Predicting gas production from future gas discoveries in the Netherlands: quantity, location, timing, quality

J. LUTGERT, H. MIJNLIEFF and J. BREUNESE

Netherlands Institute of Applied Geoscience TNO – National Geological Survey, PO Box 80015, 3508 TA Utrecht, the Netherlands (e-mail: j.lutgert@nitg.tno.nl)

Abstract: Recent policy and market developments have raised the question not only as to how much gas remains to be discovered in the Netherlands but also where and when it will be produced and of what quality. These questions are addressed by compiling a 'firm futures' database, estimating the 'potential futures in proven plays' and simulating the exploration process using the prospect portfolio.

The firm futures are assessed using an integrated play evaluation method in a combined database/GIS environment. The database includes a prospect portfolio of firm futures, well data and maps. Volumetric estimates (low, mean and high success volumes) are assigned to the prospects. The risk of each prospect is estimated using play maps and the historical success ratio for that play. Using standard GIS functionality, the distance to production nodes, the gas quality, licence status and operator are determined for each prospect. The prospect portfolio is not static but dynamic. Using the creaming curve, fractal distribution analysis and the average POS (Possibility of Success) of the play, the gas volume and the amount of structures in the potential futures category are estimated. Again, GIS functionality is used to determine possible new locations and the gas quality of the prospects. The prospect portfolio is the input for a computer program that stochastically simulates the exploration process. Assumptions about future drilling efforts, economic screening scenarios and time delays from discovery to production are made explicit. The program simulates annual drilling campaigns. A number of prospects are drilled according to an assumed exploration efficiency and the POS is used to determine the result: a discovery or a dry well. The discovered volume is economically assessed and, if economic, transformed into a production and gas-quality profile.

Keywords: prospects, future reserves, the Netherlands, North Sea, exploration simulation, GIS

A reliable estimate of remaining gas reserves is essential information for government policy-making. Insight into the reserves situation helps in defining policy strategies for future exploration. Gas purchasing and distribution companies are also interested in knowing the remaining gas reserves, but on a more technical and commercial level. Information on the expected future gas volume and quality delivered to certain entry points or pipelines supports decision-making in investments in the mid-stream system.

When it comes to reserves and production forecasting, the overall question involves much more than just 'How much?' It also includes 'When?' 'Where?' and 'What quality?'. Questions such as 'How much gas remains to be discovered in the Dutch offshore?' and 'What level and quality of gas production can we expect from the area served by a particular pipeline system?' can only be answered if reliable estimates for the remaining exploration potential are available.

Since 1992, the Ministry of Economic Affairs (MEA) has published assessments of the reserve position of the Netherlands in its annual report *Oil and Gas in the Netherlands* (MEA 2002). Figure 1 shows the published figures for the produced reserves, reserves and the future reserves. Reserves are gas volumes in proven accumulations, declared economically recoverable. The future reserves category comprises only the firm futures class (Fig. 2). The definition of firm futures is 'prospective structures, which have been identified on the basis of existing data within geological plays in the Netherlands where suitable conditions for gas accumulations exist and which have been sufficiently established by drilling'. It is remarkable that from 1992 onwards the total volume of firm futures has stayed approximately the same. This, in spite of the fact that some 220 prospects have been drilled and some $400 \times 10^9\,Sm^3$ has been discovered. Thus, the firm futures are not a static but rather a dynamic prospect portfolio. This actually follows from the definition of future reserves. The firm futures are a subset of the total futures (Fig. 2). If successful, drilled prospects migrate from the firm futures to the reserves class. Newly mapped prospects migrate from the 'potential futures in proven plays' class to the firm futures. The future class 'potential futures in proven plays' is defined as the risked gas volume of the as yet unidentified structures within proven plays. The futures classes 'potential futures in non-proven plays' and 'potential futures in hypothetical plays' have not yet been assessed.

This paper addresses the questions: How much? Where? When? and What quality? The steps taken are:

(1) compiling a prospect database in a Geographic Information System (GIS);
(2) extracting a firm futures portfolio from the prospect database;
(3) evaluating volume and risk factors of prospects in the firm futures portfolio;
(4) assigning essential prospect attributes using GIS to the prospects in the portfolio;
(5) estimating the volume of potential futures within proven plays;
(6) estimating key parameters for modelling the exploration process;
(7) running the prospect portfolio through an exploration simulation program, which results in volume and gas-quality profiles.

This chain of assessment steps is the subject of this paper.

LUTGERT, J., MIJNLIEFF, H. & BREUNESE, J. 2005. Predicting gas production from future gas discoveries in the Netherlands: quantity, location, timing, quality. *In*: DORÉ, A. G. & VINING, B. A. (eds) *Petroleum Geology: North-West Europe and Global Perspectives—Proceedings of the 6th Petroleum Geology Conference*, 77–84.

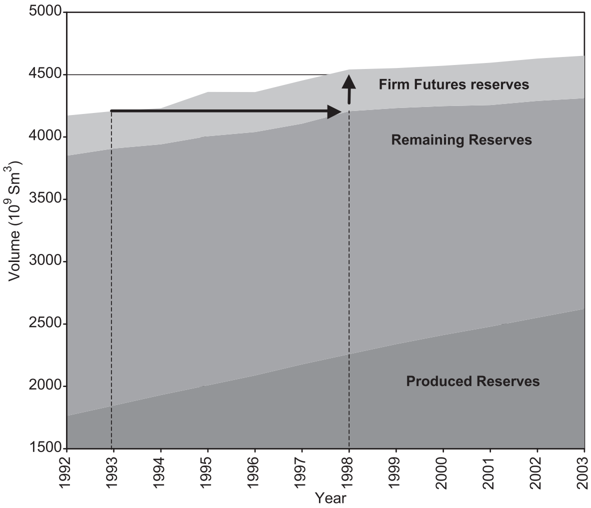

Fig. 1. Reserves and future reserves from the yearly reports of the Ministry of Economic Affairs. The Firm Futures of 1993 (some 300×10^9 Sm^3) had already been found in 1998. This means that the portfolio is a dynamic portfolio. Prospects are drilled, new prospects are added.

