XMM-LSS manuscript no.
(will be inserted by hand later)

Your thesaurus codes are:
missing; you have not inserted them

XMM-LSS
INTERNAL
REPORT n°2-Mi

May 16, 2006

Using stacking to compute rates of XMDS sources

L.Chiappetti[1]

INAF, IASF Milano, via Bassini 15, I-20133 Milano, Italy

For internal circulation

Abstract. I report on a simple method to characterize X-ray sources which are observed in the XMDS in several adjacent fields (so called *duplicated*), by stacking data from the different field images. The results should provide rates and fluxes with smaller errors and better significance, although avoiding a complete combined re-analysis at event file level. They may also be used to assess properties of undetected sources.

Key words: LSS – XMDS –

1. Introduction

We presented in Chiappetti et al. (2005) (hereafter Paper I) a catalogue of 286 XMDS X-ray sources with a significance above 4σ in the VVDS area, and I extended this in Chiappetti et al. (2006) (hereafter Report I) to 1147 X-ray sources in the entire area at any significance above the XMDS probability threshold.

While it is known that there is a large overlap between adjacent XMM pointings in the XMDS (and LSS), so far the detection and source characterization has been run independently on each field. The only action taken about *duplicated* detections in overlapping fields has been to flag them (as explained e.g. in Section 3.3 of Report I), and consider the "best" detection as primary, ignoring the other(s).

Since the database contains all information about which "secondary" detection(s) are associated with which primary one, it is however possible to try to combine the information in order to improve the statistics. This does not mean repeating the *detection*, but just doing a *characterization* which uses data from all fields involving a given source.

In section 2 I give the count of duplicated sources, a recipe to locate them, and some technical information about how they have been managed so far.

In section 3 I summarize the source characterization procedure used in Paper I and Report I, and describe the new stacking procedure.

In section 4 I illustrate the results achieved comparing rates and fluxes of combined measurements with the original ones.

In section 5 I describe the modification proposed at database level to ingest the new results.

2. Duplications in the current database

I use as starting point the *glorified correlation table* for the full XMDS catalogue described in Report I. There are 1322 detections, corresponding to 1147 primary sources. Each distinct "primary" X-ray source will have for sure, besides any other eventual entry, a *single* entry with either rank 0 or rank 1. Therefore the following `WHERE` condition will locate the entries with at least one duplication, namely the association with the "preferred" duplication :
`(rank between 0 and 1) and xmdsdup is not null`

Such condition returns 159 sources.

A few sources have more than one duplicate, i.e. they fall in the intersection of three XMDS fields (in a single case, source #704, in the intersection of four). These sources have "hidden" entries in the database with autorank=6. To return the relevant source number one has to issue a `SELECT DISTINCT xmdsepic` with a condition : `autorank=6 and xmdsdup is not null`.

One can even "or" the two conditions and produce a cumulative list. The 159 duplicated sources are 144 doubles, 14 triples and 1 quadruple, corresponding to 175 detections.

Fig. 1 illustrates the conventions used in the glorified correlation table for the entries relevant to duplicated sources. Currently the database still contains information on all detections, although only the records with non-technical ranks on primary detections are normally used. But until the historical records are not deleted, it is possible to change the status of any given detection.

3. The source characterization

I remind that the Milan pipeline is based on the procedure developed by Baldi et al. (2002), which separates the source detection stage from the source characterization.

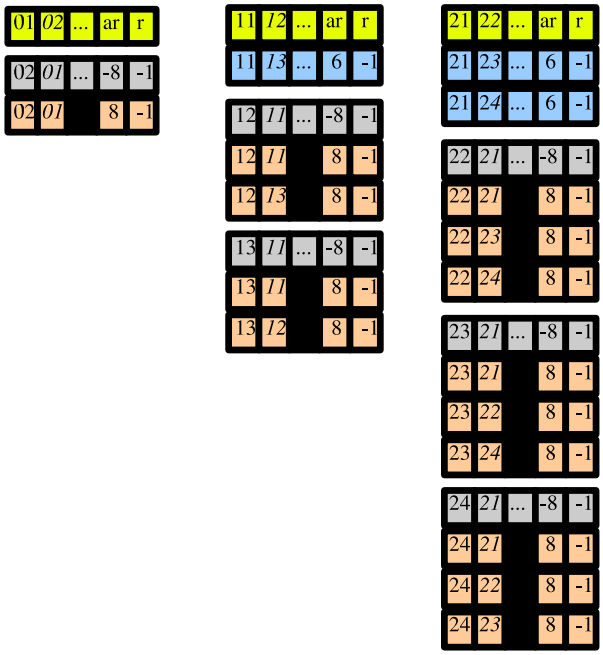

Fig. 1. Scheme of the entries present in the glorified correlation table for duplicated sources. The different columns illustrate the case of double, triple and quadruple sources.
The present panel illustrates the current arrangement of the database. The proposed modification is shown in Fig. 13.
Each family of records is represented by one of the rectangular coloured boxes with five cells. The cells represent : **(1)** the first cell is the `xmdsepic` source number of the primary source ; **(2)** the second cell is the `xmdsdup` source number of an associated duplicate detection ; **(3)** the ellipsis in the third cell indicates the body of the records with the pointers to the other optical and radio catalogues; such cell is black when all pointers are set by definition to null ; **(4)** the penultimate cell is the *autorank* ; and **(5)** the last cell is the *rank* described in Report I.
The colour coding of the cells is as follows. It also indicates how many records of a given family can be present (the *multiplicity* of the record family) : **(a)** Yellow indicates *one or more* records associated to the primary source, of which one will surely have rank=0 or rank=1, while there may be other entries with rank=2 (ambiguous) or -1 (identification rejected to be later removed) ; **(b)** For each detection which is a "secondary" duplication, there can be *one or more* records kept for historical reasons, but flagged with autorank=-8 : these are indicated in violet-gray ; **(c)** For each secondary detection there is a *single* entry which is assigned autorank=8, which "points back" to the primary; The information on the counterpart has been nullified: These are the pink entries with a black cell in third position ; **(d)** For triple and quadruple sources there are also the autorank=6 (blue) entries: there are as many blue entries as yellow entries, since they have the same content, except for `xmdsdup` which points to the other secondary (or secondaries for a quadruple).
The examples shown correspond to an hypothetic source 01 duplicated with 02, to a "triple" source 11 duplicated with 12 and 13, and to a "quadruple" source 21 associated with 22, 23 and 24.

Therefore I have not repeated the source detection, but exploited the intermediate results from the original pipeline, and just replaced the characterization stage. In particular this means that different fields are not combined together to obtain an improved detection list, and that the LogN–LogS discussion made in Section 4 of Paper I (which made reference to the full XMDS dataset for this purpose, contrary to the rest of paper, limited to the VVDS 4σ sample) is still applicable.

3.1. Summary of the standard procedure

The procedure by Baldi et al. (2002), after cleaning and merging of the event files, runs a very coarse preliminary detection which also produces preliminary background maps. It then runs a further detection, which produces a *candidate input list*, and, separately, prepares some cleaned *background maps* as well as a map of the *flux conversion factors* (CFs).

The source characterization algorithm runs for each source in the candidate input list, using the photon images, background maps, CF maps and exposure maps, and extracts counts, count rates (and fluxes and hardness ratios HR), S/N ratio and probability.

Only detections characterized by a probability $p < 2 \times 10^{-4}$ in at least one band are accepted to be ingested in the database.

Namely the source characterization program `circledetect` works as follows:

- for all 5 energy bands at the same time
- for each source position in the candidate input list in turn
  - integrates *gross* image counts in a circular area encircling 68% of the local PSF
  - gets from the precomputed background map the local background counts previously integrated in the same area
  - gets from the precomputed exposure maps the average camera exposure at the local position on the same area
  - computes the *net* counts, the associated error, the S/N and probability starting from gross and background counts
  - computes count rates dividing net counts by the total exposure (sum of the three cameras) and converts to flux using the local CF read from the precomputed CF maps
  - computes HRs from count rates

The details (with formulae) are given in Section 3 of Baldi et al. (2002); in particular the flux conversion factor is weighted on the local exposures

$$CF = \frac{\sum_i T_i}{\sum_i \frac{T_i}{CF_i}}$$

where the index i runs on the three cameras.

