



Scientific Paper

The relationship between pressure gradients and earthquake rates in Groningen

Frank P. Pijpers

June 2018

Contents

1	Introduction	3
2	Background	6
2.1	Propagation of the signal of gas extraction	6
2.2	The local gas pressure precursor for quakes	6
2.3	setting up synthetic sets	9
3	Comparison of precursors for $M \geq 1$ after Sep. 1 2014	10
3.1	The pressure	10
3.2	The averaged pressure	12
3.3	The horizontal pressure gradient	14
4	Comparison of precursors for $M \geq 1.2$ after Sep. 1 2014	16
4.1	The pressure	16
4.2	The averaged pressure	17
4.3	The horizontal pressure gradient	18
5	Conclusions	20

Nederlands

Deze rapportage is een verslag van onderzoek dat is uitgevoerd in het kader een onderzoeksproject door het CBS in opdracht van Staatstoezicht op de Mijnen (SodM) sinds 2014. Dit onderzoek is ten behoeve van een statistische onderbouwing van het meet- en regelprotocol voor optimale gasexploitatie in de provincie Groningen, met in het bijzonder de aandacht gericht op de relatie tussen de gasdruk in het reservoir, die in de tijd varieert vanwege de gas productie bij de diverse clusters, en de aardbevingen in Groningen.

Uit de analyse blijkt dat er een duidelijke correlatie is in de zin dat voorafgaand aan aardbevingen er op de locatie van die aardbeving een signaal is te zien in de gasdruk in het reservoir. Dit is een her-analyse uitgevoerd in verband met een datalevering met betere, dagelijkse, tijdsresolutie van de reservoir gasdruk. Ook is er gebruik gemaakt van een verbeterde localisatie na her-analyse van de seismische data van een aantal van de aardbevingen in de KNMI catalogus. Een interim versie van deze rapportage is verschenen in Jan. 2018 (Pijpers, 2018).

English

This report is a continuation of research, commenced in 2014, which is part 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 control of the production of natural gas in the province of Groningen. In this research the particular focus is on the relationship between reservoir gas pressure variations and the earthquakes in Groningen.


From the analysis it is clear that there is a correlation, in the sense that prior to earthquakes there is an enhanced signal in the reservoir gas pressure at the location of the earthquake. This is a re-analysis which is carried out because of the availability of improved, daily cadence, time series for the reservoir gas pressure. Also, use is made of improved localisation of a subset of earthquakes in the KNMI catalogue after re-analysis of the relevant seismometer traces. An interim version of this report was published in Jan. 2018 (Pijpers, 2018).

1 Introduction

For some decades earthquakes of modest magnitudes have occurred in and around the Groningen gas field. It is recognized that quakes can be induced by the production of gas from a field eg. (Wetmiller, 1986); (Grasso and Wittlinger, 1990); (Bourne et al., 2014). An extensive study program is in place to improve the understanding of the risk and hazard due to gas production-induced earthquakes in the Groningen field (Nederlandse Aardolie Maatschappij BV, 2013).

The causality relationship between gas production and earthquakes is difficult to establish unequivocally, purely on the basis of time series analyses. For instance, a constant rate of gas extraction might lead to a likelihood of quakes that increases with time through the compaction of the reservoir layer, cf. (Bourne et al., 2014), but a correlation analysis would not necessarily

Figure 1.1 The locations of all quakes within the zone in Groningen relevant to this analysis, recorded after 1 January 1995. The red squares are locations of the production clusters, some of which are identified by name. The production field is also shown in dark gray, overlapped on a map of the region.



provide any evidence of this. However, because in practice the rate of production does vary periodically throughout the year, as well as incidentally for operational reasons, it becomes possible to pursue a correlation analysis between variations in quake incidence and production variations. A correlation analysis between ground subsidence and production variations (Pijpers and van der Laan, 2015b) provides indications that the response time of the force balance in the reservoir layer to extraction of gas might be sufficiently short that a correlation analysis between earthquake rate and production variations would be feasible. There is no doubt that the extraction of natural gas strongly influences the balance of forces in the reservoir layers, eg. (Mulders, 2003). However, every quake is the realisation of a stochastic process, for which the balance of forces only in part determines the probability distribution function of that process. Because of this stochasticity there is an intrinsic decorrelation between variations in gas production and variations in quake incidence frequency. Conversely, even if a correlation can be established between gas production and quake incidence, it will depend on the exact nature of that correlation whether causality can be considered proven or likely.

As with the CBS reports from the previous phases of this project, starting point for the analysis of the statistics of quakes is the catalogue of quakes as reported by the Royal Netherlands Meteorological Institute (KNMI) based on their processing of the network of seismometers that they manage. Further selections from that catalogue are made as necessary, for instance in order to filter out quakes that have occurred outside of the zone of interest, or quakes from an epoch when the seismometer network was less sensitive and may suffer from more serious completeness issues. The selection/filtering criteria are chosen to be consistent with those of the previous Statistics Netherlands reports, (Pijpers, 2014b, 2015a,c, 2016b,c; Pijpers and van Straalen, 2017a,b, 2018): the spatial region of interest is the same as in those reports (see fig. 1.1). For a subset of 100 earthquakes that occurred after Jan. 1 2015, the original seismic data has been re-analysed, in order to improve the localisation. For the analyses in the interim

report (Pijpers, 2018) a set of 342 earthquakes, having occurred between Jan. 1 2014 up to the end of August 2017 is used. If an improved localisation is available from the re-analysis, this is used as data, if not then the original KNMI location is used. This is unchanged from the interim report (Pijpers, 2018), but the focus here lies on one particular subset of earthquakes (subset 4), also discussed in the interim report. For the results of the other subsets the reader is referred to the interim report.

The gas pressure inside the reservoir can be modelled by making use of detailed geophysical mapping of the reservoir and its material properties. These together determine the diffusive flow of gas through the reservoir, induced by the gas production at the various clusters. In a previous report (Pijpers, 2015c), the gas flow induced by the gas production was modelled using a very simplified model of the flow, which did not take any variation of these material properties into account. Much better models are available however, in the form of the reservoir modelling tool MoReS (Por et al., 1989). The NAM has supplied the relevant data from the MoReS code, which has been used to provide the starting point for an improved analysis. The reservoir gas pressure as a function of time at the required locations, ie. the locations of quakes and of a reference set of points, has been determined. These time series data are then combined in a way very similar to what is described in (Pijpers, 2015c) but without having to take recourse to the very simplified flow model used in that report.