The prospect database

The prospect database is a GIS-based set of files that includes a prospects map and a set of database files holding administrative, volumetric and other attribute data.

The polygons of all mapped prospects are collected and merged into one prospect map in a GIS system, such that the location and surface areas of all mapped prospects are stored. Administrative data of the prospects, such as reservoir, date and source of data and the type of map from which the prospects were digitized, are entered into database files. A number of other data items, if provided, are added to the prospect file. These include the volumetric parameters of the prospects (P90, P50 and P10 stored as the Low, Mean and High Success Volumes, LSV, MSV and HSV, respectively), Expectation (EXP) and Possibility of Success (POS), key parameters for calculating unrisked in-place volumes such as net:gross, porosity as well as estimated risk factors and structural annotations such as the type and depth of the structure. All these data items are regarded as ‘fixed prospect attributes’. The other types of prospect attributes, namely the ‘mapped attributes’ such as Gross Heating Value (GHV), are interpolated from maps using GIS functionality.

At present, the database comprises some 1200 prospective structures. These structures originate from various sources, have different vintages and often have incomplete ‘fixed prospect attributes’. Having the prospects in a GIS environment allows a variety of quality control analyses, for example combining the prospect map with regional geological maps to check or add prospect attributes and completing the missing fixed attributes that are essential for further analyses.

This prospect database forms the working set, from which the prospect portfolio of firm futures is to be extracted.

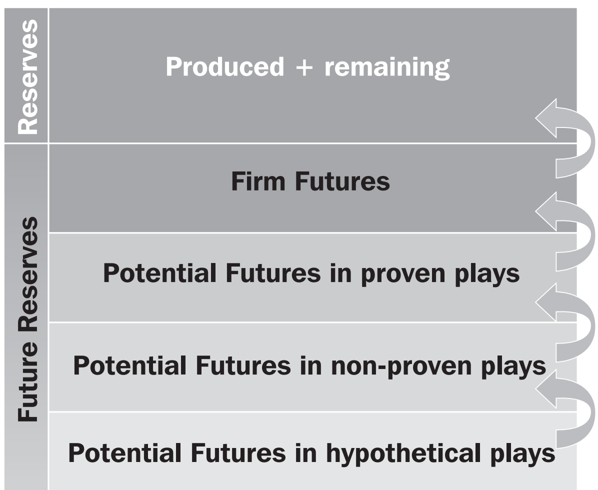

Fig. 2. Reserves classification as used in the study.

Mapping in GIS

Important play elements are reservoir, seal, source-rock distribution and source-rock maturity. These elements have been compiled and mapped mainly using well data, seismic data and various other geological studies. Within a GIS environment, a selection of play element maps can be combined into a new play map (Fig. 3). In this map, the common cross-section of the play elements describes the outline of a certain play. Another set of thematic maps is compiled to assign ‘mapped attribute’ data to the prospects. These maps are gridded using information at well locations, including parameters such as the GHV (Fig. 4), the percentage of CO_2 and N_2 and reservoir parameters such as gross thickness, net:gross ratio and porosity.

Play maps and thematic maps are easily updated when new data become available or when new concepts are developed. This procedure is much easier than storing and maintaining prospect attributes in a flat file system.

Firm futures portfolio

The prospect portfolio presently used in this study includes only the firm futures. The prospect database contains prospects both inside and outside the defined proven play areas. A selection of prospects is therefore made using standard GIS functionality. The prospect map is crossed with the play maps of the various proven plays. The only prospects selected are those lying within a play area and which have a reservoir horizon identical to that of the play reservoir with which the selection was made. The futures class is not, therefore, a fixed prospect attribute but is the result of selection using the prospects map and play maps. This reduces prospect administration when a play outline is changed as a result of new information or concepts. At present, some 1000 prospective structures from the prospect database have been identified within the proven play areas.

To address the question ‘How much?’ each prospect must be assigned an estimate of the associated gas volume and a risk factor (1-POS). For prospects lacking the required volumetric attributes, an estimate of the recoverable gas volume is derived from either a method using the surface area and reservoir property maps to estimate the LSV, MSV and HSV or a method based on the statistical analysis of fields and prospects from the same play.

The Possibility of Success (POS) is difficult to assess. Estimates vary between operators because different methods are used and because most procedures involving a notion of ‘risk’ are subjective to a degree. Another complicating factor arises from the dependency of the POS of a prospect on the result of an exploration well, which, in effect, equates to testing a prospect from a ‘prospect family’. A prospect family is a number of prospects sharing similar reservoir characteristics and risk factors. When the analysis of an exploration well shows that one of these factors is below expectation, the POS of all the related prospects will be reduced. On the other hand, when all factors prove to be positive and the exploration well is a success, the POS of related prospects should be increased. Presently the POS is estimated per play or sub-play area using the historical commercial success ratio as calibration.

