1661-8726/09/010107-14
DOI 10.1007/s00015-008-1301-1
Birkhäuser Verlag, Basel, 2009

Swiss J. Geosci. 102 (2009) 107–120

Preparing a seismic hazard model for Switzerland: the view from PEGASOS Expert Group 3 (EG1c)

Roger M. W. Musson[1,*], Souad Sellami[2] & Wolfgang Brüstle[3]

Key words: PEGASOS, Switzerland, seismic hazard, hazard model

ABSTRACT

The seismic hazard model used in the PEGASOS project for assessing earthquake hazard at four NPP sites was a composite of four sub-models, each produced by a team of three experts. In this paper, one of these models is described in detail by the authors. A criticism sometimes levelled at probabilistic seismic hazard studies is that the process by which seismic source zones are arrived at is obscure, subjective and inconsistent. Here, we attempt to recount the stages by which the model evolved, and the decisions made along the way. In particular, a macro-to-micro approach was used, in which three main stages can be described. The first was the characterisation of the overall kinematic model, the "big picture" of regional seismogenesis. Secondly, this was refined to a more detailed seismotectonic model. Lastly, this was used as the basis of individual sources, for which parameters can be assessed. Some basic questions had also to be answered about aspects of the approach to modelling to be used: for instance, is spatial smoothing an appropriate tool to apply? Should individual fault sources be modelled in an intraplate environment? Also, the extent to which alternative modelling decisions should be expressed in a logic tree structure has to be considered.

1 Introduction

A general account of the PEGASOS project will be found elsewhere (Abrahamson et al. 2002). This was the first application of the SSHAC "Level 4" methodology (Budnitz et al. 1997) outside the United States, and participating in the study was a new, and a stimulating experience for all those engaged in the capacity of experts. The first sub-project of PEGASOS (SP1) concerned the formulation of the seismic source model that would be used for the final hazard calculations. SP2 was charged with developing ground motion studies appropriate for use in Switzerland. Site response was handled by SP3, and the actual hazard calculations were conducted by SP4. This division of the project made working with PEGASOS rather different from a routine hazard study, in that each expert was to some extent isolated from the work undertaken by the other sub-projects, although the sub-projects were not sealed off from one another totally.

SP1 was further divided into four teams of three experts each. Each of these Expert Groups was charged with preparing a single model. It was understood at the outset that:

- Agreement within each Expert Group was required;
- Agreement between the Expert Groups was not required or expected;
- Each Expert Group was further expected to take into consideration the range of opinions present in the wider scientific community;
- Hence, between the four groups of three, a more or less complete perspective of relevant contemporary scientific opinions should inform the final product.

This paper presents a partial account of the modelling process as seen by Expert Group 3 (EG1c). The final product of each group was an "elicitation summary"; a document that completely describes, in text and tables, the model as defined by the group. This paper draws on the elicitation summary produced by EG1c; we feel that merely to recapitulate that

---

[1] British Geological Survey, West Mains Road, Edinburgh EH9 3LA, UK. E-mail: rmwm@bgs.ac.uk
[2] Swiss Seismological Service, ETH-Hönggerberg, CH-8093 Zurich, Switzerland.
[3] Regierungspräsidium Freiburg, Abt. 9: Landesamt für Geologie, Rohstoffe und Bergbau, Ref. 98: Landeserdbebendienst, Albertstrasse 5, D-79104 Freiburg im Breisgau, Germany.
*Corresponding author: R. M. W. Musson. E-mail: rmwm@bgs.ac.uk

report (in NAGRA 2004) would not be appropriate for the wider readership, and consequently this paper concentrates chiefly on certain issues that occasioned special discussion with EG1c, and which have wider implications for seismic hazard methodology beyond the confines of Switzerland.

2 Inclusion or exclusion?

The first point that needs to be considered as a matter of general principle is the degree to which seismic source models should explicitly incorporate competing hypotheses. It is generally considered a matter of importance that, in probabilistic seismic hazard analysis (PSHA), the full range of uncertainty should be addressed (Budnitz et al. 1997). This means addressing both issues of epistemic and aleatory uncertainty (Abrahamson 2000). Of these, one could say that issues of epistemic uncertainty are those things which are unknown but which could conceivably be known (e.g. whether a coin is weighted or fair). Aleatory uncertainty refers to those things that can never be known (e.g. what the next toss of the coin will be).

Many issues of opinion relevant to the formulation of a seismic source model are issues of epistemic uncertainty. Are tectonics in the Alpine foreland thin-skinned or thick skinned? Is a given fault active or inactive? What are the depth limits of seismicity? What are the magnitude limits? What is the completeness of the earthquake catalogue? These are all issues on which different people could express different opinions, and even a single expert might well advance different answers that could equally well be correct on present information.

A SSHAC Level 4 study is designed to encompass the full range of informed opinion on all these issues. In the first case, this enhances the scientific robustness of the study; all available information is incorporated into the model. In the second case, it also enhances the political robustness, in that the study cannot subsequently be attacked by a hypothetical “Professor X” on the grounds that the important hypothesis of “Professor X” was not taken into consideration.

However, the point remains open as to whether considering a hypothesis necessarily means including it. The SSHAC Level 4 methodology requires that all relevant studies and ideas be taken into consideration; but it does not preclude the possibility that the EG may consider a hypothesis and then reject it. This issue revolves around the question of weights. Epistemic uncertainty is generally incorporated into the final model using the principle of the logic tree (Coppersmith & Youngs 1986). Competing hypotheses may be assigned branches in the logic tree, with a weight expressing the EG’s opinion as to the probability that this hypothesis is correct. The question then comes down to this: given a hypothesis that the EG know of but have little faith in, should this hypothesis get a branch on the logic tree with a very low weight, or should it be given zero weight (i.e. excluded)?

A specific example can be given, one which generated much discussion in PEGASOS as a whole. Could the maximum magnitude in Switzerland be as high as 8 Ms? One can make a case for this as follows (A):

(i) The 1812 New Madrid earthquake was over 8 Ms (Johnston et al. 1994).
(ii) What can happen in the Mississippi Valley could equally well happen in the Upper Rhine Graben;
(iii) Therefore M_{max} could be as high as 8.

One can equally argue in a contrary fashion (B):

(i) The 1812 New Madrid earthquake was less than 8 Ms (Hough et al. 2000).
(ii) There are no structures suitable for a magnitude 8 in or near Switzerland, nor do the strain rates allow of it;
(iii) Therefore M_{max} could not be as high as 8.

Obviously these are simplified cases. The point at issue is that, from the view of the analyst favouring case B, one could follow two paths. One is to say that case B refutes case A, and therefore no weight should be given to $M_{max} = 8$. The other is to say that although the analyst does not agree with case A, some people in the wider community hold case A to be valid, therefore this case must be represented, if only at some very small weight. Therefore $M_{max} = 8$ is included in the model weighted at 0.05 or less.

There clearly exists a tendency towards inclusiveness in PSHA practice. A case often referred to, especially by detractors of PSHA, is Bernreuter et al. (1989) in which any source model submitted by an expert in the project was considered valid, even though some of the outliers could have been rejected as unrealistic (Musson 2004a, 2004b).

In fact, there is a danger in following this approach. As one considers hazard at longer and longer return periods, extreme outliers begin more and more to dominate the hazard. An early form of deterministic hazard analysis sought to model the maximum possible (worst case) ground motion. The objection to this practice was that the worst case is often so improbable that it is doubtful whether it should be considered. If one inverts this, and looks to calculate the hazard that is extremely improbable, one should not be surprised if this turns out to be the worst case.

Thus, including in the model a hypothesis in which one has no faith, “because it’s there”, at very low weights, can still exert a totally disproportionate effect on the hazard, despite the low weights.

A potential drawback to the SHAC Level 4 procedure is that the analysts preparing the source models are divorced from those doing the hazard calculations, making it hard for the former to follow a holistic approach to the modelling, in which decisions are made with full knowledge of their impact on the results (though the Level 4 procedure does provide a formal mechanism for allowing this feedback). However, the fact, known in advance, that hazard would be calculated at the very low level of 10^{-8} per annum, indicates that one should actually be careful about what is and is not allowed into the model, even at very low weights.

