Fluvial — Improved Seismic Conditioning

Arne Skorstad Ragnar Hauge Øivind Skare Lars Holden

February 18, 1996

Contents

1 Introduction

This version of Fluvial (formerly known as Fettuccini and SISA among others) has extended the use of seismic data which was included in the program in the previous version (from October 1995). The seismic data is used to condition the spatial distribution of the channels by inserting extra condition points based on seismic information only, in addition to the condition points given from well observations.

A new interaction function has also been implemented, encouraging channels to cross instead of being parallel.

For stacked channels there exist no information whether the channels have been eroded or not, and this is now considered in this version, giving more realistic realizations.

Improved output from the program has also been a subject, which resulted in a redesign of the log files given from the execution.

In Section 2 it is given a short review of the current simulation algorithm, concentrating on the parts where there has been a change from the previous version of Fluvial. Documentation of this can be found in "STORM Technical Manual", February 1996 from "Geomatic", Pb. 172, N-4033 Forus. Next, the new model file is explained in Section 3. A program for calculating the correlation between the seismic and facies from a realization in a sub volume of the reservoir has been written, and is documented in Section 4. In the following sections it is considered one major change at the time, trying to give the effect of the change on both the simulated realizations and consumed CPU-time. Section 5 results from the most extensively testing, due to the more complex nature of the changes in the algorithm when considering the seismic input. The documented border effects has been tested on the new Fluvial version and documented in Section 6. A new interaction function, allowing channels to cross was implemented, and the effect of this is shown in Section 7. The last test-section is Section 8 where the effect of conditioning on drawn uneroded thicknesses instead of the observed eroded thicknesses is documented. Finally the new log files are explained in Section 9.

2 Model and algorithm

The only change done to the model in this project is that the interaction function between families has been changed. It is still a pairwise interaction function, and its role in the model is as before. Only the function and its parameters are new. This improvement is made to allow channels to cross each other in the same plane, thus avoiding the tendency to parallelization which was seen with the old function.

The algorithm has been changed on five points, in addition to the implementation of the new interaction function. Two of them aim at improving the seismic convergence: The introduction of seismic sections, which is the main change, and an improved way to draw the direction of lines given seismic. Two others are corrections of conflicts between model and algorithm: The interpretation of well observations as possibly eroded, and the correction for line length when using seismic. The final change, the possibility of using uniform distributions when drawing the condition point for the family line is a way to provide greater flexibility when modelling.

2.1 New interaction function

The new interaction function assigns a probability to each pairwise line configuration. It is a function of the minimum distance d observed between family lines within the simulation box, and the angle α between them, and is given by

$$p(d,\alpha) = 1 - (1 - p(d))(1 - p(\alpha)) \tag{1}$$

where $p(d)$ is given by

$$p(d) = \begin{cases} 0 & d < d_{min} \\ \frac{d-d_{min}}{d_{max}-d_{min}} & d_{min} \leq d \leq d_{max} \\ 1 & d > d_{max} \end{cases} \tag{2}$$

and $p(\alpha)$ by

$$p(\alpha) = \begin{cases} 0 & \alpha < \alpha_{min} \\ \frac{\alpha-\alpha_{min}}{\alpha_{max}-\alpha_{min}} & \alpha_{min} \leq \alpha \leq \alpha_{max} \\ 1 & \alpha > \alpha_{max} \end{cases} \tag{3}$$

where d_{min} and d_{max} are minimum and maximum interaction distances specified by the user, as is the interaction angles α_{min} and α_{max}. The function is plotted in Figure 1.

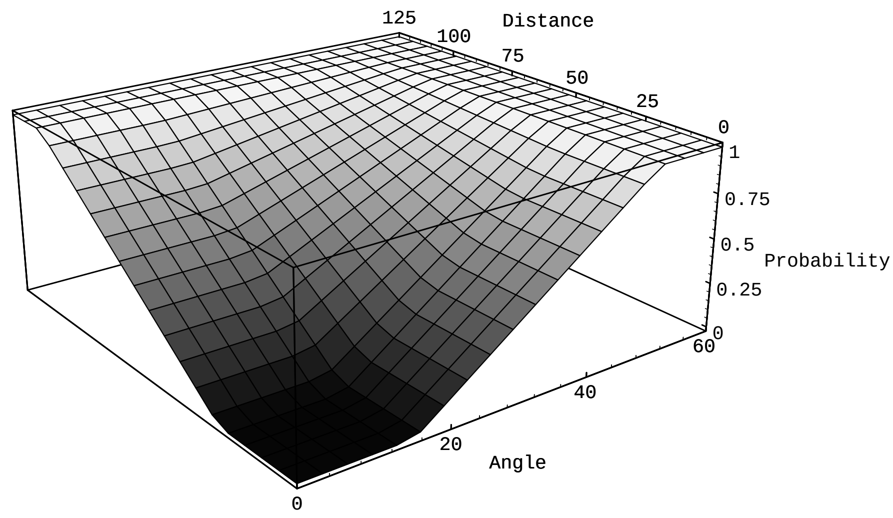

Figure 1: Interaction function where $\alpha_{min} = 15°$, $\alpha_{max} = 50°$, $d_{min} = 30$ and $d_{max} = 110$.

All parameters have default values, so there is a default interaction. This is given by

$$d_{min} = 0.25\sqrt{T^2 + W^2} \quad (4)$$
$$d_{max} = 8 \cdot d_{min} \quad (5)$$
$$\alpha_{min} = 0.25 \cdot \theta \quad (6)$$
$$\alpha_{max} = \max(\alpha_{min}, \theta) \quad (7)$$

where T is the expected thickness of a channel, W is the expected width, and θ is the pooled standard deviation of the horizontal and vertical angle of the family lines. The choice of α_{max} is made to ensure that $\alpha_{max} \geq \alpha_{min}$ also in the case when the user specifies only a minimum interaction angle.

The interaction function enters the model and simulation algorithm just as the old function did, so no major changes are done here. This new function allows more freedom in how much the user wants to restrict interaction between lines, since both angles and distances for the interaction zone can be specified.

2.2 Seismic sections

The seismic sections are planes orthogonal to the family line of a channel belt, and spaced n grid nodes apart along this line, providing there are no well sections there. If there are well sections close to seismic sections, these seismic sections are removed. Each channel belt has its own set of seismic sections.

In each seismic section, 10 different sets of values for the parameters horizontal and vertical displacement, width and thickness are drawn. The probability for each of these given seismic is computed by

$$p(x_i) = const \cdot \left(\prod_{nodes} \frac{p(f_{new}|Seismic)}{p(f_{old}|Seismic)} \right)^{\frac{a}{N}} \tag{8}$$

where x_i is the ith set of parameters, f_{new} and f_{old} are the facies type in the node before and after a channel with these parameters are placed, a is the seismic factor as specified by the user, and N is the total number of grid nodes. The product is taken over all affected nodes in the seismic section.

One set of parameters is then drawn from this distribution, and added to the observation list for this channel, as if it was a well observation. Thus, the channel has a conditioning point in each seismic section.

Note that the function given in equation (8) is the same as is used when computing the seismic potential. This ensures that the seismic conditioning points gives convergence to the same distribution as the seismic potential does.