To address the question ‘What quality?’, all prospects should carry an estimate of the expected gas composition. In the prospect database, the gas quality is not a fixed attribute. To assign gas

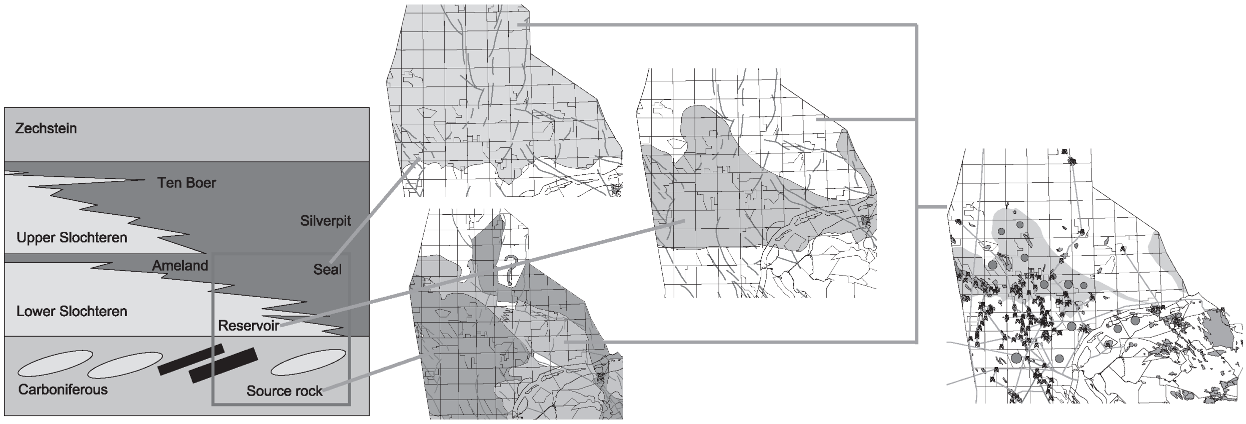

Fig. 3. A selection of play element maps used to define a play map.

properties to prospects, gas-quality maps are used. Gas-quality maps are generated using gas analyses of all the available gas samples taken in wells, not only in commercial discovery wells but also in wells with only gas shows. By gridding and contouring these observations, a variety of maps, such as GHV, CO_2 and N_2 distribution maps, are created. The interpolation of a gas-quality map at the position of a prospect is performed to obtain the desired prospect attributes. This procedure is standard GIS functionality. When new or improved data become available, for example from petroleum system analysis, the attribute maps and the dependent prospect attributes are effortlessly recalculated, resulting in up-to-date gas-quality predictions.

To address the question 'When?', all prospects should have attributes that relate to the timing of drilling a prospect. An exploration licence must be obtained before a prospect may be drilled. Later, once a discovery has been made, a production licence is required before hydrocarbon production can start. The calculation of these attributes requires that for the prospect location both the current operator of the licence and the current licence status are known, as well as the distance to the nearest production node. Using standard GIS functionality, these attributes can easily be assigned to each prospect by combining the firm futures prospect map with maps containing licence information, fields, platforms and pipelines.

Within the proven play area, the location of each prospect belonging to the firm futures class is known. However, whether a prospect at a certain location will be a success is unknown. When the question 'Where?' is addressed, the location and the POS must be taken into account. To gain meaningful predictions based on valid statistics, the evaluation should be performed on groups of prospects rather than individual prospects. To that end, the Netherlands onshore and offshore areas have been divided into 29 'source areas' that closely mirror the shape and extent of the existing infrastructure. The map of the source areas (Fig. 5) is interpolated to assign the appropriate source area attributes to the prospects.

Once all the data-processing procedures described above are complete, the prospect portfolio with all its calculated attributes is exported to a file that serves as input for the exploration-simulation program.

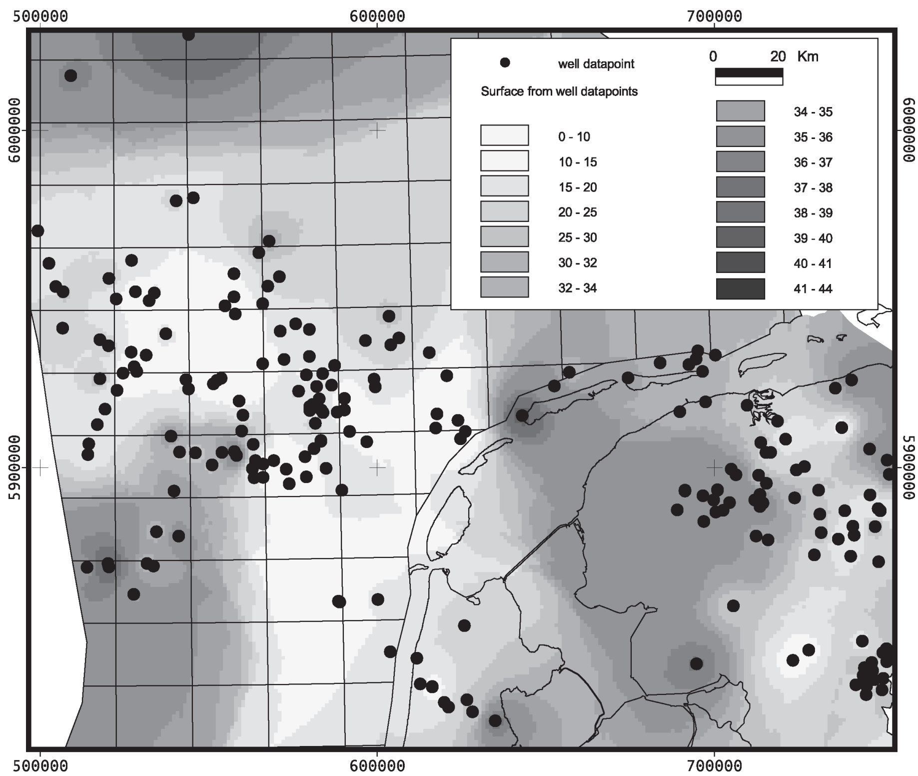

Fig. 4. Example of a thematic map: the Gross Heating Values (GHV) of gas samples are gridded to create a GHV contour map.