The time resolution that MoReS can provide is very high. In a previous analysis (Pijpers, 2016a; Sijacic et al., 2017) a relatively low time resolution of one week was used, to match the preceding analysis of (Pijpers, 2015c). In the present paper the time resolution is daily. The pressure data used are volume weighted vertical averages inside the reservoir. The horizontal resolution of the numerical grid of MoReS varies, but is everywhere sufficient for the present analysis, since spatial uncertainties of less than about 500 *m* are irrelevant given the uncertainty in localisation of quakes. The pressure data is therefore improved in accuracy. The localisation of the quakes has also been improved for a subset of earthquakes for which the original seismic data has been re-analysed. There will nevertheless still be a contribution to noise due to the uncertainty of localisation of quakes. Another source of uncertainty are the limitations of the modelling tool MoReS. While the model is calibrated with high frequency using borehole pressure data there are limitations in e.g. the resolution of the mapping of material properties of the reservoir.

The primary difference between this report and the interim report (Pijpers, 2018) lies in establishing reliable uncertainty estimates, which is achieved by repetition of all analysis steps a number of times, each time with a different realisation of a synthetic comparison set of locations and times. In the interim report, three different epochs are examined. In the present report, the focus is on the second epoch: starting September 1 2014 and ending September 1 2017, and selecting only the earthquakes with magnitudes $M \geq 1$. This corresponds exactly to subset 4 of the interim report, consisting of data for 108 earthquakes. To facilitate comparison with reporting elsewhere, a subset of this set is also reported on, for which $M \geq 1.2$ and which has data for 79 earthquakes.

2 Background

2.1 Propagation of the signal of gas extraction

In this view of the mechanism of induced earthquakes, it is not presumed that the variation in the rate of gas extraction injects large amounts of potential energy into the reservoir which is accumulated and some time later released suddenly as mechanical energy in the form of an earthquake. If that were the case a correlation analysis between variations in the production rate and the rate of energy release through earthquakes ought to be pursued.

Rather, the view is that the potential energy is built up over decades as the gas was removed from the reservoir. This potential energy is thus already present in the system, localised in particular at critically stressed fractures. The history of gas extraction and the associated (differentiated) compaction of the fractured reservoir is responsible for the accumulation of this potential energy. As when tapping the first domino stone in a display of ever increasing numbers of falling dominoes, in this situation even a minor variation in gas pressure, or a spatial gradient of gas pressure (ie. a local forcing), might then cause a cascading energy release: an earthquake. The distribution function of the energy released in any given earthquake (ie. the slope in the Gutenberg-Richter plot) is likely to be influenced more by the (fractal) distribution of fracture sizes. In this scenario there is no reason to expect a clear correlation between gas pressure variations and the amount of energy released in any given quake. Instead the scenario should produce a correlation between gas pressure variations and the rate of incidence of earthquakes (of any magnitude). For this reason the earthquake magnitude is only used as a data quality selection criterion to decide which earthquakes to include in the analysis.

It is clear that the locations of quakes do not coincide exclusively with the locations of the extraction clusters. If the gas extraction at the clusters is to have an influence at other locations, a signal has to pass from the well to those other locations. A signal occurs as gas diffuses through the reservoir layer, to replace the gas extracted at the well, or cluster of wells. A diffusive model for the signal of gas extraction has been used for instance by (Baranova et al., 1999) (their eq. 5) to model induced seismicity by hydrocarbon production in western Canada. A more general description can be found in (Russell and Wheeler, 1983). A controlling parameter for such diffusion is the diffusivity in the medium which is determined by material properties of both the medium and the material diffusing through it: the permeability of the medium and its porosity, the viscosity of the diffusing gas and its compressibility. None of these properties are constant throughout the region, and the system of larger and smaller faults also play an important role in channelling or obstructing flow. All of these factors are incorporated in a detailed finite-element model that is also calibrated on a regular basis by using a number of pressure measurement wells. For a description of the detailed modelling of the flow and pressure in the reservoir layers using MoReS the reader is referred to (Por et al., 1989) .

2.2 The local gas pressure precursor for quakes

In order to assess whether there is any kind of statistical correlation between the gas pressure throughout the Groningen field and the generation of quakes it is of interest to reconstruct the precursor gas pressure at the location of quakes in the period leading up to the time of the quake. It is assumed that the occurrence of quakes is not a fully deterministic process, but rather

that an identical set of circumstances gives an identical finite likelihood for a quake to occur. A single event is then of limited value in terms of the information that it carries concerning the likelihood and its dependence on gas pressure variations. Therefore one would wish to combine the precursor of every relevant quake to find the typical local gas pressure precursor for quake events. The dataset provided for the reservoir gas pressure by NAM has been generated using the MoReS code, with daily resolution starting on 1 January 2013. The subsequent analysis follows the same steps as is done in Pijpers (2016a); Sijacic et al. (2017).

The time $T_{quake\ i}$ at which each quake (identified with index i) occurs is the reference value which is the 0-point for the shifting in time. The reservoir pressure P as a function of time t is retrieved from the MoReS dataset for all selected quake locations $(x_{quake\ i}, y_{quake\ i})$. This is done for all selected quakes after which straightforward averaging of the N precursors for every (selected) quake achieves the goal of obtaining the typical precursor of local gas pressure, and relative pressure variation before quakes as a function of 'look-back' time from the occurrence of each earthquake. This shifting and adding procedure is sometimes referred to as co-adding, i.e.:

$$P_{co-add\ real}(t) = \frac{1}{N} \sum_{i=1}^N P_{reservoir}(x_{quake\ i}, y_{quake\ i}, (t - T_{quake\ i})) \quad (1)$$

Obtaining this averaged local pressure precursor for quakes is not yet sufficient for the purposes of the analysis. If instead of using locations and times of the selected quakes, one would pick an equal number of randomly chosen locations and times, a local pressure precursor could be obtained by applying the same technique of co-adding (i.e. shifting in time and averaging) of gas pressure precursors.

$$P_{co-add\ synth}(t) = \frac{1}{N} \sum_{i=1}^N P_{reservoir}(x_{synth\ i}, y_{synth\ i}, (t - T_{synth\ i})) \quad (2)$$

Since by definition such a randomly chosen set of locations and times is not special, it is only the differences between the precursors of true quakes and synthetic 'events' that is of interest. If the co-added pressure precursor of true quakes turns out to be statistically indistinguishable from the co-added pressure precursor of random locations and times, no causality between gas pressure and quakes can be demonstrated by this route.