As a rule, therefore, the EG1c practice was not to spare using zero weights for hypotheses which we considered were not sufficiently supportable to include in the model. Even in some cases where hypotheses were very definitely held by members of the wider community, we considered that so long as these ideas were rejected from our model as part of a thoughtful process, the principle of taking into account the wider body of opinion was satisfied.

3 Spatial smoothing

The use of spatial smoothing of seismicity as a means of modelling seismic hazard is currently a subject of much methodological discussion. The use of grids to smooth the observed earthquake pattern goes back in published studies at least as far as Jacob et al. (1994) and possibly before. It obtained some popularity through the work of Frankel et al. (1996), and the employment of kernel functions (Woo 1996) is also worthy of note.

All these methods are an attempt to remove the subjectivity involved in making decisions as to zone geometry, at the expense of abandoning any possible input from seismotectonics and geology. (Subjectivity also remains in the choice of smoothing parameters.) The use of smoothed seismicity was aptly described by Perkins (1993, pers. comm.) as providing "a good quick first approximation to the hazard". It is certainly a useful tool for computing hazard (especially hazard maps that will not be used for design purposes) in a hurry in areas where tectonic data are absent or hard to interpret. It can be viewed, for example, as a weapon of last resort.

There is also an issue of spatial stationarity. The use of a large seismic source zone implies that future seismicity can be located anywhere within that zone with equal probability. The extreme opposed view would be that historically observed epicentres will repeat themselves exactly in the same places, over and over again. This latter view would amount to supposing that seismicity was, in a spatial sense, entirely stationary. Spatial smoothing, either of an entire catalogue, or selectively within broadly delimited zones, allows an intermediate position of partial stationarity – that earthquakes in future may occur anywhere, but with diminishing probability away from locations where they have occurred in the known past.

The use of spatial smoothing is not the only way to achieve the same effect. A model used by Musson & Winter (1996, 1997) and discussed in Musson (1997) overlays a series of broad seismic source zones with a second series of small ones concentrated around past historical epicentres.

An issue that was debated in EG1c was whether it is correct to combine zoned and zoneless approaches in a logic tree. The argument against this is that a logic tree should confine itself to the expression of epistemic uncertainty (Abrahamson 2000), whereas a zoned approach and a zoneless approach represent a choice of tools, and not uncertainty about the nature of seismicity. To combine the two methods would be like attempting to combine probabilistic and deterministic hazard assessments within one logic tree.

The counter argument is that the primary epistemic uncertainty is in the conceptual models of seismicity that may be conceived by the analyst. Any procedural approach is a way of manifesting a conceptual model, and some models may be more easily expressed by one approach, and others more easily expressed by another. Hence, if one has two conceptual models of seismicity, one of which is more easily described by a set of zones, and the other by a set of smoothing parameters, then it makes perfect sense to combine both approaches in the same logic tree.

However, one does not formulate a conceptual model that seismicity follows some unknown pattern of zones, and then set out by trial and error to discover what those zones might be. Rather, one formulates a conceptual model that seismicity is higher in one place and lower in another place for certain reasons, and then uses a zone geometry as a tool to express this model. The concept of zonation has no practical meaning divorced from a real set of zones. Ultimately, a conceptual model does not follow the form of "seismicity is zoned", but rather "seismicity is zoned *in this way*".

The same is true of the zoneless approach. One's conceptual model cannot follow the form of "seismicity is smoothed" (in some unknown way), but should be "seismicity is smoothed *with these characteristics*". Spatial smoothing becomes a useful tool to express a conceptual model that already contains a smoothing shape and wavelength. It cannot be sound practice to decide on a smoothing approach without any conception of the correct form of the smoothing parameters and then try and obtain these by trial and error (and with no means of judging what was error and what was not). This would no longer be the representation of a conceptual model, but rather playing games with a computational method.

We therefore asked ourselves whether our ideas about the distribution of hazard in the study area were more faithfully represented by a set of zone geometries that we could express, or a set of smoothing parameters that we could express. We found we had no way of judging the latter, but we could express our ideas in terms of the former. We therefore chose to reject the use of smoothed seismicity models in our approach to describing the seismicity of the study area. It can also be noted, from a technical perspective, that there is nothing that can be achieved using a spatial smoothing approach that cannot be more or less duplicated, in effect, by thoughtful application of source zones. This gives a further justification in restricting the EG1c approach to the use of source zones.

This is one example of not including something in the model just "because it is there". A second, and in many ways more important case, deals with the issue of active faults.

4 Active faults in an intraplate setting

Conventional practice in PSHA includes explicit fault modelling for active faults, while seismicity for which the precise causative fault is not known is usually handled within the zone model. The explicit inclusion of individual faults has two important effects: (a) it enables the activity on this fault to be localised along its extent, possibly including a segmentation model; (b) it enables the effect of rupture dimensions to be included within the hazard estimation. The word "explicit" is used judiciously; all earthquakes modelled in the PSHA process occur on faults. If a known fault is not modelled in detail, but is encompassed within some source zone, this does not mean that the fault is presumed not to be seismogenic. It means that the fault is not considered to be more hazardous than other faults within the same zone, which may or may not themselves be mapped. Some of the earthquakes that are implied to occur within the source zones will, in fact, occur along the known fault in question. The absence of a specific fault source in the model simply implies that earthquakes will not preferentially occur on that fault rather than others in the same zone. It also implies that the rupture dimensions of earthquakes occurring on this fault are not significant to hazard. (And even this is not necessarily true, as some hazard codes – including that used in PEGASOS – treat a source as a set of pseudo-faults with known rupture behaviour.)

For a large source zone, it may well be true that rupture size is not significant to hazard. The assumption made in PSHA is that, for any zone, the epicentre of a future earthquake may be at any position with equal probability. However, the epicentre is in this context a notional point, and one could equally rephrase this as "the closest point of the rupture plane of a future earthquake projected to the surface may be at any position with equal probability". This is clearly not true for a small zone, the maximum dimension of which is only slightly larger than the projected rupture length of a large earthquake within the zone. But it may well be true for a large zone and small earthquakes.

The question now arises as to what constitutes an active fault, and which faults in the study area (Switzerland and the immediately surrounding areas) are active. Conventional definitions, largely developed in tectonic areas rather different from Switzerland, refer to any fault that has demonstrably moved in the past x years as active, where x is some large number extending certainly beyond historical times, usually back to the beginning of the Quarternary. It is common practice to examine known faults one by one, compare them to this definition, and decide if they are active or not. The number that meets this criterion indicates the number of "active faults".

We regard this process as unhelpful in a Central European context. The number of "active faults" can be more directly estimated by approaching the question from the other direction. All earthquakes occur on faults; there are approximately 2000 distinct epicentres in and around Switzerland; therefore there must be about 2000 active faults in Switzerland. Unfortunately, most cannot be identified or even guessed at.

This means that preferential modelling of those faults that can tentatively be considered as associated with specific earthquakes can have a very undesirable consequence: it effectively changes the hazardousness of faults according to whether they have been mapped or not. Consider a case of a zone within which two earthquakes of magnitude 5.8 have occurred. One of these is close to a mapped fault and consistent with having occurred on that fault; the other one cannot be attributed to a single fault (for one of various possible reasons, including uncertain location). It would be possible to identify the mapped fault as "active", model the seismicity along it on the basis of one observed event, and apportion the other event to the background of the whole zone. The effect of this is to concentrate hazard along the mapped fault, and the occurrence of the first earthquake has more impact (or at very least, a different impact) on the hazard than the second one. Yet, seismologically, the two events are not different; both occur on faults. The only difference is the state of human knowledge concerning the two generating faults: one is mapped, the other is not. It is not satisfactory to have the distribution of hazard so much affected by this rather artificial distinction.

