Earth Surf. Dynam. Discuss.,
doi:10.5194/esurf-2016-18-RC2, 2016

Earth Surface Dynamics Discussions Open Access EGU

Interactive comment on "The CAIRN method: Automated, reproducible calculation of catchment-averaged denudation rates from cosmogenic radionuclide concentrations" *by* Simon Marius Mudd et al.

G. Balco (Referee)

balcs@bgc.org

Received and published: 8 June 2016

Review of 'The CAIRN method...' by Mudd and several others.

Summary: This work is very useful and I'm quite supportive of it. The rest of the review below contains (i) some contextual remarks; (ii) a list of areas of the paper where I could not exactly understand what is happening, so clarification is needed; and (iii) a few other comments.

Contextual remarks: Basically, this paper describes a numerical scheme to compute basin-scale erosion rates (henceforth, 'E') from cosmogenic-nuclide concentrations

(henceforth, '$\bar{N}$') in stream sediment. The basic concept here is that given a number of assumptions, these two have a simple inverse relationship that is defined by the mean production rate of the nuclide in question in the basin (henceforth, '$\bar{P}$'), the nuclide's decay constant, and the apparent e-folding length for subsurface production (henceforth, 'Λ'). If we know these three things we can relate E to $\bar{N}$ very simply.

If we accept the assumptions, there are really only two problems with this method overall. The easy one is just calculating $\bar{P}$. Because P is nonlinear with elevation, the point in the basin that has the average elevation doesn't have the average P. So you can define the elevation (and topographic and other shielding, if you want to get fancy) of each pixel in a DEM of the basin, calculate P in each pixel, and take the average. If you want to reverse-engineer some code that has elevation rather than $\bar{P}$ as input, you can also represent the basin by a single point with an effective elevation and/or shielding such that P in that pixel is the same as $\bar{P}$ for the whole basin. So as long as you have a not-too-terrible DEM of the basin, or at least its hypsometry, then both of those methods are sufficient to solve the problem of estimating $\bar{P}$.

The harder problem is that the true depth dependence of production can't be represented by a single exponential with one value of Λ, which messes up the nice clean explicit relationship between $\bar{N}$ and E. In fact, it makes the formula implicit so you can't solve it analytically.

To deal with this mess in the 2008 online calculators, we used the represent-the-basin-by-a-single-pixel approach and an implicit solver for that pixel. That's efficient, but kind of lame. The correct way to do the problem is to do the implicit-solver scheme for the entire basin. That is, pick a value of E, do a full forward calculation of the nuclide concentration predicted in each pixel at that erosion rate, and take the average to obtain a predicted value of $\bar{N}$. Then repeat, using some kind of an implicit solver, to obtain the value of E that matches the observed $\bar{N}$. Besides just being the correct way to do the simple problem in which the erosion rate is assumed steady and uniform, this method has the enormously large advantage that you can implement other forward models involving nonsteady or nonuniform erosion and, potentially, use cosmogenic-nuclide measurements to constrain more complicated things than just a simple steady erosion rate. Obviously, this full-implicit scheme is much more computationally painful, because you have to do the implicit solver for all of what could be a large number of pixels. Although this has been done in the existing literature (e.g., Fox, Leith, and others in EPSL, 2015), it's pretty rare. In any case, the aim of this paper is to apply the correct and better method at a large scale, by applying some actual software engineering expertise to incorporate the capability into a GIS package.

So from the perspective of past approaches to this problem, Simon Mudd and his co-authors have basically shown up with a gun at a knife fight. Adding this tool to a topographic analysis package is really useful, and it is potentially enabling of all sorts of more complex applications of cosmogenic-nuclide measurements that haven't been thought about much so far. Although it's probably more accurate to say they have shown up with a light saber at a knife fight, because for many people using this tool will first require an extended apprenticeship with a Linux Jedi master. But the point is that the method described here is not at all lame, it is actually the correct numerical way to do this, and it is potentially enabling of new science. That's all good.

On the other hand, this work creates a new problem. This is not so much a difficulty with the paper itself as a difficulty with the method itself, but what this work focuses on is a high-precision implementation of a model calculation. The problem here is that the model calculation itself is pretty well known to be a poor representation of reality in most cases, because the geological assumptions on which the model is based are highly unlikely to be true. There are a variety of reasons for this (some of which I'll get into later), but the most important is that the model calculation implemented here assumes that the erosion rate in the basin has been constant for long enough so that the nuclide concentrations at all points (and depths) in the basin have reached an equilibrium between production and nuclide loss due to decay and erosion. Although this might be true for near-surface production at relatively high erosion rates, there is

really no possible way it can be true in the subsurface where production is mostly due to muon interactions (see more discussion below). Thus, there is somewhat of a risk of confusing the precision of the model calculation with the accuracy of the actual erosion rate estimate, and it would be incorrect to argue that the method described here, even if it executes the model calculation more precisely, will allow anyone to more accurately estimate real erosion rates. On the plus side, using the method in this paper to take the numerical difficulties in the model calculation off the table ought to allow people to focus more on the geological uncertainties, which is potentially valuable.