In terms of methodology this approach of comparing a 'treatment' population (true earthquakes) with a non-treatment control population (random locations and times) is of course common in medical research whenever the effectiveness of new treatments or medication needs to be assessed. The shift-and-add / 'co-adding' technique to extract a common behaviour or property, at as high as feasible signal-to-noise ratio, is borrowed from well-established astronomical imaging techniques (cf. Lynds et al. (1976); Christou et al. (1986)) which are still being used and developed further (e.g. Harpsøe et al. (2012)).


In the above description the pressure $P_{reservoir}$ is a volume average over a column through the reservoir. The MoReS output provides the pressure on a fixed grid of locations, so for each (real or synthetic) quake the nearest grid cell is retrieved from that output. In addition to simply selecting the nearest cell, some further indicators are also analysed using this co-adding technique. The first additional indicator is a local (horizontal) average. The reason to include this is that if a (real or synthetic) earthquake is localized very near a fracture with different pressures on either side, the localisation picks one pressure or the other. An average might then be more representative, i.e. have lower noise, than the nearest grid cell itself.

The fixed grid is adapted to some extent to the shape and features of the reservoir, and therefore it is not a completely regular rectangular grid. In order to enable forming a local horizontal average pressure the 8 cells which are the nearest neighbours of that grid cell need to be found.

$\bar{P}_i$ is the local average for grid cell i together with its 8 nearest neighbours. Since the grid is fixed in time a simple lookup table can be set up which defines for each grid cell i the 9 cells, including cell i itself, to be included in the average. This lookup result is indicated with the function $j(i, k)$ where k runs from 0 to 8: $j(i, 0) = i$ and $j(i, k)$ with $k = 1, \dots, 8$ successively provides the index for the neighbouring grid cells of cell i in counter-clockwise order (schematically shown in fig. 2.1).

$$\begin{aligned}\bar{P}_{co-add real}(t) &= \frac{1}{N} \sum_{i=1}^N \left[\frac{1}{9} \sum_{k=0}^8 P_{reservoir}(x_{quake j(i,k)}, y_{quake j(i,k)}, (t - T_{quake i})) \right] \quad (3) \\ \bar{P}_{co-add synth}(t) &= \frac{1}{N} \sum_{i=1}^N \left[\frac{1}{9} \sum_{k=0}^8 P_{reservoir}(x_{synth j(i,k)}, y_{synth j(i,k)}, (t - T_{synth i})) \right]\end{aligned}$$

Figure 2.1 Schematic diagram showing the configuration of a central cell in an irregular grid, with its nearest neighbours. The arrows indicate the base vectors used to determine a horizontal pressure gradient.



Having access to the nearest neighbours, it is also possible to calculate a finite difference approximation to the local horizontal gradient in the pressure by taking differences between pressures in the central column and neighbouring columns, and correcting for the shape factors introduced by the irregularity of the grid. With the arrangement of the cell lookup function 4 sets of local 'base vectors' can be set up, along which finite pressure differences can be calculated.

Set 1, with $\Delta P_{10} \equiv P_{j(i,1)} - P_{j(i,0)}$, $\Delta P_{30} \equiv P_{j(i,3)} - P_{j(i,0)}$:

$$\begin{pmatrix} x_{j(i,1)} - x_{j(i,0)} \\ y_{j(i,1)} - y_{j(i,0)} \end{pmatrix}, \begin{pmatrix} x_{j(i,3)} - x_{j(i,0)} \\ y_{j(i,3)} - y_{j(i,0)} \end{pmatrix} \quad (4)$$

Set 2, with $\Delta P_{20} \equiv P_{j(i,2)} - P_{j(i,0)}$, $\Delta P_{40} \equiv P_{j(i,4)} - P_{j(i,0)}$:

$$\begin{pmatrix} x_{j(i,2)} - x_{j(i,0)} \\ y_{j(i,2)} - y_{j(i,0)} \end{pmatrix}, \begin{pmatrix} x_{j(i,4)} - x_{j(i,0)} \\ y_{j(i,4)} - y_{j(i,0)} \end{pmatrix} \quad (5)$$

Set 3, with $\Delta P_{50} \equiv P_{j(i,5)} - P_{j(i,0)}$, $\Delta P_{70} \equiv P_{j(i,7)} - P_{j(i,0)}$:

$$\begin{pmatrix} x_{j(i,5)} - x_{j(i,0)} \\ y_{j(i,5)} - y_{j(i,0)} \end{pmatrix}, \begin{pmatrix} x_{j(i,7)} - x_{j(i,0)} \\ y_{j(i,7)} - y_{j(i,0)} \end{pmatrix} \quad (6)$$

Set 4, with $\Delta P_{60} \equiv P_{j(i,6)} - P_{j(i,0)}$, $\Delta P_{80} \equiv P_{j(i,8)} - P_{j(i,0)}$:

$$\begin{pmatrix} x_{j(i,6)} - x_{j(i,0)} \\ y_{j(i,6)} - y_{j(i,0)} \end{pmatrix}, \begin{pmatrix} x_{j(i,8)} - x_{j(i,0)} \\ y_{j(i,8)} - y_{j(i,0)} \end{pmatrix} \quad (7)$$

