Regime transitions and heteroclinic connections in a barotropic atmosphere

D. T. Crommelin[1,2]

3rd June 2002

[1] Royal Netherlands Meteorological Institute
P.O. Box 201, 3730 AE,
De Bilt, The Netherlands
[2] Mathematical Institute, Utrecht University
The Netherlands
E-mail: crommeli@knmi.nl

Submitted to the Journal of the Atmospheric Sciences

Abstract

By interpreting transitions between atmospheric flow regimes as a deterministic rather than a stochastic phenomenon, new insight is gained into the phase-space characteristics of these transitions. The identification of regimes with steady states should be extended with the association of transitions with nearby heteroclinic connections between steady states, as known from the theory of dynamical systems. In the context of a T21 barotropic model of the northern hemisphere, which possesses regime behaviour, steady states are found that correspond with regimes, and heteroclinic connections are approximated using a new algorithm based on adjoint modelling techniques. A 200 year dataset generated by the model is shown to posses spatial preferences in its transitional behaviour that match well with the approximated heteroclinic connections.

1 Introduction

The concept of atmospheric circulation regimes is an old idea that is still much studied in meteorology and climate science. The interest in the notion that the atmosphere can be caught, for some time, in particular flow configurations has far from subsided. Apart from the relevance for meteorological studies, the concept of regime behaviour now also draws attention from studies of climate variability and predictability. This is due to the belief that climate variability is related to changes in the probabilities of regime visits (Corti et al. 1998, Palmer 1999). Many studies have been devoted to the question if and where regimes exist in observations and in complex atmosphere/climate model data. Also, the connection of regimes to known (but still ill-understood) climate variability patterns, such as the North-Atlantic Oscillation (NAO), the Arctic Oscillation (AO) and the Pacific/North American (PNA) pattern, is a much debated issue. Surprisingly, the investigation of the dynamical origin of regime behaviour seems to have slowly moved out of sight.

The concept of global bifurcations is an instrument that may enable us to get a better grip on the phenomena of regime behaviour and low-frequency variability and their underlying dynamical nature. The notion that individual regimes can be associated with separate (simple) attracting structures in the phase space of atmospheric flow exists since Charney and Devore (1979). They identified stable equilibria (steady states, fixed points) with regimes, a hypothesis that was followed and expanded on by e.g. Reinhold and Pierrehumbert (1982), Legras and Ghil (1985), De Swart (1988a,b, 1989) and Itoh and Kimoto (1996). Even though some of these studies speculate on periodic solutions or tori to be the underlying mathematical structures of regimes, the basic idea remains the same: regime behaviour can be identified with temporary stays on or near simple attractors, followed by irregular transitions between these attracting sets.

The multiple attractors have been shown to be amenable to bifurcation analysis by e.g. Legras and Ghil (1985), De Swart (1988a,b, 1989) and Itoh and Kimoto (1996). Continuation and bifurcation software packages such as AUTO (Doedel, 1986) are able to follow steady states and periodic solutions through parameter space. They were used to shed light upon the creation and structure of the simple solutions identified with regimes in low-order atmosphere models. Much more difficult however are the mechanisms of transitions between these solutions. That may be the reason that the, from a dynamical systems point of view most natural possibility of *global bifurcations*, i.e. of deterministic connections being created between different solutions (or from a solution to itself), hasn't gotten much attention yet. The most common transition mechanism invoked is stochasticity: transitions induced by stochastic perturbations in the model equations. Introducing such perturbations is usually motivated by referring to physics and dynamics (e.g. small-scale, high frequency motions or convection) that are not explicitly resolved in the relatively simple models used for bifurcation studies in relation to regime transitions. However, introducing stochastic terms or interpreting small-scale, high-frequency behaviour in complex, deterministic atmosphere models as noise is not entirely satisfying. Ignoring the deterministic origin of such stochasticity may obscure relevant information about regime behaviour. In this paper we will hold on to this deterministic nature, and use the notion of heteroclinic

connections to get a better understanding of regime transitions.