Also note that f_{new} in equation (8) is not updated with the conditioning points of other channels in the same section. Thus, if there are more than one channel in a channel belt, their parameters in a seismic section will be drawn independently of each other. This may result in too many channels concentrating in the same area if the seismic there is favorable.

2.3 Direction of lines given seismic

Conditioning the direction of a family line on seismic observations can now be done on several grid nodes in depth, that is, grid nodes above and beneath the one the line intersects. This is reasonable, since the channels will spread out around the line. The default value is a number of grid nodes corresponding to half the expected thickness of a channel, but the user can also specify this number.

2.4 Eroded observations

When a channel is observed directly below another channel in a well, as shown in Figure 2, the possibility that the bottom channel may be eroded by the one above is considered.

When conditioning on a channel observation, it is first checked if it is possible that this observation is eroded. If this is the case, the expected thickness of the channel is drawn from the prior distribution truncated by the observed

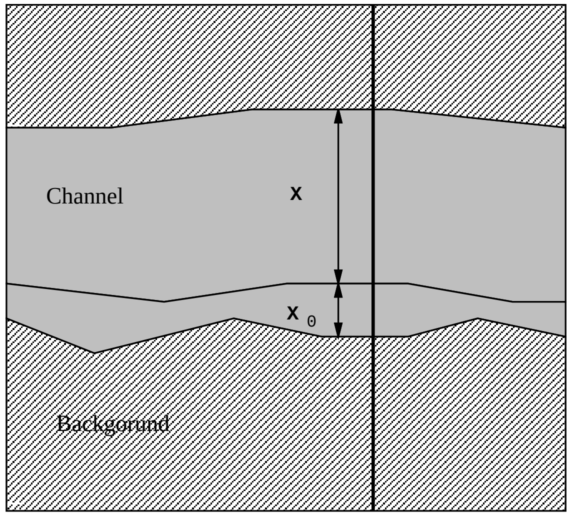

Figure 2: The bottom channel in this figure may be eroded, so the thickness in the well section is somewhere between X_0 and $X + X_0$

channel thickness as the minimum value, and the cumulative thickness of this channel and all the channels directly above it as the maximum. The thickness of the channel in the observation point is then drawn from the distribution with this expected thickness and the specified standard deviation of thickness, with the same truncation.

2.5 Correction for line length

The procedure for drawing lines when using seismic first selects a plane parallel either to the xz or yz plane, dependent on the main direction of lines. Then a point in this plane is drawn, conditioned on seismic and the prior distribution of lines. The family line is then drawn through this point.

This procedure has a greater probability of drawing lines where much of the line is inside the simulation box, since the probability of drawing a line with this setting is proportional to the observed length of the line. This factor must be included in the term describing the probability q_{ij} of drawing j as the next state given the current state i.

This is done by scaling q_{ij} with the length of the line divided by the expected length of lines when a family is added from state i to j. With this factor in place, the Metropolis-Hastings algorithm ensures that the realizations follows the distribution specified in the model.

2.6 Uniform distribution of families

The algorithm has been modified to allow uniform distribution of family lines around the reference point. This is done by drawing these parameters from a uniform distribution if this option is specified. A similar result could be obtained by specifying large standard deviations for this point in older versions of the program, but that resulted in a slower convergence since a major part of the lines were drawn outside the simulation box.

The uniform distribution of family lines ensures, in an unconditional simulation, an uniform distribution of sand/gross over the reservoir. This new option removes therefore edge effects.