Using each of these four sets, a horizontal pressure gradient is determined using finite differences in the corresponding reservoir pressures, and applying the appropriate local coordinate transform to calculate the gradient in a common (x, y) coordinate system. This process therefore provides estimates of $(\partial P / \partial x, \partial P / \partial y)$. The absolute value of these pressure

gradients is then approximated by:

$$|\nabla_h P| \equiv \sqrt{\left(\frac{\Delta P}{\Delta x}\right)^2 + \left(\frac{\Delta P}{\Delta y}\right)^2} \quad (8)$$

The 4 different determinations, from each of the 4 choices of base set, are compared and the largest value is selected. In what follows this largest value is indicated by $|\nabla_h P|_i$. The reason to proceed in this way is that a single unfortunate combination of grid cells can produce an orientation along which there is a very small and unrepresentative apparent gradient, whereas with the 4 different chosen combinations this is unlikely to occur for all 4 sets. Once a $|\nabla_h P|_i$ can be calculated, the same co-adding analysis can be applied as is done for P itself

$$\begin{aligned} |\nabla_h P|_{co-add \text{ real}}(t) &= \frac{1}{N} \sum_{i=1}^N |\nabla_h P|_i(t - T_{quake \ i}) \\ |\nabla_h P|_{co-add \text{ synth}}(t) &= \frac{1}{N} \sum_{i=1}^N |\nabla_h P|_i(t - T_{synth \ i}) \end{aligned} \quad (9)$$

2.3 setting up synthetic sets

The analysis thus requires that a set of locations and times be set up, which can be used to build a local precursor of gas pressure in the same way that it is done for the true set of selected quakes. It might not be appropriate however to generate such a set using a completely uniform distribution over space or time. It is clear from fig. 1.1 that the spatial distribution of quakes is not uniform, and it is known that the distribution over time also is not completely uniform. It is desirable to mimic at least some aspects of the real distribution over space and time in the distribution for the set of synthetic events. To this end the number of quakes in grid squares is counted and for the synthetic set the number of events in each of these squares is reproduced, but random (x, y) -coordinates uniformly distributed within each grid square are generated. This type of strategy is sometimes referred to as stratified sampling. For the synthetic set that is used to compare with the selection of earthquakes with magnitudes $M \geq 1$, the grid squares used to set up samples are 2×2 km. For the synthetic events to compare with the set with magnitudes $M \geq 1.2$ 4×4 km squares are used because for the smaller squares the numbers become too small.

For the timing of these artificial events, the full time range since Jan. 1 2014 is sectioned into intervals of length 50 d and the number of true events in each of these sections counted. The number of synthetic events must match this true number but the actual timing within each subinterval is generated from a uniform distribution.

In particular when counting earthquakes with magnitudes above $M = 1$, the detection and localisation of earthquakes has been reliable for the entire period after Sep. 2014. The much improved sensitivity over of the seismic detector network that KNMI has at its disposal, since upgrades implemented over the course of 2015, are primarily noticable for earthquakes with lower magnitudes. The detection completeness limit is now estimated to be near $M = 0.8$. Nevertheless elsewhere a limit of $M = 1.2$ is used instead, so an additional set with selection $M \geq 1.2$ is also analysed. In the interim report (Pijpers, 2018), the analysis is applied to 5 different selections from the earthquake catalogue: sets covering a shorter or a longer total length in time, and also applying different selection criteria for the magnitude of the earthquakes. In the present report the focus lies on selection 4 of the interim report: 108 earthquakes with $M \geq 1$ in the period 1/9/2014 – 31/8/2017. For this set 25 separate realisations of synthetic data are constructed and the co-adding is carried out for each realisation. After co-adding, the

average of the 25 realisations is calculated, for all quantities that are relevant for the report, and also the r.m.s. σ of the variation around that average is determined. By splitting up these 25 realisations into two disjunct subsets and examining the value of σ for each subset it is established that 25 realisations is sufficient to ensure that the relative uncertainty for the value of σ is $\delta(\sigma)/\sigma < 0.1$, i.e. σ is determined to (at least) 1 significant digit with 25 realisations.


3 Comparison of precursors for $M \geq 1$ after Sep. 1 2014

3.1 The pressure

Over the course of 2015 improvements were implemented by KNMI to the seismic detector network, installing many more detectors, spread out more evenly over the region. The localisation of earthquakes with $M \geq 1$ over that entire period was good, in particular after reanalysis of the original seismograms dating from the start-up phase. Over this same period the seasonal swing overall in production was reduced, although at individual clusters there was still considerable variation, so that in terms of variations propagating diffusively through the reservoir one would expect there to still be quite some temporal and spatial variation in the reservoir pressure as well.

Fig. 3.1 shows immediately that there is an offset between the co-added pressure for the real earthquakes and for the average of the 25 realisations of synthetic earthquakes. While an offset indicates that there is a systematic difference between the locations and timings of real earthquakes and of synthetic sets, a constant valued offset is perhaps of less interest since it does not provide a clear path of inference concerning a cause.

Figure 3.1 The co-added pressure for the real earthquakes (in blue) as well as the average for the 25 realisations of co-added pressure for the synthetic sets (red)



Given that a constant offset is perhaps of less immediate importance, the ratio of real to synthetic co-added pressure is calculated in two different ways. The first way to construct the

averaged ratio is as a straightforward average over the ratios for each realisation:

$$R_P(t) = \frac{1}{25} \sum_{j=1}^{25} \frac{P_{co-add \text{ real}}(t)}{P_{co-add \text{ synth } j}(t)} \quad (10)$$

where $P_{co-add \text{ real}}$ and $P_{co-add \text{ synth } j}(t)$ are determined using Eq. (1) and (2) respectively. In this case the offsets remain visible in the resulting ratio.