Heteroclinic orbits are solutions of (deterministic) model equations that connect one invariant set, e.g. a steady state solution, with another. Homoclinic orbits connect an invariant set with itself. Hetero- resp. homoclinic bifurcations are bifurcations in which such orbits, or connections, are created. They are mentioned a few times in the literature as a possible mechanism for transitions between regimes. Legras and Ghil (1985) use a 25-dimensional equivalent-barotropic model, perform a bifurcation analysis of its steady states and find, by numerical integration, transitions between regimes. Heteroclinic orbits as a possible explanation are mentioned, but not expanded on. In De Swart (1988a,b, 1989) a 6-component spectral model was found to possess multiple attractors, some of them with chaotic windows (parameter intervals with chaotic behaviour), that could be identified with regimes. However, transitions between these attractors (vacillation) were not observed. By adding stochastic perturbations, a transition mechanism was created (De Swart and Grasman, 1987). Also, by increasing the number of spectral components to 10, the model could be brought to vacillatory behaviour, without stochastic perturbations. But no attempt was made to find heteroclinic orbits. Itoh and Kimoto (1996,1997,1999) perform bifurcation analyses of a two-layer T15 model and of a five-layer T21 model. They find multiple attractors and transitions between them, and call this wandering between different attractor ruins an example of "chaotic itinerancy". It is not made clear if and how this chaotic itinerancy is related to the concept of homo- and heteroclinic connections.

The notion that regime transitions are not entirely random is recognised in, amongst others, Mo and Ghil (1988), Kimoto and Ghil (1993b) and Itoh and Kimoto (1996,1997,1999). These studies identify more than two regimes and point out that some transitions between regimes are more probable than others. Using this statistical approach, cycles of most probable transitions are identified; these are called "preferred transitions". Itoh and Kimoto (1997) refer to the notion of "explosive bifurcations", but they don't try to locate heteroclinic connections, nor do they mention them. In the present study, something different is meant with the notion "preferred transitions": it refers to the idea that transitions between regimes tend to follow certain routes through phase space; routes that can be associated with approximate heteroclinic connections. The focus is on preferences of the behaviour *during* the transitional episodes, not on the most probable order of regime visits without looking at characteristics of the transition phases, as in Mo and Ghil (1988), Kimoto and Ghil (1993b) and Itoh and Kimoto (1997).

The study by Plaut and Vautard (1994) shows, performing Multichannel Singular Spectrum Analysis (MSSA) on observational data, that regime behaviour is connected to the existence of low-frequency oscillations (LFOs) in the atmosphere. LFOs influence regime behaviour but not in a systematic way; rather, the interaction between two Atlantic sector LFOs (one of 70 days and one of 30-35 days) seems to be one mechanism, among others, that can produce high-amplitude anomalies (regimes). It remains unclear, though, whether these anomalies correspond to (quasi-) stationary states (thus, if the system evolution slows down). The authors hypothesize that LFOs should be associated with the presence of unstable periodic orbits in phase-space. An important conclusion is that regime transitions, at least in the Atlantic sector, are not produced randomly by synoptic transients. This

supports the choice for the deterministic (instead of stochastic) point of view taken in the current paper.

In Crommelin (2002), the appearance of homoclinic dynamics in an atmosphere model gets more detailed attention. In this paper, a two-layer model is analysed which uses Empirical Orthogonal Functions (EOFs) instead of spherical harmonics as its basis functions (see Achatz and Branstator (1999) for the derivation and formulation of the model). With 10 EOFs, some realism was still retained and at the same time it was possible to find evidence for the occurrence of homoclinic dynamics. The dynamics on long timescales turned out to be dominated by a homoclinic orbit of the *bifocal* type: an orbit attached to a steady state with two complex pairs as its leading eigenvalues. A glimpse of a heteroclinic connection between two equilibria was also found, but the role of one of the equilibria could only be made visible for small parameter regions. The other equilibrium, very close to the climatic mean state, is prominent in the model behaviour over large parameter ranges. The behaviour was therefore identified as driven by a homoclinic, not a heteroclinic connection.

In the fluid dynamics community, the possibility of heteroclinic connections as an explanation for transitions between steady states has been recognised before. Examples were found in e.g. Rayleigh-Benard convection (Proctor and Jones, 1988) and the turbulent boundary layer (Aubry et al. (1988), Armbruster et al. (1988), Holmes et al. (1996,1997)). Knobloch and Moehlis (2000) provide an overview of mechanisms that produce bursting in hydrodynamical systems (sudden transitions from and to regular, repeatedly realized states); again, heteroclinic connections play an important role. They interpret transitions between different types of behaviour as a form of bursting. Although no examples are mentioned from geophysical fluid dynamics, such a concept of bursting could equally well apply to the case of regime-transitions in atmospheric dynamics. The methods and concepts used in the study of these fluid-dynamical examples could be very useful in the investigation of regime behaviour and low-frequency variability of the atmosphere. This paper is partly inspired by these examples.