3 The model file

Here, an example of a model file which includes the parameters controlling the latest developments is listed.

```
MODEL  <model-name(string)> example ! Name of the model
<seed(string_or_int)> -11324456 ! Seed for random number
                                !   generator ;
;

VOLUME
0.0   ! x-origo
500.0 ! dx
0.0   ! y-origo
500.0 ! dy
0.0   ! angle
20.0  ! z-top, constant or file name
40.0  ! z-bot, constant or file name
0.0   ! top erosion, constant or file name
0.0   ! bot erosion, constant or file name
20    ! dz
;

FACIESTABLE
0 SHALE ;
1 CHANNEL-SAND ;
2 CREVASSE-SAND ;
3 BARRIER ;
;

FACIES
<channel-facies(string)> CHANNEL-SAND
<crevasse-facies(string)> CREVASSE-SAND
<barrier-facies(string)> BARRIER
<background-facies(string)> SHALE
;

FAMILY
<x-ref(real)> 250 ! Expected point of family interactions
<y-ref(real)> 250 ! with the simulation box
<z-ref(real)>  10 ! In local coordinates
<E(horizontal_angle)(real)> 20 ! Exp. hor. direction for family line
<E(vertical_angle)(real)>    0 ! Exp. ver. direction for family line
<sd(a)(real)> *  ! ***** NEW ***** 1
                 ! St. dev of y-value of family-line interaction.
                 | A '*' gives uniform distribution
```

```
<sd(horizontal_angle)(real)> 5
<sd(c)(real)> *  ! ***** NEW ***** 1
                 ! St. dev of z-value of family-line interaction,
                 ! A '*' gives uniform distribution
<sd(vertical_angle)(real)> 0
<etopnchn(real)> 1 ! Expected number of channels in family
                   !   at the top of the reservoir
<ebotnchn(real)> 1 ! Expected number of channels in family
                   !   at the top of the reservoir
<delta_nchn(real)> 0.1  ! The width in the uniform distribution
                        !   for the number of channels
[uniform-normal(string)] uniform ! Must be uniform or normal
[sd(delta-hor)(real)] 10 ! St.dev. of hor. displacement in normal dist.
[sd(delta-ver)(real)] 2  ! St.dev. of ver. displacement in normal dist.
[E(area-thickness)(real)] 0
[E(width-bottom)(real)] 0
[E(width-top)(real)] 0
[sd(area-thickness)(real)] 0
[sd(width-bottom)(real)] 0
[sd(width-top)(real)] 0
;

FAMEXPECTATION    !Low frequent horizontal deviation curve
<hor-sd(real)> 30.0      !Standard deviation of the amplitude
<hor-range(real)> 300.0  !Range of hor. dev. curve
;

CHANNEL
<E(mean-bot(hw))(real)> 20 !Expected mean width of a channel
                           !at the bottom of the reservoir
<E(mean-top(hw))(real)> 20 !Expected mean width of a channel
                           !at the top of the reservoir
<sd(mean(hw))(real)>     0 !Standard dev. of mean width
<E(mean-bot(vt))(real)>  5 !Expected mean thickness of a channel
<E(mean-top(vt))(real)>  5 !Expected mean thickness of a channel
<sd(mean(vt))(real)>     0 !Standard dev. of mean thickness
<E(sd(hd))(real)>       20 !Expected st.dev. for the horizontal dev.
<sd(sd(hd))(real)>       0
<E(sd(vd))(real)>        1 !Expected st.dev. for the vertical dev.
<sd(sd(vd))(real)>       0
<Corr(mean(vt),mean(hw))(real)> 0.0
<relative-var(hw)(real)> 0.2 !The st.dev of the width (pointwise) is
                             !relative-var * mean(width)
<relative-var(vt)(real)> 0.2
<Range(hd)(real)> 100 !Range of horizontal deviation from line
<Range(vd)(real)> 100 !Range of vertical deviation from line
<Range(hw)(real)> 100 !Range of width
<Range(vt)(real)> 100 !Range of thickness
```

```
[Min(hw)(real)] 0.5   !Minimum width = min(hw) * mean(width)
[Max(hw)(real)] 2.0   !Maximum width = max(hw) * mean(width)
[Min(vt)(real)] 0.5   !Minimum thickness = min(vt) * mean(thickness)
[Max(vt)(real)] 2.0   !Maximum thickness = max(vt) * mean(thickness)
[Sampling-of-channel-width(string)] yes
;

CREVASSE
<E(width)(real)> 0.5  !Expected mean crevasse width =
                      !   'E(width)' * mean(channel width)
<E(vert)(real)> 0.0   !Mean crevasse vertical position relative
                      !to the channel position, must be in [-1,1]
                      !0.0 in center of channel, -1 bottom, 1 top
<E(thick)(real)> 0.2  !Expected mean crevasse thickness =
                      !   'E(thick)' * mean(channel thickness)
<sd(width)(real)> 1.0
<sd(vert)(real)>  0.5
<sd(thick)(real)> 0.05
<range(horizontal)(real)> 200
<min-ncrev(int)> 0    !Minimum number of crevasses in channel
<max-ncrev(int)> 1    !Maximum number of crevasses in channel
;

SIMULATION
<sand-gross(real)> 0.1  ! Target channel proportion
<sg-epsilon(real)> 0.01 ! Tolerance limit for channel proportion
<nIterations(int)> 1000 ! (minimum) Number of iterations in algorithm
<sampnx(int)> 20 ! Gridnodes in x-direction
<sampny(int)> 20 ! Gridnodes in y-direction
<sampnz(int)> 20 ! Gridnodes in z-direction
[nMaxIterations(int)] 1000 ! Maximum number of iterations
[z-anisotropy(real)] 25.0  ! Scaling factor when computing the
                           !   vertical distance component:
                           !   dist^2 = hor-dist^2 + (zanis*ver-dist)^2
                           !   Used in the repulsion/interaction function
[min_interaction(real)] 20 ! Min. interaction distance for lines
[max_interaction(real)] 60 ! Max. interaction distance for lines
[min_angle(real)] 15 ! ***** NEW ***** 2
                     ! Min. interaction angle between lines
[max_angle(real)] 45 ! ***** NEW ***** 2
                     !Max. interaction angle between lines
[rapidSimulation(string)] no  ! no = use Metropolis-Hastings
[min-temp(real)] 0.00001      ! Temperature achieved after nIterations
[variogram-factor(real)] 1.5  !
[simulation-area(string)] *   !
[start-temp(real)] 0.2        ! Temperature in first iteration
;
```

```
SEISMIC
[seismic-grid(string)] ../data/test.seis   ! Seismic value grid
[g-channel(string)] ../data/test.gchn      ! Prob(sand|seismic)
![g-background(string)] is removed in present version;
!Prob(shale or crevasse|seismic) = 1 - Prob(sand|seismic)
[seismic-factor(real)] 200.0               ! Factor compensating for
                                           !   correlation in seismic values
[weight-grid(string)] ../data/test.weigth  ! Seismic weight grid
[grid-node-seperation(int)] 10             ! ***** NEW ***** 3
                                           ! Grid nodes between
                                           !   seismic sections
[number-of-seismic-layers(int)] 1          ! ***** NEW ***** 4
                                           ! How many grid nodes in depth
                                           !   are considered when choosing
                                           !   direction of family line.
;
```

The '***** NEW *****' comments indicates the parameters that controls the new features. The numbers refer to the following:

1. Uniform distribution of conditioning point. This is documented in sections 2.6 and 6.
2. New interaction function. Docmumented in sections 2.1 and 7.
3. Seismic sections. Documented in sections 2.2 and 5.
4. Seismic conditioning for line directions. Documented in section 2.3.

Note that it will still be possible to use old model files using the parameters in 'TEMP' instead of 'SIMULATION'. Many of the parameters are the same in 'SIMULATION', some are neccessary as before, and some have been optional in the this new version.

4 Correlation between seismic and facies in given windows

The program fluvial_correlation calculates the correlation between the seismic and facies from a realization in a given part of the reservoir, called a window. All input information for the program is in a model file, and the model file name is given as a command line parameter when starting the program.

The commands in the model file must be specified in the given sequence, and all commands must be specified. Otherwise, the syntax rules are the same as for all STORM model files. The commands are:

- FACIES Takes one parameter, the file to read the facies realization from. This file must be an output file from the facies_grid program.
- SEISMIC One parameter, the file to read the seismic data from. This file can be either a contsim grid or a 2D irap-file.
- RESULT One parameter, the file to write the results to. The result file will have eight columns: The first column contains the correlation, the next 6 columns the window, and the last column contains the number of points used to calculate this correlation.
- WINDOWS This command takes records of 6 parameters, each series specifying one window. There can be as many series as the user wants. The records must be separated by a ';'. The records contain x0, dx, y0, dy, z0, dz for a window. Note that the coordinates are given relative to the xmin, ymin, ztop coordinates given in the VOLUME section of the model file that generated this realization.

Example of model file:

```
FACIES example.gri ;           ! Facies realization
SEISMIC ../data/example.seis ; ! Seismic data
RESULT example.corr ;          ! Output file
WINDOWS
0       ! x0
500     ! dx
0       ! y0
500     ! dy
0       ! z0
20      ! dz
;       ! End of first record
0 250 0 250 0 20 ; ! Next record
250 250 0 250 0 20 ;
```

```
0 250 250 250 0 20 ;
250 250 250 250 0 20 ;
;
```

5 Testing the seismic option

The seismic option will affect the simulation run in almost every way, from running time and convergence speed to posterior distribution of parameters. Only the final sand/gross ratio and the correlation with well observations should be unaffected, since these are included in the annealing.

Thus, it is a vast field to test the general influence of the seismic. Here, it is focused on four main points. The first is the influence of informationless seismic, which ideally should be none. The second is how the use of seismic affects the running time of a simulation, the third is the influence of the prior distribution for sand given seismic and the last is how the new seismic intersections performs.

5.1 Influence from informationless seismic on simulation runs

In order to test if the seismic option affects the simulation more than its data content gives reason for, an informationless seismic file was made. This file consists of the same seismic observation everywhere in the reservoir, which in the current model should give no influence. Simulation runs were then done with two model files which were identical except that one of them used this informationless seismic. The results are shown in tables 1 to 6.

Seismic:	Number of families									
With	28	26	26	27	24	24	27	22	24	26
Without	26	30	25	29	28	26	23	26	27	29

Table 1: Number of families for 10 simulations with and without informationless seismic.

As can be seen from the tables, the mean width and thickness, the direction of lines and the number of families remains almost unperturbed. The only thing that really changes is the number of accepted iterations (and thus the number of adds, changes and removes). Note that these tables show only one dataset, whereas a number of different runs were done with different

Seismic:	Number of accepts (1976 iterations)									
With	620	637	638	564	661	571	636	656	613	598
Without	858	906	823	852	884	804	900	906	887	845

Table 2: Number of accepts for 10 simulations with and without informationless seismic.

