are remapped to new values to give a

flat histogram. Figure 3 shows intensity histograms after the equalisation. (The resulting histogram is not flat because of the effect of finite width binning.)

After the equalisation, a threshold cut at intensity equal to 0.5 created binary images with 0 and 1 intensity values. Binary images are more simple to work with and most morphology filters require binary input images.

3.2 Filtering

As can be seen from figure 1, the images can hold significant noise and extraneous objects. We tried several methods to filter from the images those features that could distract the assembly recognition. Morphology operations, where morphology refers to shapes and structures, have become a common technique in the field of image processing to handle such problems. Matlab offers a large set of standard morphology functions such as Dilate, Erode, Close, Open, etc. (Thompson 1995.) Morphological operations affect the shape of connected areas. Connected means that any two pixels of similar pixel value (e.g. value 1 in a binary image) can be reached without crossing a gap of one or more dark pixels. In figure 1 the BWR assemblies, in particular, have large white bars and crosses that can confuse the pattern recognition.

The Open operation, which shrinks an object and then expands it (if it still has non-zero number of pixels), removes small noise points very effectively. Getting rid of larger background objects was more complicated. In one method we first labelled all connected areas and measured their area. If their area exceeded a threshold they were eliminated from the image (e.g. set to 0). However, sometimes the binary operation resulted in a large section of the assembly being connected and then this area cut would eliminate this section of the assembly as well. This happened in particular for dim images where there is low contrast and the threshold cut would set both the rod areas and the water gaps to 1.

Finally, we wrote a *blanking out* function to specifically get rid of the bright bars. We scan the image looking for areas that matched a large rectangle, oriented horizontally in one scan and vertically in another. Those areas that matched were set to 0. Figure 4 shows a binary image before and after this *blanking out* operation. This operation required both a connected white area to be both large and rectangular. The assembly areas were unlikely to fulfil both of these requirements and were unaffected.

3.3 Assembly Location

Locating and identifying the assemblies in the images is a basic operation required by a fuel analysis system. This processing must be robust enough to find the assembly despite a wide variety of assembly types, a wide range of brightness, noise, and background objects.

Initial algorithms achieved only 50-70% accuracy in locating the assemblies. These algorithms relied on a filter, or set of filters, including the morphology operations mentioned above, to enhance the assembly area while eliminating, or, at least, drastically reducing, the backgrounds. The assemblies have a lot of *texture*, i.e. many light and dark areas, in a confined square area, while backgrounds were typically small noise points or large smooth areas. Edge detectors, which use gradients to find sharp changes in intensity, and wavelets were filters that enhanced the assembly areas while reducing the backgrounds. Projections of these filtered images onto the x and y axis would then show clusters, i.e. a section of closely packed filled bins, from the assembly pixels. The edges of the clusters would give the edges of the assembly in the projections. The overlap of the x and y edges would then locate the assembly.

However, there were several problems with this method. These filters only worked with grey scale images. Low brightness, and so low contrast, assemblies failed to be enhanced by either the edge detectors or wavelets. Also, fragments of backgrounds often remained after the filtering and managed to confuse the projections and so the assembly edges were incorrectly determined.

The most accurate method found for locating the assemblies was template matching. Templates, shown in figure 5, of three assembly types were made including two templates for one type but of different size. These ¼ scale templates were scanned across ¼ scale versions of images with a convolution operator. The white areas had value +1 and the dark areas -1 for both template and image. The position of the best match, i.e. score = number matching pixels / total pixels of template, was recorded. Initially, a scan of the SVEA64-TP templates was made. If the score was above a threshold, then that assembly type won. If below the cut, then the scanning was repeated for the 8X8 template, and so forth. If no score was above the cut, the highest scoring template was chosen as the winner.

Note that there were two templates for 8X8 assemblies. In figure 1 (b) two sizes of images for 8X8. The CLAB set gave images with larger views of the assembly than the set from the first Ringhals set due to difference in the camera to assembly distance. In practice, the SVEA64 template found the small 8X8 assembly images and so it was not necessary to scan with the small 8X8 template. However, different image sizes of the assemblies due to different distances of the camera from the assemblies must be considered in future analysis.

The template matching gave 100% accuracy in locating the assembly *close* (i.e. the centre of the template was within the assembly area) to the exact position. Figure 6 shows a sequence of images for the template matching algorithm. To obtain the exact location of the assembly, a second pass on an unscaled region around the first guess would need to be made but we have not coded this.

4. Light Distributions

There is probably little point in finding the fuel assembly if no unique information exists in the assembly light characteristics. The fuel assembly Cerenkov light, produced in the water from electrons scattered by gamma-rays, has at least one well-established distinctive feature. Because the assemblies are several meters long, the light in the narrow water gaps becomes highly collimated. The inspector scans the imager slowly across the assembly until a sharp increase in brightness appears when the camera is exactly lined up with the assembly. Another feature has been predicted by Monte Carlo (MC) studies but not numerically evaluated for real assemblies (Kulka 1995). The MC predicts that the distribution of the light for real fuel should show a peaking in the centre. Non-fuel assemblies, which receive gamma-rays from nearby fuel, or assemblies with fuel rods intermingled with non-fuel rods, should show a flatter distribution. This peaking effect, which at the 10-20% level is difficult to detect by eye, could possibly be used by the imaging system to assess the fuel, e.g. distinguish fuel from non-fuel (Chen 1996, Trepte 1996).

We therefore examined the assembly images to determine whether this effect existed in practice and if it distinguishes the fuel types. The PWR 15X15 type images, figure 1, should provide the best test of this since they do not have obstructing handles or top plates. In the analysis here, the following steps were carried out:

1. The assembly was cropped by hand from the grey scale image.
2. We calculated the ratio **R** of the sum of the pixel values, i.e. total light, in an inner square to the total light in an outer frame.

The areas of the square and the outer ring were equal. Figure 7 shows the square and ring for 3 types of assemblies studied. This R ratio is a coarse measure but for this first study, it should indicate the gross amplitude of any such effect.

4.1 PWR

The PWR 15x15 assemblies, as mentioned above, provided the best images for investigating the light distributions. There were no handles to obstruct the view and they present considerable amount of light evenly over the whole face. In addition, the PWR images came from the third test of the CVD and the experience gained from the previous two tests helped to make the images more uniform and systematic. Dark images, for example, were taken regularly during the testing. Here the dark image closest in time to the given image was used to subtract the dark currents from the images.

To determine if the light distributions were reliable and consistent, we looked at cases where there were several images taken of the same assembly. These could be several images with the same settings or with different exposure times. The light from the image is defined as the sum of the grey values (0.0 to 1.0, 8-bits). Figure 8 shows a plot for

several assemblies of the ratio of total light from the assembly face (cropped from the total image) for two images versus the ratio of their corresponding exposure times. Ideally they would lie along the diagonal line. That is, two images with exposure times of 10 and 20 seconds should have a light ratio of 2 and an exposure ratio of 2. Without subtracting the dark images the longer exposure images fall below the diagonal. Subtracting the dark currents brings them up to the diagonal.

There is also consistency in the inner-outer ratio R for images of the same assembly taken with different exposure times. In figure 9 the ratio R is plotted for several PWR 15x15 assemblies, as well as for several BWR 8x8 assemblies. The ratio is consistent within +/- 0.02 for the PWR assemblies.

Since the images seem to behave consistently, we can look at various aspects of the light distributions with some confidence. From the above discussion, it would appear that the dark image should be subtracted from the assembly image. Figures with and without dark image subtraction are given, however, since the dark images add another possible source of systematic errors.

Figures 10 and 11 shows R versus the total light from the assembly with and without dark image subtraction. The error on the total light was taken roughly to be 10%. The points rise with greater brightness and then seem to level off. Whether this plateau is a physical effect in the light distribution or from saturation in the camera isn't known.

Figure 12 and 13 shows R versus the burn-up of the fuel with and without dark image subtraction. Figures 14 and 15 show R versus the cooling time of the fuel with and without dark image subtraction. For the same cooling time of 5 years, R increases with increasing burn-up. For similar burn-up, R decreases with increasing cooling time.

Images of high-density (HD) non-fuel and the skeleton were not available with the Ultra filter. Their images were taken with the Lite filter. Images 16-21 show similar figures for the Lite filter data. There is less data than with the UV-Ultra filter and it is difficult to draw conclusions. Gamma-rays from surrounding fuel will produce Cerenkov light even in these non-fuel assemblies. However, the HD and skeleton clearly have consistently *low* R values indicating that there is little or no peaking of the light towards the centre for the non-fuel. Of course, absence of rods will be clearly be seen for the skeleton assemblies anyway.

The non-fuel assemblies produce Cerenkov light because of gamma rays from surrounding fuel. One might also expect that the light distributions of real fuel might be affected by the surrounding fuel. One assembly, I22 in Table 2, was measured in an isolated position and then moved to where it was surrounded by fuel. However, no significant change was measured in the R ratio within the error.

4.2 BWR

The BWR images come from the data taken during the two earlier tests (Ringhals May 1995 and CLAB Oct.1995). The data is less consistent and systematic since there were so many variations in the camera settings. Also, dark images were not taken often during the test and it's seems that the distributions seem to suffer from this in not being as consistent as with the PWR data. Figure 9 shows R for images of the same assembly for several cases. The scatter was more than for the PWR assemblies.

The SVEA64-TP fuel covered only a narrow range of burn-up and cooling so we discuss here only the 8X8 images.

Figures 22 and 23 show R versus the total light from the 8x8 assemblies. They display generally similar behaviour to the PWR distributions. The plateau at high values is seen. Also, displayed are points for the assemblies with missing pins, high density non-fuel, and skeletons. Again the non-fuel assemblies show low R values indicating that the peaking effect is absent.

The missing rods assemblies should give anomalous R values, for a given burn-up and cooling, since extra light will be emitted in the area where they are missing. Depending on whether they are in the outer or inner section will either decrease (if in outer) or increase (inner) the R ratio. This might be a way of identifying assemblies with missing rods if the systematic errors on the R values can be reduced enough.

Figures 24 and 25 show R versus cooling time. Figure 26 and 27 show R versus burn-up. The Ringhals data seem to show a decrease in R versus longer cooling times and lower burn-up. The CLAB data is consistent with this but is also consistent with a flat behaviour as well. As mentioned, the weight given to this data is less due to the messier quality of the data. The high density and skeleton points are included for comparison.

5. Implementation Schemes

Various image processing boards for PC expansion slots are available. These act as co-processors to the main PC processor and accelerate the execution of the image algorithms. The boards may use a single very high-performance digital signal processor (DSP) or have multiple processors running in parallel.

We have experience with the Adaptive Solutions CNAPS system (Adaptive Solutions, 1996). This system is based on a chip with multiple simple processors working in SIMD (Single Instruction Multiple Data) mode. PC accelerator cards with up to 512 processors working in parallel are available. A graphics package has also been developed. Speedups of 10-1000 can be expected over sequential processing depending on the degree to which the given algorithm can be made parallel.

Many image processor cards, usually based on DSPs are advertised. A detailed comparison of the price/performance ratios for these cards should be done to determine how they compare to the CNAPS system.

6. Discussion

As a short term feasibility study we did not try to build a working prototype of a image processing analysis system for the CVD. Instead we concentrated on a few essential items that would indicate whether one should proceed to the next step in development of such a system. We believe this study showed the following:

- Image processing techniques can equalise the images so that comparisons can reliably be made between assemblies of differing brightness.
- Recognition techniques can locate the assemblies in the image
- The light distributions from the assemblies are consistent and reliable when the images are carefully made.
- Peaking of the light, as measured by the ratio of light in inner versus outer areas, increases with the hotter fuels, whereas cool fuel show little or no peaking.
- Non-fuel assemblies consistently show flat, i.e. no peaking, light distributions

Due primarily to lack of time we have *not* shown the following:

- An automatic correction for tilting of the assemblies in the image.
- A program package that does the full analysis of reading in the image file, finding the assembly, identifying it and calculating the inner versus outer light ratio. The parts exist but have not been put together.
- The analysis could be speeded up enough in a conventional processor or even in a high-performance image processor card enough to do the image analysis in real time.
- Exactly the best way to help the inspector in assessing the fuel.

