Statistics Netherlands

2014 Scientific Paper

Phase 0 report 1 : significance of trend changes in ground subsidence in Groningen

The views expressed in this paper are those of the author and do not necessarily reflect the policies of Statistics Netherlands.

Frank P. Pijpers

Contents

Nederlands

De voorliggende tussentijdse rapportage is een verslag van het onderzoek dat tot dusverre is uitgevoerd in het kader van fase 0 van een onderzoeksproject door het CBS in opdracht van Staatstoezicht op de Mijnen (SodM). Dit onderzoek is ten behoeve van een statistische onderbouwing van het meet- en regelprotocol voor gasexploitatie in de provincie Groningen, met in fase 0 in het bijzonder de aandacht gericht op de meetbaarheid van het effect dat het sterk reduceren van productie in delen van het gasveld gehad kan hebben op de bodemdaling in het betreffende gebied.

Het is geen doelstelling van deze fase van dit project om de directe causaliteit tussen gas productievariaties en trendwijzigingen in bodemdaling te onderzoeken. Op basis van uitsluitend de analyse die hier wordt gepresenteerd kan een direct verband noch aangetoond, noch verworpen worden.

Uit de analyse blijkt dat ongeveer 2 maanden nadat de productie sterk was gereduceerd er een statistisch significante trendbreuk is opgetreden in de bodemdaling. De dalingssnelheid is een factor 2.8 lager na de trendbreuk, die optreedt in de weken tussen ongeveer 17 Maart en 15 April 2014, vergeleken bij de periode daar direct aan voorafgaand. Deze factor is vergelijkbaar met de reductie in gasproductie, gezamenlijk over de locaties waar productie sterk is gereduceerd met locaties daar direct aan aangrenzend. De trendwijziging kan zich geleidelijk gemanifesteerd hebben over een periode van enkele weken, en er is daarom een onzekerheid van ruwweg een week of twee over de centrale datum van deze overgang.

English

This is an interim report on the research that has been carried out to date within phase 0 of a research project being carried out by Statistics Netherlands and commissioned by State Supervision of Mines (SodM). This research is part of the underpinning of the statistical methods employed to support the protocol for measurement and regulation of the production of natural gas in the province of Groningen. In phase 0 the particular focus is on the measurability of the effect that the strong reduction of production in part of the field may have had on the ground subsidence in and around that region.

It is not the aim of this phase of this project to investigate a direct causality between gas production variations and trend changes in subsidence rates. On the basis of the analysis presented here by itself, such a direct relationship can neither be proved nor disproved.

From the analysis it can be determined that some 2 months after the strong reduction in production, there has been a significant change in the ground subsidence. The speed of subsidence is a facor of 2.8 lower after the break in the trend, between around March 17 and April 15 2014, compared to the period directly previous to that date. This factor is comparable to the factor by which production has been reduced, in the combined production levels of the locations where a strong reduction was implemented together with well locations directly adjacent to this where production continued. The changeover in the trend of ground subsidence can have become manifest gradually over a period of several weeks, which implies that the central date of the transition is also uncertain by roughly a week or two.

1 Introduction

For some decades earthquakes of modest magnitudes have occurred in and around the Groningen gas field. It is recognized that these events are induced by the production of gas from the field. Following an $M_L = 3.6$ event near Huizinge, and the public concern that this raised, an extensive study program has started into the understanding of the risk and hazard due to gas production-induced earthquakes.

A protocol is being established with the aim of mitigating these risks and hazards by adjusting the production strategy in time and space. In order to implement this regulation protocol and adaptively control production it is necessary to also measure the effects on subsidence and earthquakes to provide the necessary feedback.

Figure 1 The relative locations with names of each of the GPS stations from which data are available. The red diamond symbols are locations within the area of the production field `Middle', where production has been reduced. The stations indicated with blue squares are locations outside of the main field, collectively designated as `Edge', although at the southwestern station Eemskanaal there are wells in production. The two stations indicated with green triangles are also production locations in a separate part of the field `South' where production has not been reduced. Well locations where production has been reduced are shown as black crosses, with in addition three wells very close to this region (AMR, TJM, SDB). The indicated scales are in km