Seismic:	Mean width of channels				
With	696.88	715.06	710.42	691.39	711.57
Without	702.91	694.87	724.50	731.09	705.83
With	708.33	707.11	709.00	726.57	699.28
Without	728.27	721.04	714.17	698.14	708.18

Table 3: Mean width of channels for 10 simulations with and without informationless seismic (expected width 700).

Seismic:	Mean thickness of channels				
With	10.08	10.21	10.21	10.01	9.73
Without	10.05	10.14	9.87	10.29	10.12
With	10.05	9.83	10.50	9.77	9.91
Without	10.10	10.17	9.97	9.96	10.15

Table 4: Mean thickness of channels for 10 simulations with and without informationless seismic (expected thickness 10).

Seismic:	Mean hor. angle for lines				
With	3.8924	-1.4151	-2.8131	3.7547	-3.0655
Without	-1.0699	0.3210	-2.4967	-2.5991	5.5670
With	-6.3122	-0.1733	-1.3828	9.0471	0.7860
Without	3.8499	-6.1686	-2.4029	0.6880	3.4633

Table 5: Mean hor. angle of channel lines for 10 simulations with and without informationless seismic (expected angle 0).

Seismic:	Mean ver. angle for lines				
With	0.0002	0.0003	0.0001	-0.0001	-0.0004
Without	0.0000	0.0003	-0.0001	0.0005	0.0004
With	-0.0001	0.0001	-0.0003	0.0000	0.0000
Without	-0.0001	-0.0002	0.0002	-0.0002	0.0002

Table 6: Mean ver. angle of channel lines for 10 simulations with and without informationless seismic (expected angle 0).

model files. However, the observations were very much the same, and so these tables are typical.

The number of accepts decreases when the informationless seismic is used. This is because a different procedure for drawing channels is used when there is seismic. This procedure does not draw from the prior distribution, and compensates for this by scaling the accept probability. Although the number of accepts decreases, the difference between the number of adds, changes and removes remains much the same, which ensures a similar realization.

The sand/gross convergence does not seem to be influenced by the decrease in number of accepts, as Figure 3 shows. However, one can expect that the realizations with more accepts are better in other aspects, such as channel distribution, since they have been through more states and have had greater possibility to optimize other factors than just the sand/gross.

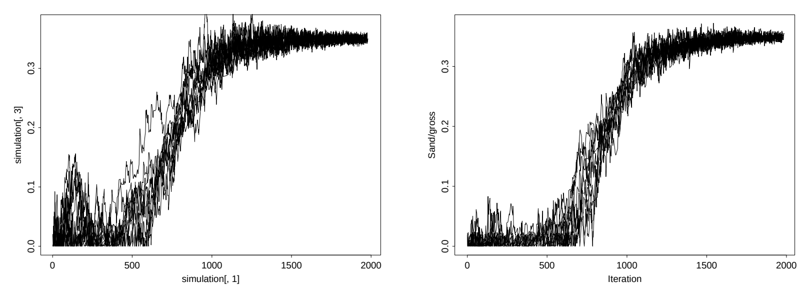

Figure 3: Sand/gross against iteration number for 10 runs without seismic (left) and with flat seismic (right).

5.2 Running time

To test how use of seismic affects the running time, different versions of a standard model file was made. They were all identical except for the seismic command, where the following versions were made: No seismic command at all, informationless seismic and seismic with information with and without seismic sections. The model file containing seismic with information was also run through the previous release of the program.

The results are shown in Table 7. As can be seen from the table, it makes no big difference whether the seismic carries information or not; the routines use the same amount of time. However, there is a big difference between using seismic and not using it, and using seismic sections increases the time spent further.

The amount of extra time it takes to use seismic varies with the number

No seismic	Flat seismic	Normal seismic	Intersections	Old version
0:26.2	2:14.3	2:13.9	3:02.0	5:41.3
0:26.0	2:17.3	2:12.1	3:06.1	5:41.4
0:25.7	2:14.2	2:22.4	3:01.1	5:42.4
0:25.9	2:08.9	2:17.8	3:07.4	6:37.7
0:25.0	2:16.9	2:23.2	2:57.6	5:42.1
0:25.5	2:13.8	2:50.4	2:59.8	5:43.8
0:26.7	2:14.2	2:16.0	3:03.9	5:43.2
0:25.8	2:17.8	2:20.6	3:15.0	5:50.7
0:25.6	2:09.0	2:19.9	2:59.0	5:41.9
0:26.9	2:14.4	2:22.8	3:01.6	5:42.3

Table 7: CPU time for ten runs in the different cases (time in minutes).

of grid nodes and the general complexity of the model file. This example was run on a 40 by 40 by 100 grid, with a simple model file (no wells, one channel per family, no crevasses). The number of iterations was low, making the overhead time spent in reading and transforming the seismic file high. When using seismic sections, these had a grid spacing of 8, which means that there was about 20 seismic sections per channel in this example.

With a more complex model, especially including wells, the time spent on seismic calculations will be relatively less. Increasing the fineness of the grid will increase this time proportion and vice versa, and so will the density of seismic sections.

It is also interesting to note that even when seismic sections were used, the program was faster than the previous version. This is due to an optimalization in the routine calculating the probability of seismic given facies, which is essential in all seismic calculations.

5.3 Influence of prior distribution

The user specified probability for sand given seismic combined with the distribution of seismic values gives an expected value for the sand/gross ratio. This value should be equal to the target sand/gross; a warning will be given if this is not so. A mismatch here does not seem to influence too much on the final realization in other aspects than the difference between prior and posterior probability for sand given seismic, which becomes larger when the two sand/gross values are unequal. This is natural, since the sand/gross potential is included in the annealing term, and thus overrules the seismic potential. However, a mismatch here influences on the convergence and thus on the number of iterations needed to get a realistic realization.

Figure 4 shows what can happen. When the expected sand/gross from seismic is too large, as in the leftmost picture, the sand/gross will overshoot during the simulation, until the annealing forces it onto target. However, when seismic and target sand/gross corresponds, the convergence is good, as seen in the rightmost picture.

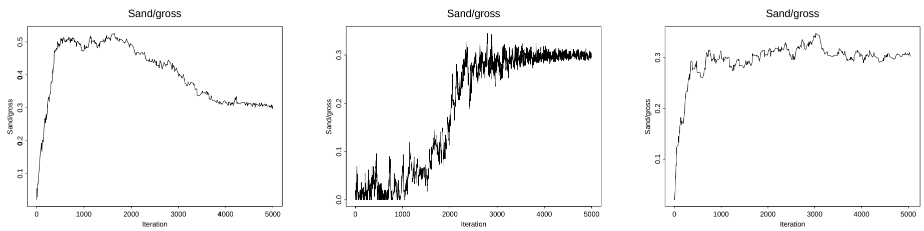

Figure 4: Sand/gross against iteration using seismic with too large expected sand/gross (left), no seismic (middle) and seismic with correct expected sand/gross (right).