A more detailed study of the light distribution is also warranted. Here the single R value gives only a coarse estimate of the peaking of the light towards the centre of the assembly. A comparison, for example, of intensity versus distance from the centre for the same type assembly but with different fuels and non-fuels has been suggested.

Methods to reduce the systematic errors could be investigated. For example, the CVD will be capable of a high image rate (approx. 15 per sec). Using the *average image* created by averaging the pixels from several images of the same assembly will reduce the noise and produce a more reliable image on which to do the analysis.

The development of a real time prototype system, although we believe it to be quite straightforward, will require another study. A number of hardware and software options are available and the trade-offs of cost, complexity, and capability will have to be analysed.

The assembly locating algorithm could certainly be simplified and the speed increased. For example, the small SVEA64 template found both SVEA64 and the small images of the 8x8 assemblies. We hope this indicates that a few *generic* templates could be used for initial scans. If a unique template must be used for every kind of assembly (> 50 types at least) then the template scanning would take much too long, especially if different sizes of each would be needed (due to different heights of the cameras above the assemblies in different storage facilities).

Once the assembly location is found, a second pass would be needed in a smaller area around this initial guess to obtain the exact location and assembly type. Then the assembly light distribution would be compared to a standard distribution for that type of assembly to determine if any anomalies were present. If so, then the inspector would be notified that something was suspicious.

However, exactly what the image analysis system should do must be determined by interaction with the inspectors. The image processing could be as simple as merely monitoring the images for good contrast and brightness, or as elaborate as giving the inspector an in-depth assessment of the fuel light distribution and the implied fuel properties.

7. Acknowledgements

We thank Lars Hildingsson for many useful discussions and advice. Also, we thank Bo Lindberg and Oliver Trepte for their help. We thank the staff at CLAB and Ringhals for their assistance in obtaining the fuel images. We appreciate Richard Ward of ORNL for his help early in this study. We also thank Anders Ingestedt, Annika Lasu, Mikael Hansen.

8. References

1. E. M. Attas, G. R. Burton, J. D. Chen, L. Hildingsson, O. Trepte and G. J. Young, "Ultraviolet digital imaging for nuclear safeguards", SPIE Vol. 2654, pp.287-298, 1996.
2. O. Trepte, L. Hildingsson, J. D. Chen, G. R. Burton, G. J. Young, E. M. Attas, "Development of a high sensitivity Cerenkov viewing device. Field test at the Ringhals 2 PWR facility, Sweden", IAEA Task A00704, SKI Report No. 96:45, 1996.
3. J. D. Chen, L. Hildingsson, O. Trepte, E. M. Attas, G. R. Burton, G. J. Young, "Development of a high sensitivity Cerenkov viewing device. SCCD field test in Sweden", CSSP report No. 88, 1996.

4. K. Kulka and A. Hallgren, "Determination of detected Cerenkov-Light intensities for verification of irradiated nuclear fuel using Monte Carlo techniques", SKI Report 95:45, 1995.
5. Matlab User's Guide August 1992, The MathWorks Inc., 24 Park Way, Natick, Mass.01760.
6. C. M. Thompson, Loren Shure, Image Processing Toolbox, 1993, The MathWorks Inc., 24 Park Way, Natick, Mass.01760.
7. R. C. Gonzalez, R. E. Woods, "Digital Image Processing", 1992, Addison-Wesley Co.
8. Adaptive Solutions Inc., 1400 N. W. Compton Dr., Suite 340, Beaverton, OR 97006, USA., www.asi.com, 1996.

Table 1. Number of images available from the three experiments for each assembly type. (Assembly types are displayed in figure 1 and discussed in Section 2.) Several *cuts* were made on the data to obtain images suitable for analysis. The analysis only used images that were taken with the telephoto lens, no binning of the CCD pixels, 2μs readout and Visilog 3 file format. Most accepted images were taken with the Ultra filter but the few with the Lite filter are given in parenthesis. Non-fuel assemblies include those with rods made of either high density (HD) steel or helium (HE). Skeletons were assemblies with no rods. Empty indicates no assembly. Dark images provide the dark noise in the pixels and can be subtracted from images to improve contrast. The accepted images may include cases of multiple images of the same assembly. Some numbers are qualitative since scans may not include a view of whole assembly in image.

Location	Date	Assembly Type	Total Number of Images	Number of Images after Cuts
Ringhals1	1995-May	SVEA64	9	2
		SVEA64-TP	87	10
		8X8	165	24
		8X8, HD	16	1
		8X8, Skeleton	13	2
		SVEA64, HE	10	0
		Empty	4	0
		Dark	6	4
		Other	7	0
CLAB	1995-Oct.	8X8	49	10(4)
		15X15	1	1
		Dark	3	2
		Other	35	0
Ringhals2	1996-Feb.	15X15	122	21
		15X15, HD	7	(3)
		15X15, Skeleton	7	(2)
		17X17, HE	9	0
		Dark	38	21
		Other	23	0

Table 2: List of **PWR 15x15** assemblies after all cuts in table 1. This data produced the data plots in figures 10-15 (except No. 5 is included only to show values for an assembly that was moved). In some cases the values arise from averages over results of several images of same assembly. Ultra filter used unless otherwise noted. **Int.** indicates the total summed intensity and **R** indicates the inner-outer ratio. . The intensity is normalised to a 10 sec exposure. **Dark Subt.** indicates that the dark image was subtracted. Data is from Ringhals Feb-1996 data unless otherwise noted.

No	Assm	Pos	Burn	Cool	Raw Data		Dark Subt.		Comment
					Int.	R	Int.	R	
1	S35	F10	42	1	---	1.08	---	1.18	Only TIF files.No abs int.
2	R29	E10	38	2	157	1.15	131	1.18	
3	R14	D10	39	3	114	1.14	85	1.19	
4	I22	L24	35	10	64	1.00	36	1.00	
5	I22	B10	35	10	63	1.01	39	0.99	Moved assembly
6	R35	D13	39	2	169	1.14	130	1.18	
7	S22	D12	41	2	160	1.15	132	1.18	
8	N15	C10	39	5	86	1.10	58	1.12	
9	N17	F08	39	5	94	1.10	66	1.11	
10	Q01	E08	31	5	89	1.05	53	1.07	
11	Q03	D08	23	5	76	1.01	41	1.05	
12	H26	A112	36	9	42.	1.08	33.6	1.08	CLAB file.
13	M02	E24	40	2	41	1.08	39	1.14	Lite filter
14	H09	D22	33	9	23	1.06	8	1.15	Lite filter
15	I22	L24	35	10	40	1.03	10	1.06	Lite filter
16	SK	B22	--	--	21	1.01	5	1.03	Lite filter, Skeleton
17	HD	F22	--	--	16	1.00	1.6	0.86	Lite filter, High Density.
18	L09	B22	41	7	26	1.05	10.	1.05	Lite filter, missing rods.

Table 3: List of **8x8** assemblies after all cuts. This data was used in the data plots and in some cases includes averages over results of several images of same assembly. **Int.** indicates the total summed intensity and **R** indicates the inner-outer ratio. The intensity is normalised to a 10 sec exposure. **Dark Subt**. indicates that the dark image was subtracted. Data is from Ringhals May-1995 data unless otherwise indicated.

No	Assm	Pos	Burn	Cool	Raw Data		Dark Subt.		Comment
					Int.	R	Int.	R	
1	9900	AH32	37	1	79	1.09	66	1.10	
2	9900	AH32	37	1	57	1.09	48	1.09	From different day
3	519	CB21	20	15	19.	0.96	5.5	0.90	
4	8235	BH21	36	7	29	1.04	17	1.06	
5	9853	CB33	35	3	40	1.05	31	1.05	
6	985	AE26	22	16	13	1.02	3.1	1.03	Missing rods
7	758	CB35	22	15	26	1.01	2.9	0.85	Missing rods
8	S003	BH23	--	--	17	0.95	5.9	0.82	Skeleton
9	S004	CB36	--	--	12	0.96	3.3	0.87	Skeleton
10	HD	BH21	--	--	11	0.97	2.2	0.82	High density non-fuel
11	80	H025	5.4	19	23.	1.01	7.3	1.01	CLAB file.
12	1885	H025	23	15	27	1.01	9.9	1.01	CLAB file
13	8514	H025	35	4	30	1.03	23	1.03	CLAB file
14	8505	H025	35	6	34	1.04	13	1.04	CLAB file
15	1059	B415	16	16	10	1.02	6.5	1.02	CLAB file,missing rods.

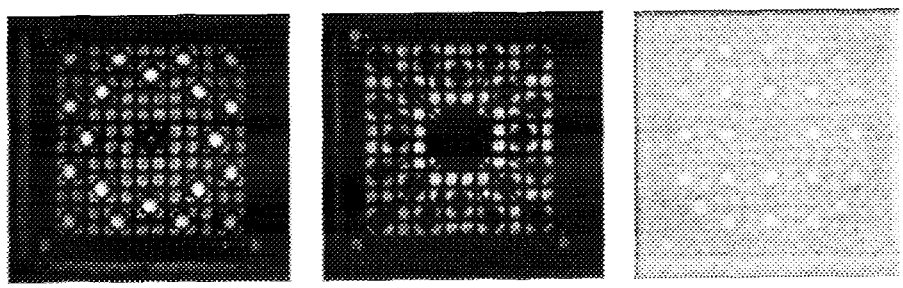

Figure 1(a) Examples of PWR assembly types found in the current image sets. 15X15-CR indicates an assembly with a control rod. The 15x15 types were made by Siemens. The 17x17 assembly, made by ABB Atom is a non-fuel assembly with Helium filled rods.

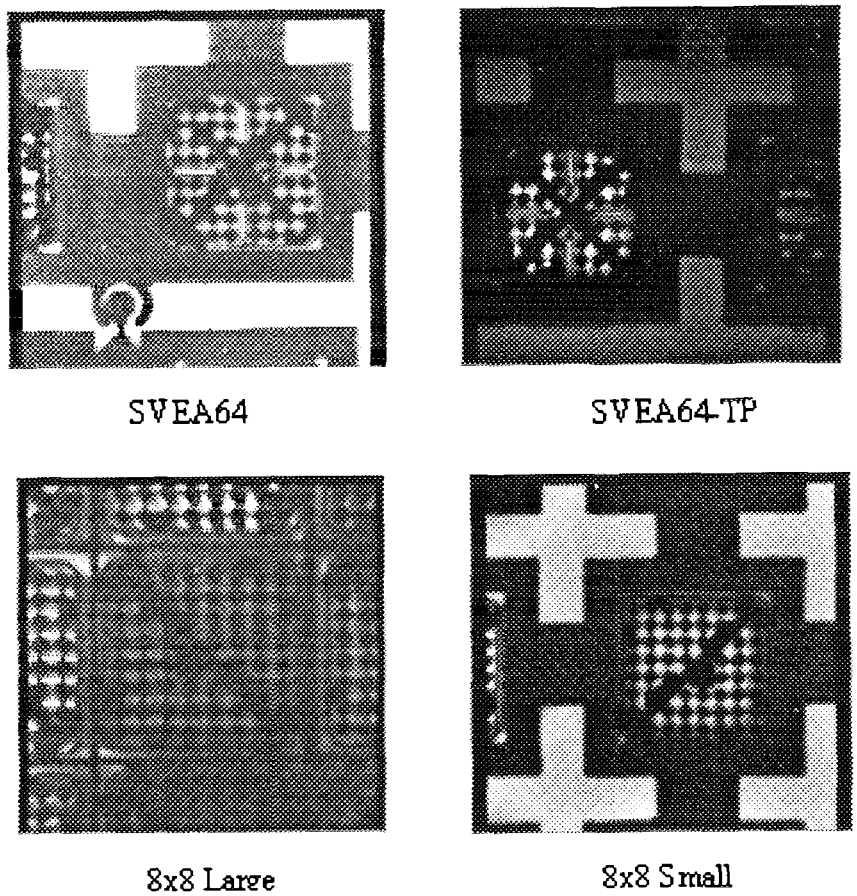

Figure 1(b) Examples of BWR assembly types, all made by ABB Atom, found in the current image sets. The SVEA64 types are same except the SVEA64-TP has a different top plate that obscures the outer pins. The 8x8 assemblies in the two lower pictures are the same type. However, in the CLAB images the assembly took up a larger area of the frame, as on the left, than in the first Ringhals images, shown on right.