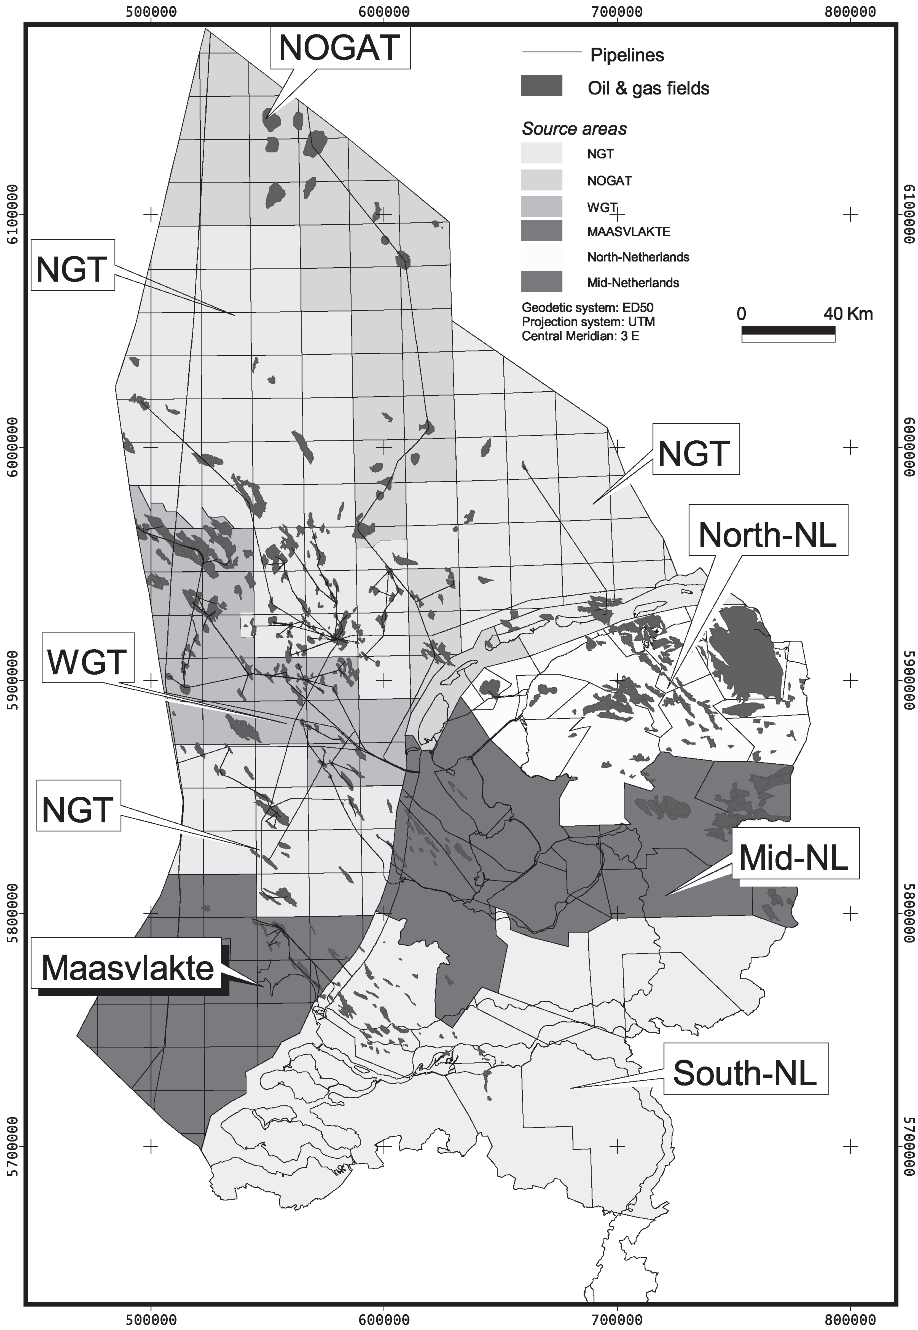

Fig. 5. Source area distribution map.

Potential futures in proven plays portfolio

Potential futures in proven plays include the yet unmapped prospects within the proven play area. But how many prospects are there in this futures class and what is their (risked) volume?

When evaluating the creaming curves (Laherre 2003) of the main plays, Rotliegend and Triassic, it appears that there is still potential in these plays (Fig. 6). This follows from the observation that the curves do not yet show any obvious signs of creaming. Estimating the remaining potential using the creaming method gives a wide spread of possible outcomes. Depending on the position of the 'regression line', fractal analysis (Barton 1997) of the same dataset results in a volume for the potential futures in this proven play of $200–600 \times 10^9\,Sm^3$ (Fig. 7). As well as giving an estimate of the remaining exploration potential of a play, the fractal distribution also gives an indication of the size class to which the future potential is expected to belong. Using this information, together with the average POS of the play, a potential futures portfolio is made that lists the number of prospects per size class (Fig. 8). The location of these future prospects is still unknown but, by definition, they lie within the proven play area. It is also known that the locations within the play outline of the existing fields and the locations that were found to be dry or uneconomic do not qualify as target locations for future prospects. If, in the GIS system, these already-occupied areas are blanked out on the play-outline map, what remains is a map indicating all possible locations of the potential futures. An estimate of the average surface area of the future prospects is found in the average surface area of all fields. Subsequently, the area of possible locations in the play-outline map is populated with grid blocks the size of the average surface area for future prospects. This results in possible locations of future prospects. To all these locations, the same attributes as those of the firm futures prospects can be assigned by combining maps and subsequently propagating the information from these maps, the gas-quality maps, field maps, production node maps and pipeline maps, amongst others, to the

