Lakshmanan,V., 2012: Image processing of weather radar reflectivity data: Should it be done in Z or dBZ?, *Electronic J. Severe Storms Meteor.*, **7** (3), 1–8.

Electronic Journal of
SEVERE STORMS METEOROLOGY

Image Processing of Weather Radar Reflectivity Data: Should It be Done in Z or dBZ?

VALLIAPPA LAKSHMANAN
Cooperative Institute of Mesoscale Meteorological Studies,
University of Oklahoma & National Severe Storms Laboratory,
Norman, Oklahoma

(Submitted 10 January 2012; in final form 24 April 2012)

ABSTRACT

It appears to be a common belief that processing, such as interpolation, clustering or smoothing, of weather radar reflectivity fields ought to be carried out on the reflectivity factor (Z) and not on its logarithm (dBZ). It is demonstrated here by means of a statistical study on a large dataset that, contrary to common belief, processing in dBZ is better for such applications.

---

1. Motivation

Image processing of weather radar reflectivity fields (the postprocessing of radar data using values at adjacent range gates) traditionally has been carried out in dBZ, e.g: Delanoy and Troxel (1993); Johnson et al. (1998); Lakshmanan (2004); Charalampidis et al. (2002); Anagnostou (2004); Steiner and Smith (2002); Kessinger et al. (2003). This author could find no successful postprocessing of radar moment data (not pulses at a single gate, but values at adjacent gates) that has been carried out in terms of the reflectivity factor measured in mmi^6 mm^{-3}, hereafter Z.

However, there appears to be a persistent belief, particularly among radar engineers (and noticed by this author in the course of having many papers reviewed), that image processing of radar reflectivity in an "artificial" quantity such as dBZ is wrong. The radar measures Z, the argument seems to go, so any interpolation or smoothing ought to be carried out in Z and not dBZ[1]. This argument is a fallacy. To see why, imagine a hypothetical instrument that measures the spatial variation of the square root of the density of a substance. Suppose that an intermediate pixel was not measured and one must estimate the value at the missing pixel. Should the value at the missing pixel be estimated by linearly interpolating measured values (the square root of the densities) at adjacent pixels or the "artificial" values (obtained by computing the squares of the measurements at the adjacent pixels, interpolating the squares and then taking the square root)? Naturally, the answer depends on the variation of density of the substance being measured, on whether the density or the square root of the density is linear. The better field for carrying out linear interpolation has nothing to do with which quantity gets measured and everything to do with which quantity varies linearly within the substance being measured.

As will be shown later, linear interpolation is closely related to operations such as clustering

Corresponding author address: Valliappa Lakshmanan, University of Oklahoma, 120 David L. Boren Blvd., Norman OK, E-mail: lakshman@ou.edu

[1]Note that this paper deals with image processing of radar moment fields, and not with pulses at a single range gate. I have nothing to say here about the relative merits of doing signal processing in Z or in dBZ.

and smoothing of radar reflectivity images. Consequently, an appropriate way to address the question of whether it is better to carry out image processing operations in Z or in dBZ is to examine a large weather radar reflectivity dataset and verify which is more linear: variation in Z or in dBZ. The author carried out such an experiment and reports those results in this note.

Does it matter whether one interpolates dBZ or Z? Yes, because the results can be dramatically different. Imagine interpolating between two values 46 dBZ and 53 dBZ. Interpolating in Z would give the mid-way value as 50.8 dBZ whereas interpolating in dBZ would give the mid-way value as 49.5 dBZ. This difference of 1.3 dB, especially aloft, can be dramatic in its effect on algorithms such as hail estimation. Interpolated Z values are consistently larger than interpolated dBZ values since an interpolated dBZ value is akin to a geometric mean of the Z values[2] and the geometric mean is always smaller than the arithmetic mean for non-negative real numbers.

2. Method

WSR-88D reflectivity data from the Little Rock, Arkansas radar from 1 January to 31 December 2007 (at 1 km $\times$ 0.95^o resolution)were analyzed in this experiment. These data were chosen solely because the author happened to have them readily available on disk. Radials of the lowest elevation scan within 230km of the radar were searched for "triads'. A triad is defined as a group of three adjacent range gates that all had valid data, i.e. data above the signal-to-noise threshold.