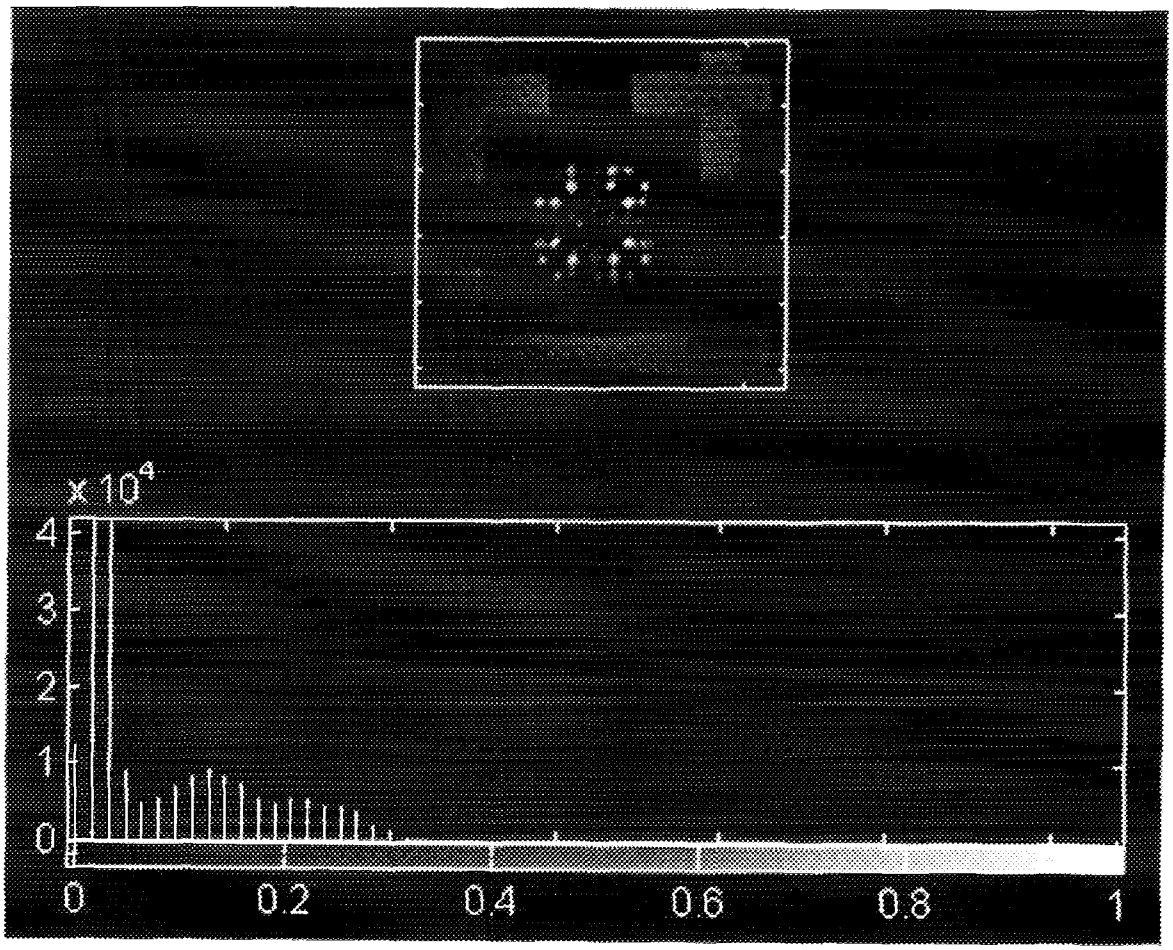

Figure 2 (a) Image of a SVEA64-TP type and a histogram of its grey scale pixel values.

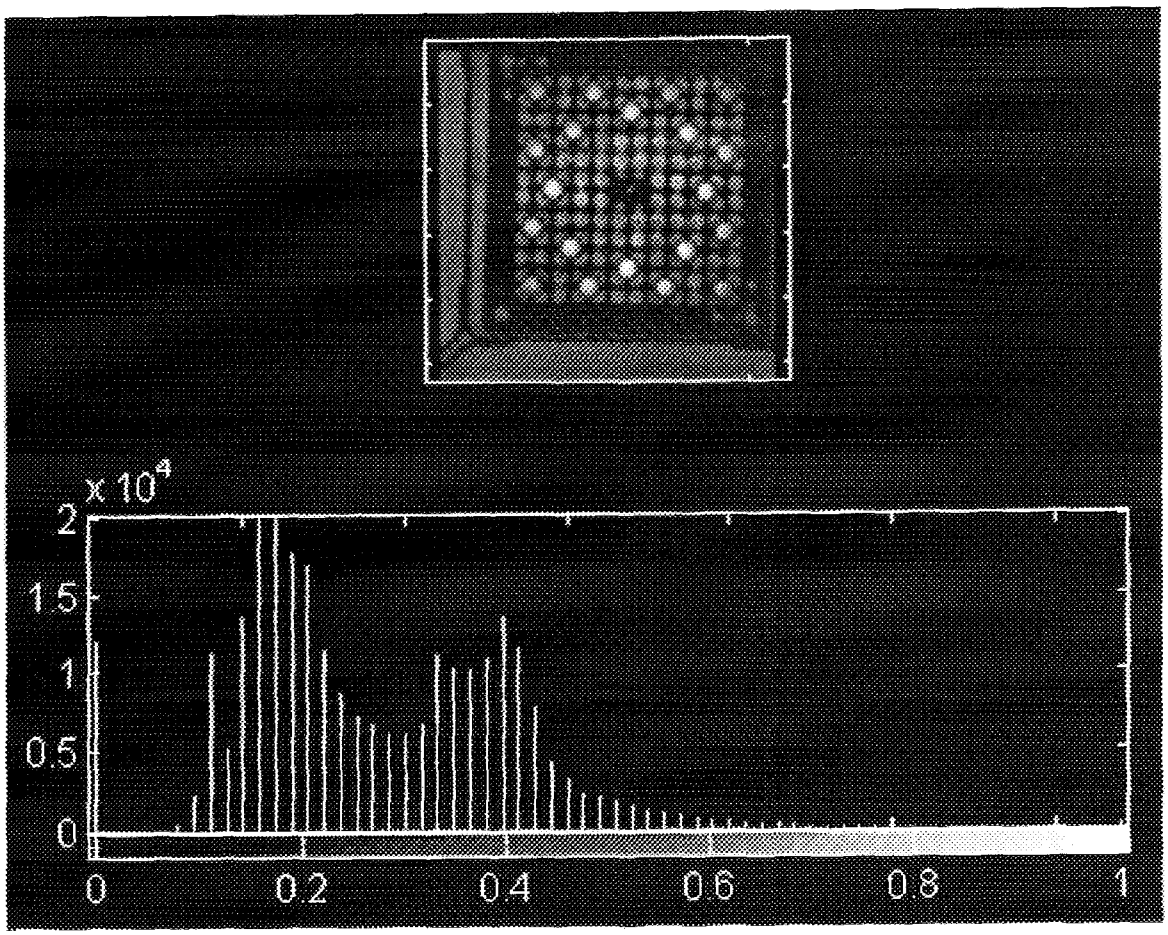

Figure 2 (b) Image of a PWR 15x15 type and a histogram of its grey scale pixel values.

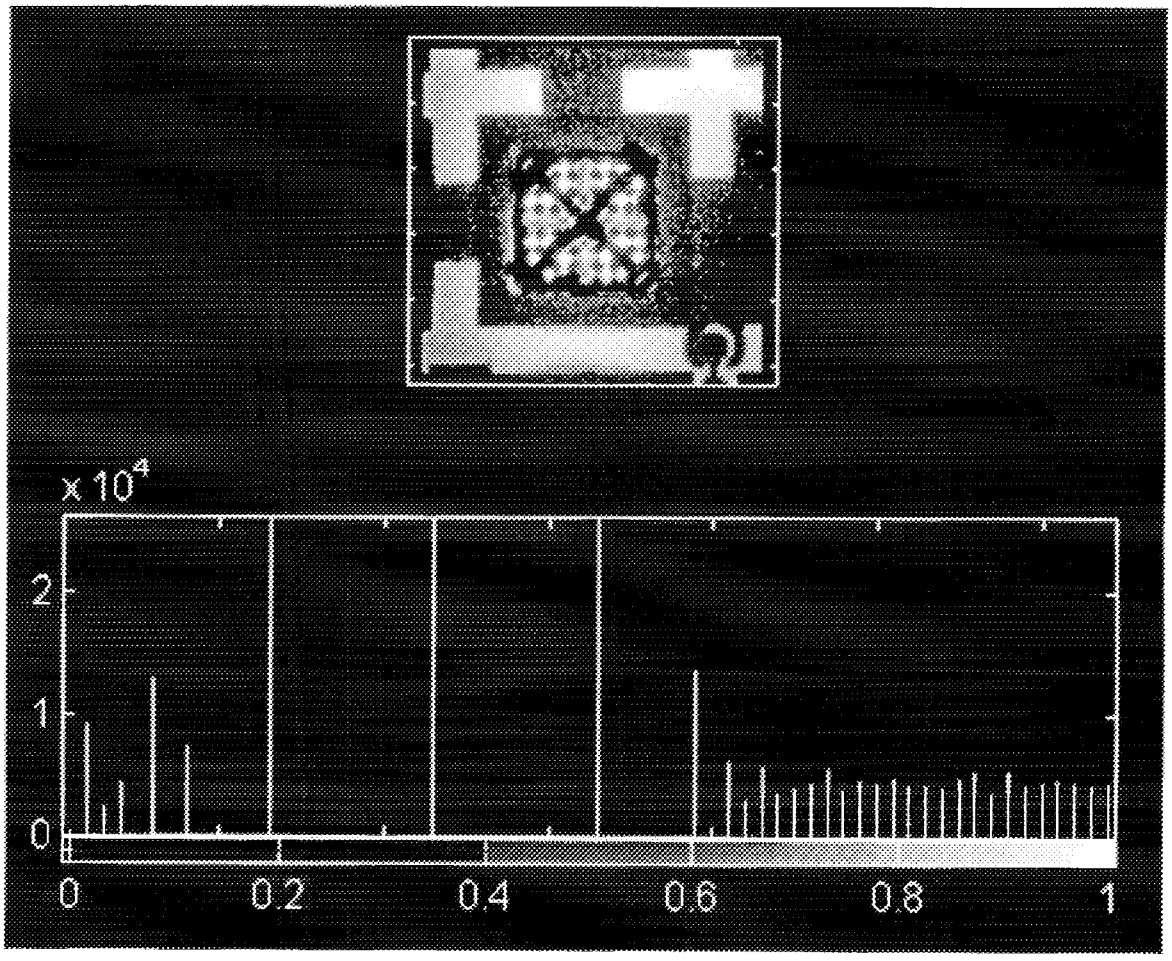

Figure 3 Equalised image and its histogram after histogram equalisation operation on image in figure 2(a).