So that is the main point of what I am trying to say here. It is great that someone has finally written nice code to do this problem properly, and incorporating the full-forward-model into a topographic-analysis software package is potentially enabling for new science, but it is important to keep in mind that the accuracy of the erosion rate estimates themselves is unlikely to have been improved. Thus, from the perspective of routine calculations of apparent basin-scale erosion rates, it is not clear that an improvement has been made over existing approximate schemes. The main value of this paper is likely to be in future work that explores how to (i) evaluate the geological assumptions in the model, and (ii) learn about unsteady or nonuniform erosion rates from cosmogenic-nuclide measurements.

End of contextual remarks. The rest of this review goes through the paper in order and highlights locations where either (i) I could not understand what was happening, or (ii) improvements could be made. Also some minor comments are included as well. Note that a weakness of this review is that I am not very fluent in C++, so it is possible that some areas of the text that I found hard to understand are clearer in the code than I think they are.

Abstract, first sentence. The first sentence can be removed – it would be more useful for the abstract to begin with the second sentence, which describes what the authors actually did in the paper.

Line 31. The word 'cosmogenic' means originating from cosmic rays. The calculators don't originate from cosmic rays, so 'cosmogenic calculator' makes no sense. Suggest just 'calculator' here.

Line 35-ish. This is a bit oversimplified. In fact, there is not a lot of variation in approaches to computing the erosion rate (really there are only two: include muons, or don't). What there is a lot of variation in is how to compute $\bar{P}$. Also, the phrase "representative parameters for the catchment" is unnecessarily complicated and also unhelpfully nonspecific. Specifically, what you are trying to estimate is just $\bar{P}$. Be more specific here.

Line 40-ish. Again, this is somewhat misleading as written. The mathematical definition of the catchment-averaged erosion rate (e.g., from Bierman and Steig, etc.) specifically has $\bar{P}$ in it, and there is no doubt about what this parameter is. Thus, it is incorrect to say that authors are trying to "choose a production rate that is 'representative' of the catchment." Instead what they are doing is trying to estimate the mean production rate in the catchment, either by computing that in some pixel-based way or by fudging the input to the online calculators to force them to compute the mean production rate internally. In any case, it would be helpful to clarify this a bit.

Line 53. I was confused by this discussion of the landslide scenario, because this scenario violates the assumptions inherent in computing the catchment averaged erosion rate from a Be-10 concentration. Thus, if you are going to compute the erosion rate using the method in Bierman, Steig, etc., you have already assumed that erosion is steady over time (although not spatially uniform), which is equivalent to stating that landslides do not occur. Thus, landsliding doesn't cause a problem with computing $\bar{P}$ (as is implied by its inclusion in this section), it basically invalidates the entire method. To clear this up, I suggest dividing this section into two subsections: (i) issues that make it difficult to compute $\bar{P}$ (e.g., shielding, snow cover, etc.); and (ii) issues that invalidate the entire method by violating its basic assumptions (landsliding, sediment storage, etc.). Alternatively (and probably better), I suggest just pending the entire discussion of landsliding to a separate section at the end of the paper, in which you can discuss more generally the point that you could potentially use this code for all sorts of nonequilibrium scenarios.

Line 60. Don't you want to continue here to mention that in addition to computing production rates, the software also does the implicit solution for erosion rate given a measured nuclide concentration? Because as written this paragraph doesn't indicate that any erosion rate estimate happens. Needs improvement. In general, also, it would be helpful for the reader at this stage if you were to explain the overall procedure of inverting a forward model for nuclide concentrations to obtain an estimate of the erosion rate, as a preview of coming attractions.

Line 80-ish. This brings up the subject of muons. In this work as in others, muons are responsible for 2% of surface production and 98% of suffering. The decision in this paper to use an exponential scheme for muon production for computational simplicity is, in fact, sensible. Unfortunately I found the explanation here to be incomplete and confusing.