A second possibility is to use a weighted average, in such a way that any constant offset between the real and synthetic co-added pressure is removed. In order to achieve this a single weight per realisation w_j can be defined:


$$w_j \equiv \frac{\sum_{k=0}^{365} P_{co-add \text{ synth } j}(t_k)}{\sum_{k=0}^{365} P_{co-add \text{ real}}(t_k)} \quad (11)$$

in which the time is regularly sampled at a cadence of once per day: $t_k = -1.0 * k$. A weighted average ratio is then

$$R_{norm \text{ } P}(t) = \frac{1}{25} \sum_{j=1}^{25} w_j \frac{P_{co-add \text{ real}}(t)}{P_{co-add \text{ synth } j}(t)} \quad (12)$$

This is exactly equivalent to normalising each of the co-added pressures to average to unity over

Figure 3.2 Top panel: The ratio $R(t)$ for the co-added pressure, averaged for the 25 realisations of a synthetic set. Bottom panel: the ratio of the normalised co-added pressures $R_{norm}(t)$ and the $\pm 1\sigma$ range (lighter grey lines).



the period from $t = -365$ days to $t = 0$ before calculating their ratio. The result is shown in fig. 3.2. Comparing the top and bottom panels of fig. 3.2 shows that the normalisation has little or no effect on the behaviour as a function of time: the ratio between real and synthetic co-added

pressure is clearly not constant as a function of time. In the bottom panel the grey lines indicate the 1σ range for the 25 realisations, from which it is clear that the deviations from unity are statistically significant. The average over the 25 realisations of the normalisation factors w_j is 0.995 with a 1σ standard deviation of 0.006. The difference from 1 of the offset, visible as a vertical shift when comparing top and bottom panels of Fig. 3.2, is therefore not statistically significant. This means that the various realisations all show essentially the same shape as a function of time, but they are simply shifted up or down as a whole between realisations.

From this result it is evident that despite the constraint on NAM to reduce the overall seasonal pattern in gas production, the variations in production rate at the level of individual clusters continue to cause variations in reservoir gas pressure. The variations in reservoir gas pressure are such that locations in the reservoir that are prone to generate earthquakes, such as fractures, continue to show sensitivity to such variations. Assuming that fractures may be in a critical state, it is not implausible that an excess effect of a few per mille, the total range in the ratio R_{norm} , is sufficient to produce an (additional) stimulus for generating earthquakes.

3.2 The averaged pressure

The local average of the grid cell nearest to a (real or synthetic) earthquake location together with its 8 nearest neighbour grid cells is subjected to the same co-adding technique as discussed in section 3.1. As one would expect, the additional local averaging over 9 near grid cells reduces systematic offsets between the co-added pressures for the real earthquakes and the synthetic earthquakes. The reason for this is that together with the cell that is genuinely closest to each earthquake, an additional 8 cells are included which are more distant and therefore cannot show as much of an effect prior to an earthquake. However, if there is some ambiguity in determining in which cell any given earthquake is located, averaging over a patch of adjacent cells should reduce the likelihood of missing a pressure precursor effect altogether. By showing and comparing both results, a measure of robustness can be obtained. Fig. 3.3 shows that the pressure precursors for real earthquakes and for the average of all 25 realisations of synthetic earthquakes are barely shifted with respect to each other.

Figure 3.3 The co-added locally averaged pressure for the real earthquakes (in blue) as well as the average for the 25 realisations of co-added locally averaged pressure for the synthetic sets (red)





Figure 3.4 Top panel: The ratio $R(t)$ for the co-added locally averaged pressure, averaged for the 25 realisations of a synthetic set. Bottom panel: the ratio of the normalised co-added pressures $R_{norm}(t)$ and the $\pm 1\sigma$ range (lighter grey lines)



In fig. 3.4 the ratio and normalised ratio of these locally average pressures are shown, defined analogously to the ratios for the pressure itself:

$$R_{P\ av}(t) = \frac{1}{25} \sum_{j=1}^{25} \frac{\bar{P}_{co-add\ real}(t)}{\bar{P}_{co-add\ synth\ j}(t)}$$

$$R_{norm\ P\ av}(t) = \frac{1}{25} \sum_{j=1}^{25} w_j \frac{\bar{P}_{co-add\ real}(t)}{\bar{P}_{co-add\ synth\ j}(t)} \quad (13)$$

where the weights now satisfy:

$$w_j \equiv \frac{\sum_{k=0}^{365} \bar{P}_{co-add\ synth\ j}(t_k)}{\sum_{k=0}^{365} \bar{P}_{co-add\ real}(t_k)} \quad (14)$$


It is interesting that the shape of this curve, the temporal pattern, is very similar to the pattern obtained in the previous section without the local averaging. Even the full range of variation is barely reduced. This demonstrates that the result is not particularly sensitive to the precise method of picking the nearest or most appropriate grid cell for each earthquake location. The deviations from unity continue to be statistically significant. The average over the 25 realisations of the normalisation factors w_j is 0.999 ± 0.006 .

This result implies that for any given realisation of a synthetic set, the entire pattern can shift up or down by a relatively large amount, but the shape overall remains the same. As argued before it is primarily the shape, i.e. the fact that the ratio is not a constant and therefore the plot not a horizontal straight line, that points to gas pressure variations modulating earthquake generation.

3.3 The horizontal pressure gradient

As described in section 2.2 the absolute value of the local horizontal pressure gradient $|\nabla_h P|$ is determined and co-added for the real and synthetic sets (cf. Eq. (9)). In this case the average values of about 105 N/m^3 and about 73 N/m^3 are sufficiently different that two separate panels are used to show their behaviour as a function of time. The full extent of the scale in both panels is the same: 4 N/m^3 . The top panel of fig. 3.5 shows the behaviour of the co-added

Figure 3.5 The co-added horizontal pressure gradient for the real earthquakes. The average for the 25 realisations of co-added horizontal pressure gradient for the synthetic sets (red)




horizontal pressure gradient for the set of real earthquakes, whereas the bottom panel shows the behaviour of the co-added horizontal pressure gradient for the set of synthetic earthquakes. It is immediately evident that the behaviour with time of the co-added time series is quite different for the earthquake locations and for synthetic locations. The development in the rest of the field is mostly flat or very slightly dropping, whereas at the earthquake locations the gradient is higher and also steeply rising.