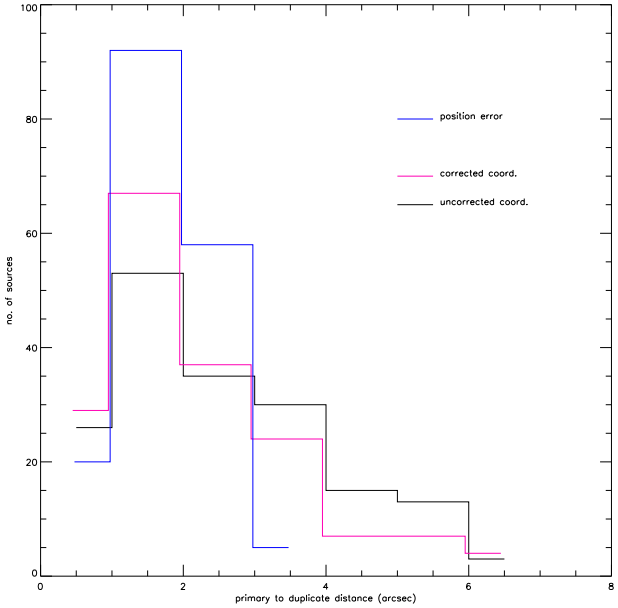

Fig. 2. Frequency histogram of the distance between X-ray coordinates of a primary detection and its duplicated detection(s), in raw coordinates (black) and in astrometrically corrected coordinates (magenta). The blue histogram shows, for comparison, the distribution of the combined position errors (nominal `emldetect` errors).

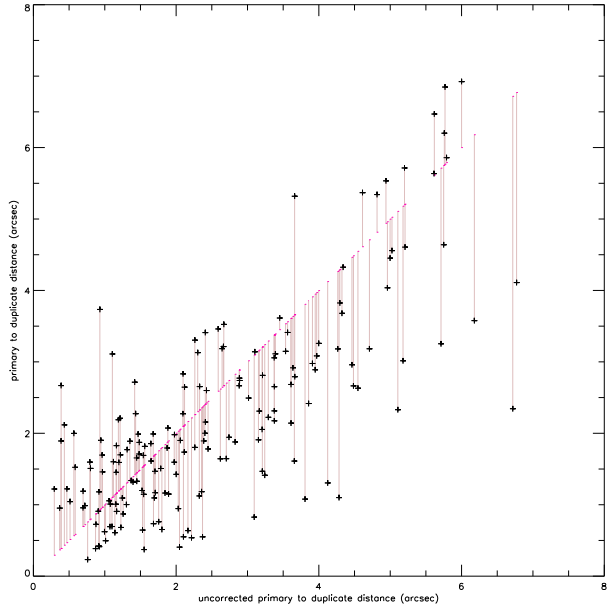

Fig. 3. Effect of the astrometric correction on the distance between the X-ray positions of primary and duplicated detections. The thin pink line connects the distance between uncorrected positions (coloured dot) to the one between astrometrically corrected positions (black cross).

3.2. The stacking procedure

I have replaced the characterization program with a modified program called `multistack`, which uses as input the same images, background maps, exposure maps and CF maps of `circledetect`.

However the program, instead of reading an input source list, is controlled by a `control.dat` file, which is simply a list of pointings (fields), in the form of a subdirectory. Each subdirectory contains all the images relevant to a given field, as well as a field-specific input list.

The program proceeds as follows :

- for all 5 energy bands at the same time
- initializes counters for gross counts, background counts, camera exposures and "partial" CFs
- for each field in `control.dat`
  - for each position in the field-specific input list
    - integrates *gross* image counts in a circular area encircling 68% of the local PSF *and adds them to the gross count counter*
    - gets from the precomputed background map the local background counts previously integrated in the same area *and adds them to the background count counter*
    - gets from the precomputed exposure maps the average camera exposures at the local position on the same area *and adds them to the exposure counters*
    - computes a T/CF value, where T is the sum of the three camera exposures in the band, and CF comes from the CF map, i.e. has already been weighted for the relative exposures in the 3 cameras ; such value is added to the "partial" CF counter
  - when the loops on positions and fields are over
- computes *one value* for the *net* counts, the associated error, the S/N and probability starting from gross and background summed counters
- computes a weighted CF (see further below) dividing the total exposure by the "partial" CF counter
- computes count rates dividing net counts by the total exposure (sum of the three cameras) and converts to flux using the weighted CF
- computes HRs from count rates

Therefore providing a single set of results per run.

The computation of the weighted CF uses a formula analogous to the one used to average on cameras (and in which the "partial CF counter" is the numerator) :

$$CF = \frac{\sum_j T_j}{\sum_j \frac{T_j}{CF_j}}$$

but the index j runs on all source positions in the input lists of all fields.

It has to be noted that `multistack` can be used in two altogether different modes.

In the first mode (used here) the control file contains entries for 2, 3 or 4 fields according to the number of overlaps (the multiplicity of the duplicated source), and

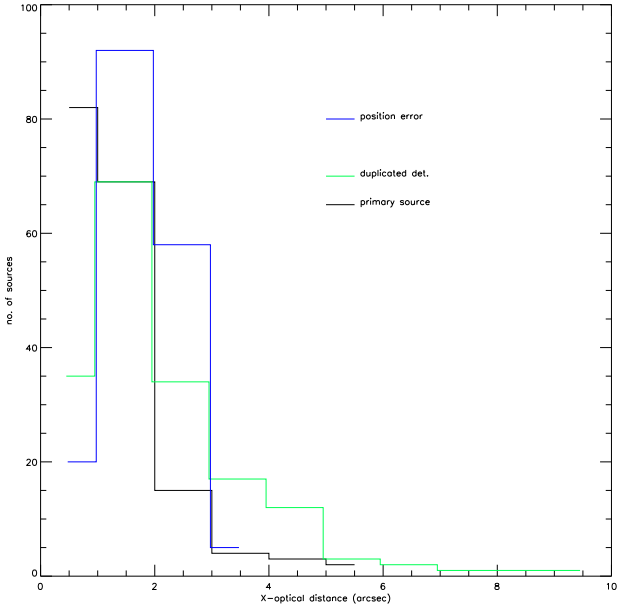

Fig. 4. Frequency histogram of the distance between the X-ray position and the position of the best optical or IR counterpart, for the primary detections (black) and the duplicated detection (green). The distribution of position errors (the same shown in Fig. 2) is also shown (in blue).

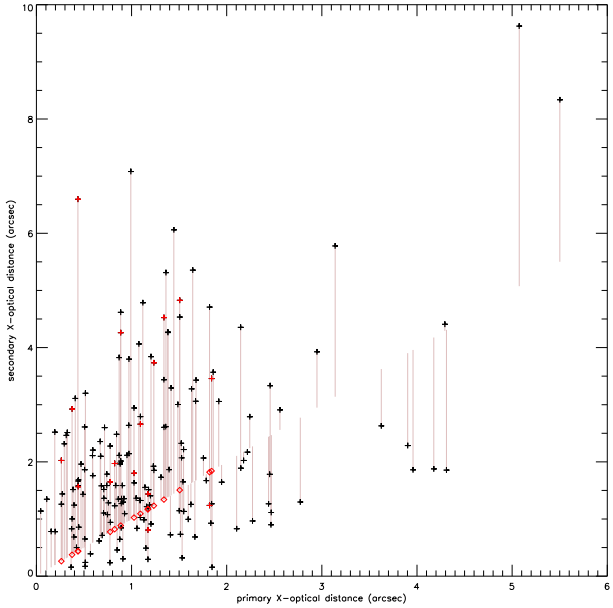

Fig. 5. X-ray to optical distance for duplicated detections as a function of the X-ray to optical distance for the primary detection. The thin pink line joins the value of the distance for the primary detection (on the diagonal) to the value (crosses) for the duplicated detections. The red crosses and diamonds mark the duplicated and primary detections with multiplicity greater than 2 (autorank=6).

each field-specific input list consists of a *single* position, which is relevant to the individual detection in such field (and is extracted from the full input list of `circledetect` used at database population time, using the stored (field-dependent) `id` as key. Otherwise said, the surroundings (within 68% of the *local* PSF) of the different detections are stacked, without repeating the detection, and irrespective of the sky coordinates to which they correspond (integration occurs in pixel space around individual field pixel positions and only integrated values are summed).