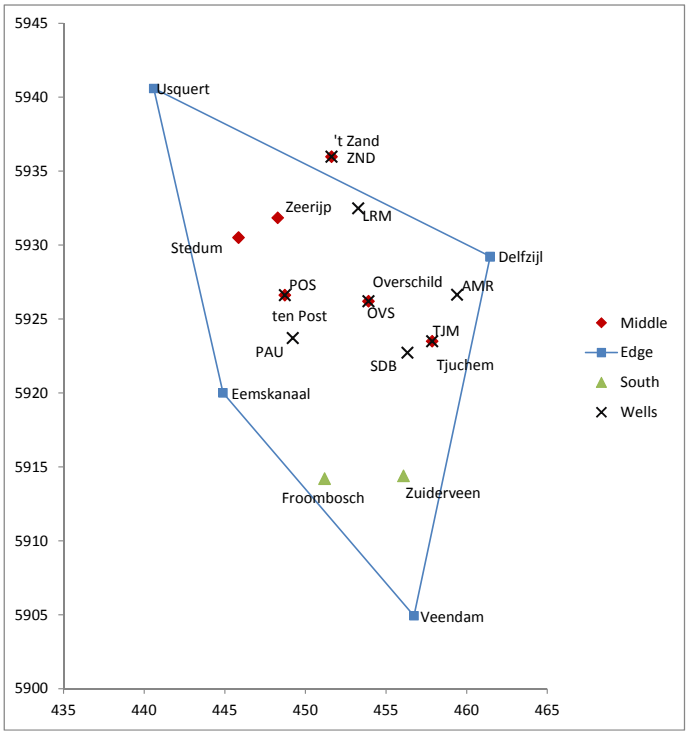

The causality of the earthquakes induced by gas production is likely to be through the interaction of reservoir compaction with existing underground faults and differentiated geology of the subsurface layers. The ground subsidence occurs because with the extraction of gas, pressure support decreases in the layer from which the gas is extracted. The weight of overlying layers then compacts that extraction layer until a new pressure equilibrium can be established.

Whether a fault will fail and cause an earthquake depends on the local stresses and friction of the fault. Pressure differences cause different local stresses on the fault, which is why part of the fault may be critically stressed and another part may not be critically stressed. This will change as depletion increases as well. The technical addendum to the winningsplan Groningen 2013 "Subsidence, Induced Earthquakes and Seismic Hazard Analysis in the Groningen Field" discusses all of these aspects in much more detail.

Within the framework of the protocol a number of geopositioning system ground stations (GPS) have been set up in addition to the ones that were already in place at Ten Post and Veendam. The height data collected from these GPS stations can be used to monitor trends in ground subsidence. This is the first link in the chain of causality between gas production and earthquakes, and therefore it is important to establish whether changes in production levels can be shown to have a measurable effect in the GPS data. In particular a strong reduction of production levels from a number of wells in the field in January 2014 is an event a signature of which can be attempted to be found in the GPS data.

In this technical note, the available GPS data are examined for a signature of breaks in trends. The analysis procedure is presented as well as the conclusions one can draw from this initial phase of the research project.

2 Background

2.1 The GPS data

The available GPS data are recorded with a cadence of once per hour. A height, as well as longitude-latitude position is available for each of the 12 locations. At two locations, Ten Post and Veendam, the GPS ground stations have existed since februari 2013. The other 10 locations have been established more recently. From the most recently established station, at Usquert, data are available only from April 5 2014 onward. In figure 1 the relative locations of these stations are indicated. The GPS stations can be divided among three groups. One group is located within the area of the production field 'Middle', where production has been reduced. A second group are locations outside of the main field, collectively designated as 'Edge', although at the southwestern station Eemskanaal there are wells in production. Two final stations are also production locations but in a separate part of the field 'South' where production has not been reduced. The five well locations where production has been reduced are shown as black crosses. In addition there are three wells very close to this region (AMR, TJM, SDB) where production has not been reduced as much.

The GPS time series of heights is of both high quality and high cadence. While this is evidently fortunate the high cadence can be problematic when attempting to establish longer term trends, because these can be masked by higher frequency signal in the time series. For this reason it is useful to apply a standard band-limited filter to each of the GPS station time series.