In the above example, it is assumed that both these faults are roughly similar ancient structures that have been reactivated in the present stress field. In such cases, there is no real reason why any neighbouring fault of similar orientation should not be reactivated. The fact that a particular fault has been reactivated once does not heavily prejudice one into believing that the next earthquake will occur on the same structure and not some other one.

However, some faults really are persistently active, such that one can state with certainty that future activity on these faults will occur at the same rate as past activity, and the slip rates can also be estimated. This occurs when faults are active because they are controlling recognisable coseismic deformation. An example is the North Anatolian Fault – it is inevitable that seismicity will continue on this structure, for tectonic reasons that are very well understood. And as a result, it is true that seismicity is *preferentially* disposed to follow this structure because it controls deformation. In which case, it is essential to model it explicitly.

The question now becomes whether any such faults can be identified in Switzerland. Lacking any major active thrust features in the Alps comparable to those of the Himalayas, there appear to be five principal candidates: the Fribourg Fault, the Vuache Fault, the Reinach Fault, and the two master faults either side of the Rheingraben. These will be considered in turn (Fig. 1).

4.1 The Fribourg Fault

As shown by Deichmann et al. (2000), there can be little doubt that the Fribourg earthquake of 14 February 1999 occurred along the Fribourg Fault. This structure is well known because

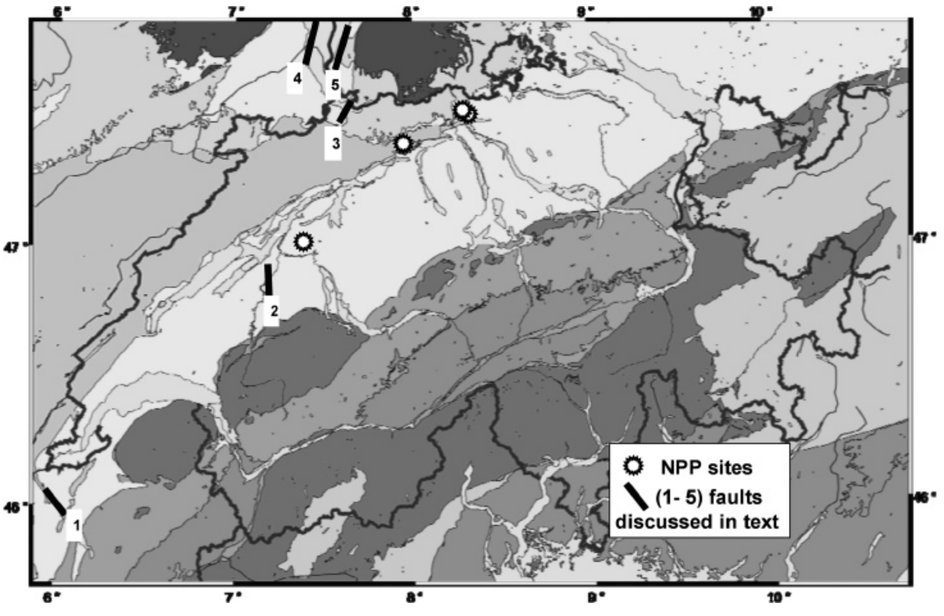

Fig. 1. Map of Switzerland showing (a) the four NPP sites of the PEGASOS study, and (b) the faults discussed in this paper (numbers as in the text).

it has been studied in detail, and can be contrasted with other faults in the same general area that have not been studied so well, cannot be characterised so well (if at all), and yet are probably equally dangerous. There is no tectonic reason why seismicity should occur preferentially along this fault, although, it may well reactivate again some time in the future. It seems inappropriate, therefore, to single this fault out for special treatment.

4.2 The Vuache Fault

This is a similar case to the Fribourg Fault; this fault produced a 5.3 ML earthquake on 15 July 1996 (Thouvenot et al. 1998). It is less significant than the Fribourg Fault because of its distance from the sites.

4.3 The Reinach Fault

A recent paper by Meghraoui et al. (2001) reports trenching across this fault south of Basel, with results that are held to show the fault rupture of the 1356 Basel earthquake as well as possible earlier events. If this evidence were correct, it could indicate a persistently active fault producing the largest regional events. However, normal faulting along this feature is an implausible hypothesis to explain the 1356 earthquake, as shown by Meyer et al. (1994). A more probable explanation for the displacement seen on the excavated Reinach Fault is slumping, which could be earthquake-triggered or not. In our view the probability that the Reinach Fault is significantly seismogenic is sufficiently small that according it special treatment is not required, and might be harmful to the model for reasons already given. (At the time of the project, our authority for this was Schmid 2002, pers. comm. – since then, Ustaszewski & Schmid (2007) write as follows: "… the evidence … for the 'Reinach fault' being an active seismogenic fault is far from convincing. The geomorphological features found in the surroundings of the 'Reinach fault' rather indicate that gravitational sliding was responsible for the 'faults' documented by the trenching …").

4.4 The Rheingraben Boundary Faults

The Upper Rheingraben is a major rift running roughly N–S, bounded by faults on either side. If the graben were still subsiding the way it probably was in earlier geological times, these boundary faults would be excellent examples of exactly the sort of fault for which special treatment as seismic sources would be required – faults the persistence of which can be assured because of their role in controlling contemporary deformation. Evidence suggests, however, that extension in the Upper Rheingraben ceased at the end of the Oligocene. The significance of the bounding faults under the current tectonic regime is debatable. We concluded that these faults are not critical in controlling current seismicity.

We were not able to identify any other candidate fault structures that required consideration, and could not conclude that any of the faults discussed above merited special treatment within the seismic source model as persistent, preferentially active sources. We therefore decided not to include any individual faults in our seismic source model. Because explicit fault modelling has a tendency to increase hazard sharply close to the fault (see, for example, Marin et al. 2004) we considered it undesirable to include explicit fault modelling just "because it was there", even with low weighting. It would have required much higher confidence on our part that any of these structures were actively controlling deformation in the present tectonic regime.

5 Source zoning

Up to now, it might seem that this paper has been somewhat negative in character, dealing more with what was left out than what was included. These issues have been emphasised largely because they are issues of general application in intraplate PSHA studies, whereas the fine details of the model that was produced are of less general interest to those working outside Switzerland. We now turn to the decisions that were taken on what was included in the model, at least in outline.

Our opinion was that the optimum way to construct a robust seismic source model follows three key stages.

The first stage involves the determination of the *kinematic model*. This is the basic element of the conceptual model of the seismic process at a sub-continental scale. The kinematic model describes, at the broadest scale, what is the relationship between large blocks in the Earth's crust in terms of relative movement. In a very simple case this might be stated as: Block A is subducting northwards under Block B, with resulting large thrust earthquakes at the interface and lesser amounts of intraplate seismicity within the two blocks in reaction to the stresses engendered by differential subduction. This describes the basic mechanisms for earthquakes that are to be expected in different parts of the area under examination.

The second stage refines the kinematic model into the *seismotectonic framework*. In this part of the process, the very broad divisions used in the kinematic model are looked at in more detail, with the aim of dividing them up into volumes of crust that are sufficiently structurally distinct that it is improbable that seismicity could be considered to be uniform across the boundaries of such divisions. At this stage, the key elements to be assessed do not include seismicity except in a rather broad sense. Rather, one is seeking to characterise areas that have a similar style of faulting, are experiencing a similar pattern of crustal stresses, and so on. In drawing up the seismotectonic framework one may start drawing basic crustal divisions that will ultimately form the outline of the seismic source model itself.

The third stage is the final construction of the *seismic source definition*. Here the final partition of the seismic source model is made from analysis of the seismotectonic framework together with the detailed pattern of observed seismicity and local geological structure.

By following this sequence of steps, we believe that an informed basis can be found for decision making about the detail of the model. How those decisions were taken in this study will now be described.

5.1 The Kinematic Model

In looking at the kinematics of crustal deformation in the area of study, a few key questions immediately arise. Firstly, what is the relationship between Alpine tectonics and seismicity? Can we see active seismicity along Alpine thrust structures indicating that mountain building is still in progress? Secondly, and clearly related, is the nature of the boundary between the Adriatic Plate and the European Plate. Thirdly, there are a number of contentious questions about the tectonics of the Alpine Foreland. We will take these topics in turn.