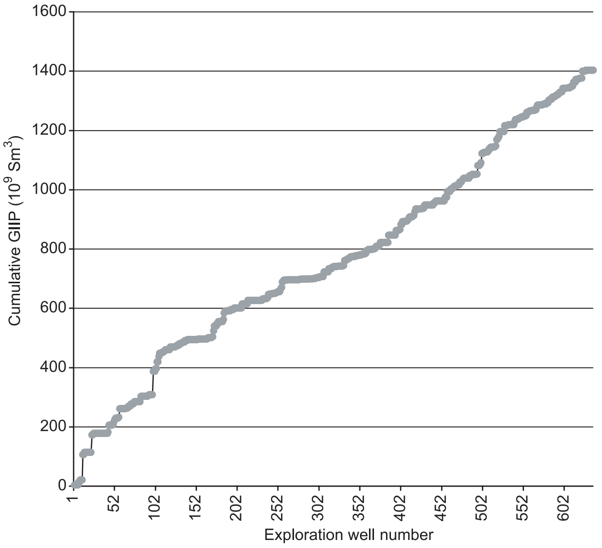

Fig. 6. Creaming curve of the Rotliegend play.

possible locations of future prospects. The final product of this method is a file containing the possible location of prospects with the necessary attributes. The possible volumetrics are listed in a separate file.

Combining data from these two files using random sampling will result in a realistic estimation of the volume, location and quality of the 'potential futures in proven plays' portfolio, up to the total assumed play potential.

The Exploration Simulator

The Exploration Simulator (EXPLOSIM) is a computer program that models future exploration using the firm futures portfolio. The output of this program includes future-production profiles and gas-quality profiles. The exploration process is not just a question of drilling away prospects in a certain order but is a process that is influenced by numerous factors. Within EXPLOSIM, the program's input parameters are factors of prime importance to modelling the exploration process. These parameters are the drilling

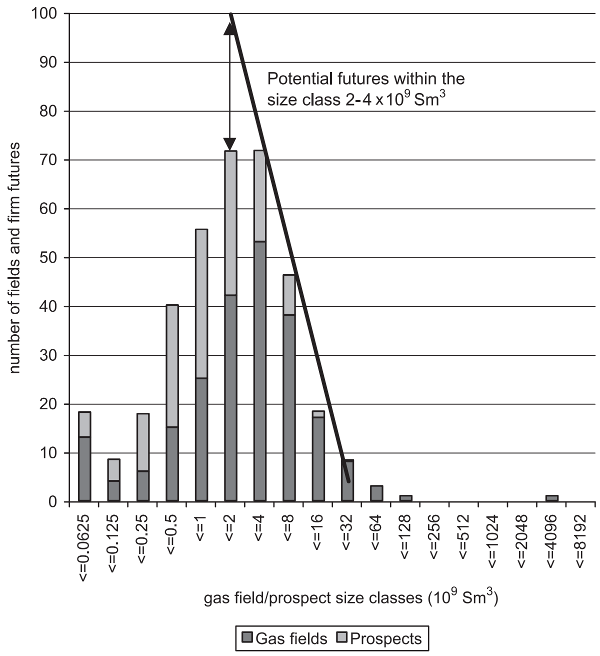

Fig. 7. Fractal analysis of fields and firm futures of the Rotliegend play.

Upper Boundary size class (10^9 Sm^3)	predicted total volume	number successes to be explored *	number of structures left to be explored **	expectation volume of potential futures ***
<=0.0625	213.75	196	723	6.11
<=0.125	191	183	696	11.41
<=0.25	168.25	150	536	28.21
<=0.5	145.5	105	383	39.55
<=1	122.75	67	249	50.42
<=2	100	28	104	42.63
<=4	77.25	6	20	16.62
<=8	54.5	8	29	49.74
<=16	31.75	13	51	161.28
<=32	9	1	2	15.84

Fig. 8. The number of prospects per size class deduced from the fractal analysis. * total predicted for the class − number of fields and prospects; ** number of successes divided by the POS is the number of structures to be explored; *** number of successes to be explored multiplied by the mid class is the approximation of the expectation of the potential futures in proven plays.

efficiency and economic screening parameters, assumed drilling efforts and time delays related to changes in the licence status.

Prospect ranking and drilling sequence

The prospect portfolio of firm futures as described above is the input for the program. At present, this portfolio contains some 1000 prospects, including very small (low risk/low reward) prospects that will not pass any economic screening and very large (high risk/high reward) prospects. The future drilling sequence is compiled by ordering the prospects in the portfolio in such a way that the most attractive prospect is at the top of the list and the least attractive at the bottom. The attractiveness of prospects is influenced by numerous factors. In practice, all factors that influence project economics – for example, the geological risk, the expected volume of gas (MSV), the gas price, the distance to existing infrastructure, the current licence status, tax regimes, contract and strategic considerations – play a role in determining the attractiveness of a prospect and, thus, the position in the drilling sequence. In EXPLOSIM, ranking is done using a parameter that specifies the relative degree of 'attractiveness' of drilling a prospect. Examples include the expected gas volume (MSV) and the Expected Monetary Value of the project (EMV).

Drilling efficiency

All the fields presently known were once prospects. If all these fields had been drilled in descending order of actual reserves (i.e. the largest field first and the smallest field last), maximum efficiency would have been achieved (max efficiency line in Fig. 9). However, in practice, the exploration process is not that efficient. Conversely, if an algorithm were to draw samples at random from the prospect database, the outcome would mimic random exploration (random exploration line in Fig. 9). Thanks to the time and effort invested in exploration studies, historical drilling efficiency exceeds that of random exploration. To derive a measure of historical drilling efficiency, the ratio of the surface area between the actual recorded efficiency and the random curve to the surface area between the random curve and the maximum efficiency curve is used. Thus calculated, the drilling efficiency for historical exploration in the Netherlands is approximately 30%.