The program reports the expected sand/gross from seismic, and gives a warning if this falls outside the target area. In order to match the sand/gross expectations, two things can be done. The easiest is to move the target sand/gross to the value expected from seismic. However, this value is most likely to be better founded than the probability for sand given seismic, so this is where the corrections should be made.

A value from seismic above the target area means that the probability of sand is too large. This can be corrected by moving the entire probability for sand to the left, that is, decreasing the values for seismic in the input distribution. A value below the target area, on the other hand, indicates that the probability of background is too large. To correct this, increase the seismic values in the file containing the probability of channel for given seismic values.

Even if the expected sand/gross from seismic matches the target sand/gross, overshooting of sand/gross during simulation and mismatch between prior and posterior probability of sand given seismic may occur. This is due to a too large seismic factor. There will also be mismatch between prior and posterior distribution if the seismic factor is too small.

A too small seismic factor will tend to flatten the probability for sand given seismic, whereas a too large factors will turn it into a Heaviside function. The middle point will depend on how large the correlation in the seismic grid is. If there were no correlation, the seismic factor should be equal to the number of grid nodes, but since there normally is rather high correlation, a much smaller factor should be used.

The seismic factor must be tuned from example to example; it depends on the number of grid nodes, the correlation in the seismic, the prior probability of sand given seismic, and expected width and standard deviation of horizontal displacement for channels. Before an attempt to tune the seismic factor is done, the expected sand/gross from seismic must match the target sand/gross.

When this is done, a too large seismic factor will reduce the variance in the posterior distribution of seismic given facies as compared to the prior, whereas the opposite is true for a too small factor. Thus the way to adjust the seismic factor can be determined from looking at these distributions. They are logged in the .seisini and .seissim files. An almost perfect match between prior and posterior distribution of seismic given facies should be possible if there is no strong interaction and not too many well observations.

Another way to get a rough idea of whether the seismic factor used is far off mark is to look at the correlation between seismic and facies for different values of the seismic factor. If the correlation does not increase significantly with increased seismic factor, this means that the seismic factor being used is too large.

Using a too large seismic factor is equal to having a binary seismic - seismic values above a certain level indicates background, below indicates sand. If the specified prior probability for sand given seismic is to have any effect, the seismic factor must be near the correct value; however, there is some robustness, so a perfect match is not required. Using a too small seismic factor is similar to reducing the weight of the seismic observations; the prior still has influence, but it is weaker than it should be.

5.4 Seismic sections

The idea behind the seismic sections is that these should make it easier to generate channels that matches the seismic observations. To test this assumption, an artificial seismic file was made, containing one straight channel of seismic with high sand probability, whereas the seismic everywhere else indicated very low probability. The model file was set up so that there would be space for only one channel in the good seismic area; there was also no wells and no target sand/gross, so only the seismic potential and the prior distribution affected the result. The prior distribution of channels was set to be uniform, with direction parallel to the seismic channel, and with expected width and thickness equal to the width and thickness of the seismic channel.

This model file was run with and without seismic sections. The runs using seismic sections performed better both with respect to how many iterations they needed to hit the channel and how well they hit it. The average number of iterations before a hit with seismic sections spaced 5 grid nodes apart were 3.2, half of what was observed without. The correlation between sand and seismic averaged -0.85 (in this setting, -1.0 was theoretically possible) compared to -0.59 without. Increasing the distance between the grid nodes to 10 did not increase the number of iterations needed to hit, but decreased the correlation to -0.83. In this synthetic example with binary seismic, a correlation of -1.0 is desirable. This is normally not the case, since a high correlation indicates a binary interpretation of the seismic.

As can be seen from Figure 5, the general shape of the channel changes when using seismic sections. With these relatively close to each other (5 grid nodes apart in the figure) the channel exhibits a more high frequency, low amplitude oscillation. The reason for the high frequency is that the channel parameters in one section is conditioned only on the previous section, whereas a variogram is used between sections and when there are no sections. The low amplitude is of course due to the fact that the channel must hit the seismic with high probability.

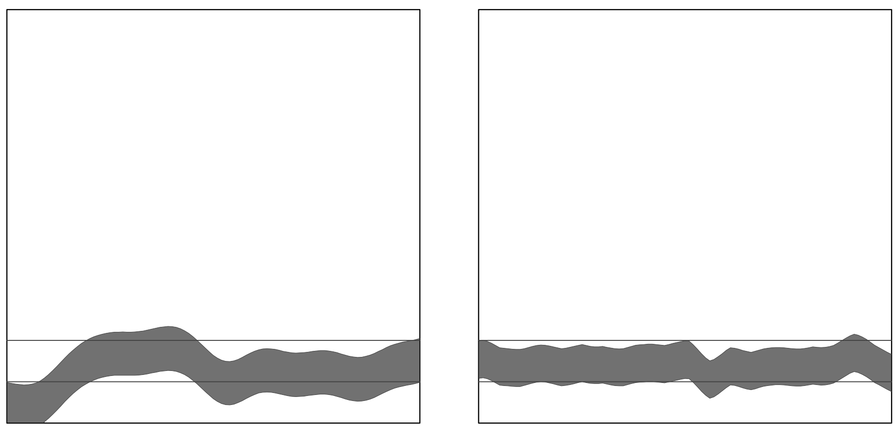

Figure 5: Realization without seismic sections (left) and with (right). The horizontal lines show where the seismic channel is.

Note that placing a seismic section in every grid node (setting the grid-node-separation to 1) does not necessarily provide the best results. This is due to an approximation error; when drawing the position, width and thickness of a channel in a seismic section, it is assumed that this is done from the real distribution. However, what is done is that 10 different realizations are generated, and then one of them is drawn according to their probability. The error made here takes effect only when there are seismic sections in all grid nodes, since there otherwise is calculated a seismic potential which makes this effect insignificant. This is also the case with a similar error which is done when drawing the condition point and direction of the family line.

Seismic sections were also used in an example with a more realistic seismic file. Here, high correlation is not necessarily a positive sign, since it may indicate a perturbation towards a Heaviside probability function for sand given seismic. In the examples run, the seismic factor was tuned, so no such perturbation should take place.

In addition, convergence was considered, instead of looking at the realization. The simulation runs were aborted 100 iterations (out of 1000), and the amount of sand/gross and correlation between facies and seismic was considered. The simulations done with seismic sections averaged a sand/gross of 0.21 and a correlation of -0.29, whereas those without averaged 0.20 in sand/gross and -0.26 in correlation. The same simulation was run with a very fast cooling so it converged in sand/gross after 150 iterations; the correlation between seismic and facies was then -0.21 in mean with seismic sections, compared to -0.17 without. All this indicates a faster convergence when using seismic sections.

In general, seismic sections will be more useful the less smooth the seismic is. The realistic seismic used here was rather smooth, and so the improvement from using seismic sections was limited. However, as the example with binary seismic shows, the seismic sections help the channel to find a good path through the reservoir when there are large differences in the seismic. Also, it helps channels avoid small regions where the seismic is unfavorable.