5.1.1 Alpine Tectonics

Although one might superficially expect the Alps to be similar to the Himalayas with respect to tectonics and seismicity, this is clearly not the case. With the Himalayas, the major active thrust planes (such as the Main Boundary Fault and Main Central Thrust) are easy to identify, and their activity is incontrovertible (Chandra 1978; Singh et al. 1990) even if the details are still subject to discussion.

The same is evidently not true in the Alps. There are no great earthquakes, and there is no apparent correlation between seismicity and major structural features. The rate of convergence and seismicity in the Himalayas is an order of magnitude greater than in the Alps.

While the Alps are still a young mountain range, it is clear, that the active orogenesis such as is seen in the Himalayas cannot be used as a model for the Alpine region.

5.1.2 The Adriatic-European Plate Boundary

Nevertheless, it is clear that the southern Alps mark a plate boundary between the Adriatic Plate and the European (Eurasian) Plate. The Adriatic Plate has been considered either to be part of the African Plate or an entity of its own; either way, it is not part of the European Plate and its interaction with that plate is clearly significant.

It has generally been accepted that the European Plate is subducting or has subducted under the Adriatic Plate, though the situation is far from clear. The "lithospheric root" discussed in Mueller (1997) can be interpreted as a broken-off slab, now almost vertical, though this seems to be controversial.

The role of the Adriatic Plate as a rigid indenter is significant and will be returned to frequently in the course of this report. On the one hand, the fact of continental-continental collision taking place (albeit at a relatively low rate) can be viewed as a driving force in terms of seismicity. Secondly, the nature of the collision, and the rotational movement involved (Meletti et al. 2000), can be seen to be strongly influencing the local stress field. The change in direction of maximum compressive stress from the Western Alps to the Eastern Alps is well documented, and has been interpreted since the work of Pavoni (1961, 1975) as due to this cause. The precise stress pattern in the Alpine region is suggested to be an interaction between the prevailing continental stress direction and the radial pattern resulting from the rotational collision of Adria and Europe (Kastrup et al. 2002).

Since the seismicity of Switzerland appears not to be interpretable in terms of active tectonic deformation along new features, the consensus of opinion is that the dominant cause of seismicity is the reactivation of old features in a typical intraplate manner. In this case, the radial stress pattern is rather important, as the distribution of seismicity is likely to be related to the interaction of stress direction and the availability of suitably oriented structures for reactivation.

5.1.3 Alpine Foreland Tectonics

The recent tectonic history of the Alpine Foreland, including the Molasse basins and the Jura, has long been a subject of controversy amongst geologists. The debate has been summed up recently by Sommaruga (1997). The basic dichotomy is between the schools of "thin-skinned" and "thick-skinned" tectonics. The former model states that northward movement has taken place on a crustal decollement over the Alpine Foreland area without involvement of the underlying basement. The thick-skinned model holds that the basement of the Foreland region has also been affected by northward thrusting and displacement.

The weight of contemporary geological opinion, we believe, generally favours the thin-skinned school. However, it is also clear from the earthquake catalogue that seismicity is far from being concentrated within the upper crustal decollement; nor is there any concentration of seismicity at the sedimentary cover-basement interface. In fact, seismicity appears to be more significantly concentrated within the basement.

The implication is that, whatever the dominant mode (or modes) of deformation have been in the Alpine Foreland region in geologically recent times, currently, seismicity in the basement is more important. From this, it would appear that most geological data about the structure of the decollement is hardly relevant to the analysis of the seismicity. Most geological studies are principally concerned with the topmost 4 km of crust, where the seismicity is generally low. The structure of the basement between 5–15 km is not much known, yet this is where the majority of seismic energy is being released, even allowing for poor constraints on depth in the earthquake catalogue.

And yet, allowing for the fact that more seismicity seems to occur in this depth band than any other, the geographical pattern of seismicity in the basement and the detachment do not seem to be greatly different, apart from a pronounced absence of even small seismicity in a central part of the Molasse Basin in Switzerland at shallow depth (<5 km). It seems paradoxical that, if there is no involvement of the basement in the shortening processes that occur in the decollement, there should nevertheless seem to be patterns of seismicity that persist in both, and perhaps even have some correlation with formations (such as the Jura) that exist only in the decollement. One senses that the domination of the thin-skinned school of explanations may not be wholly justified.

5.2 Seismotectonic Framework

Having established the broadest outline of the model, we now proceed to look for subdivisions within this, where we can discern reasons for supposing that the seismicity within certain volumes of crust is similar, and different from neighbouring volumes of crust. A number of different data sets were considered as potentially giving information on which decisions could be based.

Obviously, the earthquake catalogue is a data set of primary importance, and to some extent the usefulness of other data sets is proportional to the extent to which they shed light on variations in seismicity. The depth distribution of seismicity is an important feature that is a function of the earthquake catalogue; particularly in the present case where major differences in depth distribution can be discerned from one area to another.

Maps of seismicity are useful both in showing the broader variations in seismicity that help to form the seismotectonic framework, and also in shaping the precise boundaries of individual zones later.

Other data sets considered included topography, geological units (considered in three dimensions), faulting, fault plane solutions, stress inversion, depth to Moho, depth to Mesozoic, subsidence/uplift measurements, in situ stress measurements, palaeoseismic data, gravity data, magnetic anomaly data, heat flow, thermal spring distribution, and others, including derived parameters such as regional variations in magnitude-frequency *b* value, derived from the earthquake catalogue.

Not all of these were found to be useful. In many cases it was considered advisable to inspect a class of data in case it seemed to shed light on the distribution of seismicity (which often it didn't), rather than because we expected it *a priori* to make a significant contribution.

The role of faults as individual sources has already been discussed; but even when approaching the modelling process entirely with a zoned approach, the use of fault data is important. In particular, one is interested in distinguishing areas where the faulting has similar characteristics. One would not normally group an area with a strong N–S trend of faulting with an adjacent area where the faulting was principally on a NW–SE trend; because it is unlikely that seismicity would be similar in both areas. In the model, observed faulting patterns (either mapped or from fault plane solutions) were used to determine expected orientation of future earthquake ruptures. Observed faults were not used directly to restrict estimates of expected maximum rupture dimensions; these were derived from estimates of maximum magnitude. However, decisions on maximum magnitude were informed by the absence of suitable structures to host exceptionally large earthquakes.

The practice of deriving stress inversions from groups of fault plane solutions overcomes any problems rising from vagaries within fault populations. Some excellent studies by Eva et al. (1997), Eva & Solarino (1998) and Kastrup (2002) show that conducting inversions for homogeneous earthquake populations can indicate the characteristic faulting type(s) and local stress field. Attempting the same thing with heterogeneous sets of earthquakes gives conspicuously poor results. The very ap-