The second possible mode of usage is when the control file lists an arbitrary number of fields, and each field-specific input list contains several positions of arbitrary objects. This could be e.g. a list of optical objects of a given class, undetected in X-rays, whose positions are converted into pixel positions in a field. In this case the program stacks the surroundings of all positions, counting them one or more times according to where they fall in various fields, and returns a sort of average X-ray property of the class. We plan to use this mode for a study of the X-ray properties of X-ray undetected VVDS sources (Garilli et al., 2006).

4. Results

I have extracted from the database the list of the 175 detections of overlapped sources as described in 2, and generated from it the list of 159 `control.dat` files and 334 field-specific input lists. I have then run `multistack` on each of them, and merged the output files, obtaining a list of 159 "merged measurements" of counts, rates and flux.

Details of ingestion of such data in the database are deferred to the next section (5) while here I compare the results.

4.1. Primary to duplicated distance

Before of the illustration of the results of the stacking procedure, let us have a look at the distance between the X-ray positions of the primary detection (which will be used as *the* position of the source) and of the duplicated detections (this information predates and is independent of the stacking procedure).

From Fig. 2 one sees that for the 175 duplications (including 15 cases with multiplicity greater than 2) the distance between the X-ray positions of the two detections peaks between 1 and 2″. If the distance is computed applying to each field the appropriate astrometric correction (see Report I) the distribution is even better peaked. The X-ray to X-ray distance is comparable with the nominal position error (computed as quadratic combination of the `emldetect` position errors of the two detections).

Fig. 3 gives a different view of the same fact, namely indicating that the application of the astrometric correction in the majority of cases puts the detection in different fields closer between each other, and usually within less than 4″.

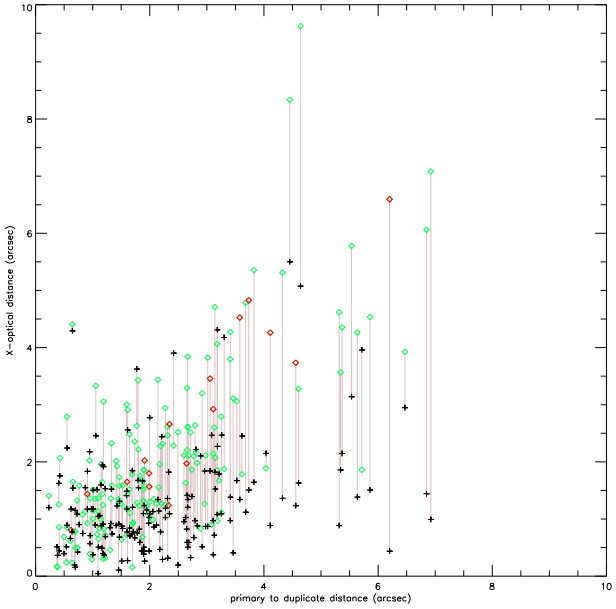

Fig. 6. Relationship between the X-ray to optical distance and the X-ray primary to duplicated distance for primary detections (crosses) and duplicated detections (diamonds). Autorank=6 cases are shown in red.

This latter value is also a safe value for the X-ray to optical distance (compare Fig. 5 of Report I). In fact I also report in Fig. 4 the distribution of the distance between the X-ray (corrected) position and the position of the best counterpart (D1, W1, VVDS, `sacphot`, SWIRE as in Report I) for the primary and duplicated detections.

This, as well as Fig. 5, indicates also that the choice of the primary detection is in general such to select the best X-ray to "optical" distance. Also, with the exception of 3 cases, at least one or the other X-ray to optical distances are below $4''$.

Finally Fig. 6 shows the X-ray to optical distance as a function of distance between primary and duplicated X-ray positions.

4.2. Changes in rates, fluxes etc.

I have produced 5 figures from Fig. 7 to 11, one for each energy band, which summarizes the changes in count rates, fluxes, S/N and chance detection probability. A further Fig. 12 gives the changes in the two hardness ratios (HRs) included in the database.

I also compared the number of net counts, and the exposures. In this case the difference is trivial, the value for the combined measurement resulting from the stacking procedure is just the sum of the values for the 2 (or 3 or 4) individual duplicated detections already contained in the database. Therefore the count rates (as ratios of counts to exposure) could have been computed straight from the database. However the error bars, the S/N and the probability require the full stacking procedure, since they also depend on the background maps, and the conversion to flux also depends on the CF maps (and ultimately on the exposure maps).

The summary figures are all arranged in 9 panels in three rows, as explained in the caption to Fig. 8. Let us examine each panel in turn.

It can be seen that the count rates (panel 1) and fluxes (panel 4) of the combined measurement agree quite well with the old individual measurements (most of them lie on or around the diagonal identity line). Large changes occur only for faint objects.

However, as apparent from the error bars in panels 1 and 4, and in more detail from panels 2 and 5, the size of the error on rate and flux is now significantly smaller. In this respect the stacking procedure provides an improved information.

Panels 3 and 6 provide a different way of looking at the rate or flux change, i.e. they show the new/old ratio vs the old value. One sees that the values are in general the same, or slightly higher, clustering around the horizontal unity line. Only some old low values get a substantial increase.

Panel 7 illustrates the actual flux conversion factor, derived a posteriori taking the flux to count rate ratio. This is combination of the local CFs used for each detection, which in turn are a combination of individual camera CFs as described in Section 3.2. The values cluster around a band specific value, between extrema which are due to detections in single cameras. The fluctuations give an idea of the effects due to position in the FoV, chips gaps, borders etc.

Panel 8 is more interesting, as it compares the S/N ratio. One can see that it always improves, though generally not a lot. In particular I have noted the cases where the individual source measurement was below 3σ in the given band, and the combined measurement is above (green crosses), and those who improve going above 4σ (purple diamonds). Of course some fall in both categories.

Panel 9 compares the chance probability detection as defined in Baldi et al. (2002). I remind that the database includes only objects which are below a probability threshold (shown by the horizontal dashed line in the figure) in *at least one* band (they may be above the threshold in other bands). I indicate in colour the sources for which individual measurements where above the threshold, but combined measurements have improved and go below.

Fig. 12 shows the changes in hardness ratios (as defined in Paper I). I remind that we store in the database two HRs, a soft one involving bands B and C, and a hard one involving C and D.

From the left hand panels in the figure one can see that most HR values are confirmed by the stacking procedure. The error bars are however large, specially for the hard HR. I remind also that asymmetric error bars are more suitable for HRs (termed "low" and "high" error bars).