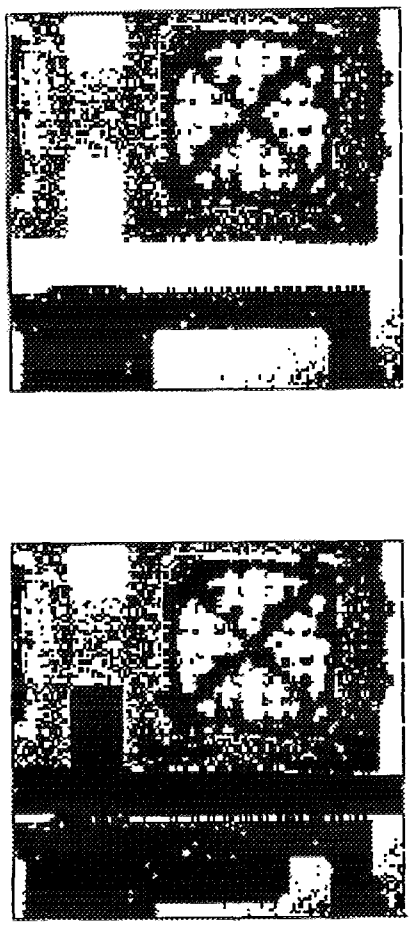

Figure 4 Blanking out operation. Top image shows a SVEA64-TP image after equalisation and binary operation. Lower image shows result of blanking out with a rectangle operator that looks for matching white areas.

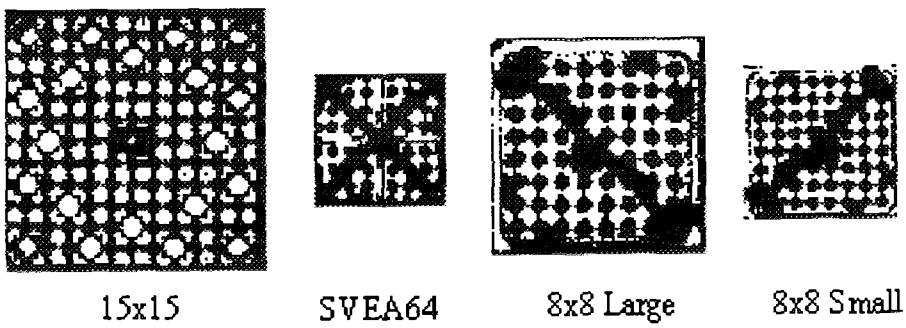

Figure 5 Templates to find assemblies in the images.

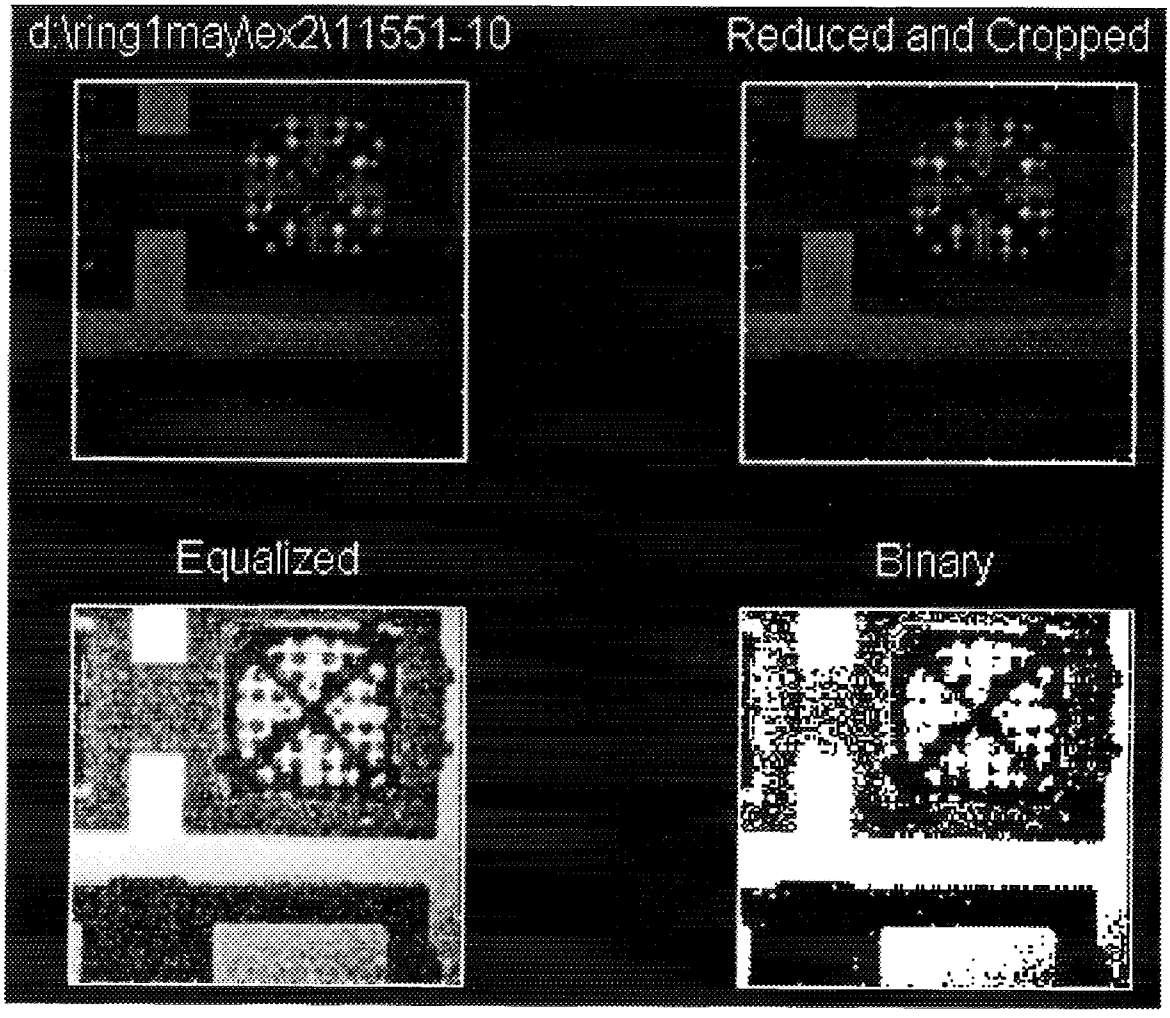

Figure 6 (a) SVEA64-TP assembly finding. Initial image is reduced by ¼ and cropped to eliminate black bands around edges. Histogram equalisation is then followed by a threshold cut at 0.5 (on 0-1 scale) to produce a binary B/W image.

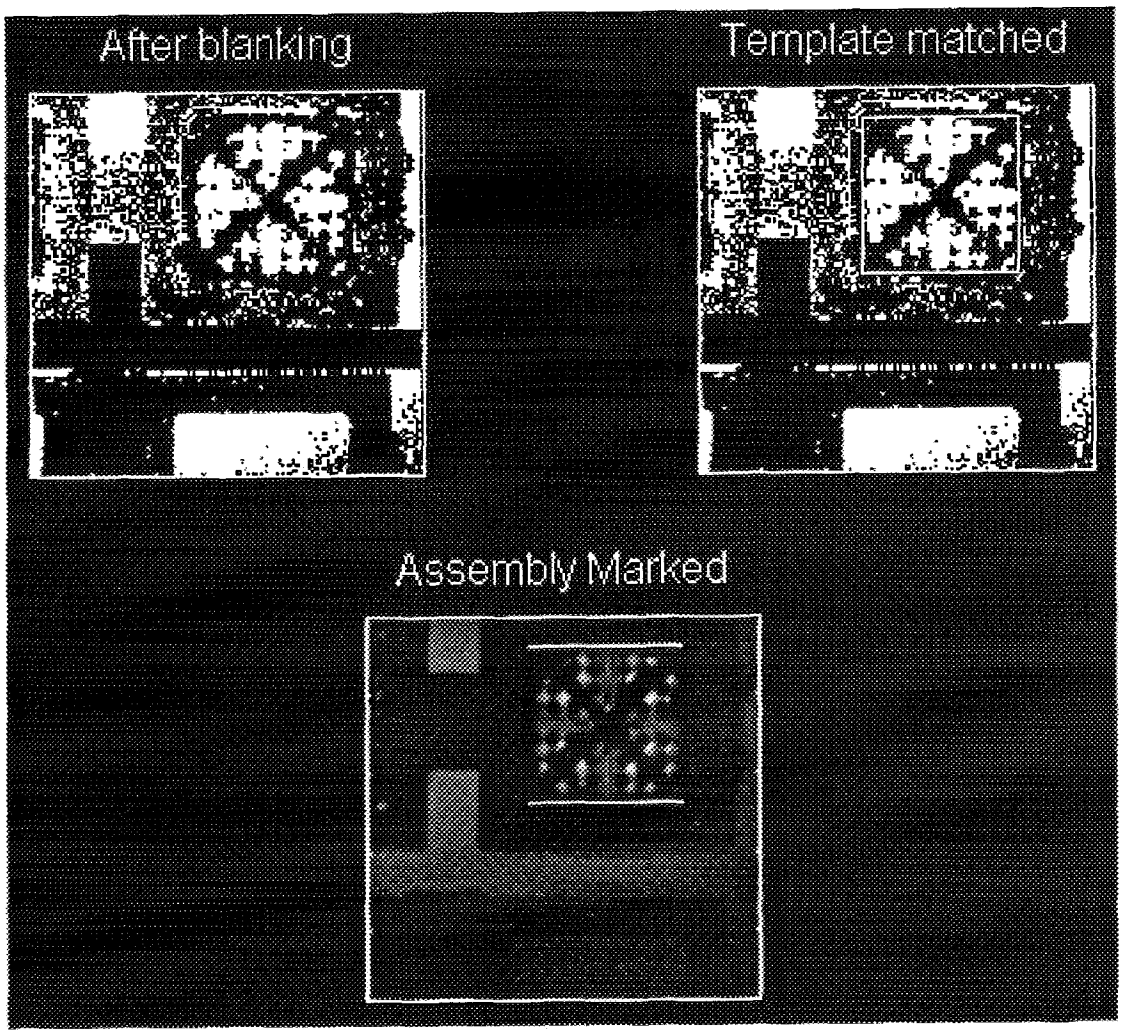

Figure 6 (b) Blanking out operation on binary image is followed by template matching. Matched co-ordinates are determined for the original image and indicated here.