If the variance in seismic is on a too large scale compared to the allowed displacement and size variation of the channel, seismic sections does not make much difference. Another thing to keep in mind is also that large displacements and variations in size may reduce the effectivity of the sections, since the sample of ten parameter sets made in each section will become less representative.

When using seismic sections, it is important to give the channels enough freedom so there is a real choice of where to place the channel. Thus, small standard deviations for horizontal and vertical displacement, thickness and width will reduce the effectivity of the seismic sections.

6 Testing border effects

In case studies performed by Statoil, border effects were observed near the edges of the reservoir box. Sand/gross was here lower than expected from the seismic input probabilities.

As explained in sections 2.5 and 2.6, the line drawing algorithm undersampled lines in the edges of the reservoir. This explains most of the documented border effects.

However, if the simulation is not run a sufficient number of iterations, border effects may still be observed. This is due to:

1. Channels which contribute much to sand/gross have larger accept probabilities early in the simulation.
2. The seismic line drawing algorithm generates rather seldom lines at the edges, therefore the simulation must be run a sufficiently number of iterations to ensure that edged lines may be generated.

A testing of the present version of Fluvial has been effectuated. 50 independent realizations are summarized in averaged sand/gross projections (xy-plane) for the following cases: unconditional simulation, conditioning on information-less seismic and conditioning on synthetic seismic (Statoil example). The present and previous versions were compared in the unconditional case. The previous version did however contain a defect that gave inverse border effects rending a comparison in the seismic cases not possible.

Values for vital model parameters influencing the sand/gross distribution:

```
VOLUME
0.0    ! x-origo
4000.0 ! dx
0.0    ! y-origo
4000.0 ! dy
100.0  !dz
;
<sd(a)(real)> *
<sd(c)(real)> *

<sand-gross(real)> 0.43
<sg-epsilon(real)> 0.005

<nIterations(int)> 5000

<sampnx(int)> 40 ! Gridnodes in x-direction
```

```
<sampny(int)> 40 ! Gridnodes in y-direction
<sampnz(int)> 100 ! Gridnodes in z-direction

[seismic-factor(real)] 800
```

In the previous version of Fluvial some of the corresponding values differed:

```
<sd(a)(real)> 2000
<sd(c)(real)> 50

<temp-rate(real)> 0.998
```

The '*'-mark specifies an uniform prior distribution of family lines. 5000 iterations corresponds to the temp-rate of 0.998 in the old model file.

Figure 6 shows the results from unconditional simulation for the new (left) and previous (right) version of Fluvial. The characteristic bell shape of the normal distribution is seen in the previous version. The sand/gross distribution is centered in the reference point, and decreases from this reference point. The present version of Fluvial gives better result, but sand/gross is slightly centered in the middle of reservoir.

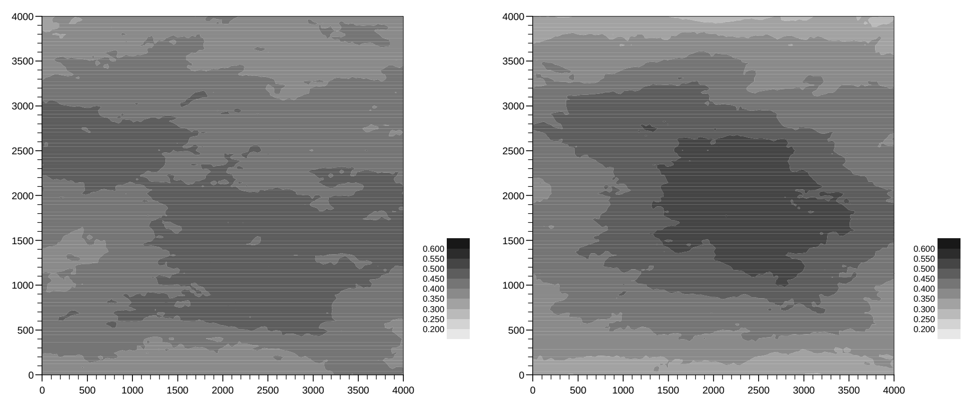

Figure 6: Sand/gross maps (xy) averaged over 50 independent unconditioned realizations. Results from present (left) and previous (right) version of Fluvial.

Information-less seismic should give in expectation an uniformly distributed sand/gross, it should therefore be a good reference for testing the simulation program for border effects. Figure 7 shows the results from the present version of Fluvial. The left figure is iterated 5000 times for each run, while the right is iterated 20 000 times. No significant border effects are observed in the edges, however sand/gross seems to be slightly centered in the middle of reservoir. There are a clear improvement in increasing the number of iterations from 5000 to 20000. This indicates that the non-uniformed sand/gross distribution is caused by a lack of convergence.

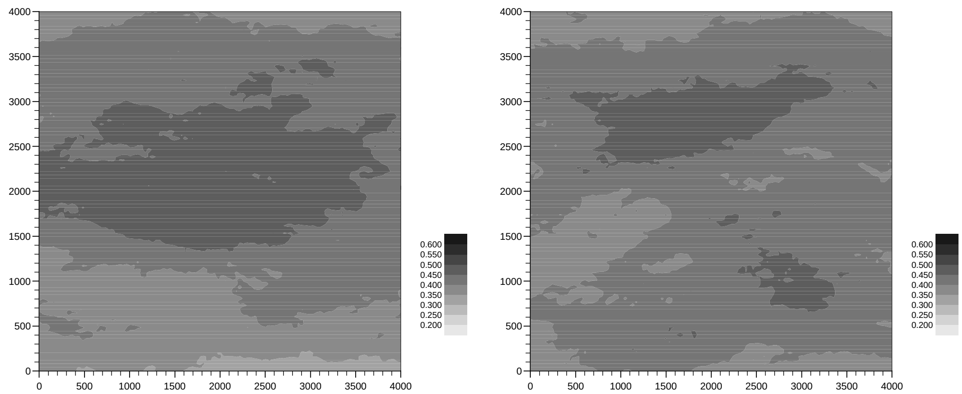

Figure 7: Projected sand/gross map (xy) averaged over 50 independent realizations conditioned on information-less seismic. Results from present version of Fluvial with 5000 (left) and 20 000 (right) iterations for each run.

Figure 8 shows the result from the present version of Fluvial using synthetic seismic (left), and the expected sand/gross distribution according to the input seismic (right). No significant border effect is observed, however the sand/gross distribution is thresholded. This may be explained by a too high seismic factor.

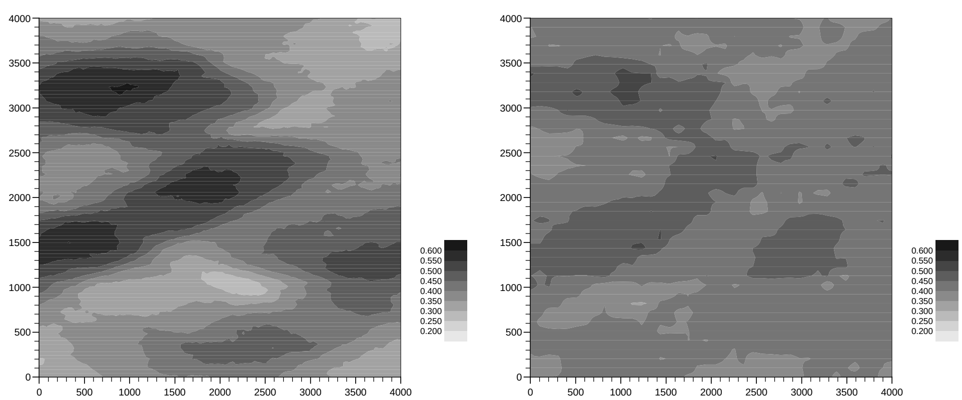

Figure 8: Sand/gross map (xy) using synthetic seismic. Average over 50 independent realizations (left) and expected sand/gross from seismic (right).