From the central panels one can see that combined measures allow to reduce the size of both error bars. This

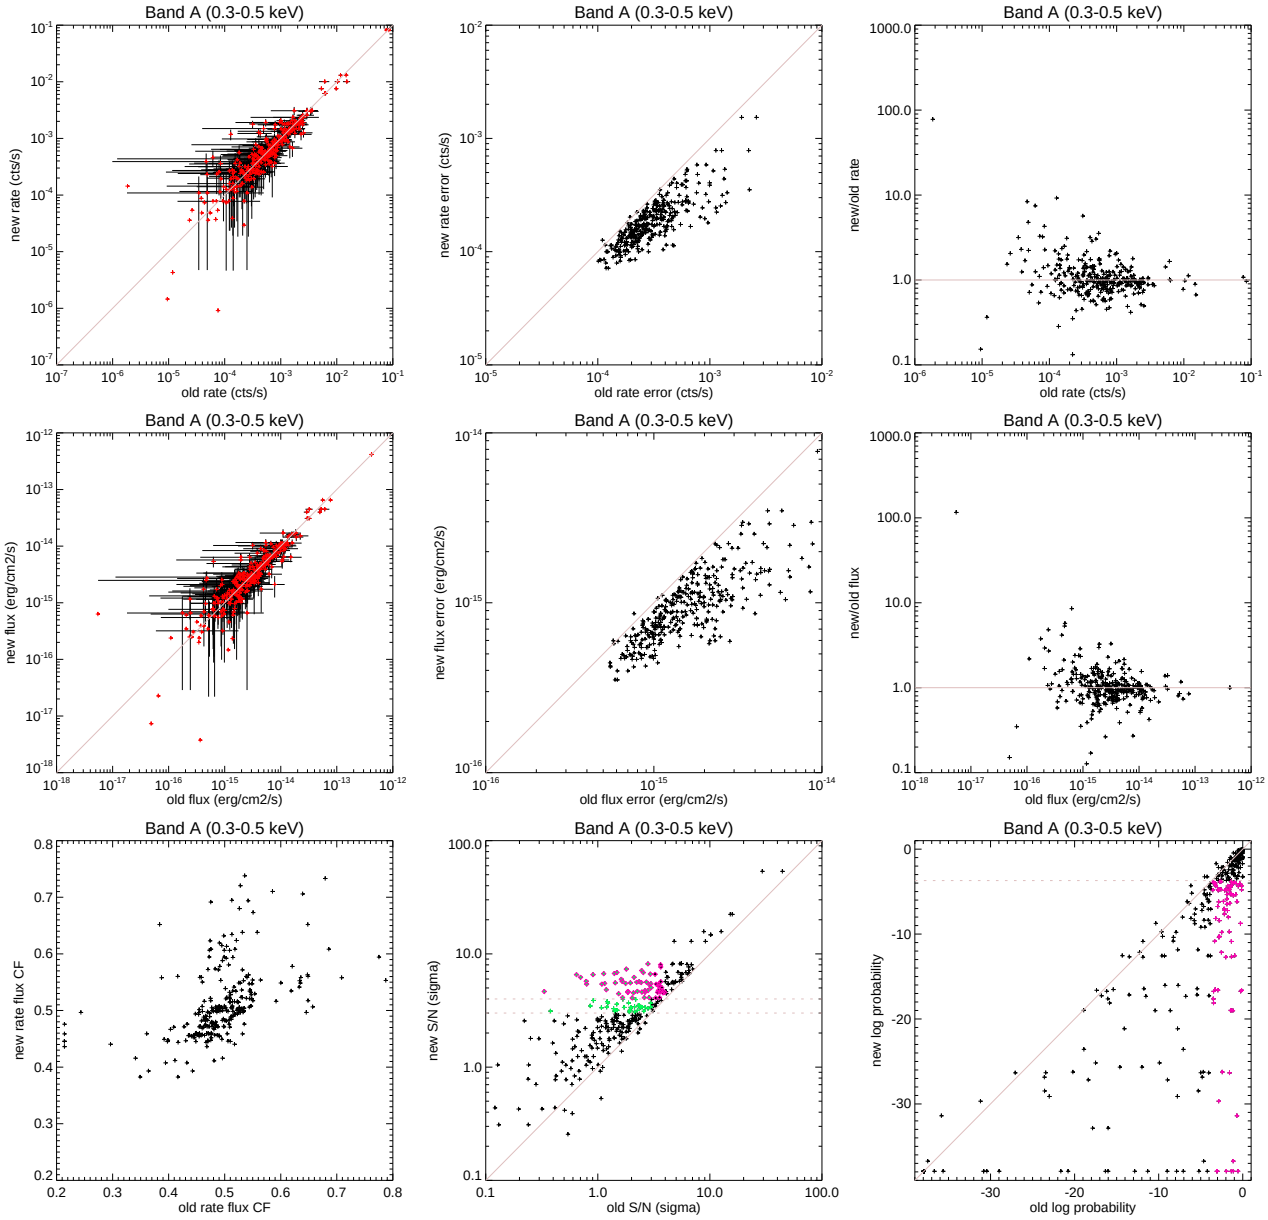

Fig. 7. Summary plot of changes for band A (0.3-0.5 keV). See Fig. 8 for further details.

is more evident for the soft HR where the error bars are about halved. Unfortunately the effect is smaller for the hard HR, where a quantity of large error bar remains (note that, although the HR ranges from -1 to 1, the scale for the hard HR errors goes from 0 to 2, i.e. there are some hard HRs which are totally undefined and the error bar size is simply the entire range).

The right hand side panels show instead how the difference between low and high error tends to become more symmetric for the combined measurements, and at the same time the error size decreases. This is quite apparent for the soft HR, while the hard HR has still a number of quite bad values.

5. Changes to the database

The cumulative output file with the merged measurements described in the previous section (4) has the same format of the output of `circledetect` used for normal database population, with just a few differences in content :

- the `id` of each record corresponds to the `seq` of the primary source of an n-uple
- the positional information (RA, Dec and error and pixel positions) are *copied* from those of the primary detection
- the remaining information (counts, count rates, fluxes, HR) are relevant to the result of `multistack`