The order of things to come in this paper is as follows. In section 2 the barotropic model that was used will be described. The model shows bimodality and regime behaviour, treated in section 3. The regimes can be associated with steady states; this is described in section 4. In section 5 the role of heteroclinic connections will be explained. A new algorithm, based on adjoint modelling techniques, to calculate approximations of such connections will also be presented in this section, together with its results. Section 6 contains a discussion and summary.

2 A barotropic model

For this study a barotropic spectral model was used. It is the same model that was used in Selten (1995), a standard model with realistic orography and a forcing that was calculated from observations in order to get a realistic climate mean state and realistic low-frequency variability.

The model is based on the barotropic vorticity equation, and reads:

$$\frac{\partial}{\partial t}\xi = -\mathcal{J}(\psi, \xi + f + h) + k_1\xi + k_2\Delta^3\xi + \xi^* \tag{2.1}$$

This partial differential equation describes the temporal evolution of the relative vorticity field $\xi(t, \lambda, \phi)$, a scalar field, on the sphere. The spatial coordinates are the longitude λ and latitude ϕ. The streamfunction $\psi(t, \lambda, \phi)$ is related to the vorticity via $\xi = \Delta\psi\ (= \nabla^2\psi)$. The terms multiplied by k_1 and k_2 are damping terms (Ekman friction resp. scale-selective damping); f denotes the Coriolis parameter and h is the (nondimensional) orographic height. $\mathcal{J}$ denotes the Jacobi operator, which reads $\mathcal{J}(A, B) = \frac{1}{\cos\phi}\left(\frac{\partial A}{\partial\lambda}\frac{\partial B}{\partial\phi} - \frac{\partial A}{\partial\phi}\frac{\partial B}{\partial\lambda}\right)$. The (time-independent) vorticity forcing ξ^* is based on observations:

$$\xi^* = \mathcal{J}(\psi_{\rm cl},\, \xi_{\rm cl} + f + h) - k_1\xi_{\rm cl} - k_2\Delta^3\xi_{\rm cl} + \overline{\mathcal{J}(\psi',\, \xi')}. \tag{2.2}$$

This procedure was proposed by Roads (1987). It involves the use of the observed climate mean states at 500 hPa of the winter season $\xi_{\rm cl}$ and $\psi_{\rm cl}$. The last contribution to the forcing, due to the transient eddy forcing, is calculated using (winter) anomalies of the 10-day running mean at 500 hPa, ξ' and ψ'. The overbar denotes, as usual, time average.

Equation (2.1) is projected onto spherical harmonics and triangularly truncated at wavenumber 21 (i.e. T21), yielding a coupled set of nonlinear ordinary differential equations for the time-dependent spectral coefficients. By discarding all modes with zonal wavenumber plus total wavenumber being even, a model of hemispheric flow is obtained with a total number of 231 variables. It will serve to describe the (barotropic) atmosphere of the northern hemisphere. For more detailed information on the formulation of the model the reader is referred to Selten (1995).

The model has been integrated for 73 000 days (200 years), using a Runge-Kutta fourth-order routine with a 30 minutes timestep and daily output. The initial state was a realistic flow field, so no spin-up time was needed. The timescale of the Ekman damping was set to 15 days; the scale-selective damping was tuned such that the damping timescale is 3 days for wavenumber 21. Figure 1 contains the time-mean and variability (root mean square, or rms) of the thus obtained 200 year dataset. Shown are the mean and rms of the 500 hPa geopotential height Z_{500}, which is calculated from the streamfunction (at 500 hPa) using the so-called linear balance equation $g_0\,\nabla^2 Z_{500} = \nabla\cdot(f\,\nabla\psi)$, $g_0 = 9.80665\,\mathrm{m\,s^{-2}}$ (see Holton, 1992). The mean resembles the mean field of the real atmosphere quite closely. The rms is not as close to the observations, but still quite realistic considering the model resolution. Most notable are the unrealisticly high variability above the north pole and the location of the rms-maximum over Europe, which should be over the north-east Atlantic. For a more detailed comparison with observations, see Selten (1995). The 200 year dataset will be analysed in order to get a better understanding of its low-frequency aspects.