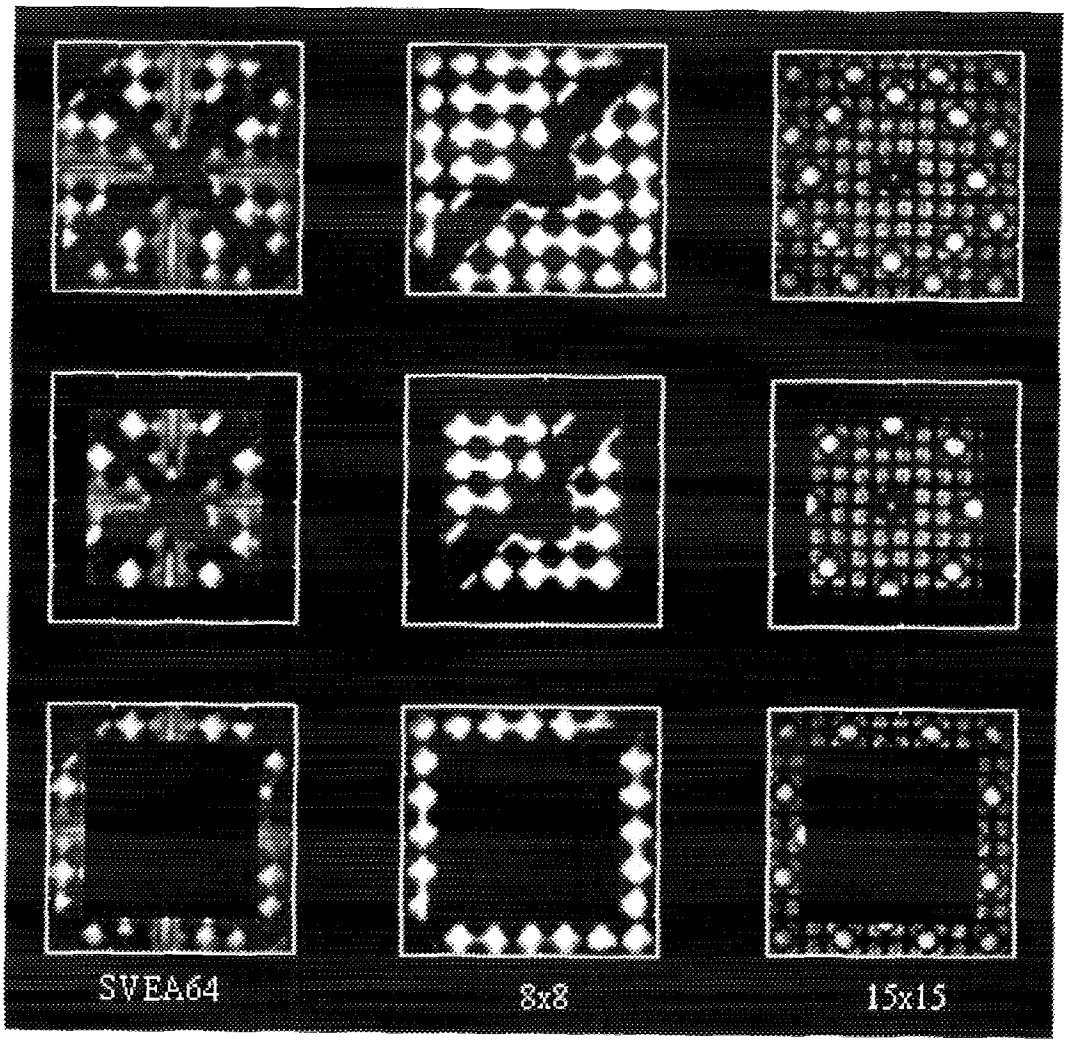

Figure 7 Cropped assembly faces are divided into inner and outer sections of equal areas. Ratio - R - of the inner to outer total light sums (i.e. sum of the grey scale values of 0 to 1 per pixel) is used to look for peaking in the light distributions.

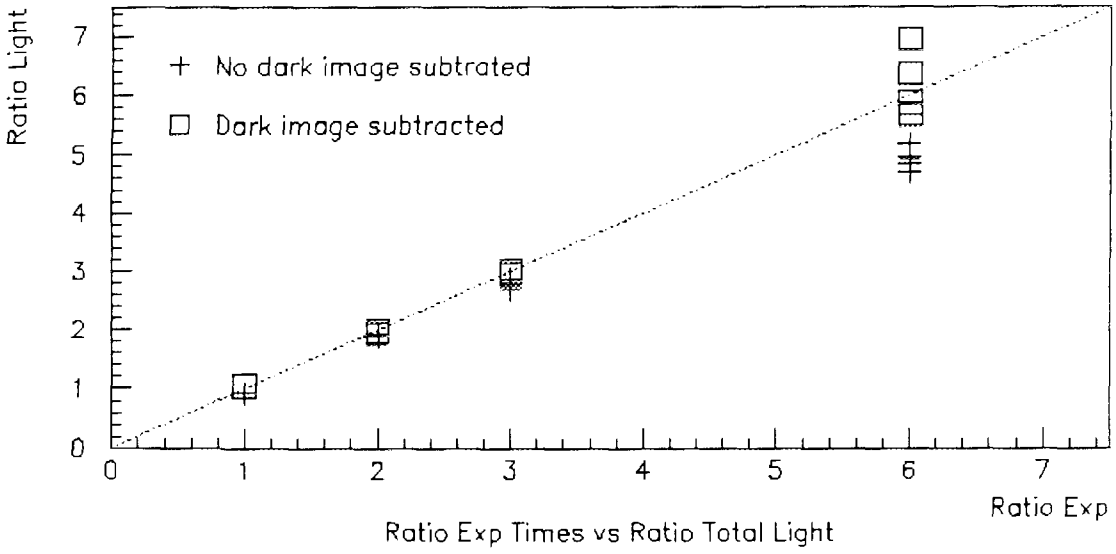

Figure 8 PWR 15X15 Ratio of the exposure times to ratio of total light for pairs of assembly images.

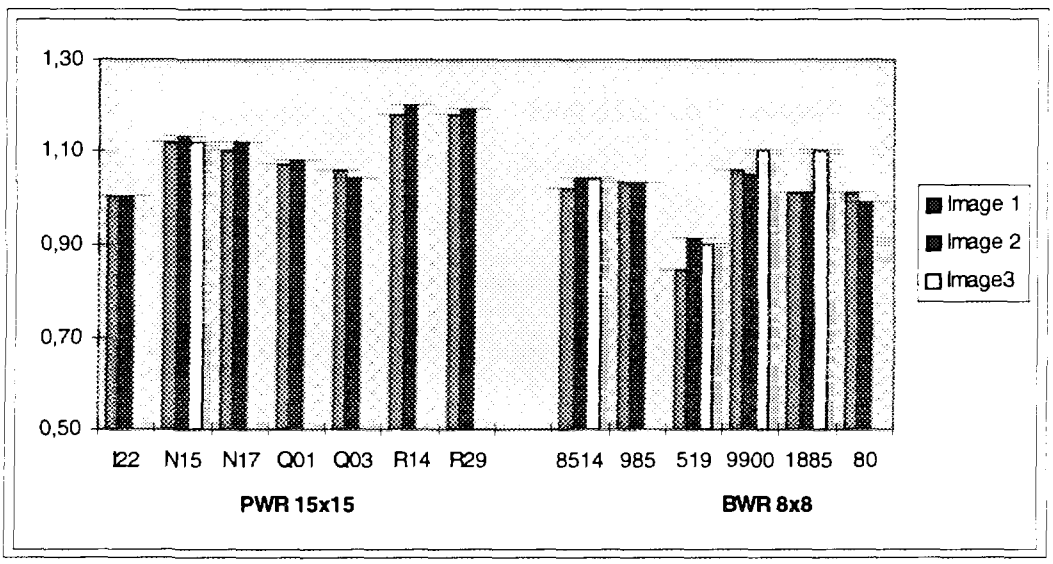

Figure 9 The ratio R for different images of the same assembly for seven PWR 15X15 assemblies and seven BWR 8x8 images (after dark image subtraction for both types.)

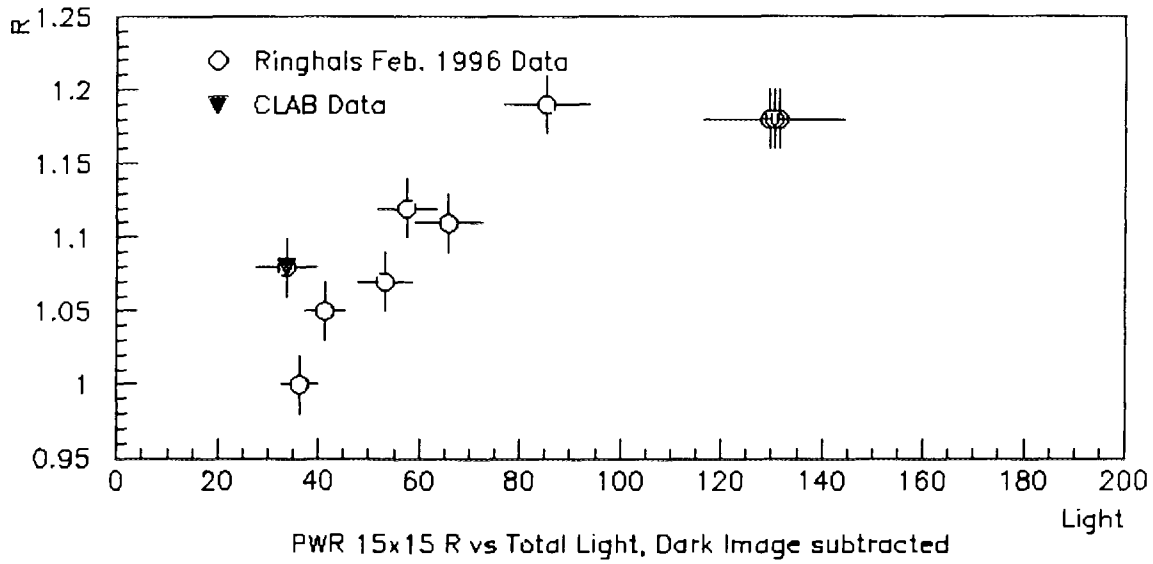

Figure 10 PWR 15X15: The ratio R *versus* total light (sum of grey scale pixels) with dark image subtraction.

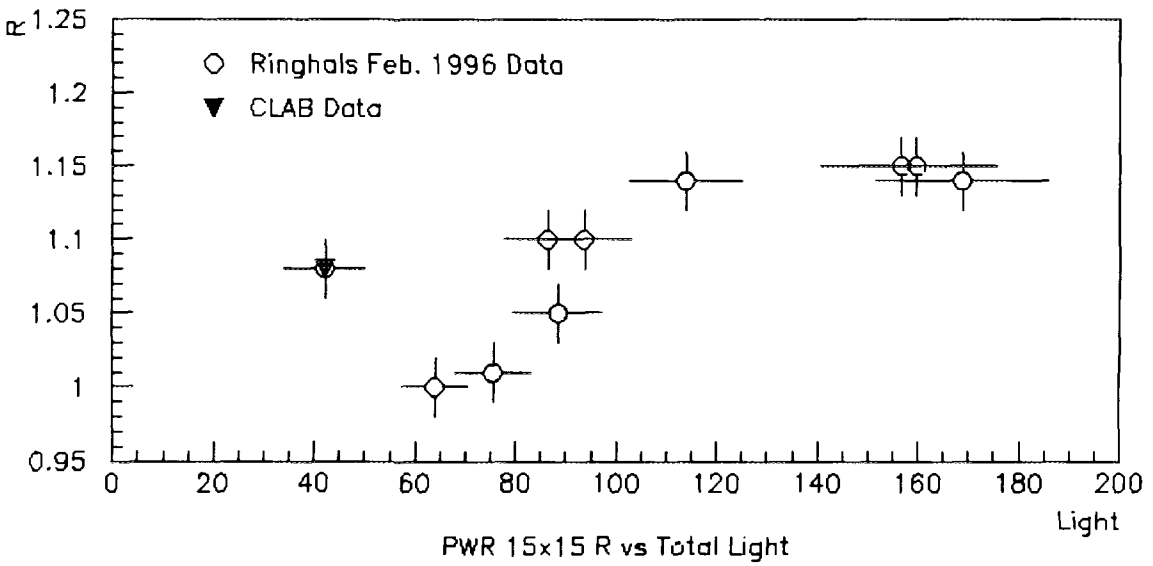

Figure 11 PWR 15X15: The ratio R *versus* the total light (sum of grey scale pixels) with NO dark image subtraction.