2.2 filters in the time domain and the Fourier domain

For a thorough discussion of the effects that a general linear filter has on a time series, a valuable diagnostic is the frequency response of that filter. This frequency response is the Fourier transform of the filter factors. In general the Fourier transform $Y(\omega)$ of a function of time $y(t)$ is:

$$Y(\omega) \equiv F(y) = \int dt\, e^{i\omega t} y(t) \tag{1}$$

in which ω is the frequency of a signal, related to the period P by:

$$\omega = \frac{2\pi}{P} \tag{2}$$

The Fourier transform is often also referred to as the Fourier spectrum. For this Fourier transform there is also an inverse operation:

$$y(t) = F^{-1}(Y) = \frac{1}{2\pi} \int d\omega\, e^{-i\omega t} Y(\omega) \tag{3}$$

If the function $y(t)$ is only known at discretely sampled times t_i which are regularly spaced, a discrete version of the Fourier transform can be defined as well, where the integrals are replaced by summations with appropriate weights for each of the samples $\{y_i\}$. Since there are particularly efficient algorithms (fast Fourier transforms or FFTs) for computing discrete Fourier transforms if the number of samples is exactly a power of 2, normally time series are truncated to an appropriate length or padded out symmetrically with appropriate average values in order to get the number of samples to satisfy this criterion. A useful reference for FFT implementations and their properties is Press et al. (1992).

One property of Fourier transforms, whether discrete or continuous, that will be useful for the further discussion, is that they are linear operations, and therefore so are the inverse operations. Linearity implies that the Fourier transform of the sum of functions $y(t)$ and $z(t)$ is the sum of their Fourier transforms (cf. Bracewell, 1965):

$$F(a\, y + b\, z) = a\, F(y) + b\, F(z) \tag{4}$$

where a and b can be any constants.

Of particular importance when carrying out discrete Fourier transforms is a limitation that is imposed by the sampling cadence. A fundamental aspect of this is that with a finite cadence of sampling, with time interval Δt, it is impossible to detect any periodic signal with a frequency that is higher than the Nyquist frequency. This frequency is:

$$\omega_{Nyq} = \frac{\pi}{\Delta t} \tag{5}$$

A discretely sampled, finite, time series can be represented without loss of information by its discretely sampled Fourier transform. Since the Fourier transform or spectrum is in general complex valued, the number of sampling points required is only half the number of points in the time series. If a standard FFT is used, these sampling points are uniformly spaced over the frequency interval $[0, \omega_{Nyq}]$.

The reason for discussing Fourier transforms in the context of filtering is that filtering in the time domain corresponds to a particularly simple operation in the frequency domain. In general, filtering in the time domain of a regularly sampled time series can be written in the form of a weighted (moving) average:

$$\tilde{y}_i = \sum_{k=-m}^{n} w_k y_{i+k} \tag{6}$$

where the weights w_k are the filter factors. This is a general form, allowing for asymmetric filtering. For instance in causal filters $n = 0$ so that only information from the past and present of point i is used and none from its future. While m and/or n could in principle be infinite, this has no practical purpose in the current context. In the context of seasonal filtering, of which the present application is a variant, it is more usual to employ symmetric filters so that not only $m = n$ but $w_{-k} = w_k$. Often an additional property that is imposed is that $\sum w_k = 1$, but in the applications at hand this is not always used.

The equation (6) above is known as a convolution in the time domain of the function y and the function w, in which the latter is simply the set of averaging weights interpreted as (samples of) a function of time. It can be demonstrated that the Fourier transform of the convolution of two series in the time domain is identical to the product of the Fourier transforms of the two time series (cf. Bracewell, 1965). Therefore the Fourier transform of the filtered time series $\tilde{y}$ can be calculated as the product of the Fourier transform of the original time series and the Fourier transform of the filter function w.

$$F(\tilde{y}) = F(y * w) = F(y)F(w) \tag{7}$$

in which the $*$ indicates the convolution of functions. Multiplication is a more straightforward operation than calculating the above moving averages of (6), and in the frequency domain it is also more straightforward to see the effect of the filter on signals with a particular frequency or period. However, there are advantages to carrying out the filtering in the time domain, for instance for time series that are continually being updated with more recent data.

Figure 2 The band-pass filter designed for extraction of signal at intermediate frequencies. Removal of such signal is achieved by applying (1-Response) to a time series.