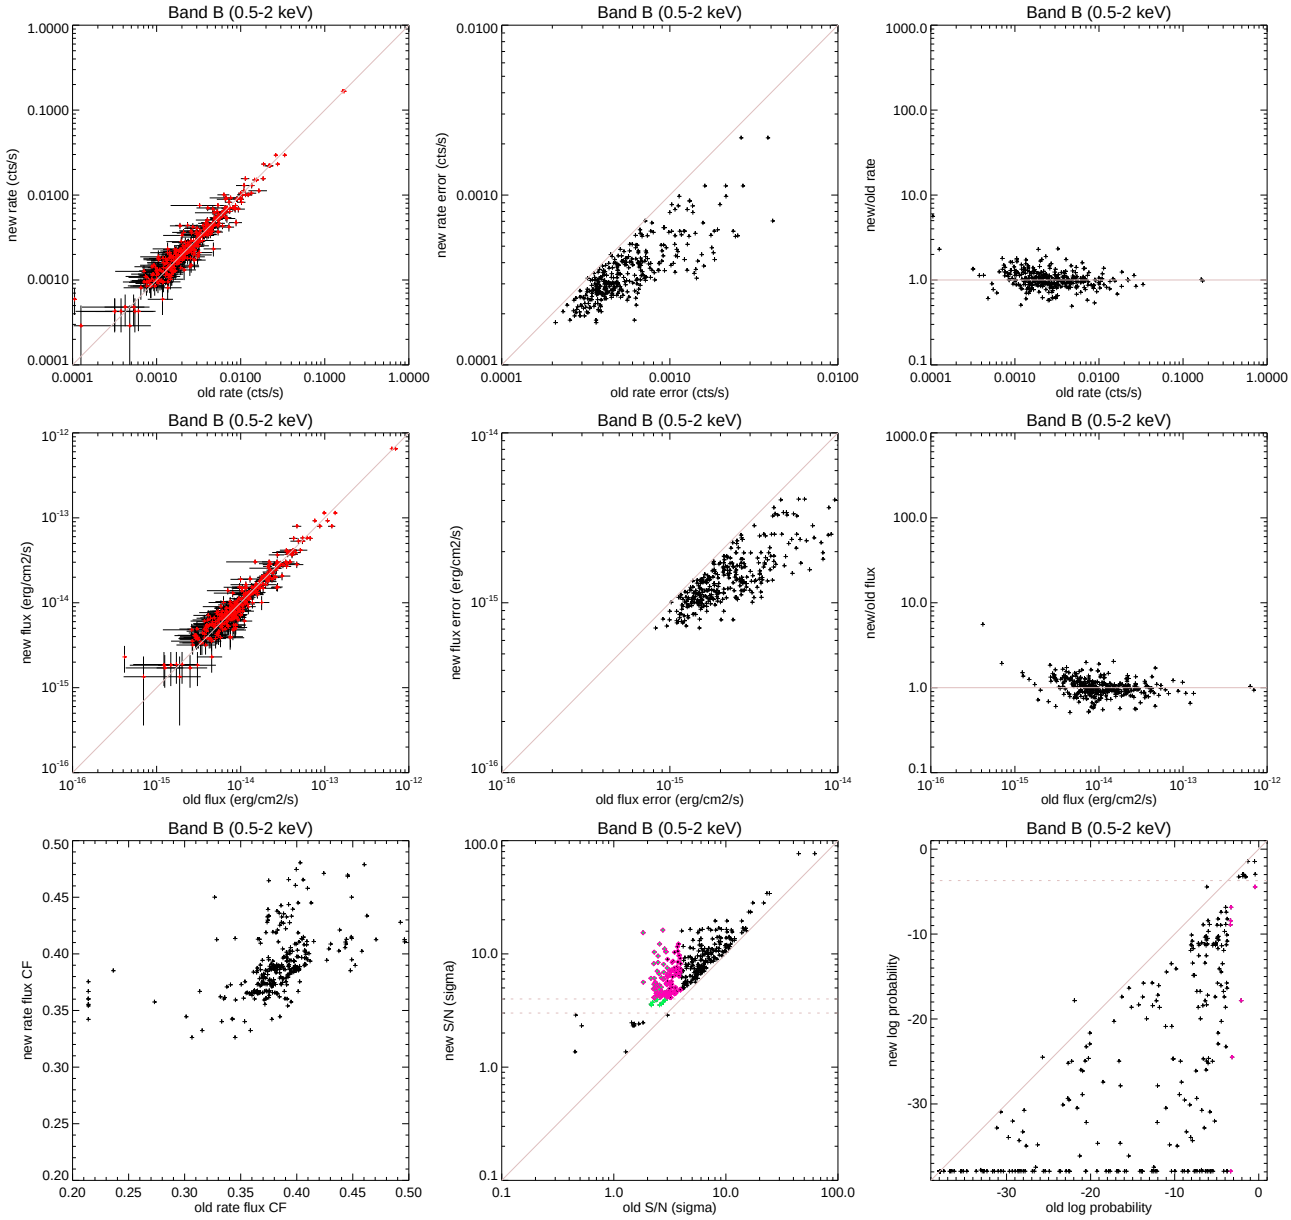

Fig. 8. Summary plot of changes for band B (0.5-2 keV). The top row (panels 1-3) is relevant to changes in count rates and associated errors. The middle row (panels 4-6) is relevant to changes in flux, and uses the same conventions as the top row. The bottom rows (panels 7-9) contains miscellaneous information.
Panels 1 and 4 show the new rates (or resp. fluxes) resulting from the stacking procedure vs the old values for individual detections. For clarity data values are in red over error bars in black. Here and in all other panels a diagonal pink line is the locus of equal new and old values.
Panels 2 and 5 show the new rate (or resp. flux) errors vs the old values (they are clearly systematically lower).
Panels 3 and 6 show the ratio of the new rates (or resp. fluxes) to the old ones as a function of the old rates or fluxes. The horizontal line is for unitary ratio, i.e. the locus of equal new and old values.
Panel 7 compares the actual flux conversion factor (scaled by 10^{11}) i.e. the flux to count rate ratio.
Panel 8 compares the new S/N ratio (in σ) to the old one. The horizontal fiducial dashed lines are for values of 4σ and 3σ. The coloured symbols indicate sources with a noticeable improvement in S/N. Namely for green crosses new S/N go above 3σ, and for purple diamonds above 4, while old S/N were below.
Panel 9 compares the probability of the detected counts to be a chance fluctuation of the background. The horizontal fiducial dashed line is for the probability threshold of 2×10^{-4}. The purple crosses show points for which the new values goes below the threshold while the old one was above, i.e. the improvements.