A conclusion of these tests is that no significant border effects are observed for the present Fluvial version although sand/gross is slightly centered around the reservoir centre. This seems to be a convergence problem, the results for the information-less seismic indicates much better results when the number of iterations increase. A future improvement of Fluvial should be a removal of the effect that channels contributing much to sand/gross have larger accept probabilities.

7 Testing the new interaction function

In the present version of Fluvial, the interaction function has been simplified compared to the old version. However, this simplification also provides greater flexibility when modelling the interaction between channels in a reservoir because of the new angle parameters.

To illustrate a main aspect of the new interaction function, a model was made where all channels had the same z-coordinate. Thus, if they crossed, they had to do it in the same plane. Figure 9 shows two realizations with different interaction functions.

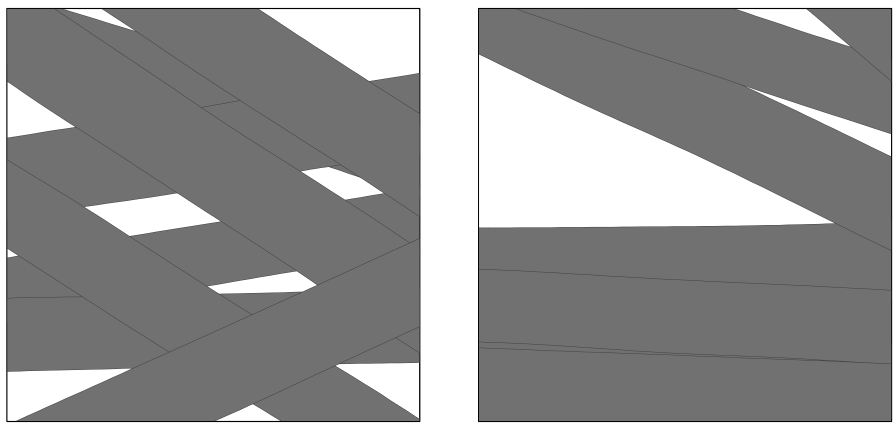

Figure 9: Realization of the same model with the new interaction function(left) and the old (right).

This figure shows clearly the difference. With the old interaction function, no channels cross each other. They do cross with the new function, but only if the angle between them is sufficiently large. Here, the minimum interaction angle was set to 10 degrees, the maximum to 30.

A comparative study of the interaction functions was also done using a more realistic model. The interaction angles were still set to 10 and 30 degrees. The results are summarized in Table 8.

The leftmost table shows the mean number of families in a realization, the middle shows the standard deviation of the mean direction of lines between the realizations, and the rightmost shows the mean of the standard deviation for the direction of lines within one realization. All numbers are calculated from series of ten realizations.

The tables show that the number of families show no significant difference between the old and new interaction function. This is no surprise, since the difference between the interaction function should work both ways. On one hand, the old interaction function should have a tendency so squeeze

Number of families

S/g	Old	New
0.15	10.6	9.4
0.35	25.0	25.5
0.55	41.9	41.1
0.75	65.2	63.5

Sd between real.

S/g	Old	New
0.15	4.49	2.97
0.35	1.91	1.35
0.55	2.10	0.95
0.75	2.32	1.20

Sd inside real.

S/g	Old	New
0.15	9.11	9.11
0.35	9.29	9.83
0.55	9.32	10.00
0.75	9.44	10.23

Table 8: The number of families and st. dev. for direction of lines inside and between simulations with various sand/gross and interaction function.

channels out along the edges, so their contribution to the sand/gross would be smaller, and more channels would be needed. On the other hand, the new function allows them to intersect, and thereby reduces the sand/gross contribution from each individual channel. These effects seem to cancel out.

What distinguishes the new and old function in these tables are the difference in standard deviation of direction of families within and between realizations. As was expected, the new interaction function gives larger standard deviation within the realization, and smaller between.

This is due to the fact that the old function had a tendency to make the channels parallel, in a direction decided by the early drawn channels. Thus, the variance inside a realization would be lower than without this tendency. Furthermore, since different realizations would have different directions of the channels drawn early, the variance of mean direction between realizations would be larger.

This shows that the new interaction function fulfills its purpose. It is also faster to calculate than the old function; however this has no measurable effect on the total running time of the program.

8 Testing the conditioning on uneroded thicknesses

This section presents some tests on conditioning on the uneroded thicknesses. First ten unconditional realizations from a model file used at Statoil in earlier tests were simulated. The model had one channel for each family, no crevasses, no correlations between the channel parameters and long variogram ranges. The target sand/gross was 0.35.

The two conditional versions 'Eroded' and 'Uneroded', were then generated by adding ten synthetic wells to one of the unconditional realizations. These wells were distributed as shown in Figure 10. The channel observations in these wells were used as inputs to the conditional versions.

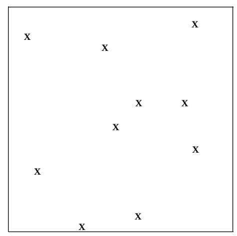

Figure 10: Distribution of synthetic wells.

Now, ten realizations with conditioning on both eroded and uneroded thicknesses were run, and the results for some vital parameters are given in Table 9. The numbers in the parenthesis are standard deviations for the means of the 10 realizations given in the table. Note that the directional values are the horizontal values.

s/g = 0.35	Prior	Unconditional		Eroded		Uneroded	
Mean Thickness	10.0	10.0	(0.1)	9.3	(0.1)	9.6	(0.1)
Mean Width	700.0	710.6	(10.5)	699.6	(12.2)	694.4	(18.4)
# Channels		28.0	(2.3)	27.9	(2.5)	25.9	(1.9)
Mean Direction	0.0°	-1.1°	(5.5°)	-1.5°	(3.7°)	-2.7°	(4.5°)
std(Direction)	30.0°	28.8°	(2.3°)	28.4°	(1.9°)	28.3°	(2.6°)

Table 9: Test results, first synthetic example.

From Table 9 it is seen that the mean thickness increases when conditioning on the uneroded thickness instead of the eroded thickness as expected. The mean thickness was less than for the unconditional case. In this dataset, some of the observed channels which could not have been eroded had significantly lower thickness than the expected thickness, something which reduced the mean. The number of channels was also reduced as expected, this follows from the thicker channels and the constant sand/gross target. The mean width in this case is however lower for the 'Uneroded' case, but not enough to fully compensate for the increased thicknesses. The differences in the directional values are all within one standard deviation, and gives no knowledge about a change in these values due to the change in thickness conditioning.