Within each triad, an interpolated value was computed from the two end values and compared to the value at the center. The interpolation was carried out in two ways: 1) by simply averaging the dBZ values and 2) by converting the dBZ values to Z, averaging the Z values, and then converting the averaged Z value back to dBZ. Using this procedure, one can measure the interpolation error when interpolating in dBZ and in Z. If the interpolation error is measured in dB, the interpolation error when interpolating in dBZ is given by:

$$e_{dBZ}(dB) = z_0 - (z_{-1} + z_1)/2 \tag{1}$$

where the reflectivity values in the triad are (z_{-1}, z_0, z_1) with z_0 being the reflectivity value of the center gate in dBZ. The interpolation error (in dB) when interpolating in Z is given by:

$$e_Z(dB) = z_0 - 10\log_{10}((10^{z_{-1}/10} + 10^{z_1/10})/2) \tag{2}$$

Similarly, one can compute the interpolation error in mm$^6 i$ mm^{-3} by first converting dBZ values to Z and then computing the error. The interpolation error in mm^6 mm^{-3} when interpolating in Z is given by:

$$e_{dBZ}(Z) = 10^{z_0/10} - 10^{(z_{-1}+z_1)/20} \tag{3}$$

whereas when interpolating in dBZ is given by:

$$e_Z(Z) = 10^{z_0/10} - (10^{z_{-1}/10} + 10^{z_1/10})/2 \tag{4}$$

3. Results

Statistics of the interpolation errors on all triads from a year's set of WSR-88D data are shown in Table 1. The root mean square error (RMSE) is shown first with errors in dB and then in mm^6mm^{-3}. As can be observed readily from the first row of that table, interpolation using dBZ is better whether we consider the RMSE, the bias, the variance or a pair-wise comparision. All these statistics demonstrate that linearly interpolating the dBZ values is better. The total number of triads considered and the finite values of the variances make most statistical tests here superfluous. For example, the significant t-value for a paired difference test at 10^8 degrees of freedom and significance level of 0.001 is about 3.3 whereas the computed t-value is on the order of 10^4. Thus the differences in this table are all statistically significant.

Do the results change if one were to consider only low reflectivity values, only high reflectivity values, or only triads where the endpoints have large reflectivity gradients? The answer to all of these questions, as shown in the remaining rows of Table 1, is "no". Whether one considers the overall set of triads, or slices and dices them based on the reflectivity values or based on reflectivity gradients,

[2] This is because:

$$\begin{aligned} dBZ_{mean} &= 0.5(dBZ_1 + dBZ_2) \\ &= 0.5(10log(Z_1) + 10log(Z_2)) \\ &= 5log(Z_1 Z_2) \\ &= 10log(\sqrt{Z_1 Z_2}) \end{aligned}$$

Table 1: Comparing two forms of interpolation. The first row shows the statistics over the complete year of data. The remaining rows show statistics when the data are divided in different ways based on the values of the endpoints of the triads: cases where one endpoint is greater than some threshold, cases where both endpoints are greater than some threshold and cases where the gradient between the endpoints is greater than some threshold. The penultimate column is the fraction of triads where interpolating in dBZ is better (ties are not counted) and the last column provides the total number of triads in each category.

Triad Category	Interpolating Z		Interpolating dBZ		Fraction where	Number of
	RMSE (dB)	RMSE (Z)	RMSE (dB)	RMSE (Z)	dBZ better	triads
overall	4.17	6795.74	3.8	6713.8	0.54	2.2e+09
one $>$ 0 dBZ	4.08	7234.26	3.65	7147.04	0.55	2.0e+09
one $>$ 20 dBZ	4.29	16962	3.44	16757.5	0.57	3.6e+08
one $>$ 40 dBZ	4.63	62600.1	3.66	61835.2	0.55	2.6e+07
both $>$ 0 dBZ	3.49	7840.36	3.26	7747.59	0.55	1.7e+09
both $>$ 20 dBZ	2.78	20480.6	2.64	20356.9	0.54	2.4e+08
both $>$ 40 dBZ	2.77	91436.3	2.73	93215.5	0.52	1.0e+07
gradient $>$ 2.5 dBZ/km	6.35	10218.4	5.33	9848.66	0.6	5.4e+08
gradient $>$ 5 dBZ/km	9.15	14129.8	6.99	12656.3	0.65	1.6e+08

the result is consistent: interpolating in dBZ results in a lower interpolation error. Physically, then, the quantity being measured (the hydrometeor content of the atmosphere) by a weather radar is more linear in dBZ than in Z. Hence, based on the data, interpolation across range gates is, on average, more accurate when done in dBZ than in Z.