3 Bimodality

EOFs of the streamfunction data are calculated; in figure 2 the streamfunction patterns corresponding to the leading four EOFs are shown. For the calculation, the time mean

state was subtracted from the data; the norm chosen was the kinetic energy norm (see appendix A). The amount of kinetic energy variance represented by EOFs 1–4 is 18.3, 9.8, 6.8 and 6.0 percent, respectively; together these four represent 40.9 %. The principal components (PCs) were calculated by projecting the (spectral) data onto the EOFs.

The shape of the two-dimensional Probability Density Function (PDF), made from the projection of the full 73000 days dataset onto the PC1,PC2–plane, reveals a bimodal structure with a main maximum at PC1 ≈ 0.0055 (MAX1) and a second maximum at PC1 ≈ -0.0089 (MAX2). The PDFs made from projections onto the PC2,PC3–, the PC2,PC4– and the PC3,PC4–plane (not shown) yielded unimodal distributions; the bimodal structure mainly projects onto EOF1. Figure 3 (dashed lines) shows the PDF in the PC1,PC2-plane; the various symbols in the figure will be explained later. The PDFs were calculated by collecting the projected datapoints in 15×15 bins and calculating the corresponding histogram.

The locations of the two maxima of the two-dimensional PDF were estimated in all 231 PCs by first identifying the PC1 values of the two maxima. Then, two sets of datapoints were created by selecting all datapoints with their PC1 value in a small band around the two maxima. For all higher PCs two one-dimensional PDFs and their maxima were calculated, thus giving estimates for the location of the two maxima in all PCs. The Z_{500} patterns corresponding to the maxima are shown in figure 4 (note that these patterns consist of contributions of all 231 PCs). The distances to these two maxima during the model run were calculated, using again the kinetic energy norm. For a segment of the data, the timeseries of these distances are shown in figure 5. They show persistent periods close to the maxima as well as rapid transitions between them.

The PDF-maxima MAX1 and MAX2 can be interpreted as flow regimes, the transitions between them as vacillatory behaviour. The flow-pattern (figure 4) of MAX1 represents a situation of intensified zonal flow, the pattern of MAX2 shows a blocked flow over Europe (a split jetstream and a positive anomaly, or "high", over Europe). Both regimes have previously been found to correspond with steady states in low-order models (Charney and DeVore (1979), Legras and Ghil (1985)). With this correspondence in mind, one might expect the presence of (saddle) equilibria at or near MAX1 and MAX2 in our T21 model as well. In the next section, the presence of equilibria in the model will be investigated. The clear transitions between MAX1 and MAX2 suggest the influence of two heteroclinic connections between these assumed equilibria. After all, the model is completely deterministic; it does not contain any stochastic process that could have induced the transitions. Moreover, from the swiftness of the transitions and the persistence of the regimes it can be seen that the system evolution slows down near MAX1 and MAX2. The possibility of unstable periodic orbits being present in phase space that pass both MAX1 and MAX2 still leaves the question why the motion slows down near these maxima. The influence of heteroclinic connections can provide an explanantion, since such connections always end with a slowing down of the system evolution, due to the approach of a fixed point at the end of each connection.

The difference between MAX1 and MAX2 is largely spanned by EOF1, which resembles the AO pattern (Thompson and Wallace, 1998). From a phase space point of view, the AO

pattern in this model seems to be the direction, or phase space vector, connecting the two regimes. The blocking-character of the MAX2-regime is a regional aspect of a hemispheric pattern that resembles one of the phases of the AO. Long timescale variability of the AO should be related to the dynamics created by, or associated with, the near-presence (in parameter space) of a heteroclinic cycle (connection back and forth). In this perspective, variability of the AO has its dynamical origin in the (ultra-)low-frequency aspects of the dynamics of regime transitions. A better understanding of the dynamical mechanism(s) of regime transitions will help to better understand climate variability patterns such as the AO.