In fig. 3.6 the ratio and normalised ratio of these horizontal pressure gradients are shown, defined analogously to the ratios for the pressure itself:

$$\begin{aligned}
 R_{nabla P}(t) &= \frac{1}{25} \sum_{j=1}^{25} \frac{|\nabla_h P|_{co-add \text{ real}}(t)}{|\nabla_h P|_{co-add \text{ synth } j}(t)} \\
 R_{norm \nabla P}(t) &= \frac{1}{25} \sum_{j=1}^{25} w_j \frac{|\nabla_h P|_{co-add \text{ real}}(t)}{|\nabla_h P|_{co-add \text{ synth } j}(t)}
 \end{aligned} \tag{15}$$

Figure 3.6 Top panel: The ratio $R(t)$ for the co-added horizontal pressure gradient, averaged for the 25 realisations of a synthetic set. Bottom panel: the ratio of the normalised co-added horizontal pressure gradient $R_{norm}(t)$ and the $\pm 1\sigma$ range (lighter grey lines).



where the weights now satisfy:

$$w_j \equiv \frac{\sum_{k=0}^{365} |\nabla_h P|_{co-add synth j}(t_k)}{\sum_{k=0}^{365} |\nabla_h P|_{co-add real}(t_k)} \quad (16)$$

Note that the full extent of the scale in fig. 3.6 for the ratio is now 0.08 rather than 0.004 : the variations in this ratio are larger by a factor of roughly ~ 20 . As is indicated by the grey $\pm 1\sigma$ lines in the bottom panel of fig. 3.6 with time this ratio is clearly rising and this result is statistically significant. The average over the 25 realisations of the normalisation factors w_j is 1.5 ± 0.3 . This means that the slight overall deviation from 1 of the average of this ratio is not statistically significant at a 90% confidence level: the variation between realisations is large. However, this variation between realisations is mostly a straightforward shift up or down of the entire pattern. For this reason it can be determined that this ratio is not constant and this result is statistically significant.

4 Comparison of precursors for $M \geq 1.2$ after Sep. 1 2014

4.1 The pressure

While the localisation of earthquakes with $M \geq 1$ over that entire period was good, in particular after reanalysis of the original seismograms dating from the start-up phase, it has been argued that the localisation accuracy is such that it is preferable to select earthquakes with magnitudes $M \geq 1.2$. For this reason in the literature this higher cut-off is often still used. It is therefore worthwhile to restrict the dataset to these higher magnitudes, although of course that does reduce the numbers of earthquakes in the set from 108 to 79, which will have an adverse effect on the signal-to-noise ratio of the various pressure ratios shown in the present analysis.

Figure 4.1 The co-added pressure for the real earthquakes (in blue) as well as the average for the 25 realisations of co-added pressure for the synthetic sets (red). Only quakes with $M \geq 1.2$





Fig. 4.1 shows that there is still an offset between the co-add pressure for the real earthquakes and for the average of the 25 realisations of synthetic earthquakes. The ratio of real to synthetic co-added pressure is again calculated in the two different ways described in the previous section. Either as a straightforward average over the ratios for each realisation or as a weighted average, in such a way that any constant offset between the real and synthetic co-added pressure is removed. The result is shown in fig. 4.2. Comparing the top and bottom panels of fig. 4.2 shows that the normalisation has little or no effect on the behaviour as a function of time. By selecting only earthquakes with $M \geq 1.2$, there is as expected a slight loss in terms of signal-to-noise ratio, but there is nevertheless statistically significant evidence that the ratio between real and synthetic co-added pressure is not constant as a function of time. In the bottom panel the grey lines indicate the 1σ range for the 25 realisations, from which it is clear that the deviations from unity are statistically significant. The average over the 25 realisations of the normalisation factors w_j is 1.001 with a 1σ standard deviation of 0.007. The difference from 1 of the offset, visible as a vertical shift when comparing top and bottom panels of Fig. 3.2, is therefore not statistically significant.

Figure 4.2 Top panel: The ratio $R(t)$ for the co-added pressure, averaged for the 25 realisations of a synthetic set, $M \geq 1.2$. Bottom panel: the ratio of the normalised co-added pressures $R_{norm}(t)$ and the $\pm 1\sigma$ range (lighter grey lines).



4.2 The averaged pressure

The local average of the grid cell nearest to a (real or synthetic) earthquake location together with its 8 nearest neighbour grid cells is subjected to the same co-adding technique as discussed in section 3.2. As one would expect, the additional local averaging over 9 near grid cells reduces systematic offsets between the co-added pressures for the real earthquakes and the synthetic earthquakes. Fig. 4.3 shows that the pressure precursors for real earthquakes and for the average of all 25 realisations of synthetic earthquakes are barely shifted with respect to each other.

Figure 4.3 The co-added locally averaged pressure for the real earthquakes (in blue) as well as the average for the 25 realisations of co-added locally averaged pressure for the synthetic sets (red)





Figure 4.4 Top panel: The ratio $R(t)$ for the co-added locally averaged pressure, averaged for the 25 realisations of a synthetic set. Bottom panel: the ratio of the normalised co-added pressures $R_{norm}(t)$ and the $\pm 1\sigma$ range (lighter grey lines)



The ratio is shown in Fig. 4.4, where for the top panel the ratio is calculated without any weighting, and for the bottom panel with a weighting as shown in Eqs. (13). As is seen also in section 3.2 the shape of this curve, the temporal pattern, is very similar to the pattern obtained in the previous section without the local averaging, although the deviations from unity are a little smaller. This demonstrates that the result is not particularly sensitive to the precise method of picking the nearest or most appropriate grid cell for each earthquake location. The uncertainty margins are also slightly smaller, so that the deviations from unity are just statistically significant. The average over the 25 realisations of the normalisation factors w_j is 1.003 ± 0.007 .


This result implies that for any given realisation of a synthetic set, the entire pattern can shift up or down by a relatively large amount, but the shape overall has some common behaviour. As argued before it is primarily the shape, i.e. the fact that the ratio is not a horizontal straight line, that points to gas pressure variations modulating earthquake generation.

4.3 The horizontal pressure gradient

As described in section 2.2 the absolute value of the local horizontal pressure gradient $|\nabla_h P|$ is determined and co-added for the real and synthetic sets (cf. Eq. (9)). In this case the average values of about 110 N/m^3 and about 85 N/m^3 are sufficiently different that two separate panels are used to show their behaviour as a function of time. The full extent of the scale in both panels is the same: 4 N/m^3 .