Basically, the difference between the Heisinger integration scheme and a simple exponential approximation à la Braucher is that the latter is incorrectly representing the physics. What is really happening physically is that as depth increases, the mean energy of the remaining muons increases, so the instantaneous e-folding length for muon production continually increases with depth. You can't represent that with a finite sum of exponentials. If you have a bunch of summed exponentials, you can do pretty well at shallow depths, but there is some depth below which you are quite wrong. So that is what the actual difference is.

However, there are two reasons the exponential approximation is OK here.

The first is just that only a very small fraction of the nuclide concentration you are measuring was produced by muons at greater than a few meters depth, so accurately estimating production at that depth doesn't matter much. And when the erosion rate is high enough to care about this approximation, you have other problems to do with the model assumptions.

The second is that if you include muons in the production rate calculation, you are also including muon production in the basic assumption of the method that the surface has been eroding at a steady rate long enough for the Be-10 concentration to reach production-erosion equilibrium. For muon production, this takes a really long time. It is easiest to think of this as a half-life for approach to equilibrium, which is $-\ln(0.5)/(\lambda + E/\Lambda)$ where λ is the decay constant of the nuclide in question (4.99e-7 for Be-10); E is the erosion rate (g/cm$^{2]}$/yr), and Λ is the e-folding length for the depth dependence of production (g/cm$^{2]}$). Take a landscape that is eroding at 50 m/Myr. For spallogenic production ($\Lambda \simeq 150$), this half-life is 13,000 years. It is already a bit of a stretch to assume that a typical surface has had an unchanging erosion rate for several times 13,000 years, i.e. from well before the LGM until now. However, for muon production (which has $\Lambda > 1500$) this half-life is at least one order of magnitude longer (> 120,000 years, increasing with depth). It is highly unlikely that any landscape has experienced steady erosion for many times 130,000 years. Because of these different averaging times, tt is also highly likely that the erosion rate recorded by muon-produced nuclides is very different from the erosion rate recorded by the spallogenic inventory (see Dethier and others, Geology, 2014 for a nice example). So including an accurate representation of muons in the calculation would, in principle, improve the accuracy of the erosion rate calculation if the steady state assumption is true. However, it also decreases the likelihood that the steady state assumption is true. For this reason, there is no point getting too worried about how exact the muon production estimate is in this calculation.

To summarize, it is a sensible move to limit the complexity of the muon production calculations because any increase in precision from making them more elaborate is probably illusory. However, this issue should be explained more clearly in the text.

A final point here is that there would potentially be lots of other ways to speed up the computation whilst still using the Heisinger scheme, if you wanted. Mainly this is

because the production rate due to muons will not be very different between adjacent cells. Thus, it is a big waste of time computing muon production separately in each pixel. You can probably get away with doing the muon calculations in a very small minority of cells in a typical watershed and extending those results to the other cells simply by a regression formula in elevation, without loss of accuracy. Or do muon production on a much coarser grid than spallogenic production. Of course, this would be a big rewrite of the code, but if you really want it to run maximum fast it would be the next obvious strategy. If muons are 2% of surface production, why give them more than 2% of processor time?

Line 105. Again, equations are not cosmogenic.

Line 120. Note again the implications of assuming that muon production is in steady state. Not likely to be true.

Line 130, "self-shielding." I am not sure I understand what is going on here because under normal circumstances, the sediment leaving the catchment would be assumed to be from an infinitesimally small surface layer, so no integration in depth would be required. So this appears to me to be overly complex. My understanding of what is happening in this part of the paper is that the authors have just put in this capability to facilitate later use of the same code for a patchy-erosion model where finite thicknesses of sediment are removed at once (e.g., landslides). And then the discussion of steady state is confusing here as well, because, of course, if landslides are occurring then there is by definition not a steady state. Overall, more explanation needed here. As noted above, I think this would be clearer if all discussion related to the landsliding issue was deferred to a separate section.

Line 160 and below. Topographic shielding. This is another example where the calculation is an precise representation of simplified physics, so gives illusory precision. Specifically, this code includes a quite precise calculation of topographic shielding under the assumptions that (i) the cosmic-ray flux is totally attenuated below the apparent horizon, and (ii) the zenith angle dependence of the cosmic-ray flux responsible for production is a cosine to a constant power. Neither (i) or (ii) is actually true. Because secondary particle production takes place throughout the atmosphere (including that part of the atmosphere that is between you and the apparent horizon on the other side of the valley), a nonzero amount of production will actually be due to cosmic rays originating below the visible horizon. In addition, the cosine-to-a-power dependence is highly approximate. See this paper:

Argento, D.C., Stone, J.O., Reedy, R.C. and O'Brien, K., 2015. Physics-based modeling of cosmogenic nuclides part II–Key aspects of in-situ cosmogenic nuclide production. Quaternary Geochronology, 26, pp.44-55.