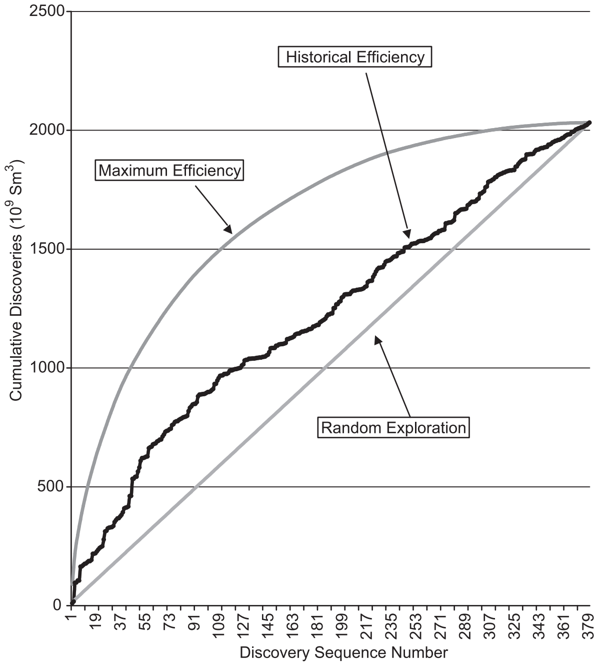

Fig. 9. Analysis of the historical drilling efficiency.

The drilling efficiency is implemented in the program by limiting the prospect sample population from which the simulator is allowed to draw its targets to a subset containing only the *n* most attractive prospects. By varying the size of this subset (sampling window on the prospect database), the drilling efficiency of the simulation can be calibrated against the known historical drilling efficiency.

The impact of changing the length of the efficiency-sampling window of the resulting production profiles is shown in Fig. 10. In this graph, the cumulative volumes for each of the curves are constant. The black dashed curve, representing random exploration, was calculated by allowing the program to sample the entire prospect population at random. The black solid curve, representing maximum exploration efficiency, was calculated by always picking the top-ranking prospect in the database. It is clear that by decreasing the size of the sampling window, the production profile is accelerated.

The simulation algorithm

The prospect volumes LSV, MSV and HSV are recalculated into a triangular distribution of volume versus POS (Fig. 11). Subsequently, this distribution is transformed to a Probability Density Function (PDF – Fig. 12).

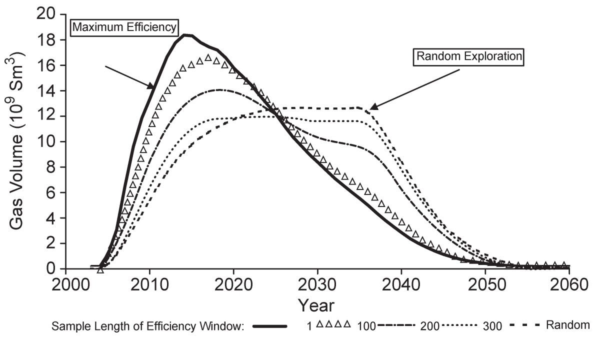

Fig. 10. The impact of the efficiency sampling window on production profiles (ranking on EMV assuming 15 exploration wells per year).

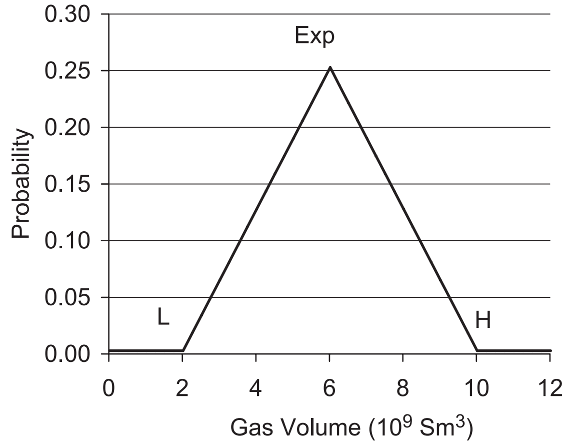

Fig. 11. Example of a triangular volume distribution of a prospect.

The program uses Monte Carlo simulation to mimic the exploration process, sampling the prospect database as it is executed. For every prospect, a random number x ($0 \leq x \leq 1$) is generated. If the value of x falls in the range $0 < x \leq$ POS, the simulated drilling of the prospect results in a success. To obtain the associated reserves, the value of x is used to interpolate the PDF (Fig. 12). If the value of x falls outside the success range ($x >$ POS), the result of the simulated exploration well is a dry hole and the prospect is labelled 'dry'.

Economic screening criteria

Next, if the draw is a success, the resulting volume is verified against a cut-off criterion for economic development. If the volume does not meet the economic constraint, the draw is labelled 'sub-economic'; otherwise, it is labelled 'economic'.

In the program, two economic screening methods are implemented. The first is a simple cut-off on success volume. An example for the Dutch situation as used by TNO-NITG is a volume cut-off at $2 \times 10^9\,Sm^3$ for an offshore prospect and at $0.5 \times 10^9\,Sm^3$ for an onshore prospect. The second method is more sophisticated. This method involves a pre-defined economic scenario and screening criterion, i.e. screening oil price at $US20 BBL^{-1} and EMV > 0. The economic scenario and criterion can be represented by one hyperbolic curve in a POS–MSV plot (Fig. 13). During execution of the program, the POS value and the associated gas volume determined by Monte Carlo simulation are verified against this hyperbolic curve. If the combination of POS and MSV plots above the curve, the project is economic, otherwise it is sub-economic.