The top panel of fig. 4.5 shows the behaviour of the co-added horizontal pressure gradient for the set of real earthquakes, whereas the bottom panel shows the behaviour of the co-added

Figure 4.5 The co-added horizontal pressure gradient for the real earthquakes. The average for the 25 realisations of co-added horizontal pressure gradient for the synthetic sets (red)



horizontal pressure gradient for the set of synthetic earthquakes. It is immediately evident that the behaviour with time of the co-added time series is quite different for the earthquake locations and for synthetic locations. The development over time in the rest of the field is mostly flat, whereas at the earthquake locations the gradient is higher and also steeply rising.

In fig. 4.6 the ratio and normalised ratio of these horizontal pressure gradients are shown, defined analogously to the ratios for the pressure itself:

$$R_{nabla P}(t) = \frac{1}{25} \sum_{j=1}^{25} \frac{|\nabla_h P|_{co-add real}(t)}{|\nabla_h P|_{co-add synth j}(t)}$$


$$R_{norm nabla P}(t) = \frac{1}{25} \sum_{j=1}^{25} w_j \frac{|\nabla_h P|_{co-add real}(t)}{|\nabla_h P|_{co-add synth j}(t)} \quad (17)$$

where the weights now satisfy:

$$w_j \equiv \frac{\sum_{k=0}^{365} |\nabla_h P|_{co-add synth j}(t_k)}{\sum_{k=0}^{365} |\nabla_h P|_{co-add real}(t_k)} \quad (18)$$

Note that the full extent of the scale in fig. 4.6 for the ratio is now 0.08 rather than 0.004, just as in section 3.3: the variations in this ratio are larger by a factor of roughly ~ 20 just as is the case in section 3.3. As is indicated by the grey $\pm 1\sigma$ lines in the bottom panel of fig. 4.6 with time this ratio is clearly rising and this result is statistically significant. The average over the 25 realisations of the normalisation factors w_j is 1.4 ± 0.5 . This means that the overall deviation from 1 of the average of this ratio is not statistically significant at a 90% confidence level: the variation between realisations is large. However, this variation between realisations is mostly a

Figure 4.6 Top panel: The ratio $R(t)$ for the co-added horizontal pressure gradient, averaged for the 25 realisations of a synthetic set. Bottom panel: the ratio of the normalised co-added horizontal pressure gradient $R_{norm}(t)$ and the $\pm 1\sigma$ range (lighter grey lines).



straightforward shift up or down of the entire pattern. For this reason it can be determined that this ratio is not constant and this result is statistically significant.

5 Conclusions

There is a strong indication, already from (Pijpers, 2016a; Sijacic et al., 2017), that the likelihood for earthquakes to occur reacts causally to gas production in the sense that variations in gas pressure in the reservoir, induced by variations in production rates, leads to an increased earthquake likelihood. Previous research (Pijpers, 2016a; Sijacic et al., 2017) already indicated that pressure variations need to arrive at specific quake generating sites such as a faults or fractures. The improved time resolution of the reservoir gas pressure data available allows for an improved view of the precise behaviour of the reservoir gas pressure in the period leading up to a quake.

In the reports (Pijpers, 2014b), (Pijpers, 2015a), and (Pijpers, 2015b) it is demonstrated that a strong reduction in gas production at a number of clusters in the beginning of January 2014 was followed during that year by a reduction in frequency of quakes in the vicinity of those clusters, which is consistent with the time scale reported now. In work of (Pijpers, 2014a), (Pijpers and van der Laan, 2015a), (Pijpers and van der Laan, 2015b) a rapid response of the subsidence to

production variations with a delay of about 9 weeks has been seen. Thus it would appear that production variations once propagated to the location of faults or fractures are followed by subsidence and the generation of quakes.

It is already pointed out in Pijpers (2016a) that proof of causality requires the ability to exclude or identify the influence of all other possible factors that could affect the likelihood of quakes. In the case at hand an, as yet unidentified, external process that might explain the timing of quakes instead of the gas pressure would need to introduce a modulation that would have to be phased just right for the gas production variations propagated to quake sites to coincidentally line up in the period leading up to quakes, as well as phasing these to increase just before quakes. The present analysis does not give any reason to reject a hypothesis of (partial) causality through reservoir pressure variations induced by the gas production.

This behaviour is consistent with eg. the discussion in (Bourne et al., 2014), where it is the interaction of reservoir compaction, because of pressure depletion, with pre-existing faults and fractures that generates quakes. The locations of those pre-existing faults or fractures is then the relevant spatial characteristic of the quakes. The timing is important because it plays a role in 'lining up' in time gas pressure variations at the fracture locations. Within the context of critical systems, cf. (Main, 1995), one may hypothesize that the pressure depletion over long time scales has put the system of faults into a (nearly) critical state. The additional short term effect of production variations in this critical systems context might be that they act to push the system temporarily but repeatedly to a critical state and act as a trigger for individual earthquakes or quake avalanches.

Recently a simulation study was conducted (Postma and Jansen, 2018) which looked in particular at the effect of pressure transients. That study concludes that there would be a transient effect of production increases, which will become larger as the overall reservoir pressure decreases, but in their simulations the magnitude of such stresses remain very small compared to ambient stresses at least where it concerns zero-throw faults. In faults with non-zero throws the effects do become larger but Postma and Jansen (2018) point out that at such faults stresses will increase over time under any production strategy, so that slower changes in production might delay but probably not prevent earthquakes. The implication is that slower changes in production rate could help in reducing the earthquake rate but would not bring that rate down to 0.

The magnitude of the effect on pressure gradients of a few percent over a year, shown in the present report, are certainly larger than the magnitude of transients that the simulation study produces for zero-throw faults. Postma and Jansen (2018) point out that their study focusses on short time scale effects (within days). Investigation of seasonal fluctuations, with multiple wells and a realistic mixture of fault properties would require further investigation, and also requires incorporating additional physical effects.