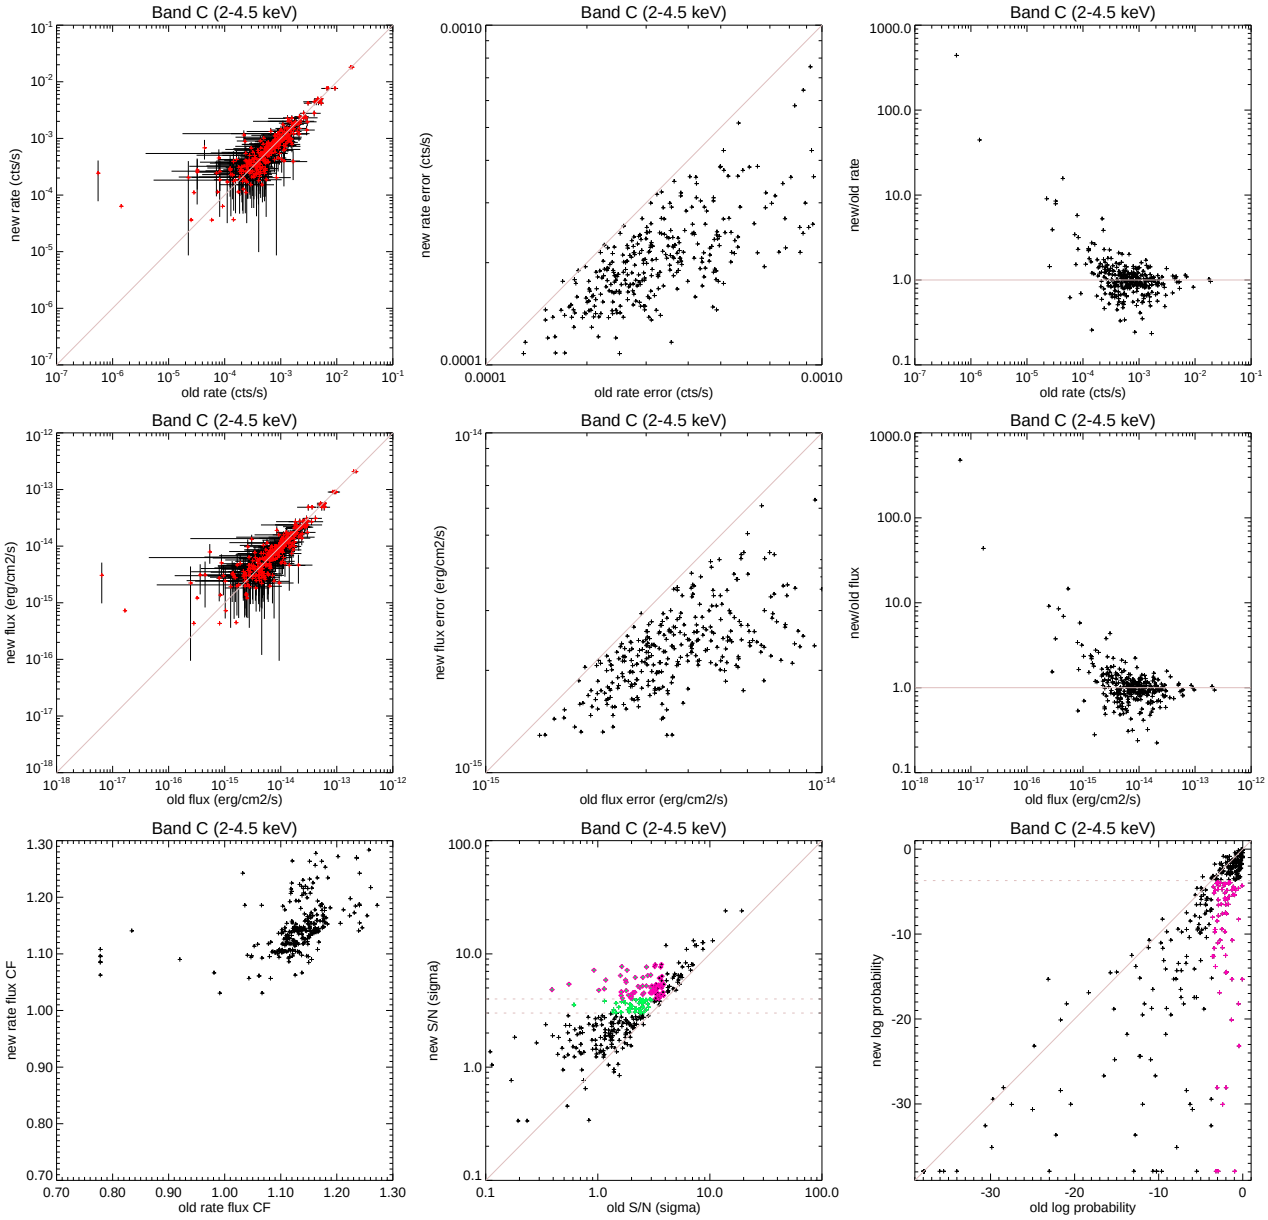

Fig. 9. Summary plot of changes for band C (2-4.5 keV). See Fig. 8 for further details.

Such data could be in principle ingested in the `xmdsepic` database table, but some missing fields will have to be redefined :

- the new measurements will receive new `seq` numbers (in the range 1325-1483)
- the `field` identification will be set to `1000` indicating "merging of the XMDS G fields" (since fields `Gnn` are already conventionally identified as `10nn`)
- the corrected RA and Dec, and the FITS thumbnail flag are *copied* from the value for the primary detection
- the simplified `gapflag` will receive the *worse* of the `gapflag`s of the duplicated detections contributing to the merged measurement

I have currently not yet ingested the data but in a temporary table, however a shell script to do it is available. Once ingestion into `xmdsepic` will be done, the merged measurements could be selected with a condition
`field=1000`
or excluded, restricting oneself to the original measurements, with
`field<>1000`

At present *and consistently with current database practice* I do not foresee a mechanism to *exclude* duplicated detections, be they primary or secondary, from the main table `xmdsepic`. They can be located and/or selected and/or excluded only using the technical ranks in the glorified table or using hand-made correlation with `xmdsdup`. I won-

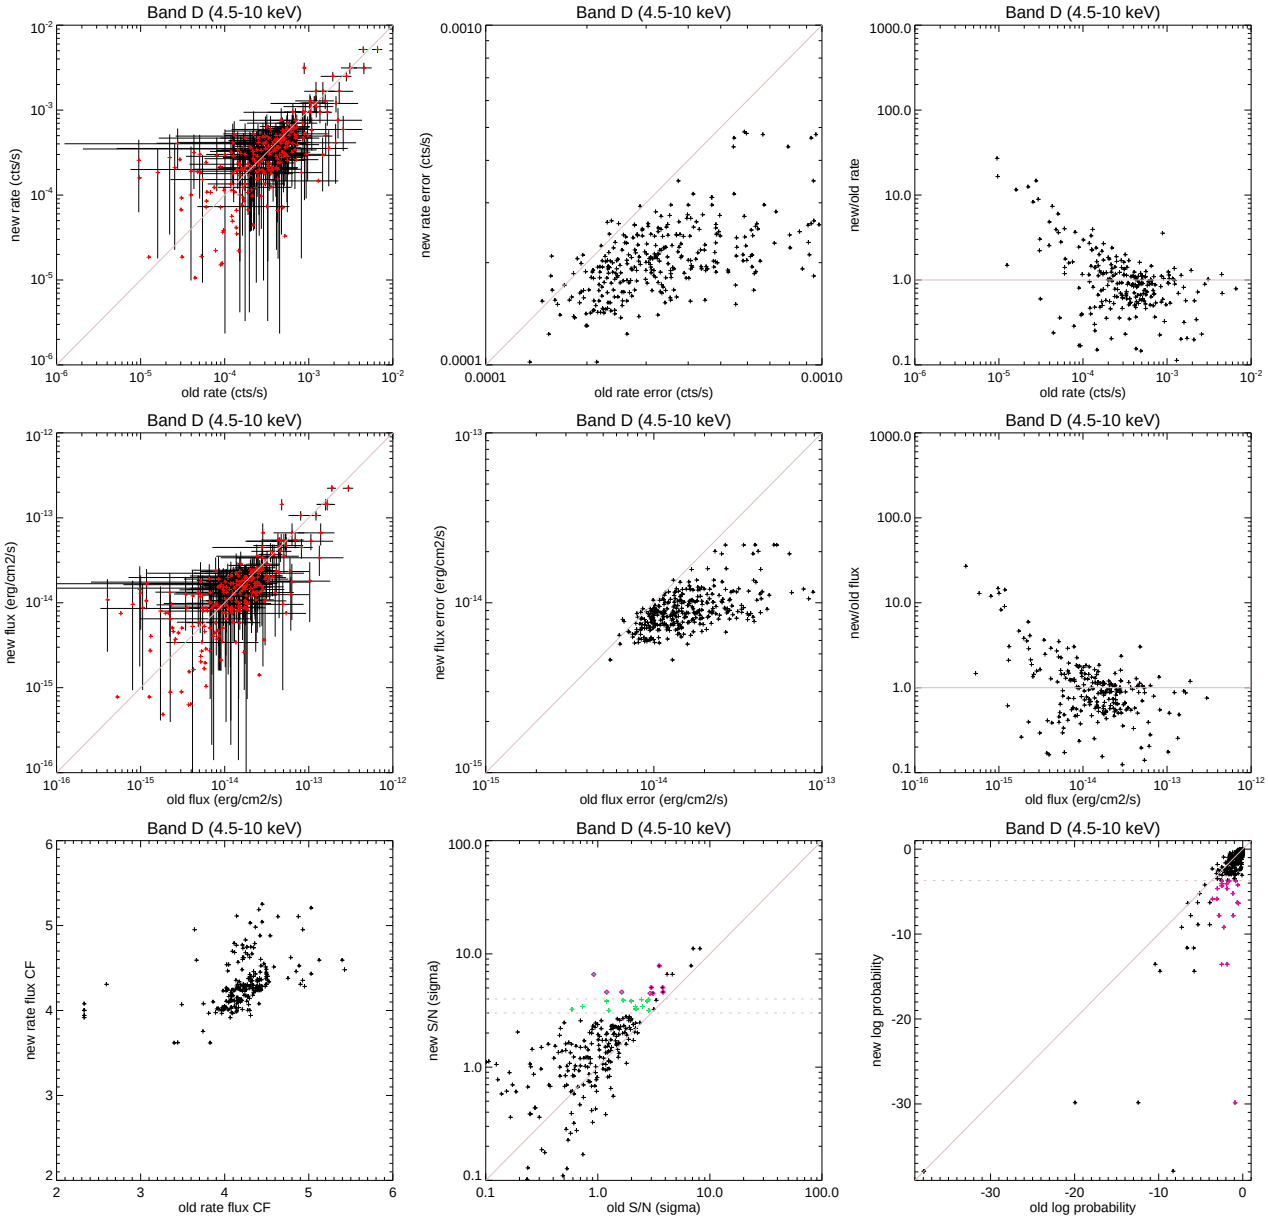

Fig. 10. Summary plot of changes for band D (4.5-10 keV). See Fig. 8 for further details.

der whether other mechanisms are desirable (e.g. add a flag column to `xmdsepic`, or a predefined correlation with `xmdsdup`). *This is probably unnecessary if duplicated detections are "removed" at the level of virtual tables.*

One further simple and immediate maintenance step (for which the code is ready in the same shell script quoted above) would be to update all correlation tables between `xmdsepic` and any other table adding entries for the new merged sources. This is obtained simply copying the entries for the primary source (as the position has not changed) and repointing it to the new `seq`.