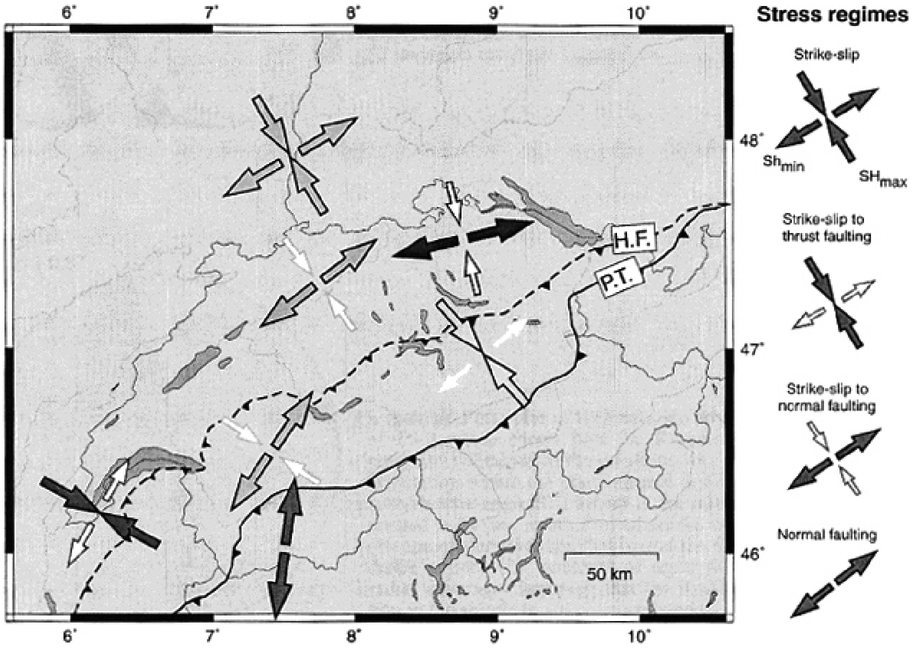

Fig. 2. Stress regimes in Switzerland, showing rotational pattern due to Alpine indentation. HF = Helvetic Front; PT = Penninic Thrust (after Kastrup 2002).

plication of this method, therefore, serves to delineate areas where the seismicity is generally consistent, and then arrives at descriptions of that consistent character. We therefore accorded this data a high priority in establishing the seismotectonic framework.

Even a simple inspection of focal mechanisms in Switzerland shows differences in deformational style by area. The study by Kastrup (2002) found that acceptable stress inversions could be obtained by making about eight local groupings of events. The resulting analysis showed a range of stress regimes (Figure 2).

In the first place, it is clear that the direction of the maximum compressive stress shows a radial pattern, going from WNW–ESE around Geneva to NNW–SSE in the Zurich area. This confirms the pattern already discussed. This pattern, which follows the curve of the Alps, is only observed close to the mountain chain. As one moves away to the north, it gradually fades into the general European stress field arising from North Atlantic ridge push (Kastrup 2002).

In the second place, the predominant styles of faulting also change, from normal faulting in the Penninic Alps to strike-slip in the Rheingraben, with areas of strike-slip to thrust (e.g. Graubünden) contrasting with strike-slip to normal regimes (e.g. Zürich).

From Kastrup's (2002) study it also appears that lines marking major crustal boundaries are also significant in dividing different faulting regimes: in particular, the Helvetic Front and Peninnic Thrust appear to have this role. To these we would add the Insubric Line, which appears to play a similar role in Eva et al. (1997).

The significance of the Helvetic Front as a divider between areas of quite different seismogenic properties is underlined by the dramatic contrast in seismicity depth profiles to south and north of it, as shown by Deichmann et al. (2000). This boundary clearly separates volumes of crust with very different earthquake populations, and therefore cannot be ignored in the formulation of a seismic model.

It is rather striking, when considering the radial Alpine stress pattern, to note that precisely where the direction of maximum compressive stress is perpendicular to the Helvetic Front (and also the important Hercynian trend) and parallel to the dip direction of the Moho, seismicity is conspicuously low both north and south of the Helvetic Front. Either side of this NW–SE band of low seismicity are bands of rather higher seismicity, again trending NW–SE, and beyond these, seismicity decreases again.

This may be a reflection of interaction between the direction of maximum compressive stress and the distribution of faults of given orientation. Whatever the explanation, the fact is that we do see a seismotectonic pattern in Switzerland that forms the basis of a seismic source partition, and, in simple terms, this consists of a series of dividing lines following the curve of the Alps, marking major crustal divisions, and a second series of dividing lines perpendicular to the first, following the direction of maximum compressive stress. The resulting pattern contrasts most strongly in seismicity rate as one moves from west to east, and contrasts most strongly in depth distribution (and possibly also b value) when moving from north to south.

Any seismic source model reflects the purpose for which it was constructed, and in this case the model was designed specifically for calculating hazard at four sites in Switzerland. The margins of the study area are relatively unimportant from the point of view of hazard and we felt it appropriate to treat them in less detail. For Northern Italy south of the Alps, our only concern was to construct the model so that the right number of earthquakes appear at roughly the right distance. Most of the territory of France that appears within the study area has been treated only in a very broad way.

The northern part of the study area is treated, so far as the seismotectonic framework is concerned, as three areas: a stable area in the east, a stable area in the west, and the Rheingraben separating them. The latter is a major structural weakness and is clearly the locus for most of the seismicity in that part of the study area north of the Alpine Foreland. The evolution of it is discussed comprehensively by Schumacher (2002), illustrating the shifting of the principal depocentres in geological time. The current regime is characterised as predominantly left-lateral strike-slip faulting with the central graben acting as a restraining bend. Young pull-apart basins are presumed to be forming in both the northern and southern graben segments. The detailed treatment of this in the zone model is discussed in the next section.

It is particularly noteworthy that there appears to be a rather higher level of seismic activity where the southern end of the Upper Rheingraben meets the Alpine Foreland. This area includes the largest earthquake in the study area (the 1356 Basel earthquake).

5.3 Source zonation

The zonation itself is based on a generalised kinematic schema, in which Italy is seen as a rigid indenter creating a radial stress pattern in the Alpine region. This creates a pattern of rings and sectors. The "rings" are progressively further from the Italian collision zone and are separated by major structural divisions (Penninic Thrust, Helvetic Front). The "sectors" are due to the rotation in the local stress regime, which interacts with the general structural grain (SW-NE) to produce a pattern of alternate zones of high and low seismicity.

The basic seismic source model consists of area source zones defined by simple polygons. Seismicity is assumed to be spatially homogeneous within these sources, with the exception of the distributed Basel source described below. The bulk of the model consists of a single set of unvarying polygonal source zones, defined in the conventional PSHA manner, the boundaries of which are firm divisions between different sets of seismicity parameters (i.e. most of the zone boundaries are not "soft" in the sense of allowing seismicity to percolate from one zone to adjoining territory).

In some source models, "soft" boundaries are used to indicate uncertainties in the location of the edges of source zones; particularly in cases where the model involves active sources surrounded by a matrix that is either discounted as aseismic or is modelled as a low-activity background area. Historically, this approach was first used to eliminate sharp "cliffs" from hazard maps (Bender & Perkins 1987), and its carry-over into site studies has not always found support (to judge from some informal conversations in the wider hazard community).

In this model, the source zones constitute a complete tessellation of the area of study, and we consider such an approach of less relevance, as it has the general effect of scattering earthquakes in both directions across any source zone boundary. To some extent many of the source boundaries have uncertainty inasmuch as we could postulate numerous minor variations in geometry; indeed, many such refinements were made in the course of the development of the model. Such variations would better express the uncertainty in the boundary positions than an ad hoc application of soft boundary methods; however, our consideration was that incorporating such uncertainties would add hugely to the complexity of the model with very little actual benefit in terms of results.

The basic model is shown in Figure 3, together with the seismicity. It is apparent to casual inspection that the zones satisfy fairly well the requirement of PSHA that zones are homogeneous as regards earthquake occurrence; this was further tested and checked using nearest neighbour analysis (Musson 2000, 2004b).

5.4 Zone models and logic trees

The last point leads on another general issue, the use of alternative source zone models in the logic trees used in PSHA to represent uncertainty and imperfect knowledge. Exactly how multiple source models should be prepared is not something much discussed in the literature (in fact, as has often been commented, there is not much guidance in the literature in preparing source zone models at all).

It is clear from experience that if one gives the same data to two hazard analysts, two different source models will result. There is nothing necessarily wrong with this; it means that more than one different interpretation of the data is possible. It was interesting to see, in the course of this project, the different models prepared by the different Expert Groups from the same data. Points of similarity but also contrast can be found in all. We would regard the EG1c model as the most "classical" in style, consisting as it does of relatively simple polygons. One can debate as a matter of practice how abstracted source zones should be. At the one extreme, using a very few, highly geometrical zones, risks oversimplifying the seismicity to such an extent that the result is completely unrealistic. At the other extreme, elaborate zone boundaries based on the surface expression of geological units may have no real correspondence to seismogenic processes in the lower crust.