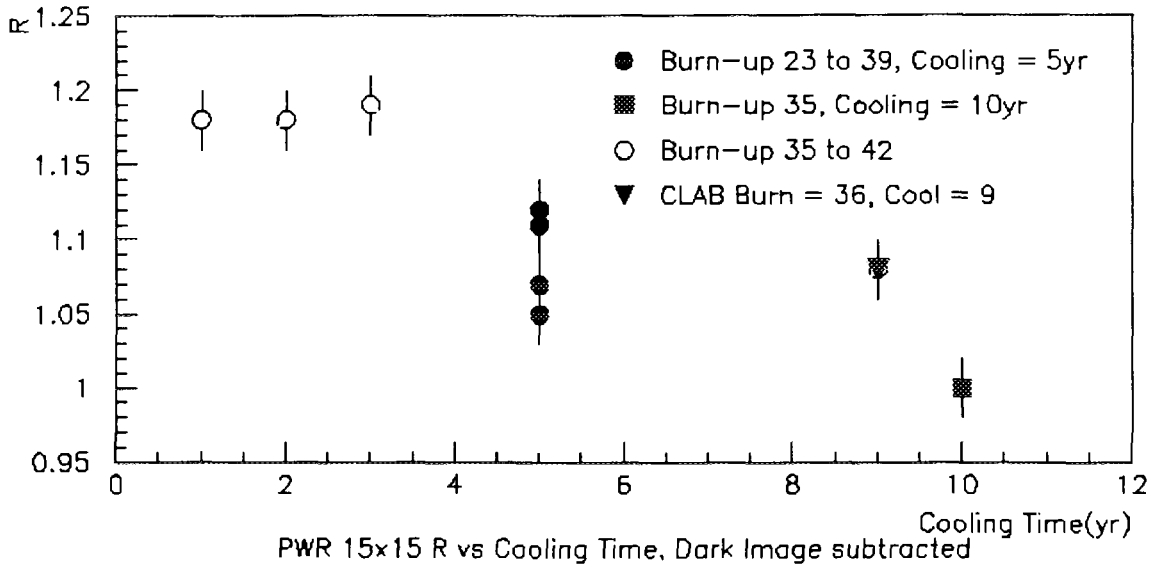

Figure 12 PWR 15X15: R *versus* cooling time with dark image subtraction.

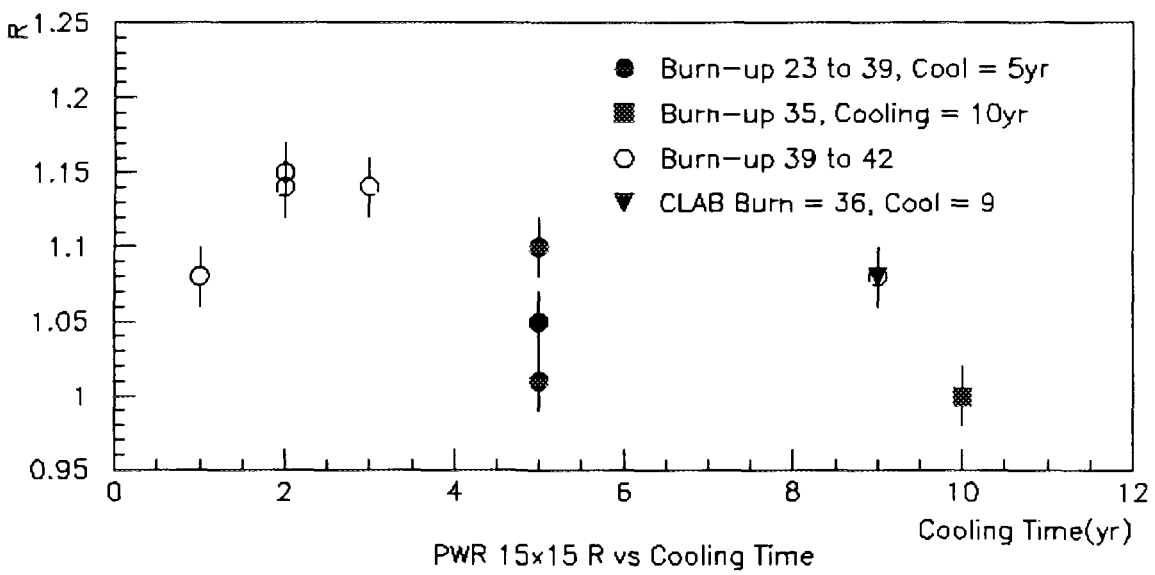

Figure 13 PWR 15X15: R *versus* the cooling time with NO dark image subtraction.

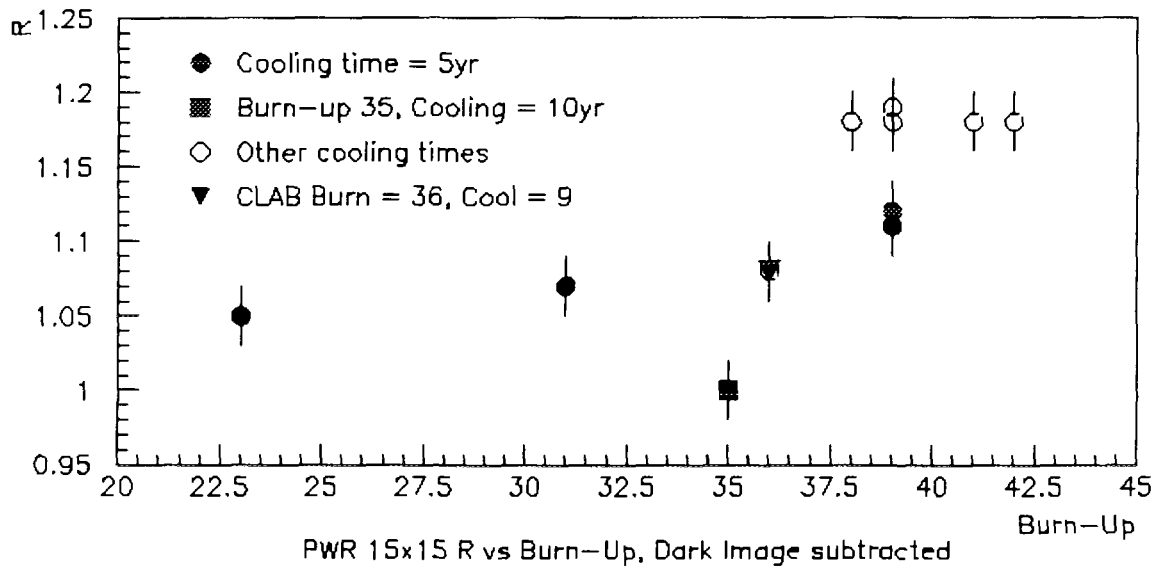

Figure 14 PWR 15X15: R *versus* Burn-up with the dark image subtraction.

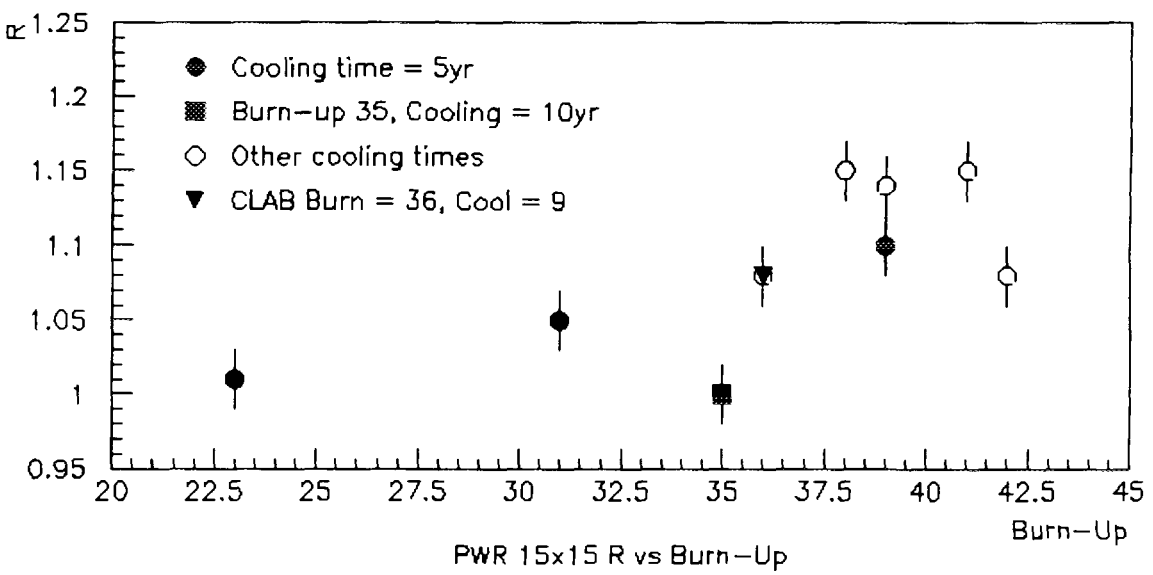

Figure 15 PWR 15X15: R *versus* burn-up with NO dark image subtraction.

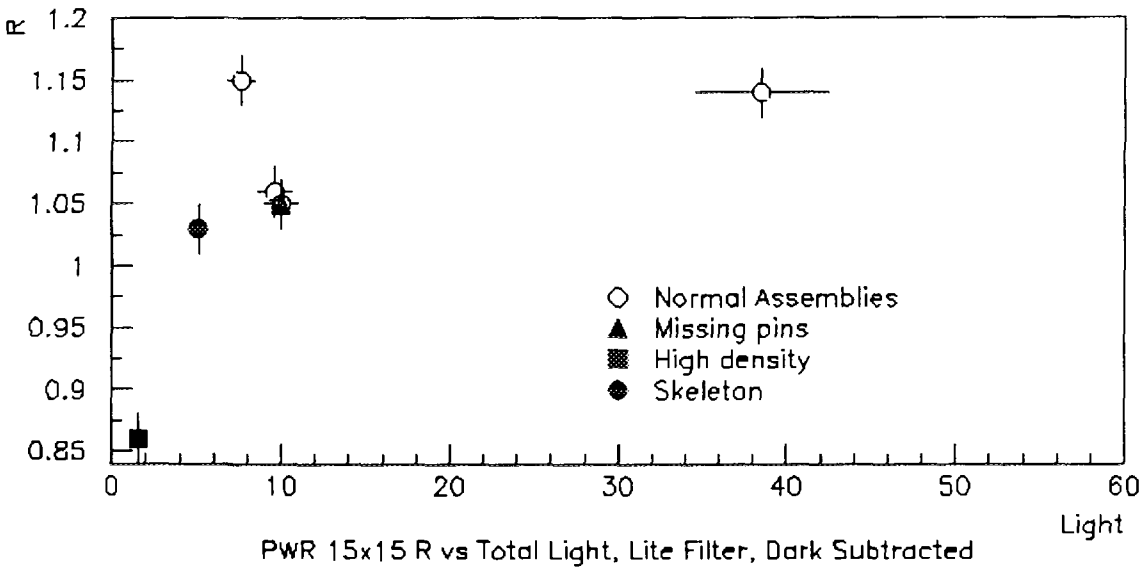

Figure 16 PWR 15x15 UV-Lite Filter: R *versus* total light with dark image subtraction.

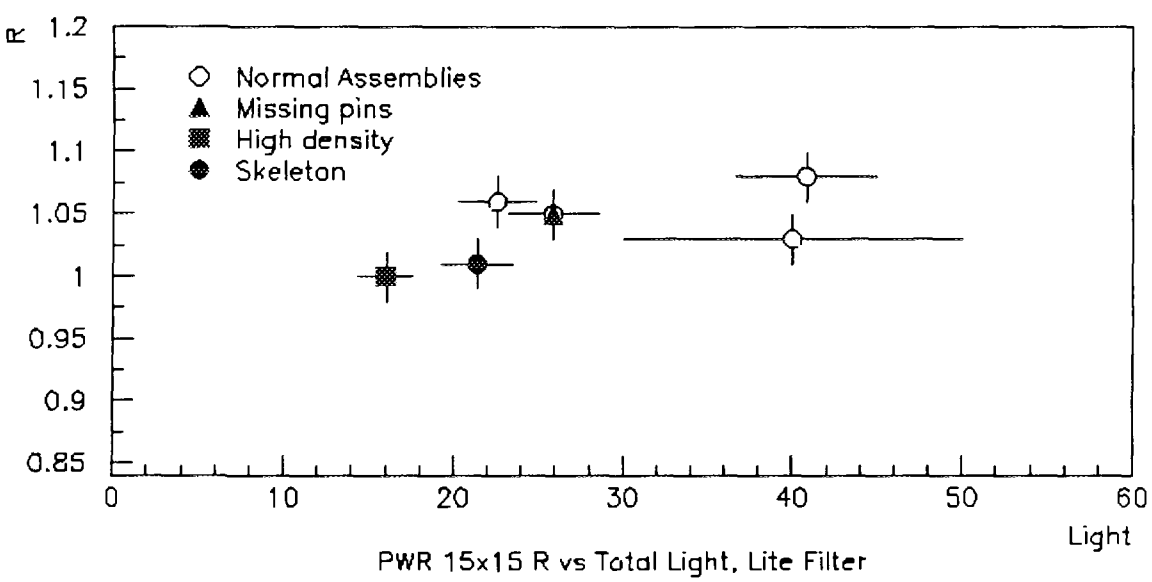

Figure 17 PWR 15x15 UV-Lite Filter: R *versus* total light with NO dark image subtraction.