The next step would be to somehow replace the reference to the old "duplicated primaries" to the new merged measurement in the glorified correlation table (so that the virtual table corresponding to the catalog will actually use rates and fluxes resulting from the stacking instead of the original values for the primary detection).

At the same time I would like to keep track of the original entries, so that one could easily make reference to them. This could be obtained assigning appropriate technical ranks.

A minimal proposal in this respect is presented in the top part of Fig. 13. More complex schemes are possible, which could mantain a more closer simmetry with the conventions for "normal" duplicates of Fig. 1 (like introducing more autorank 6 and autorank -8 entries), but I feel they are unnecessary, since such entries will almost never be

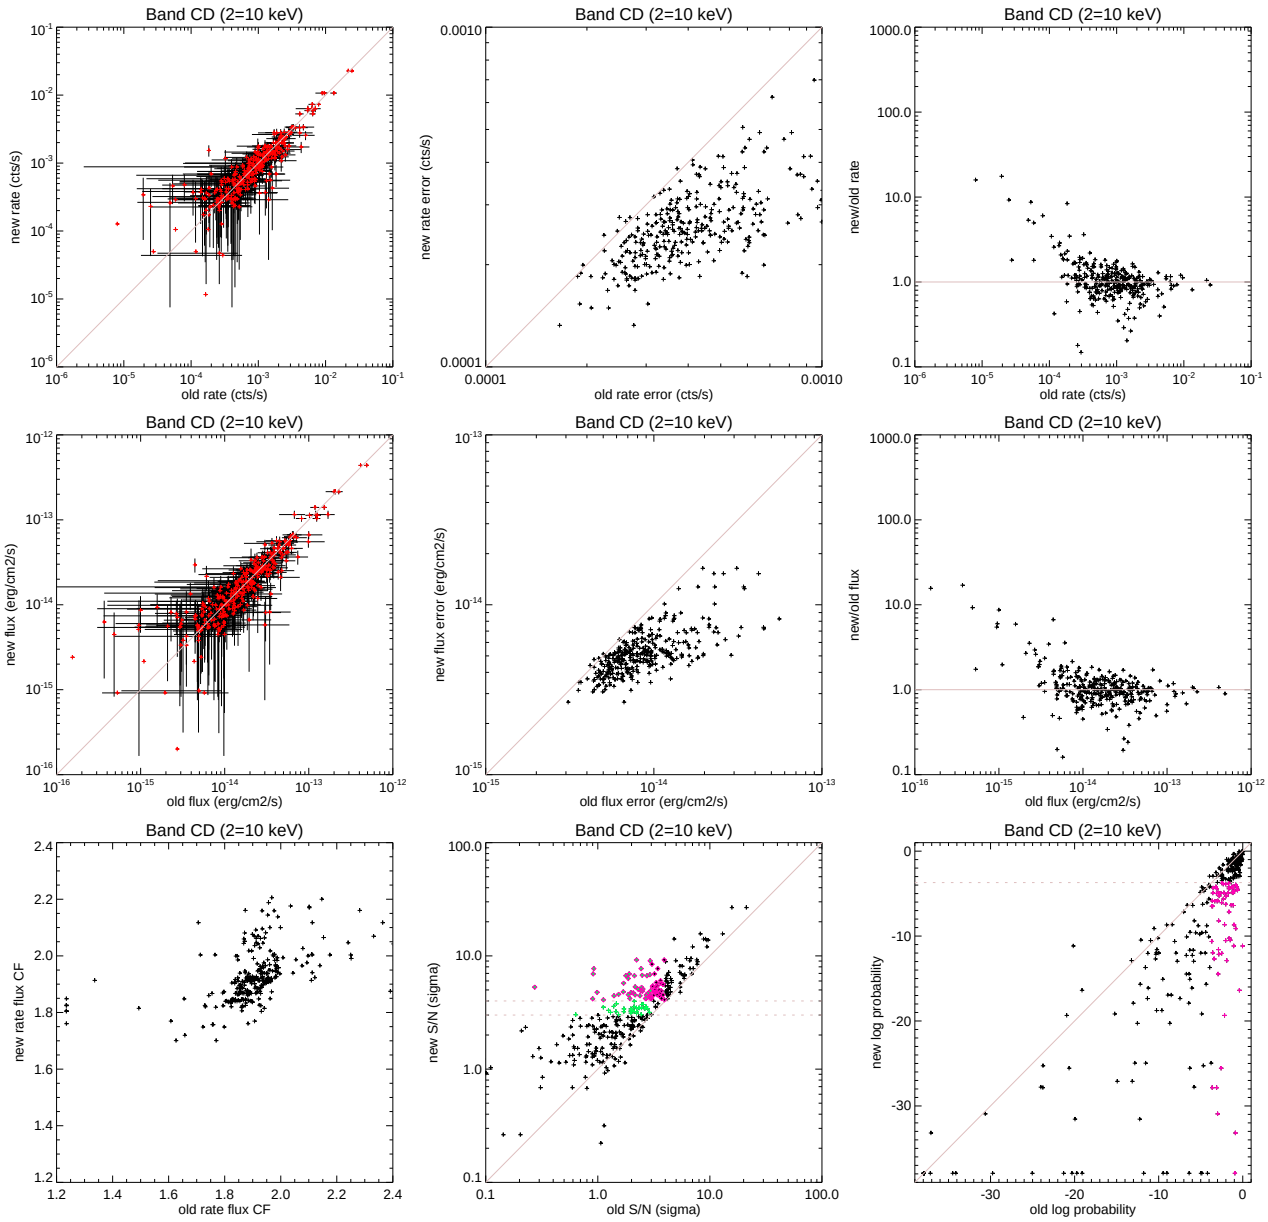

Fig. 11. Summary plot of changes for band CD (2-10 keV). See Fig. 8 for further details.

accessed, and would just duplicate the content of the real, main entries.

Essentially I would repoint the entries for the primary detection to the new merged measurement, and provide backward pointers between all duplicated detections and the merged measurement.

But not all links might be explicitly indicated. The current scheme presented in Fig. 1 is equivalent, in terms of graph theory, to *directed graphs* of order 2, 3 and 4 in which the in-degree and out-degree of each vertex is the same (resp. 1, 2 and 3), and each edge is bidirectional, i.e. each detection points to any other and is pointed to by any other.

In the proposed scheme of Fig. 13 (top) each graph order increases by 1 due to the addition of the new vertex corresponding to the combined measurement. However the in-degree and out-degree of the vertices is now different. In particular : secondary duplicated detections have one less links in entrance (there is no link pointing to them from the primary detection); primary detections have a single outgoing link pointing to the new measurement, while they are linked to from all secondary detections (but not from the new measurement); the new combined measurements have links only to the secondary detections, while are pointed to by all detections (one less outgoing link).

I am not in a condition to tell whether this is of any relevance, since I do not know how users will access the