One can think of five bases for including variant model geometries in a source model.

The first would be the attempt to include totally different interpretations; as if, say, EG1c were to try to guess how EG1a were thinking, and produce a completely different set of source zones on different principles. This is difficult, and certainly re-

Fig. 3. The basic zone model of EG1c; shown also is the seismicity ≥ 3.5 Mw since 1750.

dundant in PEGASOS where in any case, four different models will be combined from the four EGs.

The second would be to try and include all the minor variations possible within one interpretation. For example, in the development of the EG1c model, the exact configuration of the zones around the Jura was redrawn several times. One could regard each redrawing as a variant with its own validity, but by and large, we considered the successive variants to be improvements rather than alternatives.

The third case is where specific questions arose which could not be settled in discussion. For example, is the Rhine Graben a single homogeneous seismogenic unit or is it not? If this is a question without a clear answer, it can be handled by devising variant models where in one variant the Rhine Graben is unified and the other it is divided.

The fourth case is where the use of multiple models is used to obtain an effect that can be achieved by overlaying different geometries. For example, overlaying a set of broadly defined zones and a second set of smaller zones is a way of partially smoothing the seismicity within a region with recognised centres of activity (Musson 1997).

Fifth is a set of circumstances where different model geometries are applied to different subsets of the seismicity. The obvious example is where seismicity varies with depth, and one set of zones models shallow seismicity, with one or more ancillary sets of zones beneath. In this case the zone geometries merely overlap; they are not weighted alternatives. Another example would be the use of one set of zones for moderate magnitude seismicity and another one for larger earthquakes, if one believed that the spatial distribution of larger events was distinct. According to Woo (1996) zone models are necessarily independent of magnitude; this is not the case at all.

In the EG1c model for PEGASOS, we used variants for the model geometry only for the third reason. There were three instances, the first of which was the Rhine Graben partitioning already mentioned. The second involved the curious concentration of seismicity in the Swabian Jura, which has seen intense seismicity since 1911 but with no previous historical precedent. There is no good geological reason for this localisation of activity (it coincides with the Hohenzollern Graben, but is entirely beneath this rather shallow feature). One could posit either (a) there is some unknown seismogenic feature at the location of the 1911 earthquake that periodically becomes active in the same place, or (b) such a concentration of seismicity is not tied to one geological feature and could recur at almost any place in the broader Swabian area. This leads to two model variants, one where the Swabian Jura seismicity is represented by a small separate zone, and one where the 1911 and post-1911 seismic sequence is merged into the surrounding zone. The last case involved the Permo-Carboniferous grabens on the Swiss-German border which arguably are a significant seismic source, and arguably are not, leading to two variant zone configurations. These alternatives are shown in Figure 4. They combine to give eight possible configurations in all.

6 Assessing model parameters

We now discuss some issues relating to assessing the parameters defining the seismicity in each zone. The conceptual framework we adopted was one intended to combine a series of approaches ranging from the essentially general, in which divisions between the zones are minimised, and the particular, in which the seismicity parameters in each zone are heavily dependent on data local to that zone.

This developed into a system where the logic tree was based around three main branches, most easily categorised according to the procedure used for assessing maximum magnitude (M_{max}). The logic tree we used combines methods for estimating M_{max} with methods for assessing seismicity rates; each branch contains one M_{max} technique and one seismicity rate/b value technique. These are shown in Figure 5.

The first approach is to assess M_{max} in a very general way. One can show a number of cases worldwide (especially in an intraplate environment) where approaches to estimating magnitude M_{max} have failed (or would have failed), because recent earthquakes have occurred with magnitude larger than what might have been previously assumed using geological or seismological indicators. To choose an example at random, a study by Al-Tarazi (1999) estimated a single maximum magnitude value for the Gulf of Aqaba of 5.7 ML, based on statistical

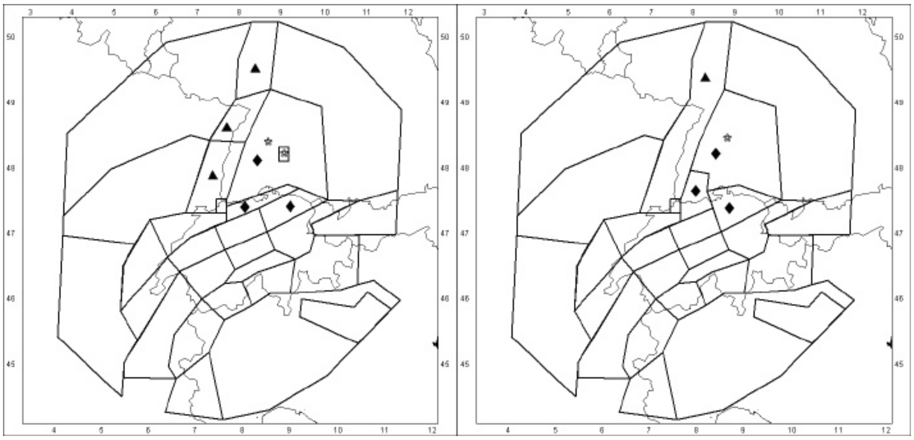

Fig. 4. Variations in the EG1c zone model. Symbols show the zones that can be varied independently in different versions of the model: triangles = zones affected by Rhine Graben partition; lozenges = zones affected by Permo-Carboniferous graben inclusion; open star = zones affected by Swabian Jura modelling.

analysis of an earthquake catalogue closing in 1993. In 1995 an earthquake of magnitude 7.2 Ms occurred in this locality.

We therefore created a branch in the logic tree with allowance for a set of "global" values for M_{max} being 6.5, 7.0 or 7.5 Mw, these values applying to all zones equally. So in this global branch it is believed that M_{max} is most likely 7.0 Mw anywhere in the region (i.e. 7.0 has the highest weighting), without taking into account the local features, and with a smaller probability even as high as 7.5 Mw. The presence of this branch in the logic tree is intended to cover the pessimistic possibility of an anomalously large event on some unknown feature that might strike anywhere.

It will be noted that one part of this branch supposes that nowhere in the study area will any earthquake exceed 6.5 Mw, which may seem strange when the ECOS (2002) catalogue includes an earthquake larger than this. We are taking into consideration the fact that the magnitude values of medieval earthquakes are inherently uncertain, and that the largest historical Swiss earthquake may perhaps not have been larger than 6.5 Mw whatever the "best-estimate" value in the catalogue is.

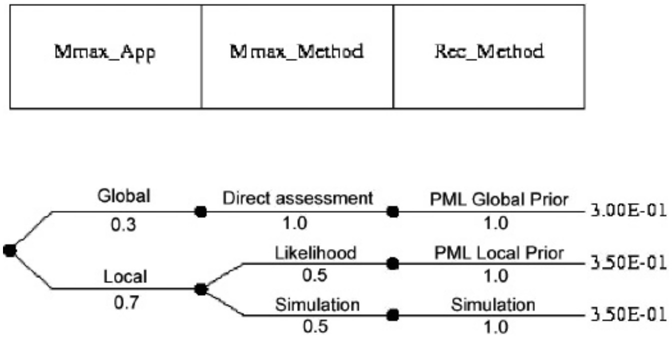

Fig. 5. Part of the overall logic tree for the study showing the weights assigned to the maximum magnitude approach, the method used for maximum magnitude assessment, and the method used for assessing magnitude recurrence (figure courtesy of RR Youngs).

This approach to M_{max} was combined with a penalised maximum likelihood approach to seismicity parameters (Weichert 1980; Johnston et al. 1994). For all zones in all parts of the model we used a truncated linear Gutenberg-Richter model, with truncation applied as a sharp cut-off. Examination of seismicity data did not suggest that there was evidence for seismicity in any zone not to follow this model. Since this first branch of the model was intended to treat the hypothesis that seismicity across the whole of Switzerland is essentially similar, in this branch, although seismicity parameters were determined for each zone individually, they were all determined using the same prior, which was derived from the entire catalogue. In some PSHA studies it is the practice to use the same *b* value derived from the total catalogue for every zone and vary only the activity rate; this we regard as too great a level of abstraction and liable to lead to unrealistic cases. This combination of M_{max} and seismicity parameters is the top path in Figure 5.