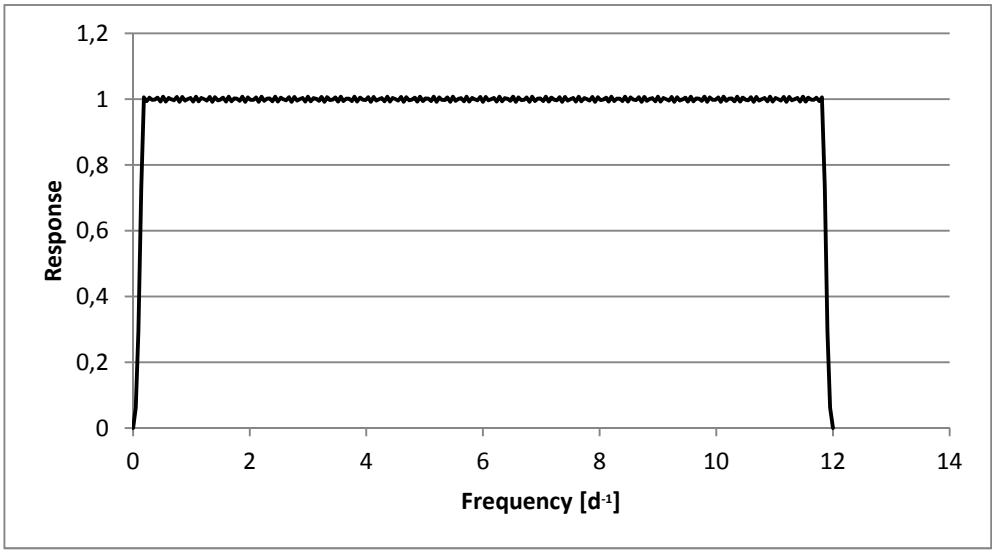

By removing the signal that the band-pass filter $F(w)$, shown in figure 2, lets through, all signal in a band between frequencies $\omega/2\pi$ of between $\sim 1/6$ cycle/day and $\sim 71/6$ cycle/day is taken out of the time series. From equation (4) it follows that this subtraction is equivalent to multiplying in the Fourier domain by $(1 - F(w))$ of the filter $F(w)$ shown in fig. 2. The signal at lower frequencies is in essence the trend in the time series with in addition some remaining low frequency variations. The signal in the range of higher frequencies up to the Nyquist frequency of 12 cycle/day, in principle could be real periodic signal but in that frequency range the

stochastic component is dominant. For convenience therefore the term 'noise' is used. It is the power in this range that will be used to establish confidence limits on parameters to be derived.

For the filter shown in figure 2, the weights w_k in the time domain for all odd values of k are identically equal to 0, as are the weights for $k > 212$. The property that the weights are zero for all odd values of k produces an additional advantage: the high frequency section of the spectrum can be determined in a very simple second step, which is evidently statistically independent. If $\tilde{y}$ is the time series from which the intermediate timescales have been removed with the above filter then the high frequency component h is obtained by taking:

$$h_i = (2\tilde{y}_i - \tilde{y}_{i-1} - \tilde{y}_{i+1})/4 \quad (8)$$

whereas the trend time series c (without high frequency contributions) is obtained from :

$$c_i = (2\tilde{y}_i + \tilde{y}_{i-1} + \tilde{y}_{i+1})/4 \quad (9)$$

With the weights w_k in the time domain, obtained from an inverse Fourier transform of the response function, the filtering can be carried out entirely in the time domain. No further Fourier transforms are necessary.

2.3 The filtered time series for each station

The first processing step of the time series of the GPS stations is to apply the filtering procedure with the filter factors as derived in the previous section. The filtering is done in the time domain. The average height of each station is of less interest for the analysis at hand. For presentational purposes as well, it is more convenient to subtract, for each station, the average value of its height over a common period which here is taken to be day 50 to day 270, where day 1 starts at time 0^h on Jan. 1 2014.

From figure 3 it is clear that even in the filtered data there is still considerable variation with time. For a subset of stations the vertical displacement appears to show correlation between the stations as well. In principle it is possible that some of this correlation is due to a bias in the measurements, which at this stage is impossible to distinguish from genuine ground displacement. Not all of these variations are of immediate interest for the present analysis however. It is appropriate to use the geographical groupings of section 2.1 to construct, by straightforward averaging, a time series for each of these groups: 'Middle', 'South', and 'Edge'. In addition, the collective movement of the entire province (ie. all stations) with respect to a fixed external reference frame is of subordinate interest. The reason for this is that this collective displacement does not cause deformations or the build-up of stress energy, and is therefore less likely to have any association with earthquakes.

The quantity that is more of interest is the 'sagging' of the stations in group 'Middle' and in group 'South' with respect to those in group 'Edge'. It is this differential displacement over a spatial scale of the order of, or smaller than, the field that will lead to the build-up of stresses. Therefore the displacements, averaged for each group, are subtracted in the sense 'Middle' minus 'Edge' and 'South' minus 'Edge'. The result is shown in figure 4. An additional reason to construct such a differential measurement between several groups is that this can compensate to some extent for any unknown bias that has produced a spurious correlation between all GPS stations. For instance in figure 3 around day number 175, ie. towards the end of June 2014, by eye a systematic downward drift appears to set in for all GPS stations. In the differential measures shown in figure 4 this effect disappears. Reprocessing of the original GPS data by the

Figure 3 The time series of the GPS height, after filtering out intermediate time scale variations. For presentational purposes for each station the average of its time series over the period from day 50 to day 214 is subtracted. The vertical scale is in m. Top Panel: GPS stations near the centre of the production field, bottom panel: GPS locations in groups edge and south.