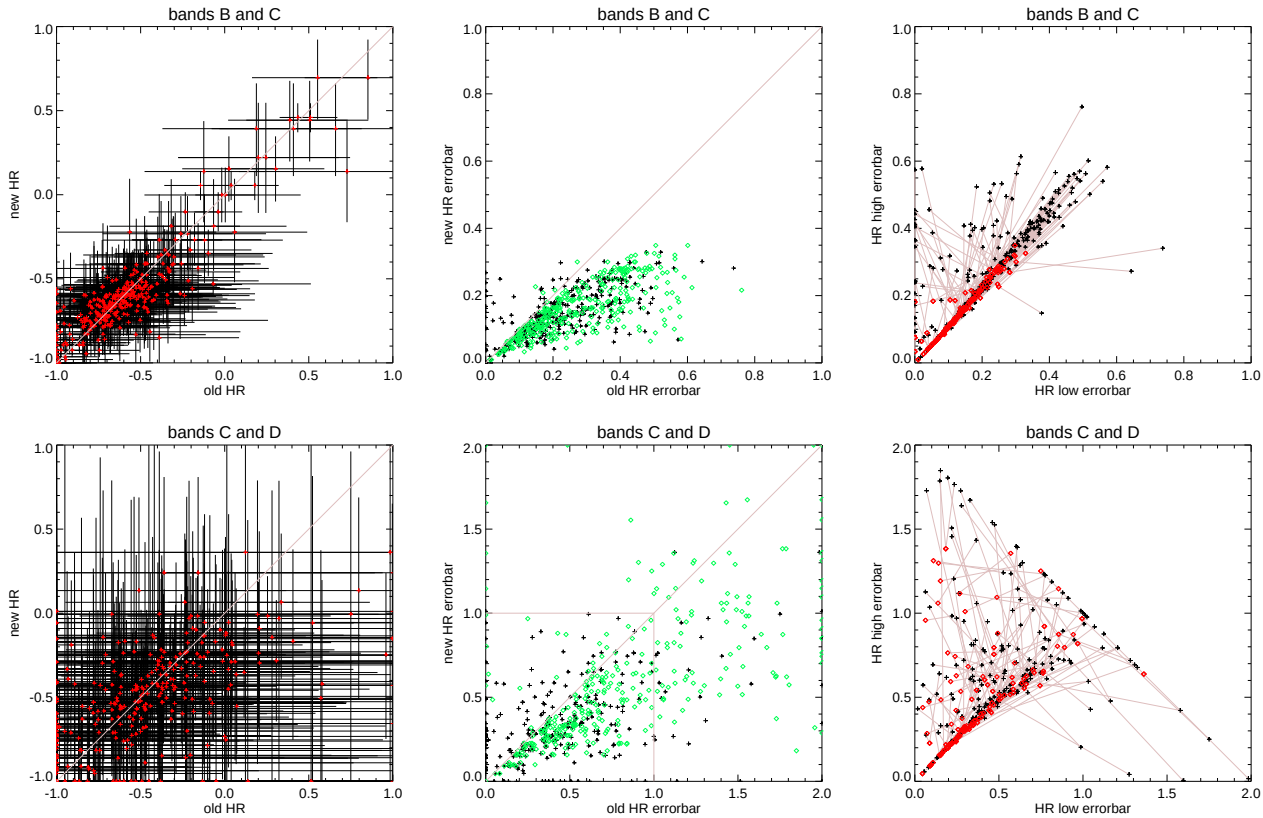

Fig. 12. Summary plot of changes for hardness ratios. The top three panels are for the soft HR (between bands B and C), the bottom three panels for the hard HR (between bands C and D).
The left hand side column plots the new HR resulting from the stacking procedure vs the old HR for the individual (primary and secondary duplicated) detections. The data values are plotted in red, with their asymmetric error bars in black.
The central column show the new HR errors vs the old ones. The black crosses are the asymmetric errors on the low side, the green diamonds the asymmetric errors on the high side. In both cases the diagonal pink line is the locus of equal old and new values.
The right hand side columns gives another comparison of the errors, plotting the high side vs low side asymmetric errors. The old values are in black, the new values in red, and measurements relevant to the same source are joined by a pink line.
Note that the two lower error panels for bands C and D have a scale double than those above, because the nominal errors are quite larger. The pink square in panel 5 is equivalent to the entire area of panels 2 and 3. Undefined HR values are excluded from the last plot (panel 6).

information about original detections (specially the old primary ones), however the unequality of the in-degree and out-degree looks inelegant.

Alternative solutions could be

– introduce a missing link at least from the merged measurement to the old primary detection. This could be a pointer devoid of other information, similar to an autorank=8 entry, using a new technical rank (e.g. the discontinued autorank-5). This way one could at least know all duplicated detections associated to new measurements (not directly all those associated to the old primary detection).
The bottom part of Fig. 1 inserts both this backward reference from merged to primary detection (with autorank-5) as well as autorank-8 links from the primary detection to the secondary detections. All other technical-rank entries are considered uninteresting and not duplicated.
– duplicate the glorified correlation table. This way one will have de facto two alternate catalogues. In one the user will access only the new combined rates and fluxes, in the other the old individual ones from the primary detection.
– duplicate the glorified correlation table *and also* the `xmdsepic` table. One version remains the current one, in the other the old primary detections will be altogether replaced by the new combined results, so that even users of the X-ray table alone can choose which dataset to use.

I plan to apply the first two items during the process of construction of the complete XMDS catalogue.

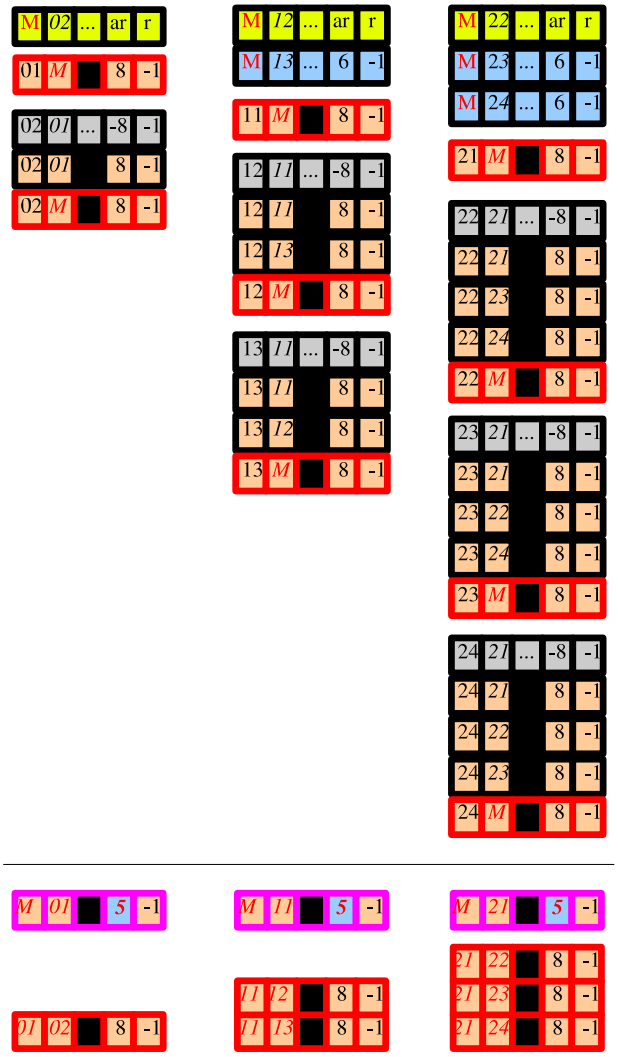

Fig. 13. This panel indicates the proposed changes to the glorified correlation table. The colour conventions used are described in the caption for Fig. 1. New records are indicated by a red or magenta border. Edited fields are indicated by red characters.
The basic idea (above the horizontal line) would be to reassign the primary source entries (yellow; and also the autorank 6 entries) as "belonging" to the merged measurement (indicated by a `seq=M`), and to introduce new autorank 8 entries (red bordered) to provide backward pointers between the duplicated and the merged measurement.
A more complete scheme (below the horizontal line) could also insert further autorank 8 entries to provide full pointers among primary and secondary duplicated detection, as well as one autorank 5 link between the merged measurement and the primary detection (magenta-bordered).

6. Conclusions

I have presented a simple data stacking procedure, which allows to produce easily and in a reproducible way count rates and fluxes from the combination of multiple detections in different fields. The procedure can also be exploited to assess global properties of undetected X-ray sources stacking data at the relevant positions.

While (for the case of duplicated detections) the result values are not different from those which can be obtained by plain addition of the existing database content, the stacking procedure allows improvements in the estimate of error bars and S/N.

It is therefore appropriate to replace in the database the original values of each primary duplicated detection with the results of the stacking procedure.

References

Baldi, A., Molendi, S., Comastri, A., et al. 2002, Ap.J, 564, 190
Chiappetti, L., et al. 2005, A&A, 439, 413 (Paper I)
Chiappetti, L., et al. 2006, XMM-LSS Internal Report N. 1-Mi (Report I)
Garilli, B., et al. 2006, in preparation