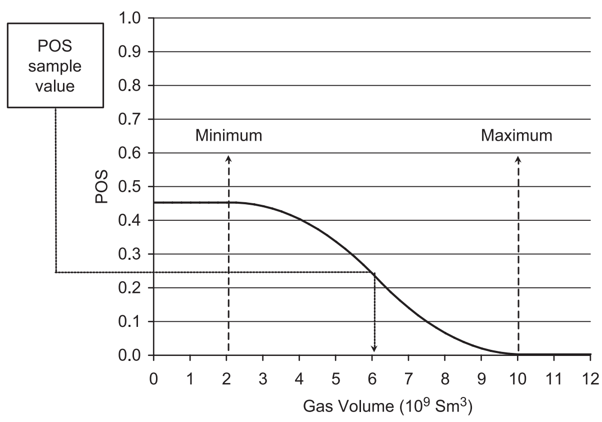

Fig. 12. Principle of volume distribution sampling from a PDF.

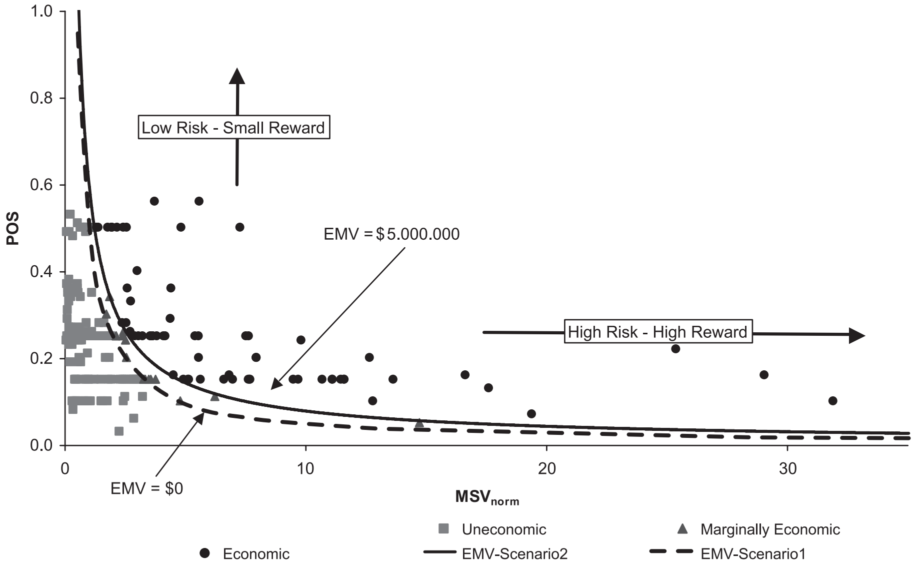

Fig. 13. Hyperbolic curves representing EMV scenarios applied to the economic screening of prospects (POS–MSV_{norm} plot for $US20 BBL^{-1}, unit of production and tax efficient company).

Drilling effort-related time delay

The rate at which new fields will be discovered depends on the drilling effort. In EXPLOSIM, the expected drilling activity level must be defined. The future drilling activity may be estimated using a combination of extrapolated historical drilling activity and the known plans of the various operators. The timing of the draw, i.e. simulated exploration well (see section entitled 'The simulation algorithm'), can now be calculated using the sequence number of the draw and drilling effort per year.

An additional option is that the drilling activity per operator can be set. The ownership of a prospect is one of the prospect attributes, namely the operator of the licence in whose area the prospect lies. The timing of a discovery can now be determined by recording the number of prospects owned by an operator and drawn during simulation, and dividing that number by the drilling effort of the associated operator. The resulting number is the delay-time in years related to that operator's drilling effort.

Licence status-related time delay

A prospect can only be drilled if it lies within the scope of an exploration or a production licence. The acquisition of acreage, the evaluation of the petroleum geology and the identification of drillable prospects within the acreage all take a lot of time. An analysis of historical data has shown that the time delay – from getting a prospect in an open area through exploration to the production licence stage – may take as long as ten years. For prospects that are already in an area under exploration or covered by a production licence, the average time required is five and two years, respectively.

Thus, the overall time delay to production for a prospect is a combination of two factors, the 'drilling effort', which is dynamically determined during simulation, and the 'licence status', which is a user-defined constant.

Production profile

Once the discovered commercial volume of gas and the time delay from discovery to production have been established, the simulated discovered reserves are converted into a production profile. Each profile is defined in terms of a start-up period, a plateau phase, a decline phase and abandonment. The profiling procedure can easily be adjusted to comply with various depletion schemes.

The production profiles of all the economic prospects resulting from one simulation run are placed at their appropriate places along the time axis. They are then aggregated to form the cumulative production profile of that run. Each curve in Fig. 14 represents such a cumulative production profile and is the result of one simulation run. Next, the results of all the aggregated Monte Carlo runs themselves are averaged to give the expected future production profile together with an uncertainty range expressed as one standard deviation. Next, volumetric averaging of the GHV associated with the economic discoveries is used to generate a