The gradient-wise segregations point out something very critical (Figure 1). Based on the statistical evidence, as the absolute value of the gradient (between the two values being interpolated over) increases, interpolation in dBZ becomes the clearly better choice. Indeed, if the two end points are separated by more than 10 dBZ (so that the gradient is more than 5 dBZ/km), interpolation in dBZ is better twice as often (65% to 35%) than interpolation in Z. The higher the gradient, i.e. the less uniform the field, the more important it is that interpolation be carried out in dBZ.

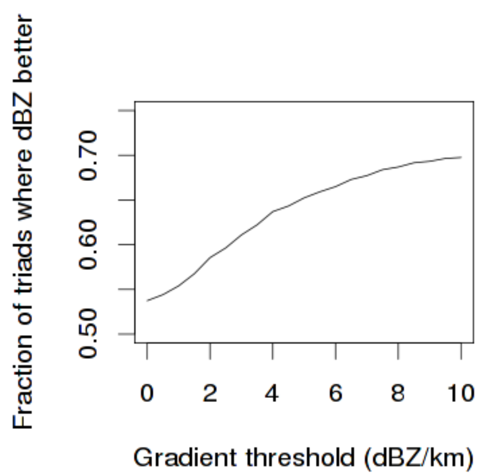

Figure 1: The odds favoring dBZ interpolation over interpolation in Z as the gradient exceeds a specific threshold. The higher the gradient, the more important it is that interpolation be carried out in dBZ.

4. Discussion

The implications of this result go beyond just interpolation and play into the choice of how to carry out any linear operations on the radar reflectivity values at adjacent gates. Nearly all image processing filtering operations are linear. For example, smoothing a field is commonly done by replacing a pixel by the average value of its neighbors. This is the same as subtracting an error estimate from each pixel's value, with the error estimate being $x-\bar{x}$, where x is the value at that pixel and $\bar{x}$ the average pixel value in the pixel's neighborhood. Expand out $\bar{x}$ in terms of the neighboring values and it is easy to see that smoothing a dBZ field implies subtracting a fraction of the dBZ values of neighboring pixels. Hence, smoothing of radar reflectivity images is better done in dBZ than in Z.

The literature on image processing of weather radar data has gotten this right, carrying out processing of adjacent range gates in dBZ and not Z. A median filter is different because $\bar{x}$ is independent of monotonic transformations in x, but techniques that use a weighted average (Delanoy and Troxel 1993; Lakshmanan 2004; Anagnostou 2004) to compute $\bar{x}$ involve subtracting dBZ values. Similarly, parameters such as the texture of the dBZ field used by Kessinger et al. (2003) boil down to arithmetic operations on dBZ values, not Z. Methods of storm identification from reflectivity also operate on constant-dBZ increments (Johnson et al. 1998; Rosenfeld 1987; Crane 1979) because increments in Z are numerically problematic. Such constant dBZ increments, not constant Z increments, are widely used whenever adjacent pixel values need to be compared, for example by Steiner and Smith (2002) to compute their textural feature.

Pixel-wise operations such as Z-R relationships can afford to operate in Z, but any clustering, smoothing, segmentation or feature extraction technique that relies on the reflectivity value at adjacent pixels is better done in dBZ than in Z.

ACKNOWLEDGMENTS

Funding for the author was provided by NOAA/Office of Oceanic and Atmospheric Research under NOAA-University of Oklahoma Cooperative Agreement #NA11OAR4320072, U.S. Department of Commerce. I wish to thank John Cintineo for providing the year's archive of data, and acknowledge helpful discussions with Kurt Hondl and Rodger Brown of NSSL.

REFERENCES

Anagnostou, E., 2004: A convective/stratiform precipitation classification algorithm for volume scanning weather radar observations. *Meteor. Appl.*, **11**, 291–300.

Charalampidis, D., T. Kasparis, and W. Jones, 2002: Removal of nonprecipitation echoes in weather radar using multifractals and intensity. *IEEE Trans. Geosci. Remote Sensing*, **40**, 1121–1132.