The two other branches of the logic tree shown in Figure 5 present source-specific approaches to maximum possible magnitude; M_{max} was estimated in each source zone individually. Some general limits are set for both branches (Figure 6). For each zone, M_{max} was never considered to be less than 5.5 Mw or the largest historical earthquake in the zone (rounded up to the nearest half-magnitude), whichever was the larger. Also M_{max} was never allowed to be more than 7¼ Mw (it is regrettable to write this as 7.25; the decimal system has disadvantages when the inherent inaccuracy of data makes working in quarter-units necessary). This limit was based on a combination of a slight increase on the largest observed historical event in the whole catalogue, general judgement on maximum observed earthquakes in corresponding areas, and a lack of significant structures that would be reasonable to expect very large earthquakes to occur on. The slight increase on this limit in the first branch is specifically designed to provide an extra margin of conservatism in the overall model.

Within these limits, in the second branch of the logic tree in Figure 5, M_{max} was calculated using a simple maximum likeli-

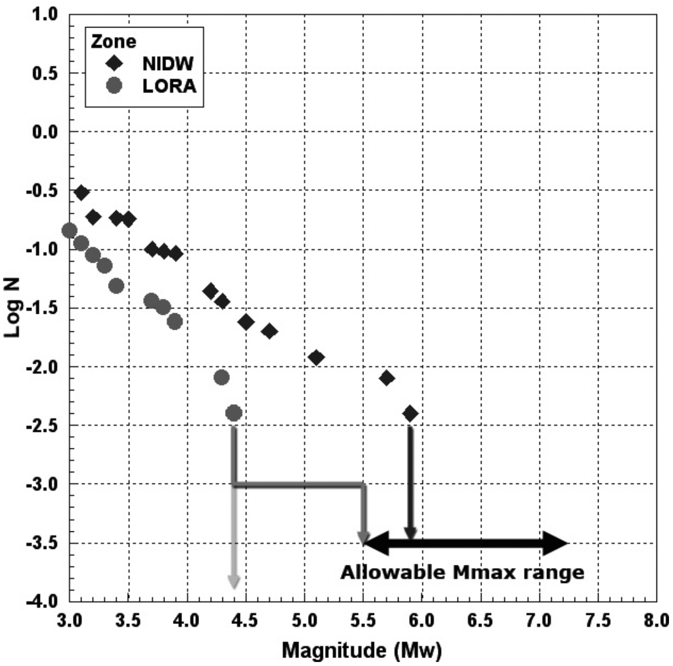

Fig. 6. Limits on zone-specific maximum magnitudes; example showing two zones. For NIDW (Nidwalden) M_{max} may not be less than the historical maximum (5.9 Mw). For LORA (Lorraine), the historical maximum (4.4 Mw) is less than the globally allowed smallest value of M_{max}, set to 5.5 Mw. Therefore M_{max} is considered to vary between 5.5 and 7.25 Mw for LORA, and between 5.9 and 7.25 for NIDW.

hood approach with no prior, taking into account the historical completeness values for the zone. Other studies that have used a maximum likelihood approach (e.g. Wahlström & Grünthal 2001) have usually restricted the results by using a prior derived from observations from similar crustal types. We preferred not to follow these examples, but to rely entirely on the local data and accept the degree of uncertainty in the results that this decision entails. The imposition of an upper bound keeps the distribution of results within desired limits and preserves the desired shape.

In this second branch, seismicity parameters were again estimated by the maximum likelihood method, but, in contrast to the first branch where a single global prior was used, in this case local priors were used for each zone, emphasising the differences rather than the similarities. These local priors were derived from least-squares solutions to the magnitude-frequency data for each zone.

The final branch provides a joint determination of M_{max}, activity rate and b value. This method is explained in detail in Musson (2004c), and will not be gone into in great detail here. The method relies on selecting possible seismicity parameters for a zone at random and attempting to generate synthetic earthquake catalogues (subject to the same historical completeness constraints) that match the real earthquake catalogue within an acceptable tolerance level. Values that give successful matches are noted, and the logic tree is ultimately compiled from this distribution.

Considering only M_{max}, a simplified illustration can be given. Suppose that activity rate and *b* value are known. Choose a value for M_{max} at random, generate a synthetic catalogue subject to historical constraints, and check whether an earthquake occurred larger than the maximum historical observed earthquake. If the historical maximum was not exceeded, note the M_{max} value. Proceed until 5,000 successes have been recorded. The distribution of M_{max} values for the logic tree is constructed from the distribution of 5,000 successful values.

Conceptually, this method for M_{max} is similar to the maximum likelihood approach with a flat prior (the same as no prior). It has the advantage that it estimates all the seismicity parameters at the same time. Both our zone-specific approaches to maximum magnitude attempt to answer the same question, "Given the historical catalogue and the constraints upon it, what is the likelihood that events that are x magnitude units larger than the observed magnitude are possible, yet never happened in historical times?" One approach seeks to address this analytically, the other approach experimentally.

In terms of activity rate and *b* value, the attraction of this method is that it is entirely driven by the data within each zone and requires no assumptions or subjectivities. The output determines even the number of logic tree branches to be used, as well as their values and weights. Also, there is no danger that the values derived are not entirely consistent with the historical data.

The output of the joint inversion is a series of branches that relate to triplet combinations of M_{max}, activity rate and *b* value. One could use these directly, but in practice, although *b* value and M_{max} are correlated, the correlation is weak, and it is adequate to separate out the weights for the different M_{max} values in order to reduce the total number of logic tree branches.

7 Conclusions

This paper has presented some aspects of the EG1c seismic source model for PEGASOS. Partly for reasons of space, many issues and details are not given here, but also because most of the details, while critical for calculating the hazard results in this project, are not so interesting to the general reader. We have instead concentrated more on the philosophical aspects of the way in which the model was constructed over a three year process, and the thought processes at work. These broader issues are relevant to a wide range of PSHA situations, and whether or not the reader agrees with the decisions that were adopted by us in this study, as seismic hazard is still a relatively young discipline, we consider that exposing methodological issues to discussion in the general literature is helpful to the development of practice.

Acknowledgements

The authors would like to express their thanks to the PEGASOS TFI team for their support in preparing the model, especially to Kevin Coppersmith and Bob Youngs. We would also like to thank Jim Farrington for all the hard work he put into the day-to-day management of the project, and also to his colleagues in Proseis, Zurich. The PEGASOS project was funded by swissnuclear, whose support, and that of NAGRA, is acknowledged. Two informal project meetings were hosted by British Geological Survey and Landesamt für Geologie, Rohstoffe und Bergbau, Baden-Württemberg. The contribution of Roger Musson to this paper is by permission of the Director of the British Geological Survey (NERC). The contribution of Wolfgang Brüstle to this paper is by permission of the Director of the Landesamt für Geologie, Rohstoffe und Bergbau. It should be noted in passing that although this paper is published in 2008, the work reported on was completed in 2002 and this paper was drafted at the beginning of 2005. It has not been attempted to report extensively on developments postdating the original PEGASOS project. Reviews by Stefan Schmid, Jean-Pierre Burg and Dario Slejko improved the text of this paper.

REFERENCES

Abrahamson, N. 2000: State of the practice of Seismic Hazard Evaluation. Proceedings of GeoENG. 19–24 Nov 2000.

Abrahamson, N.A. Birkhäuser, P., Koller, M., Mayer-Rosa, D., Smit, P., Sprecher, C., Tinic, S., & Graf, R. 2002: PEGASOS – A comprehensive probabilistic seismic hazard assessment for nuclear power plants in Switzerland. 12th European Conference on Earthquake Engineering, London, 8 pp.