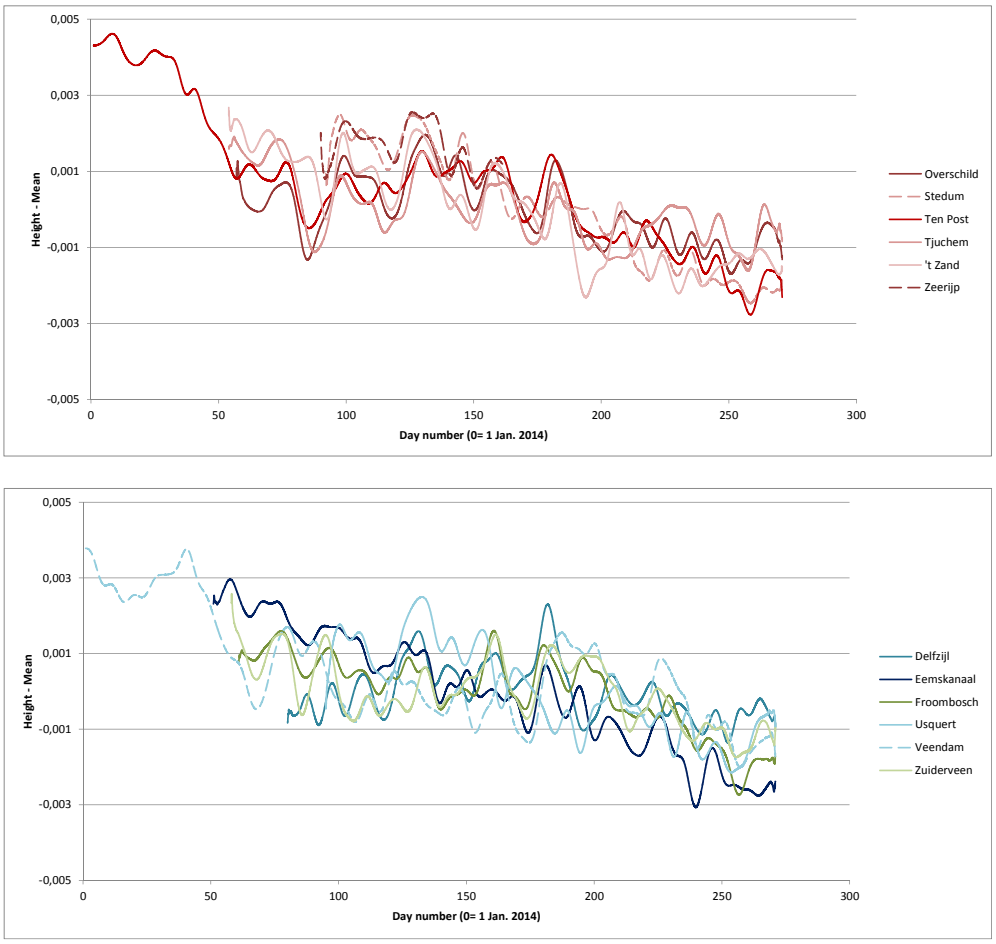

supplier of these data might shed light on whether there is an issue to resolve, but the conclusions of the present analysis appear not to be affected.

The differences between the group averages are shown in figure 4. To indicate the point-to-point uncertainty due to the high frequency noise in the time series an example error bar corresponding to the 99% confidence interval is drawn as well, at day number −200 and displacement 0. From this diagram it becomes more clear what the systematic differential movement is of the two production fields compared to the points external to them. The more 'ragged' character of the time series near the beginnings and ends of the series are due to the fact that the moving average filter would require points for which no measurements are available, so the quality of the result suffers. Because the filter factors decrease very rapidly from the central point in the averaging window, the extent of this effect is limited.

To assess the extent to which correlations might still be present in the differential displacement, two additional differential time series are shown in figure 4. One concerns the differential displacement between stations that are all within group 'Middle', in the sense (Zeerijp+ten Post) − (Stedum+Overschild) which is shown as a dashed black line. The other concerns the

Figure 4 The time series of the filtered GPS height for the differences between the group averages `Middle' – `Edge' (in red), and `South' – `Edge' (in green). The initial section of the series `Middle' – `Edge' is given a different color (blue), because for that section there are only data available for one of the stations in the group `Middle' (Ten Post) and also only one station in the group `Edge' (Veendam). The 99% confidence interval for each point is indicated by the example error bar in black plotted at day -200 and displacement 0. Solid black line : differential measurements between stations all within group `Edge'. Dashed black line : differential measurements between stations all within group `Middle'.