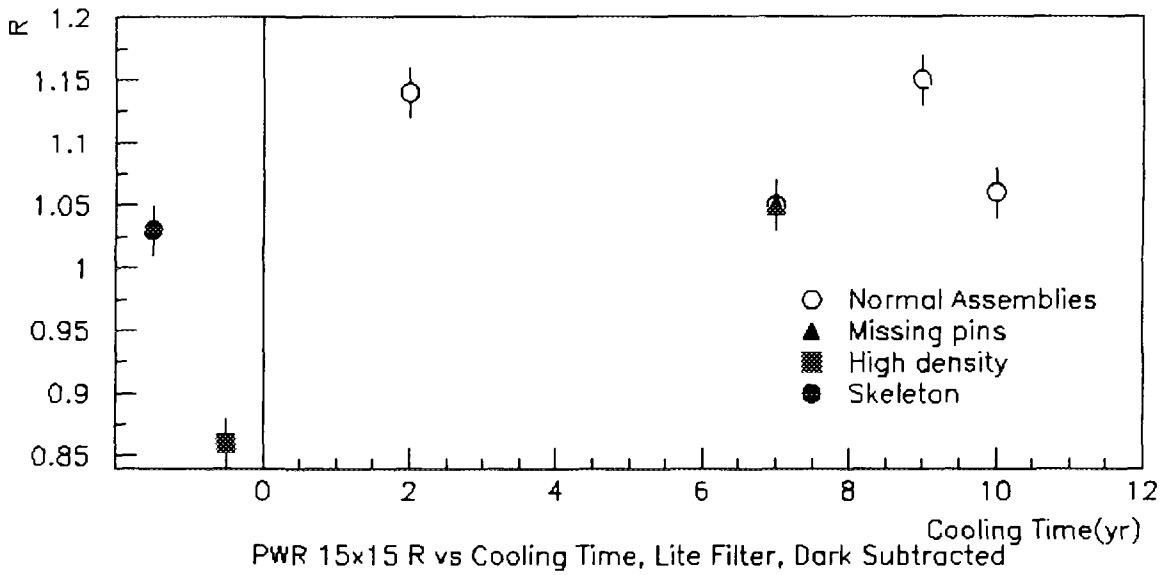

Figure 18 PWR 15x15 UV-Lite Filter: R *versus* cooling time with dark image subtraction.

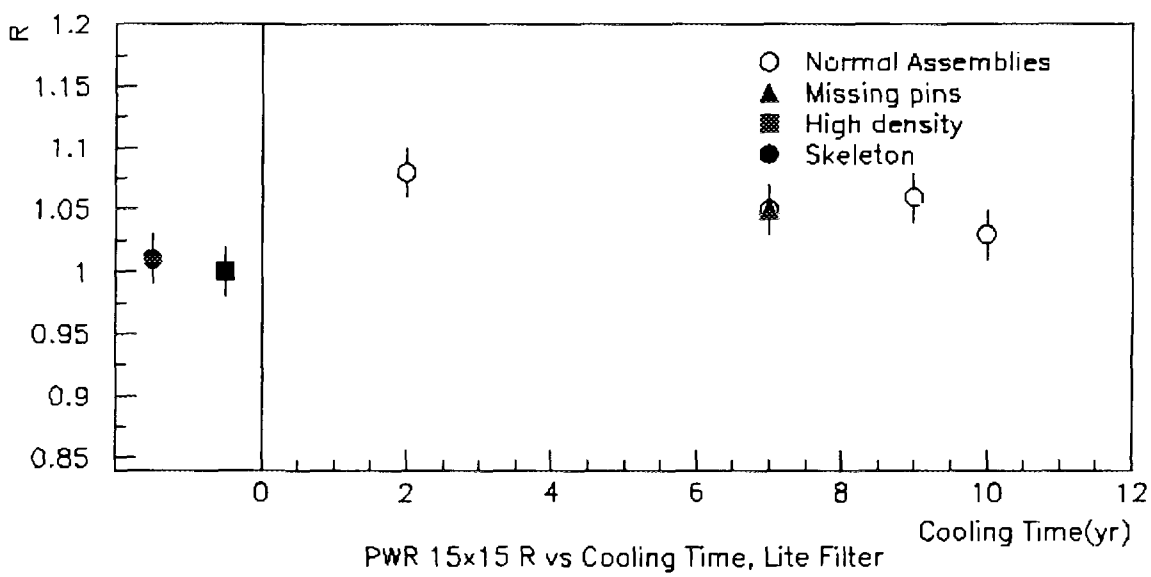

Figure 19 PWR 15x15 UV-Lite Filter: R *versus* cooling time with NO dark image subtraction.

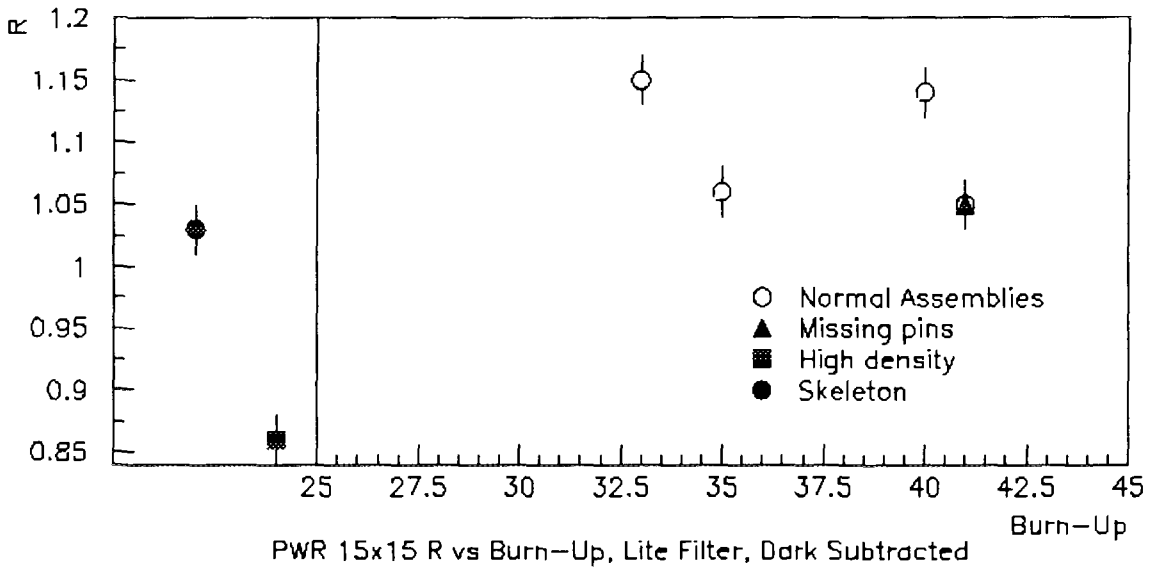

Figure 20 PWR 15x15 UV-Lite Filter: R *versus* burn-up with dark image subtraction.

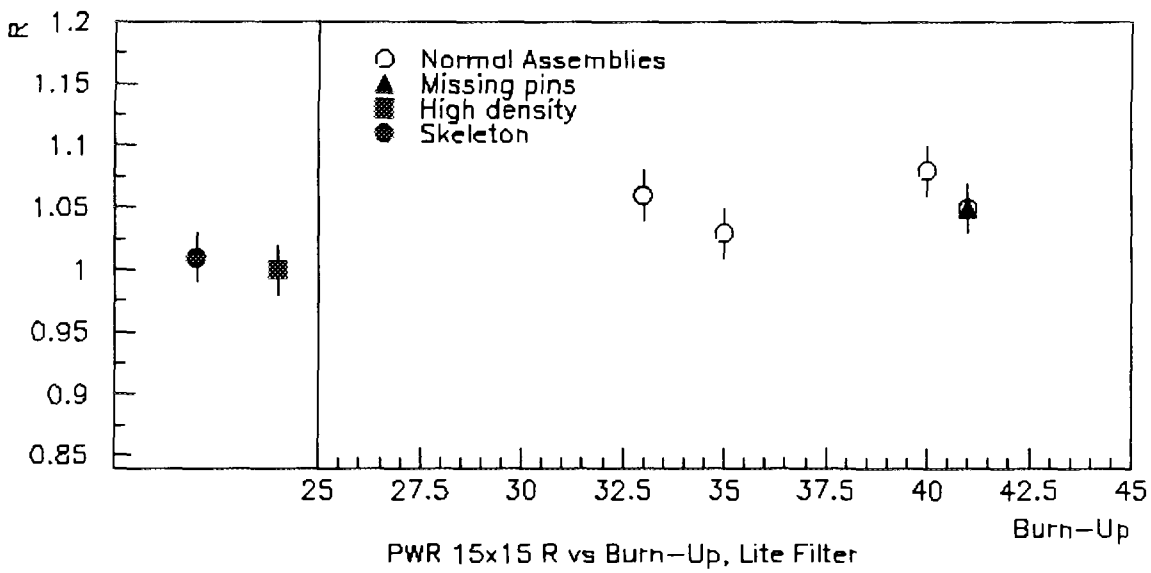

Figure 21 PWR 15x15 UV-Lite Filter: R *versus* burn-up with NO dark image subtraction.

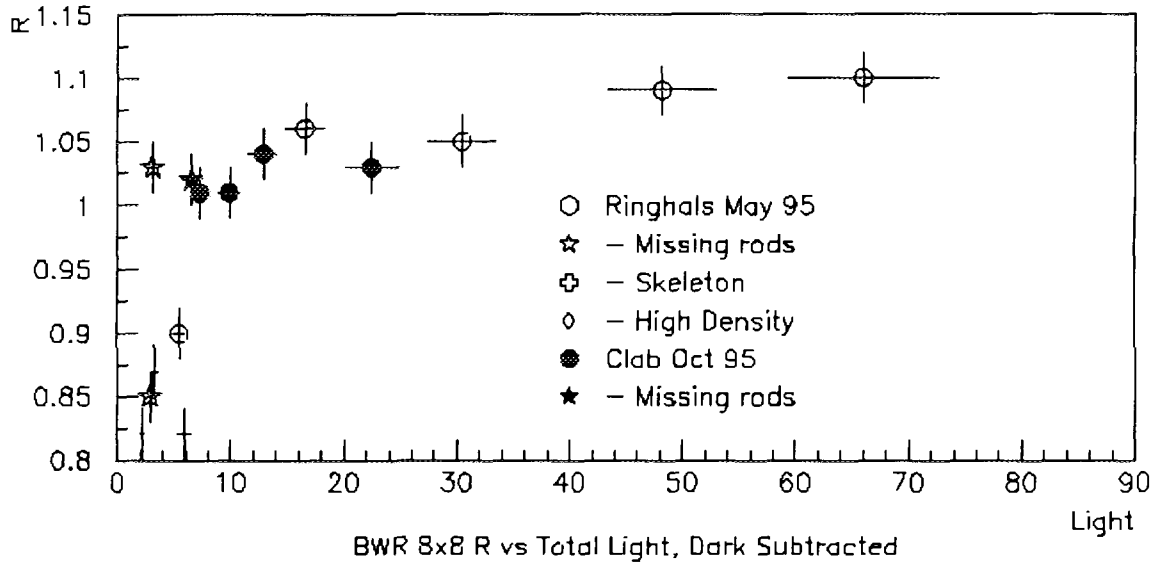

Figure 22 BWR 8X8: R *versus* total light with dark image subtraction.