As is pointed out in (Pijpers, 2018) and further confirmed by the sensitivity analysis presented in the present report, earthquake locations appear to have a quite different precursor gas pressure, and in particular a quite different precursor in terms of a horizontal gradient in this pressure. Although the overall production rate has become more stable as a function of time, the reservoir gas pressure variations have not ceased. This is due to the variations in production rates at individual cluster locations. Such variations propagate into the reservoir and therefore there will continue to be variations in gas pressure and in its horizontal gradient throughout the reservoir which could act as earthquake triggers. The previous analyses (Pijpers, 2015b, 2016a; Sijacic et al., 2017) suggested that such variations might act as a trigger on a critically stressed system, so that sensitive locations (fractures) have an enhanced likelihood of producing an earthquake if

they are subjected to the variations, even if these might appear to be relatively minor in terms of relative change.

In a recent ministerial decision the intention has been expressed to cease gas extraction in Groningen altogether, as quickly as feasible within operational boundaries and (inter-)national commitments. To obtain a further reduction in quake incidence more quickly, this research implies that it might be beneficial to operate with less variation in production rate at each individual cluster while also implementing the reduction overall in production.

References

- Baranova, V., A. Mustaqeem, and S. Bell (1999). A model for induced seismicity caused by hydrocarbon production in the western Canada sedimentary basin. *Can. J. Earth Sci.* 36, 47--64.
- Bourne, S., J. Oates, J. van Elk, and D. Doornhof (2014). A seismological model for earthquakes induced by fluid extraction from a subsurface reservoir. *J. Geophys. Res.* 119, 8991--9015.
- Christou, J. C., E. K. Hege, J. D. Freeman, and E. Ribak (1986). Self-calibrating shift-and-add technique for speckle imaging. *Optical Society of America, Journal A* 3, 204--209.
- Grasso, J. and G. Wittlinger (1990). 10 years of seismic monitoring over a gas field area. *Bull. Seismol. Soc. Am.* 80, 450--473.
- Harpsøe, K., U. Jørgensen, M. Andersen, and F. Grundahl (2012). High frame rate imaging based photometry Photometric reduction of data from electron-multiplying charge coupled devices (EMCCDs). *Astronomy and Astrophysics* 542, A23 1--9.
- Lynds, C., W. S.P., and H. J.W. (1976). Digital image reconstruction applied to alpha orionis. *Astrophysical Journal* 207, 174--180.
- Main, I. (1995). Earthquakes as critical phenomena: Implications for probabilistic seismic hazard analysis. *Bull. Seismological Soc. Am.* 85, 1299--1308.
- Mulders, F. (2003). Modelling of stress development and fault slip in and around a producing gas reservoir. Technical report, Delft University Press, Delft University, Netherlands.
- Nederlandse Aardolie Maatschappij BV (2013). A technical addendum to the winningsplan Groningen 2013 subsidence, induced earthquakes and seismic hazard analysis in the Groningen field. Technical report.
- Pijpers, F. (2014a). Phase 0 report 1 : significance of trend changes in ground subsidence in Groningen. Technical report, Statistics Netherlands.
- Pijpers, F. (2014b). Phase 0 report 2 : significance of trend changes in tremor rates in Groningen. Technical report, Statistics Netherlands.
- Pijpers, F. (2015a). Phase 1 update may 2015 : significance of trend changes in tremor rates in Groningen. Technical report, Statistics Netherlands.
- Pijpers, F. (2015b). A phenomenological relationship between gas production variations and tremor rates in Groningen. Technical report, Statistics Netherlands.

- Pijpers, F. (2015c). Trend changes in tremor rates in Groningen : update november 2015. Technical report, Statistics Netherlands.
- Pijpers, F. (2016a). A phenomenological relationship between reservoir pressure and tremor rates in Groningen. Technical report, Statistics Netherlands.
- Pijpers, F. (2016b). Trend changes in tremor rates in Groningen : update may 2016. Technical report, Statistics Netherlands.
- Pijpers, F. (2016c). Trend changes in tremor rates in Groningen : update nov 2016. Technical report, Statistics Netherlands.
- Pijpers, F. (2018). Improved time resolution relationship between pressure and earthquake rates in Groningen. Technical report, Statistics Netherlands.
- Pijpers, F. and D. van der Laan (2015a). Phase 1 update May 2015: trend changes in ground subsidence in Groningen. Technical report, Statistics Netherlands.
- Pijpers, F. and D. van der Laan (2015b). Trend changes in ground subsidence in Groningen update November 2015. Technical report, Statistics Netherlands.
- Pijpers, F. and V. van Straalen (2017a). Trend changes in tremor rates in Groningen (Jun 2017). Technical report, Statistics Netherlands.
- Pijpers, F. and V. van Straalen (2017b). Trend changes in tremor rates in Groningen (October 2017). Technical report, Statistics Netherlands.
- Pijpers, F. and V. van Straalen (2018). Trend changes in tremor rates in Groningen (May 2018). Technical report, Statistics Netherlands.
- Por, G., P. Boerrigter, J. Maas, and A. de Vries (1989). A fractured reservoir simulator capable of modeling block-block interaction. Technical report, SPE Annual Technical Conference and Exhibition, 8-11 October, San Antonio, Texas.
- Postma, T. and J. Jansen (2018). The small effect of poroelastic pressure transients on triggering of production-induced earthquakes in the Groningen natural gas field. *Journal of Geophysical Research-Solid Earth* 123(1), 401--417.
- Russell, T. and M. Wheeler (1983). *The Mathematics of Reservoir Simulation*. SIAM Frontiers in Applied Mathematics.
- Sijacic, D., F. Pijpers, M. Nepveu, and K. van Thienen-Visser (2017). Statistical evidence on the effect of production changes on induced seismicity. *Netherlands Journal of Geosciences* 96, 27--38.
- Wetmiller, R., . (1986). 'earthquakes near rocky mountain house, alberta, and their relationship to gas production. *Can. J. Earth Sci.* 23, 172--181.

Colophon

Publisher

Statistics Netherlands
Henri Faasdreef 312, 2492 JP The Hague
www.cbs.nl

Prepress

Statistics Netherlands, Grafimedia

Design

Edenspiekermann

Information

Telephone +31 88 570 70 70, fax +31 70 337 59 94
Via contact form: www.cbs.nl/information

© Statistics Netherlands, The Hague/Heerlen/Bonaire 2018.
Reproduction is permitted, provided Statistics Netherlands is quoted as the source