The point being that the very comprehensive analysis of discretization errors in the shielding calculation here clouds the fact that there exist larger systematic errors due to simplified physics. This issue has basically no practical relevance to the erosion rate calculation overall (because it is still much less important than violations of the basic method assumptions). However, the authors should note here that they are concerned with the precision of a representation of simplified physics, which may or may not be the same as the precision of the calculation relative to real life.

In addition, this could also be sped up a lot if you really wanted to. The shielding factor doesn't change much between adjacent pixels. In addition, most of the pixels in a landscape don't actually shield anything; nearly all the shielding is due to ridgelines and smaller things that are near you. There is some code that has been kicking around since 2001 here:

http://depts.washington.edu/cosmolab/oldweb/P$_{by_G}IS.html$

that uses this simplification (note that I think the editor's comment on this subject gives a bad link). But although of historical interest, that is beside the point here.

A final important issue here is that it was not clear to me whether telescoping of the

mean free path length on dipping surfaces (see Dunne, also Fig. 5 in Balco, 2014 in Quaternary Geochronology) is included in this calculation.

Line 200 et al. The issue of non-time-dependent vs. time-dependent scaling is actually more important than described here. The reason for that is that production rates are calibrated using data mostly from the last 20,000 years, and the Earth's magnetic field has been stronger than its long-term average during that time. Thus, at erosion rates low enough that the residence time of material in the soil profile is much longer than this (e.g., most normal erosion rates), the use of a non-time-dependent scheme likely creates a systematic error due to an underestimate of the long-term production rate. Basically this is yet another violation of the steady-state assumption. Again, this is a non-issue compared to much bigger issues in the application of this method, but the text is somewhat inaccurate here as written.

Line 215, section 2.6. Unfortunately, I simply don't understand why the calculation described in this section is necessary. I didn't look back at the Vermeesch paper, but if I am remembering correctly this whole procedure was just needed to make the equation relating erosion rate to concentration explicit so it could be solved analytically?? Here you don't need to worry about that, because you are only doing the forward calculation, so why are you doing this? I may not be remembering this correctly, but in other words, it seems to me that all the S's and F's needed for Equation 13 are known a priori, or should be. Typically one would compute scaling for spallogenic production and muons separately (for example, this is in the Stone (2000) scheme as published), and because they are different, that should take care of the fact that muons are less important at higher elevation. Each pathway has already been assigned its own attenuation length, so you can compute mass shielding for each pathway. Then it seems like all you need to do is decide how to compute topographic shielding for muons (I don't know the answer...in the 2008 online calculators it is just disregarded), and you are done. What am I missing here?

In any case, this needs to be better explained.

Line 270-ish. I should point out (in response to the editor's comments) that the issue of asymmetry of uncertainty distributions is really a total non-issue from the geological perspective. Anything with an e in it will have an asymmetric uncertainty distribution, of course, but it's hard to think of any cases where it's actually important from the geological perspective.

Line 280-ish. Numerical partial differentiation by repeatedly doing the full calculation is almost certainly overkill (especially because you've already linearized it). I would do this simply by assuming that the basin has one pixel with the effective $\bar{P}$ derived from the whole basin. I agree that it is interesting to do it once, though.

Line 300. In physical science, 'conservative' is typically used to mean that something is being conserved, e.g., mass or energy. This use in the context of uncertainty analysis is common but incorrect. Instead one should state that the uncertainty estimate is supposed to be an upper bound.

Line 315. This is an excellent point, that the divergence among various theoretical expectations of how snow shielding works is much less important than the practical difficulty of actually measuring the mean snow depth distribution throughout the year. Frankly, in my view it is not even necessary to mention the various models here, because that issue is pretty much totally unrelated to this paper. As an aside, I found the Delunel paper to not be persuasive because, as far as I can tell from reading the paper, we don't know the effect of snow cover on the energy dependence of the neutron monitors (snow is basically like changing the amount of polyethylene on the outside, which affects the spectral response). This would imply that it is likewise unknown whether or not the variation in monitor count rate with snow cover is applicable to cosmogenic-nuclide production at all. But that is totally off topic.

Line 320-ish. Again, I suggest moving all discussion related to landsliding and non-steady/nonuniform erosion to a separate section at the end. First, implement the basic model; then, at the end, introduce the abilities to deal with complications that are half-

baked at present, but potentially useful in future.

Line 345-ish. It is probably overkill to generate separate effective elevations and average shielding factors for input to the 2008 online calculator. I know that technically it's required because the elevation affects the muon proportion of total production and the shielding factor doesn't, but this issue is well down in the noise.