In order to better see the effect of the uneroded conditioning, the target sand/gross was decreased to 0.15, the expected channel thickness increased to 20.0, and the correlation between width and thickness set to 0.4. We also had only one well, listed below, with many observations of possible eroded channels. A repetition of the whole procedure gave Table 10.

s/g = 0.15	Prior	Unconditional		Eroded		Uneroded	
Mean Thickness	20.0	19.9	(0.5)	17.1	(0.9)	19.9	(0.2)
Mean Width	700.0	711.0	(29.1)	705.6	(22.8)	695.2	(39.3)
# Channels		4.8	(1.9)	5.6	(0.7)	5.2	(0.4)
Mean Direction	0.0°	1.6°	(8.3°)	1.8°	(11.5°)	−4.5°	(16.8°)
std(Direction)	30.0°	24.2°	(10.6°)	27.2°	(8.6°)	24.4°	(8.2°)

Table 10: Test results, second synthetic example.

Again the mean thickness increased while the number of channels decreased. For the 'Uneroded' case the mean width decreased compared to the unconditional, the 'Eroded' case is not comparative as an error in the previous version was discovered which influenced this value. The increase in mean widths between the realizations is due to the way the widths are drawn, after the thicknesses have been drawn truncated, the mean width are drawn. This makes the mean width dependent on an earlier drawn value with large variation. The differences between the mean directions of the realizations is shown to be large from the standard deviation value of 16.8°. The low number of channels in this example is probably the reason for this high value. The differences between the standard deviations within the realizations is quite comparable with the 'Eroded' case.

The well facies file was:

```
well01
 TVD   FACIES  FAMILY  CHANNEL
  0.0     2       1       1
 18.0     2       2       1
 32.0     2       3       1
 47.0     1    -999    -999
 52.0     2       4       1
 72.0     2       5       1
 82.0     1    -999    -999
100.0     1    -999    -999
```

This means that the uppermost channel was observed to have thickness equal to 18, the next was a possible eroded channel observed to be 14, but with a maximum thickness of 32. Family/channel 3 was a new possible eroded channel, with uneroded thickness somewhere in $[15, 47]$. The fourth channel was observed to be equal to the expected value, 20, while the bottom channel was possible eroded, with uneroded thickness between 10 and 30. Here only one of the observed channels which could not have been eroded had a thickness below the expected mean, and the relative difference between the average value for uneroded thickness conditioning is reduced compared to Table 9. The raw data had 3 realizations above the expected level and values ranged in $[19.5, 20.4]$. Although a minor bias from the observations existed, this did not give a systematic error in the output, which leads to the conclusion that the new conditioning behaves as hoped and expected.

The difference in CPU time used by the two different conditioning versions was not significant. Although each conditioning will take a little longer in the new version, this is negligible compared to other effects. Note that the simulations may run faster due to faster convergence, depending on the model. This will especially occur when the uncertainty in the thickness in the model is low while a channel is observed in different wells with very different thicknesses.

9 Log files

There are three levels of output files from the program fluvial_facies. The bottom level contains the essential files, that is, the raw output data and a general log of the current simulation run. The data files will not be discussed here.

The second level of output contains a number of files with more detailed information of the simulation run, and some processed data. The files from the third level are intended mainly for debug purposes. By specifying the command line switch -l (number) where (number) is 0, 1 or 2, the user can set the output level. Default value is 0, that is, only the essential files.

9.1 Level 0 file:

<mod-name>.log
This file is a log of the simulation run. It is divided into 5 parts: The first part shows some input parameters for the simulation. These are the seed for the random number generator, the global volume simulated (xmin, xmax, ymin, ymax, zmin, zmax) and the channel observations made from wells.

The second part shows how the iterations ran. For every 30 iterations, the number of channels added, changed and removed is shown, as well as the number of accepted proposals. The total sum for each of these is also given.

Part three is present only if the current simulation included seismic data. Then the mean and standard deviation for seismic values in sand and shale is given, as well as the correlation between facies (sand) and seismic.

Part four goes sequentially through all families, and all channels within families. For each family, the direction of the family line is given by the horizontal and vertical angle. The mean depth is also given, as is the maximum, minimum and actual number of channels in this family. For each channel, the mean width and thickness is given.

The last part contains some summary information about the channels. It shows the number of families, mean number of channels per family, mean number of crevasses per channel, mean width and thickness of the channels, mean and standard deviation for the direction of the lines, final sand/gross ratio, and sand/gross ratio as expected form the seismic data and interpretation. The numbers in parenthesis are the expected values, as given by the user. The last number in the parenthesis after the sand/gross ratio is the acceptance deviation from the target sand/gross ratio.

9.2 Level 1 files:

<mod-name>.simulation
This file shows how the sand/gross ratio and the seismic potential varies throughout the iterations. The first column is the iteration number, the second is the temperature, the third is the sand/gross ratio, and the fourth is the value of the seismic potential from last state. Data is written to this file only when a proposed state is accepted, since the two last values change only then.

The following level 1 files are generated only if seismic is used:

<mod-name>.condseis
This file contains the simulated probability of sand given a seismic value. The first column contains interval midpoints, and the second contains the probability that the observation is sand given a seismic value in this interval. Occurrences of NaN means that there were neither sand nor shale observations with seismic in this interval.

<mod-name>.seisini
This file contains the a priori probability of finding seismic in an interval given the facies. The first column contains the interval midpoints, the second the probability of getting a seismic value in this interval given that the facies is sand, and the third is as the second, but with shale instead of sand.

<mod-name>.seissim
As the file above, except that these values are calculated from the simulation.

9.3 Level 2 files:

<mod-name>.acceptprobs
This file contains detailed information about the iterations. The three first columns are the probability of choosing remove, change or add as action. The four next columns show the change in different potentials between current and proposed state; column 4 is the sand/gross, 5 is the number of uncorrelated observations, 6 is the interaction between families and 7 is the seismic. Column number 8 contains the accept probability for the proposed state, 9 shows a letter for which action was chosen ('a' for add, 'c' for change and 'r' for remove'), and the last column contains four stars if the proposed state was accepted, a minus otherwise. Data is written to this file for each iteration.

<mod-name>.cond
This file contains information about the channel-well intersections, and is created only if there are well observations. It has two parts. The first

part tells where the top and bottom of the well observation and channel is. Column 1 is the family, 2 is the observation number, 3 the well number, 4 and 5 the x- and y-position of the well, 6 and 7 the top and bottom of the observation in the well, 8 and 9 the top and bottom of the channel, 10 is an indicator telling if the channel intersects the well, and should be 1, whereas the last column holds the faciestype.

The second part tells about the composition of the wells. The first column is the well number, the second is the sand/gross ratio in the well, whereas the third and fourth are the parts from channel and crevasse.

<mod-name>.famPot
This file contains the potential of each family. The family number is in the left column, the potential in the right.

<mod-name>.dist_<well-name>
These files are created only if there are well observations, and the command WELLTEST is given in the <mod-name>.MOD file. They contain the shortest distance from the well to the wall of a channel passing through the well. One file is generated for each well. The first column contains family number, the second channel number, and the third the distance.