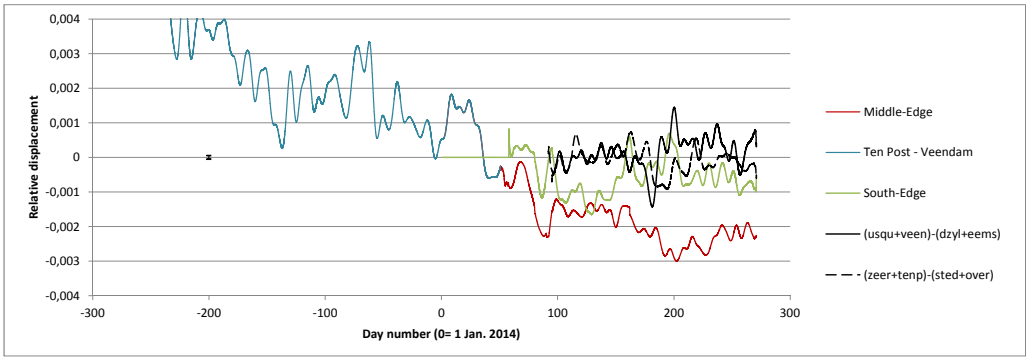

differential displacement between stations that are all within group 'Edge', in the sense (Usquert+Veendam) – (Delfzijl+Eemskanaal) shown as a solid black line. In the absence of remaining biases these differential measurements should not show any trends. From figure 4 it can be seen that the data are consistent with this. There are no visible trends in relative vertical displacements between stations separated by smaller horizontal distances as is seen from the measurements between GPS stations all within the group 'Middle'. Also, there are no trends between stations separated by somewhat larger horizontal distances as is seen from the measurements between GPS stations all within the group 'Edge'. It is difficult to envisage a remaining source of bias that would affect the 'Middle'- 'Edge' measurements so as to mimic a trend, but that at the same time does not affect differential measurements between stations within each of these groups. Most likely the trends seen in the 'Edge'-'Middle' graph are genuine rather than due to unknown biases.

3 Drift trends

In figure 4 the very long-term trends are already more easily identified by eye than in figure 3, even though there remains some variation visible at intermediate time scales. While it would in principle be possible to design a filter that removed even more of the intermediate timescale variations, at this stage it is more straightforward to fit a straight trendline to the time series of 'Middle' minus 'Edge'. The time series for 'South' minus 'Edge' is of less interest at this stage, because while there have been variations in the gas production of wells near these locations, there has not been the strong reduction in production that has taken place at the wells at/around the 'Middle' group stations. Also, the filtered differential time series does not appear to show a very systematic downward drift compared to the 'Edge' stations.

Since the reduction of the gas production at some locations near stations in the group 'Middle' has been rapid, it is of interest to explore whether a break can be found in the linear component

of the downward drift for the time series for 'Middle' minus 'Edge'. To this end one can fit not only a single straight line, using standard least-squares fitting, but also introduce a break-point with a different straight line fit before and after that point. Given that the latter fit has more degrees of freedom than the former, this should be accounted for in assessing whether the two-section fit with a break is genuinely better than a single straight line. To this end the χ^2 measure is calculated for each of the subsections, and for each of the sections the χ^2 is divided by the numbers of degrees of freedom. These two *reduced* χ^2-values are then added. This is the appropriate value of the χ^2 for the full fit, including the break, which is to be compared with the reduced χ^2 of the fit where a single linear trend line is fitted to the full time series. Note that care should be taken with determining the numbers of degrees of freedom. Since the data have been filtered the number of degrees of freedom is much smaller than the number of points in the series . The factor by which this number is reduced can be determined from the properties of the filter (cf. Yan et al., 2003) and in the present case is about 0.019.

Figure 5 The reduced χ^2 measures for the fits for each subsection, shown as a function of the location of the break point. On the vertical scale, 1 corresponds to the value of the reduced χ^2 for a single linear fit of the entire time series `Middle'− `Edge'.