Crane, R., 1979: Automatic cell detection and tracking. *IEEE Trans. Geosci. Electronics*, **17**, 250–262.

Delanoy, R. and S. W. Troxel, 1993: Machine intelligent gust front detection. *Lincoln Laboratory J.*, **6**, 187–212.

Johnson, J., P. MacKeen, A. Witt, E. Mitchell, G. Stumpf, M. Eilts, and K. Thomas, 1998: The storm cell identification and tracking algorithm: An enhanced WSR-88D algorithm. *Wea. Forecasting*, **13**, 263–276.

Kessinger, C., S. Ellis, and J. Van Andel, 2003: The radar echo classifier: A fuzzy logic algorithm for the WSR-88D. *3rd Conf. on Artificial Applications to the Environmental Sciences*, Long Beach, CA, Amer. Meteor. Soc., P1.6.

Lakshmanan, V., 2004: A separable filter for directional smoothing. *IEEE Geosci. Remote Sensing Letters*, **1 (3)**, 192–195.

Rosenfeld, D., 1987: Objective method for analysis and tracking of convective cells as seen by radar. *J. Atmos. Oceanic Technol.*, **4**, 422–434.

Steiner, M. and J. Smith, 2002: Use of three-dimensional reflectivity structure for automated detection and removal of non-precipitating echoes in radar data. *J. Atmos. Oceanic. Technol.*, **19**, 673–686.

REVIEWER COMMENTS

[*Author's responses in blue italics.*]

REVIEWER A (Mark Askelson):

Initial Review:

Recommendation: Major revision

Substantive comments:

p1, P2: I believe that this is a good point, but note that, as indicated by the questions at the end of this paragraph, it is based upon the assumption that a linear interpolation technique is going to be used to estimate the value at the missing pixel. Of course, this is only one objective analysis approach that one may use and oftentimes will produce undesirable results (e.g., discontinuities in derivatives at observation locations). Moreover, objective analysis fields commonly replicate nonlinear characteristics of input fields, which is, of course, desired. Thus, please consider either indicating earlier in the paragraph that the estimation method being used in this example is linear interpolation or making the argument more general.

Yes, I have now explicitly used "linearly interpolating" so that it is clear where I am headed. Also, as stated later in this review, the point of this paper is not objective analysis but image processing. Thus, we are concerned mainly with linear operations over short distances. I use interpolation as the machinery to demonstrate linearity.

p1, P4: This is an important point and I am glad you have included it. However, I think it would be helpful if you show how the two means are related so that the reader can have a better understanding of the arithmetic-vs. geometric-mean statement (here, capital letters indicate logarithmic units and lower-case letters indicate linear unitsplease see technical item 1 for why I follow this convention):

$$A_1 = \frac{Z_1 + Z_2}{2} = \frac{10\log z_1 + 10\log z_2}{2} = \frac{10\log(z_1 z_2)}{2} = 10\log(z_1 z_2)^{1/2}$$

and

$$A_2 = 10\log a_2 = 10\log\frac{z_1 + z_2}{2}$$

Thus, averaging two logarithmic radar reflectivity factors produces 10log of the geometric mean of the linear-units values while averaging of linear-units values, once converted into a result having logarithmic values, produces 10log of the arithmetic mean. From a theoretical standpoint I believe that it is important to show this as it illustrates why the two approaches will not generally produce the same value.

Done.

In addition, I believe that the 52 dBZ result presented is incorrect. When I compute the average radar reflectivity value by converting 46 dBZ and 53 dBZ into linear values, averaging them arithmetically, and then converting the result back into logarithmic values, I obtain 50.78 dBZ. Thus, I obtain a difference of 1.28 dB (the proper unit for dBZ - dBZ is dB).

The reviewer is correct; I have now fixed this.

I have a concern here — it is that testing using logarithmic units inherently favors the arithmetic logarithmic average. I believe that to be fair, one should also compute errors in linear units in the following way, with $e_{a\log a}$

meaning error of the arithmetic logarithmic average and e_{lina} meaning error of the arithmetic linear average:

$$e_{a\log a}(mm^6 mm^-3) = z_0 - 10^{\frac{z_{-1}-z_1}{2}/10} = z_0 - (z_{-1}z_1)^{1/2}$$

and

$$e_{alina}(mm^6 mm^{-3}) = z_0 - \frac{z_{-1}+z_1}{2}$$

where again capital letters imply logarithmic units and lower-case letters imply linear units.