Al-Tarazi, E. 1999: Regional seismic hazard study for the eastern Mediterranean (Trans-Jordan, Levant and Antakia) and Sinai region. Journal of African Earth Sciences 28, 743–750.

Bender, B.K. & Perkins, D.M. 1987: SEISRISK III: a computer program for seismic hazard estimation. USGS Bulletin 1772.

Bernreuter, D.L., Savy, J.B., Mensing, R.W. & Chen, J.C. 1989. Seismic hazard characterisation of 69 nuclear power plant sites east of the Rocky Mountains. NUREG/CR-5250, US Nuclear Regulatory Commission.

Budnitz, R.J. Apostolakis, G., Boore, D.M., Cluft, L.S., Coppersmith, K.J., Cornell, C.A. & Morris, P.A. 1997: Recommendations for probabilistic seismic hazard analysis: guidance on uncertainty and use of experts. U.S. Nuclear Regulatory Commission, NUREG/CR-6372.

Chandra, U. 1978: Seismicity, earthquake mechanisms and tectonics along the Himalayan mountain range and vicinity. Physics of the Earth and Planetary Interiors 16, 109–131.

Coppersmith, K.J. & Youngs, R.R. 1986: Capturing uncertainty in probabilistic seismic hazard assessments within intraplate tectonic environments. Proceedings of the Third US National Conference on Earthquake Engineering, Charleston, 301–312.

Deichmann, N., Baer, M., Braunmiller, J., Ballarin Dolfin, D., Bay, F., Delouis, B., Faeh, D., Giardini, D., Kastrup, U., Kind, F., Kradolfer, U., Kuenzle, W., Roethlisberger, S., Schler, T., Salichon, J., Sellami, S., Spuehler, E. & Wiemer, S. 2000: Earthquake in Switzerland and surrounding region during 1999. Eclogae Geologicae Helvetiae 91/2, 395–406.

ECOS 2002: ECOS: earthquake catalogue of Switzerland. PEGASOS technical report EXT-TB-0032, SED.

Eva, E., Solarino, S., Eva, C. & Neri, G. 1997: Stress tensor orientation derived from fault plane solution. Journal of Geophysical Research 102 (B4), 8171–8185.

Eva E. & Solarino S. 1998: Variations of stress directions in the western Alpine Arc. Geophysical Journal International 135/2, 438–449.

Frankel, A., Mueller, C., Barnhard, T., Perkins, D., Leyendecker, E.V., Dickman, N. Hanson, S., & Hopper, M. 1996: National Seismic Hazard Maps: Documentation June 1996. U.S. Geological Survey Open-File Report 96–532, 44 pp.

Hough, S.E., J.G. Armbruster, L. Seeber, & J.F. Hough 2000: On the Modified Mercalli intensities and magnitudes of the 1811–12 New Madrid earthquakes. Journal of Geophysical Research 105, 23839–23864.

Jacob, K., J. Armbruster, N. Barstow, & S. Horton 1994: Probabilistic ground motion estimates for New York: comparison with design ground motions in national and local codes. In: Proceedings of 5th U.S. National Conference on earthquake Engineering, Chicago 3, 199–128.

Johnston, A.C., Coppersmith, K.J., Kanter, L.R. & Cornell, C.A. 1994: The earthquakes of stable continental regions. TR-102261-V4, Electric Power Research Institute, Palo Alto.

Kastrup, U. 2002: Seismotectonics and stress field variations in Switzerland. PhD. Thesis ETHZ 14527, 153 pp.

Marin, S., Avouac, J.-P., Nicolas, M. & Schlupp, A. 2004: A probabilistic approach to seismic hazard in Metropolitan France. Bulletin of the Seismological Society of America 94: 2137–2163.

Meghraoui M., Delouis, B., Ferry, M., Giardini, D., Huggenberger, P., Spottke, I. & Granet, M. 2001: Active Normal Faulting in the upper Rhine Graben, Science 293, 2070–2073

Meyer B., Lacassin, R., Brulhet, J. & Mouroux, B. 1994: The Basel 1336 earthquake. Which fault produced it? Terra Nova 6, 54

Meletti, C., Patacca E. & Scandone P. 2000: Construction of a Seismotectonic Model: The case of Italy. Pure and Applied Geophysics 157, 11–35

Mueller, S. 1997: The lithosphère-athenosphere System of the Alps. In: Pfiffner, O.A. et. al. (Eds.): Deep Structure of the Alps. Birkhäuser Basel, 338–347.

Musson, R.M.W. 1997: Seismic Hazard studies in the U.K.: Source Specification Problems of Intraplate Seismicity. Natural Hazards 15, 105–119.

Musson, R.M.W. 2000: Evaluation of seismic hazard source models. In: Lapajne, J., & Vidrih, R. (Eds): Seismicity modelling in seismic hazard mapping. Slovenian Geophysical Survey, Ljubljana, 53–66.

Musson, R.M.W. 2004a: Comment on "Communicating with uncertainty: A critical issue with probabilistic seismic hazard analysis". Eos 85(24): 235–236.

Musson, R.M.W. 2004b: Objective validation of seismic hazard source models, 13th World Conference on Earthquake Engineering, Vancouver, paper 2492.

Musson, R.M.W. 2004c: Joint solution of seismicity parameters for seismic source zones through simulation. Bolletino di Geofisica Teorica ed Applicata 45: 1–13.

Musson, R.M.W. & Winter, P.W. 1996: Seismic hazard of the UK. AEA Technology Report No AEA/CS/16422000/ZJ745/005.

Musson, R.M.W. & Winter, P.W. 1997: Seismic hazard maps of the U.K.. Natural Hazards 14, 141–154.

NAGRA 2004: Probabilistic Seismic Hazard Analysis for Swiss Nuclear Power Plant Sites (PEGASOS Project) prepared for Unterausschuss Kernenergie der Überlandwerke (UAK). Final Report Vol. 1–6, 2557 pp. To be obtained on request at swissnuclear by writing to info@swissnuclear.ch.

Pavoni, N. 1961: Faltung durch Horizontal Verschiebung. Eclogae Geologicae Helvetiae 54, 515–534

Pavoni, N. 1975: Zur Seismotektonik des Westalpenbogens. Vermessung Photogrametrie, Kulturtechnik III/IV-75, 185–187.

Schumacher, M.E. 2002: Upper Rhine Graben: Role of pre-existing structures during rift evolution. Tectonics 21, 1006, doi:10.1029/2001TC900022.

Singh, R.P., Li, Q., & Nyland, E. 1990: Lithospheric deformation beneath the Himalayan region. Physics of the Earth and Planetary Interiors 61, 291–296.

Sommaruga, A. 1997: Geology of the central Jura and the molasses Basin. Mémoires de la Société des Sciences Naturelles de Neuchâtel 12, 176 pp.

Thouvenot, F., Fréchet, J., Tapponnier, P., Thomas, J.-Ch., Le Brun, B., Ménard, G., Lacassin, R., Jenatton, L., Grasso, J.-R., Coutant, O., Paul, A. & Hatzfeld, D. 1998: The ML 5.3 Épagny (French Alps) earthquake of 1996 July 15: a long-awaited event on the Vuache fault. Geophysical Journal International 135, 876–892.

Ustaszewski, K., & Schmid, S.M. 2007: Latest Pliocene to recent thick-skinned tectonics at the Upper RhineGraben – Jura Mountains junction. Swiss Journal of Geosciences 100, 293–312.

Wahlström, R. & Grünthal, G. 2001: Probabilistic Seismic Hazard Assessment for Fennoscandia using the logic tree approach for regionalisation and non-regionalization models. Seismological Research Letters 72, pp. 33–45.

Weichert, D.H., 1980: Estimation of the earthquake recurrence parameters for unequal observation periods for different magnitudes. Bulletin of the Seismological Society of America 70, 1337–1346.

Woo, G. 1996: Kernel Estimation Methods for Seismic Hazard Area Sources. Bulletin of the Seismological Society of America 86/2, 353–362.

Manuscript received June 5, 2008
Revision accepted September 9, 2008

Editorial Handling: Stefan Bucher