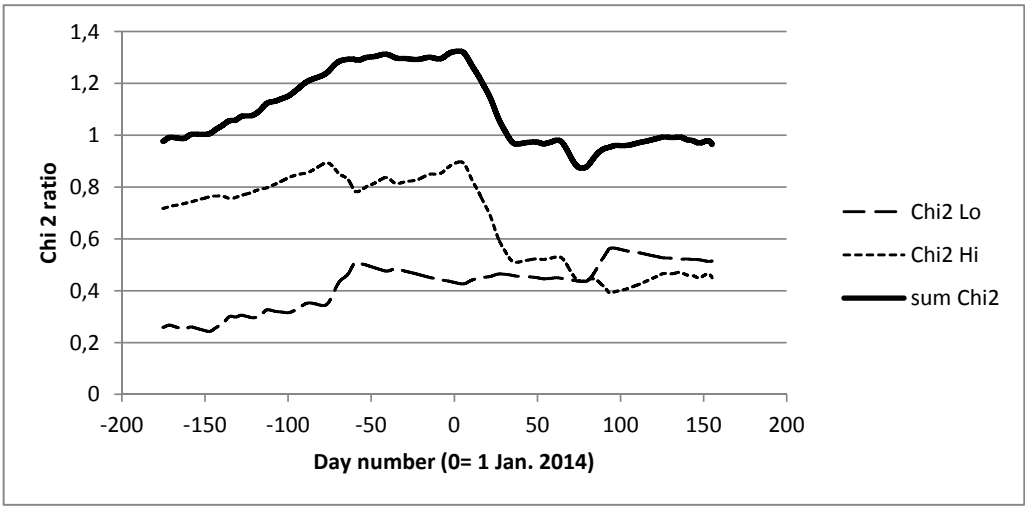

It should further be noted that the differential time series before day 51 are based on only two GPS stations, whereas towards the most recent end of the time series, all GPS stations are operational. This must be accounted for by appropriate weighting of each data point. There is no a priori knowledge about the time at which a break should be located, therefore a fit is constructed separately for a break located at any of the sampling times between around day -175 (where 0 corresponds to January 1^{st} 2014) and day 160. The reduced χ^2 of the fit for each section is shown in figure 5, divided by the reduced χ^2 of the single linear fit, as a function of the location of the break point. Also shown in figure 5 is the sum of these two. No attempt is made to produce fits with a break before day -175 or after day 160 because the length of the shorter section of the series is then too small to produce a valid fit.

From figure 5 it can be seen that for a choice of breakpoint in most of the range the reduced χ^2 is not lower than the reduced χ^2 of a single linear fit. However, there is a small range for which the χ^2 ratio does dip below 1, reaching a minimum for a breakpoint located at day 76.8. Even with the reduced number of degrees of freedom compared to unfiltered data, the reduction in

reduced χ^2 compared to a single linear fit is statistically significant at a confidence level in excess of 99%.

Figure 6 The time series of the filtered GPS height for the differences between the group averages `Middle' − `Edge'. A single linear trend fitted to the data (dashed line) is also shown, as well as a two-section trend with the best choice of a jump point (solid line), and a two-section trend additionally constrained to be continuous (dash-dot).