My concern regards the logarithmic operations that occur in different relative locations in (1) and (2) and what impact they have on relative error values. Putting (1) and (2) into linear units results in

$$e_{a\log_1} = \frac{z_0}{(z_{-1}z_1)^{1/2}} = \frac{1}{1+\epsilon_1/z_0}$$

and

$$e_{alin_2} = \frac{z_0}{\frac{z_{-1}+z_1}{2}} = \frac{1}{1+\epsilon_2/z_0}$$

where $(z_{-1}z_1)^{1/2}$ = $z_0 + \epsilon_1$ and $\frac{z_{-1}+z_1}{2}$ = $z_0 + \epsilon_2$. If logarithmic values of radar reflectivity factor have a linear spatial variation over three consecutive range resolution volumes, then $e_{a\log a} = 0$ and $e_{a\log a_(1)} = 1$. Alternatively, if linear values of reflectivity factor have a linear spatial variation over three consecutive range resolution volumes, then $e_{alina} = 0$ and $e_{alin_(2)} = 1$. Thus $e_{a\log a}$, e_{alina}, $e_{a\log a_(1)}$, and $e_{alin_(2)}$ provide the desired error characteristics for the situation of purely (spatially) linear fields (for $e_{a\log a_(1)}$ and $e_{alin_(2)}$ the logarithmic operation produces 0 dB in both cases).

However, because the error values computed using (1) and (2) involve taking the logarithm of $e_{aloga_(1)}$ and $e_{alina_(2)}$, ϵ_1 (or ϵ_2) values that have the same magnitude (in linear units) but opposite signs will produce errors that have different magnitudes. If, for instance, ϵ_1/z0 = +/-0.1, one obtains -0.4139 dB from (1) for ϵ_1/z0 = +0.1 and 0.45757 dB for ϵ_1/z0 = -0.1. Moreover, considering that the geometric mean is always smaller than the arithmetic mean and that arithmetically averaging logarithmic reflectivity values is equivalent to taking 10 log of the geometric mean of the linear-units values whereas the other approach involves taking 10 log of the arithmetic mean of the linear-units values, it seems possible that ϵ_1/z0 is slightly negative more often than ϵ_2/z0 is which, owing to the logarithmic operation could skew the results. This argument, of course, leverages the idea of linear-units reflectivity values that are of the same magnitude but different signs, which of course will not generally be the case. However, I believe that it does illustrate the potential impact of the logarithmic operations in (1) and (2) and, thus, motivates the computation of errors using $e_{a\log a}$ and e_{alina} for, at the very least, comparison purposes.

Fair point: a new column has been added to the table, with the RMSE measured in mm^3 mm^{-6}. The results remain the same. In any area with some structure, the dBZ values are more linear. Also, the pair-wise tests (the fraction when dBZ is better) is independent of the units being used.

My most significant concern regards the linearity-implies-which-field-to-use-argument. While I agree that if the underlying field is spatially linear one would do well to apply an analysis technique that retains its spatial linearity, I argue that the situation is much more complicated than this. The examples given (linear interpolation and boxcar averaging) are both relatively simple objective analysis schemes, the first of which enforces linearity between interpolation points and the second of which does not generally produce spatially-linear analysis fields (and both of which have significant shortcomings like discontinuities in derivative fields and ringing in the response function). Objective analysis schemes used in the atmospheric sciences can, generally, be expressed in terms of Distance-Dependant Weighted Averaging (DDWA; e.g. Askelson et al. 2005), which is a linear operation. However, just because DDWA is a linear operation, this does not mean that DDWA will produce linear spatial (or temporal) variations. On the contrary, most geophysical fields contain significant nonlinearity, the degree of which depends upon the scales being considered. Thus, DDWA schemes generally retain at least some of this nonlinearity. Radar reflectivity factor fields, whether viewed using logarithmic or linear units, undoubtedly contain nonlinear variations. Here, values were analysed over a 3 km distance. Of course, even a highly nonlinear field can be decomposed into linear changes if these linear changes occur over small enough distances.