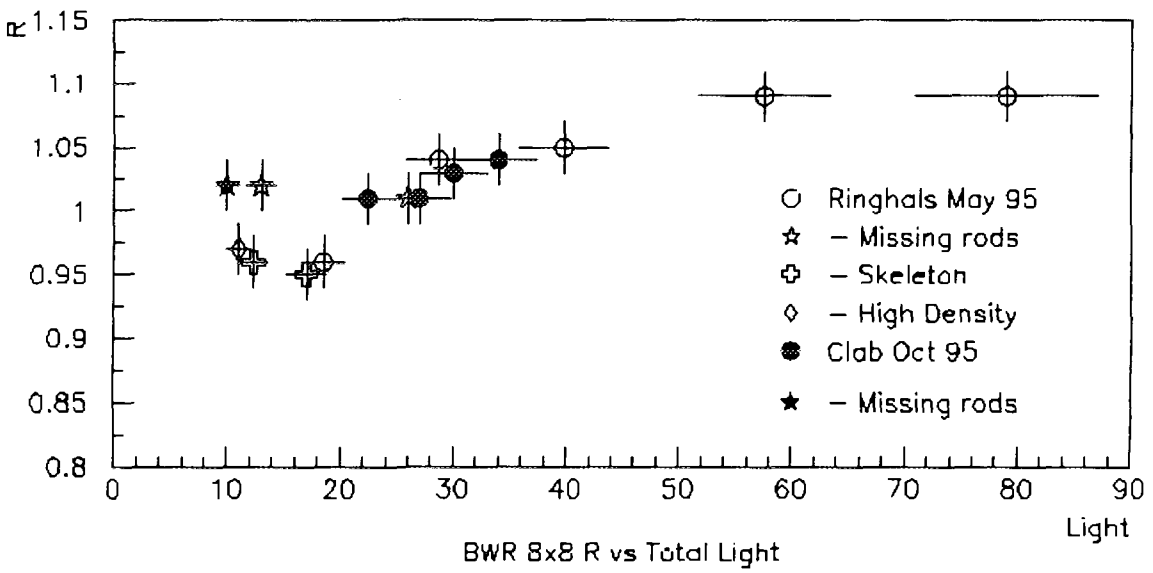

Figure 23 BWR 8X8: R *versus* Total light with NO dark image subtraction.

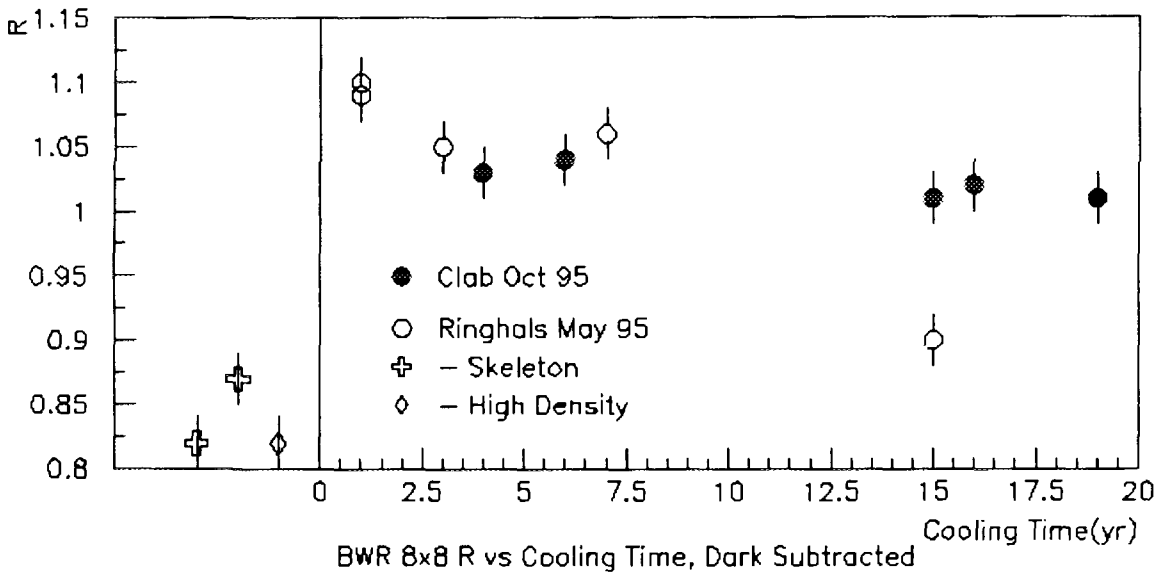

Figure 24 BWR 8X8: R *versus* Cooling time with dark image subtraction.

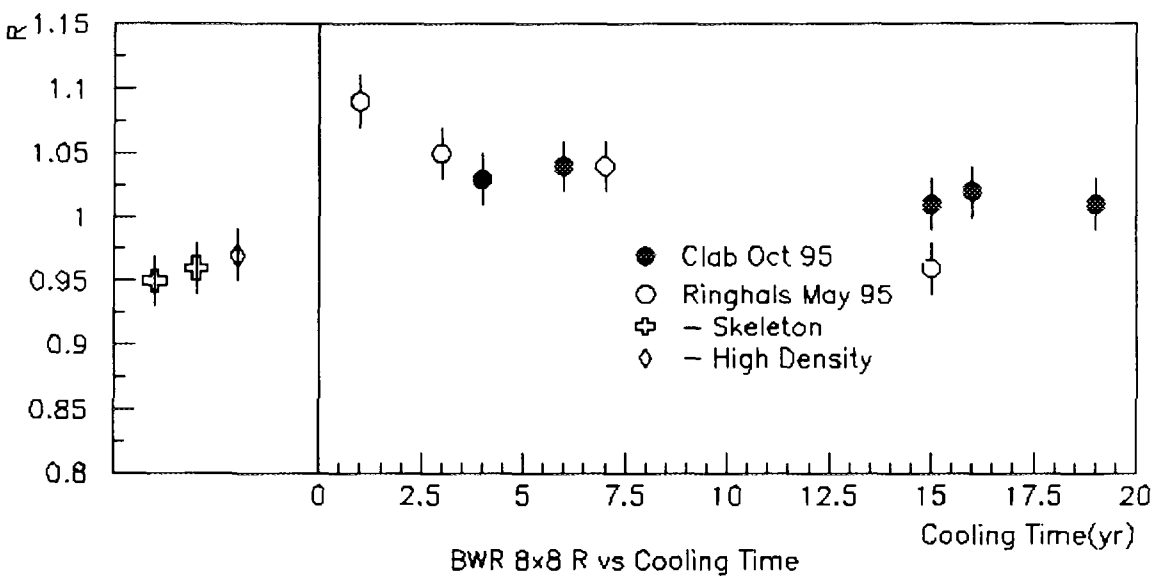

Figure 25 BWR 8X8: R *versus* Cooling time with NO dark image subtraction.

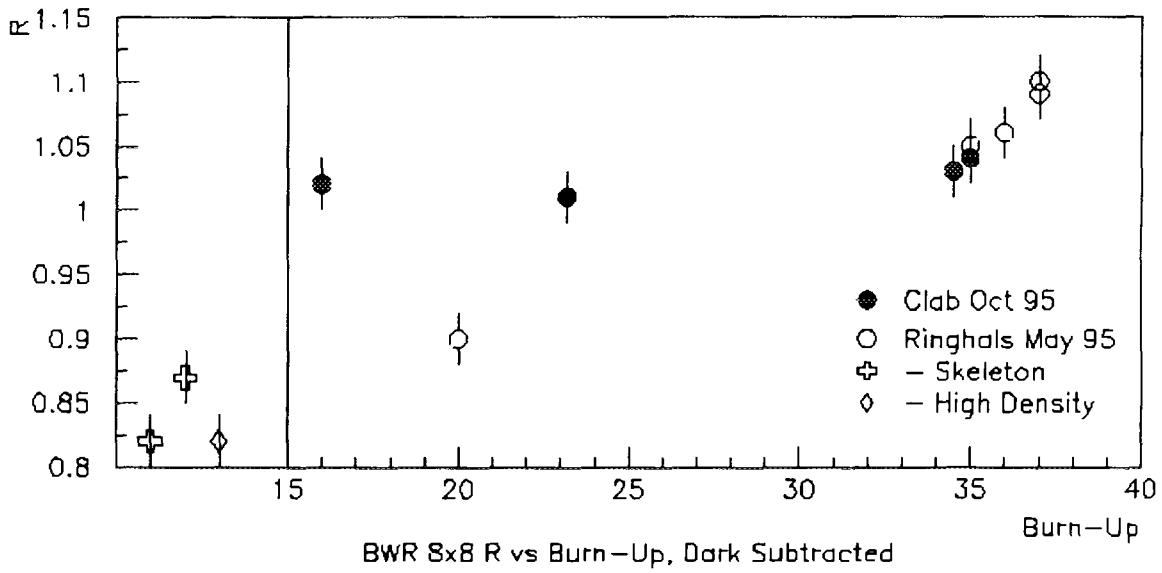

Figure 26 BWR 8X8: R *versus* Burn-up with dark image subtraction.

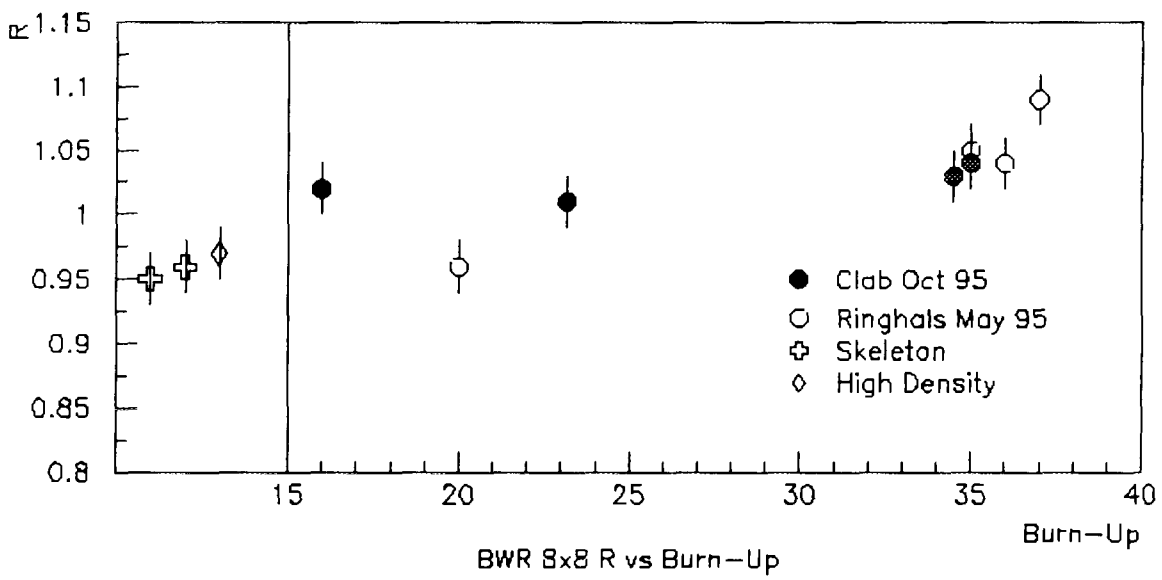

Figure 27 BWR 8X8: R *versus* Burn-up with NO dark image subtraction.

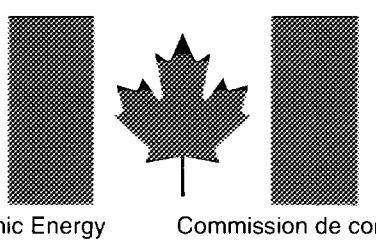

Atomic Energy Control Board

Commission de contrôle de l'énergie atomique

Postal address
SKI
SE-106 58 Stockholm
Sweden

Postal address
AECB
P.O. Box 1046, Station B
Ottawa
Canada K1P 5S9