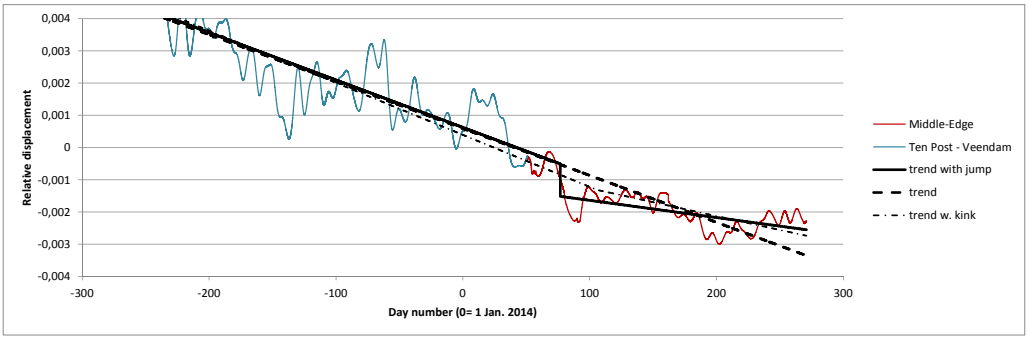

In figure 6 the fit lines corresponding to a single linear trend for the entire series, and a two-section trend with a break are shown. The slope of the fit for the initial section, shown in figure 6, is -0.015 ± 0.0002 mm/day whereas the slope for the final section is -0.005 ± 0.0004 mm/day. The difference in slopes is highly statistically significant at a level well in excess of a 99.9% confidence level. The location of the jump at day 76.8 corresponds to the afternoon of March 17, 2014. Also shown as a dashed line is a fit where the two sections are constrained to match at the breakpoint, so that there is a kink but the fit remains continous. This is also a statistically significant better fit that a single linear trend. In this case the location of the kink is at day 105.9 which corresponds to the afternoon of April 15, 2014. For the discontinuous fit the ratio of the slopes before and after the jump is 2.8, for the continous fit the ratio of the slopes before and after the kink is 1.9. These linear fits are intended to enable a measurement of changes in average downward speeds. The filtered data show that there is not an abrupt break, but rather a gradual change in average speed that manifests itself over a period of a few weeks. Hence the 'break'-dates associated with each fit serve as a guide between which dates the changeover takes place.

The location in time of the break is very close to the calendar date from which GPS-data from the most recently added station, Usquert, start to be available, which is part of the 'Edge' group of stations. To assess whether this in some way has influenced the fit, and the location of the break, the analysis has been repeated removing all GPS-data from the Usquert station from the analysis. After removing these data the result of the fitting process is that there is still a preference for a two-section fit with a break rather than just a single linear trend, where the break is located a few hours later than was the case before. This implies that the existence of the break is not an artifact of adding in the Usquert GPS data.

4 Conclusions

From the GPS data it can be concluded that there is continued subsidence of the ground in the area of the wells where production was reduced in the month of January of 2014. At the time of

writing of this report, the production data of the wells are available with a cadence of once per month. On January 17 of 2014 the production was reduced by around a factor of 20, at wells at 5 locations ('t Zandt, Leermens, Ten Post, Overschild, De Paauwen). The location of the break is around 9 weeks later, but there is a fair margin of uncertainty (some 2 weeks) around the exact value of the time gap. Also, the location of the gap shifts somewhat depending on whether the fitting function is continuous or discontinuous.

The ratio of the speeds of subsidence before and after the break around the beginning of April imply that the speed is reduced by a factor of around 2.8 after the jump, or a factor of 1.9 if a continuous fitting function is used instead. This is much less than the reduction of the production levels by a factor of 20 at the locations mentioned above. However, at locations directly adjacent, and well within the region relevant to the GPS stations in the group 'Middle' for which the subsidence is measured, the production has not decreased. If the production at the locations Amsweer, Tjuchem, and Siddeburen is also taken into consideration, together with the 5 well locations mentioned above, the combined rates of production have decreased only by a factor of around 2.5.

The combination of the GPS and production data at this stage is consistent with the hypothesis that a reduction in production levels can produce a similar reduction in subsidence speeds after a delay of between 7 and 11 weeks. However, this hypothesis is based on a single event ie. the strong reduction in production levels in January of 2014 at 5 well locations. A causal connection can therefore neither be demonstrated nor be disproved by this analysis.

While there is clear statistical evidence that a break has occurred in the subsidence rates, and the reduction in subsidence speeds is measurable with a high degree of precision, it is desirable to perform a more systematic analysis of the correlation between the production time series and the GPS height time series.

As next steps in this research it is recommended to:

- Cross-correlate time series for gas production with the time series for the GPS data, to establish a more firm basis for and better measurement of the time delay between production variations and subsidence rates. To this end it is necessary for well production data to be available with a much higher cadence, such as daily. Such an analysis is necessary to assess whether and to what extent a chain of causality can be established between production variations and subsidence variations.
- Explore the relationship between subsidence rates, or breaks in subsidence rates, and the frequency, location and intensity of earthquakes. This is the next link in the chain of causality towards assessing gas-production induced risk and hazard in Groningen.

References

Nederlandse Aardolie Maatschappij BV, 2013, A technical addendum to the winningsplan Groningen 2013 Subsidence, Induced Earthquakes and Seismic Hazard Analysis in the Groningen Field.

Bracewell, R.N., 1965, The Fourier transform and its applications, 5ff, 24ff, McGraw-Hill (paperback edition: 1999).

Yan, H.M., Zhong, M., Zhu, Y.Z., 2003, The determination of degrees of freedom for digital filtered time series, *Acta Astronomica Sinica* 44, 3, 324-329.

Press, W.H., Flannery, B.P., Teukolsky, S.A., Vetterling, W.T., 1992, Numerical Recipes in FORTRAN 77 2^{nd} Ed. : the art of scientific computing, 490ff, 551-556, 569-576, Cambridge University Press.