The finding that logarithmic-units values are more linear over small distances (up to 3 km) than linear-units values only, in my opinion, has relevance for variations having scales of up to 3 km. Because DDWA schemes are commonly applied over distances greater than 3 km when analysing radar data and because they may not retain spatial linearity on a 3 km scale, I believe that statements should be qualified in a manner similar to "...if one is using an objective analysis method that retains spatial linearity over short distances, then using logarithmic values is preferable relative to using linear values." Thus, my concern is that given the evidence presented here, the statement that logarithmic reflectivity values should be used in all objective analysis schemes (all smoothing) is too broad. Before closing it is noted that if a field is linear over certain spatial scales and an objective analysis does not replicate that linearity, then this would of course be a detriment. However, such a scheme may sacrifice linearity on those scales to more accurately represent variations on other scales. Consequently, trade-offs at different scales may occur that may be good or bad, depending upon an analysts preferences.

The goal of the article is to determine the field on which image processing operations ought to be carried out. It is not about determining how to carry out objective analysis. In other words, we know that most image processing operations (e.g: interpolation, smoothing, clustering) are linear and work only when the data values that are being operated on vary linearly. I use linear interpolation at known values to see if the field varies linearly over the scale of a few pixels. It turns out that operating in dBZ is better than operating in Z if we want local linearity. Therefore, I argue, it is better to carry out image processing in dBZ.

The reviewer is absolutely correct that over larger distances, atmospheric fields are non-linear. But that does not affect our argument which is solely about image processing operations (not objective analysis).

There are two reasons why such nonlinearity over large distances does not affect our argument (again our argument is about image processing, not objective analysis: the title of the paper makes our focus quite clear). Firstly, linear interpolation is used in this paper just as the machinery to verify a Hilbert space — a system of units within which spatial variation at the scale of a few image pixels is linear. Thus, we are not concerned with linear interpolation over large distances or multiple pixels. Secondly, most image processing operations (smoothing, texture, subsampling, etc.) model an image as a Markov random field by which we mean that they assume a pixel's value is correlated only on its immediate neighbors. Thus, non-linear, sharp spatial changes will be poorly handled anyway by image processing operations. This is regardless of what space we operate in.

I think it is clearer to refer to reflectivity factor, regardless of the units, as reflectivity factor and then to specifically indicate whether its units are linear or logarithmic. Rinehart (2010) keeps these distinct by using a lower case z to indicate that the radar reflectivity factor units are linear and an upper case Z to indicate that the radar reflectivity factor units are logarithmic. While I am not suggesting that you have to do this, the use of the symbol dBZ to indicate logarithmic units is a bit awkward, in my opinion, since dBZ indicates units, not the variable itself. One would not, for instance, label an east-west distance variable as ft.

This is something I went back-and-forth on. I decided to use the terminology in this paper because it is much clearer (when spoken, for example) than the z-Z dichotomy. Somewhat equivalently, we could talk about working in standard-atmospheres or hPa and I believe most people would understand whether we are working with ratios or millibars.

REVIEWER B (Mario Majcen):

Initial Review:

Recommendation: Accept with minor revisions

General comments:

The paper is acceptable with minor revisions and no further review is requested unless major changes are made in accordance with other reviews. I really enjoyed reading this manuscript. It reveals a fallacy behind a commonly used data manipulation technique.

REVIEWER C (David Dowell):

Initial Review:

Recommendation: Accept with minor revisions

General comments: This short, focused article summarizes results of an important study for radar-data analysis. The writing style is somewhat informal, which mostly works for this article. A few instances where more precise and/or formal wording would help are mentioned below. Also, some suggestions are provided to avoid overstating results, resulting in a stronger manuscript.

Substantive comments: In paragraph 2, the motivation for the study relies on hearsay. Can we instead refer to more concrete evidence (preferably publications) that the community often favors Z processing over dBZ?

As stated in the manuscript, the literature on the topic has it right — published work in this area involves processing dBZ, not Z. I would hypothesize that even believers in Z quickly discover that it doesn't work well, and resort to working in dBZ. Thus, I was unable to find any publication doing anything in Z.

p. 3, column 2, paragraph 1: If retained, the discussion This is the same neighboring pixels. belongs earlier in the manuscript and requires more explanation. Or maybe this discussion can be deleted entirely.

This is an important point. I've retitled this section "Discussion" and added the link text that we are going beyond linear interpolation to addressing any linear filtering operation.