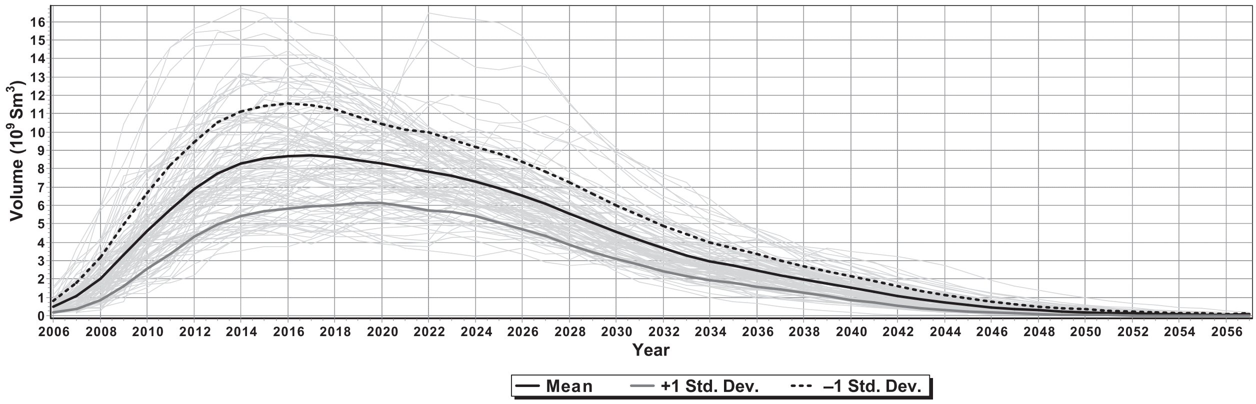

Fig. 14. Simulation results for gas volumes for a certain source area (100 runs).

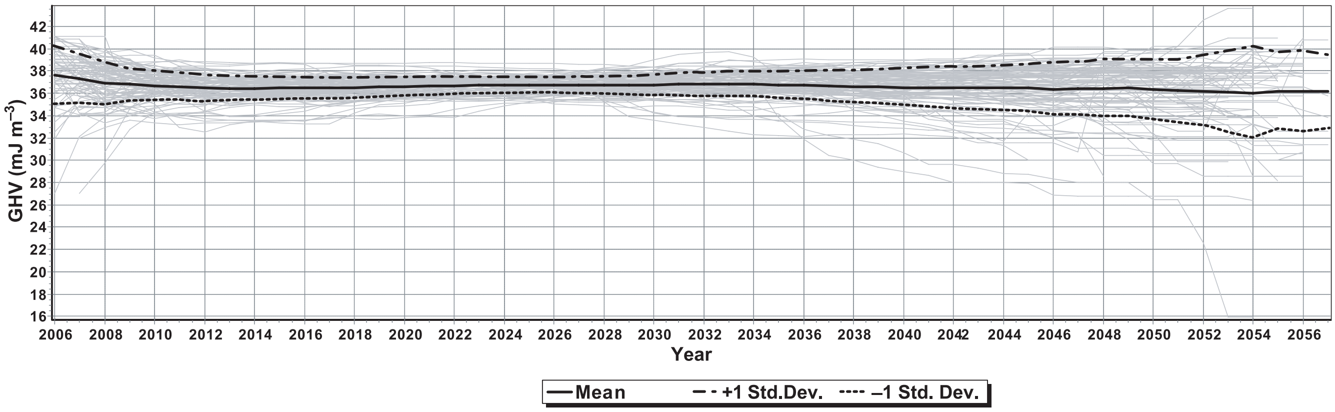

Fig. 15. Simulation results for the gas gross heating value for a certain source area (100 runs).

gas-quality profile. One line in Fig. 15 shows the volumetrically averaged GHV profile associated with one of the volume profiles in Figure 14. Profiles of percentage CO_2 and percentage N_2 are generated in a similar way. The resulting profiles can be displayed for the whole prospect portfolio, for the prospects in a certain source area or for combinations of source areas.

Simulation results

Figures 14 and 15 depict the modelling results of 100 simulation results for a certain source area in the Netherlands. The profiles show only the contribution of simulated economic prospects in that source area, although the full portfolio of firm prospects was used as input for the simulation. The full portfolio is used to address properly the ranking and the timing issue.

The simulation results may be used, among other things, to study the identification of potential growth areas for infrastructure, in pipeline ullage studies and to predict trends in gas quality in the evacuation system.

Future developments and conclusion

The simulation program will always be a simplification of 'real-life' exploration but various modifications can be made to improve its functionality. At present, a fixed number of prospects, the firm futures portfolio, is used. As described in the section 'Potential Futures in proven plays portfolio', the firm futures portfolio is replenished by re-evaluations of known structures resulting in updated volume and risk factors, while replenishment of the firm futures portfolio is realized by the addition of prospects from the potential futures in proven plays. This dynamic updating procedure of the prospect portfolio has not yet been taken into account but will be implemented in future versions of the program.

Other improvements may be achieved by incorporating dependencies on pipeline ullage, by observing the lifetime of the infrastructure and through the incorporation of more advanced production-related economic estimators. Future developments aside, even in its present state, this simple program clearly illustrates the extra information and added value that can be generated by taking a different approach to analyse exploration data. However, the methodology described in this paper critically depends on two conditions: all the data pertaining to exploration must be properly archived and continuously maintained; and a GIS system must be available that is capable of interfacing with electronic databases and of being used as a spatial analysis tool.

References

Barton, C. C. 1997. *A New Approach for Estimating Hydrocarbon Resources*. US Geological Survey Open-File Report 97-xxx. World Wide Web Address: http://www.coastal.er.usgs.gov/barton/pubs/hydrocarbon.pdf.

Laherre, J. 2003. Forecast of oil and gas supply to 2050. *In*: *Petrotech 2003: Hydrocarbon Resources*; New Delhi. World Wide Web Address: http://www.oilcrisis.com/laherrere/Petrotech090103.pdf.

Ministry of Economic Affairs 2002. *Oil and Gas in the Netherlands, Exploration and Production*. World Wide Web Address: http://www.nitg.tno.nl/oil&gas/downloads/jb2002eng.pdf.