Line 370. Again, no need to get into the details of snow shielding here. The point remains that it is unclear whether one can estimate the snow depth accurately in any case.

In general, in this part of the discussion (i.e., all of section 5) I think it would be easier to understand if you break down the discussion into two parts: things that are linear with respect to the production rate (e.g., topographic shielding), so can be pixel-averaged by themselves; and things that are nonlinear with respect to the production rate (e.g., elevation, latitude, snow shielding), that have to be converted into production rates, averaged, and then unconverted into an effective summary value. At present these two things are mixed up and it's hard to understand what is happening. Overall, this section could be made more clear.

Section 6. This sort of comparison is a terrible mess because of the need to sort out the differences between inherent properties of the algorithm (e.g., point approximation vs. full-basin calculation) and the input parameters (mainly the production rate and the muon interaction parameters). Overall, however, it is accomplished fairly well in this paper; I like the approach of selecting a few representative data sets rather than trying to show a global comparison over all of scaling and erosion rate space, and the explanation is quite clear as regards which errors apply where in which comparisons. I only have a couple of comments about this section.

Line 470. It is quite interesting that there is a systematic difference in shielding factors vs. those originally reported. Do you think this really is because of the averaging-a-nonlinear-thing effect? But in any case, as discussed above, the precision of this measurement is overstated in any case due to simplification of the physics.

Line 480. I think production rate differences are much more likely responsible here than anything to do with shielding calculations.

Line 486. The snapping issue is by far the biggest problem I can think of in wholesale automation of this process. Especially potentially disastrous for literature data.

Line 500+ and Figures 7-8. As noted by the other reviewer, this effect is nearly all due to differences in the input parameters (production rate and muon interaction cross-sections) and very little due to the spatial-averaging issue or any other aspects of the various algorithms. In large part this is my fault because I have been too much of a slacker to update these parameters in the online code (that is, make v 2.3), which is, frankly, embarrassing. Sorry. However, the need here for the purposes of this paper is to clearly separate these issues. I can think of two ways to do this. One, change the parameters in the CAIRN code to increase spallogenic production by a factor of (4.5/4) and muon production by a factor of (1/0.44). That will very nearly account for the various input parameter differences. Then do the comparison on that basis. Two, leave these figures unchanged but add in the background some lines showing the expected effect of those changes in the parameters (for elevations vaguely resembling the input data). In any case, this would be extremely helpful in distinguishing the various effects of differences in the algorithm itself vs. differences in the calibrated parameters. This would also make the discussion in lines 520+ more clear as well.

Besides the above comments on the MS, I have a couple of comments on another review ("interactive comment" by anonymous reviewer 1).

1. Reviewer's comment 3. It is not correct to say that 'the model of Heisinger overestimates muon production.' The model accurately estimates muon fluxes; the problem is that the cross-sections for production of Be-10 and Al-26 by muon interactions, when applied to those flux estimates, overestimate Be-10 and Al-26 production. In other words, it's not the model that's wrong, it's the cross-section measurements needed to

convert the model prediction to a production rate. I think the present paper is mostly correct on this point.

2. Reviewer's comment 7. The reviewer is correct here, and this is important. The options here are (i) to get all numbers properly standardized in this discussion, or (ii) preferably, to not get into the details here and simply note that best estimates of production rates are about 10% lower than they were 10 years ago because of improved calibration data.

Finally, some final remarks:

Use of acronym 'CRN.' Do you really want to exclude stable cosmogenic nuclides? Because this code would work for them too. Perhaps, having written the paper, the authors could just globally search and replace 'CRN' with 'cosmogenic nuclide?'

'CAIRN' acronym. This is really beside the point with respect to the scientific content of the paper, but when you have to remove important words ('cosmogenic') and randomly select letters from the middle of other words to eventually wind up with a word that appears in the dictionary, I question whether things have really been improved. If we are allowed to select randomly from all the letters in 'Catchment-averaged denudation rates from cosmogenic nuclides' as long as we maintain their original order, why not 'CRIMES,' 'EVIL,' 'CAGE-FREE,' or 'ACNE' (I favor 'CAGE-FREE')? Certainly the authors are entitled to call their code whatever they want, but why such a desperate bid for an acronym at any cost? Especially as it's already a part of something with a different name (LSDTopoTools). Perhaps we could just call it the cosmogenic-nuclide module of LSDTopoTools?

Interactive comment on Earth Surf. Dynam. Discuss., doi:10.5194/esurf-2016-18, 2016.