4 Equilibria

For models like the one under investigation there are two different (numerical) techniques available to find equilibrium solutions (fixed points). One is continuation: starting from a known steady state at specific (often unrealistic) parameter values, it is in principle possible to follow this steady state as parameter values are changed, toward realistic values. For an atmosphere model, one may think of starting at zero flow in the absence of forcing, and gradually turning on the forcing to a realistic intensity. A drawback of the continuation method is that it is hard, sometimes impossible, to find different (disconnected) branches of fixed points. Using continuation one easily (if not: generally) misses equilibria.

The other, less elegant method, used in the current study, is a way of "brute force calculation": bluntly calculating zeroes of the set of model-ODEs. For low-order models one can think of a Newton-like rootfinding algorithm; for larger models some sort of minimization-technique can be used (direct rootfinding in large models often doesn't work because of the presence of many local minima in phase space that are not solutions). For this study, a method was chosen that minimizes a scalar function, given its gradient. It was used before by Branstator and Opsteegh (1989). The squared norm of the vorticity-tendency $\partial\xi/\partial t$ served as the scalar function (denoted F); it is clear that a zero of this function is both a necessary and sufficient condition for the existence of a fixed point. The precise definition of F and the derivation of the gradient of F with respect to the model variables can be found in appendix B. The algorithm used comes from the NAG software library (routine E04UCF); it is a quasi-Newton method.

Many different states, sampled from all over the model attractor, need to be used as initial guesses for the minimization routine, in order to find as many different steady states as possible. This was done by starting the minimization from 5000 datapoints taken from the 200 year model integration, each 10 days apart. This set of 5000 points is assumed to give a reasonable coverage of all regions of the attractor. Not all 5000 minimizations converged to points with acceptably small F: about 1000 converged to points with $F \leq 10^{-14}$, giving 68 different fixed points in total. For comparison: the values of F before minimization were of order 10^{-4}.

It must be stressed that we do not claim to have found all the equilibrium solutions of

the model. However, it appears we have found the most relevant equilibria, namely those fixed points located in or near the two maxima of the PDF. We will apply two (dependent) criteria to assess the importance of the various fixed points for the large-scale, low-frequency dynamics of the barotropic model. One is the number n_u of unstable eigenvalues of the fixed points: fewer unstable eigenvalues suggests that the equilibrium is "more attracting" and therefore more likely to attract and influence the system as it evolves in time. We have calculated the eigenvalue-spectrum of each fixed point by evaluating the Jacobian matrix (see Appendix B) and calculating its eigenvalues. The result is shown in figures 6 and 3, which will be discussed later.

If the fixed points with the least unstable eigenvalues are indeed most likely to influence the dynamics, they can be expected to project strongly onto the leading EOFs. This should be visible using the second criterion: the extent to which each equilibrium state projects onto the leading EOFs. The behaviour of the dominant EOFs is most likely to be influenced by equilibria that strongly project onto them. Fixed points that lie far outside the subspace of the leading EOFs, and hence project relatively strongly onto the trailing EOFs, are less likely to be very relevant for the large scale dynamics of the model. Projecting each equilibrium onto the EOFs gives a vector $(a_1, a_2, \ldots, a_N)^{\mathrm{T}}$ of coefficients a_i. The total number of EOFs, N, equals 231. By calculating the ratio

$$R_M = \sum_{i=1}^{M} a_i^2 \; / \sum_{i=1}^{N} a_i^2 \tag{4.1}$$

for low values of M, one gets an impression of the contribution of the dominant EOFs to the equilibrium flow patterns. In fact, since the kinetic energy norm was used in the calculation of the EOFs, R_M equals the fraction of the turbulent energy contained in the first M EOFs.

Figures 6 and 3 contain the results of the calculations. Figure 6 shows a scatter plot of R_2 versus n_u, the number of unstable eigenvalues. There is a clear correlation between the two, which confirms the intuitive picture described earlier: the lower n_u, the more the equilibrium influences the dynamics of the system and hence the stronger it projects onto the leading EOFs (as was pointed out by one of the reviewers, figure 6 is in fact a test of the validity of applying the first criterion). Figure 3 shows the location of all the fixed points when projected onto the PC1,PC2 plane. Different symbols have been used in order to show the locations of the equilibria with low n_u (top) and high R_2 (bottom). Also shown are the contour lines (dashed) of the PDF. Applying the two criteria, low n_u and high R_2, gives a selection of five steady states: the ones with $n_u \leq 6$ (they all belong to the group with highest R_2 as well). One of the equilibria is located near MAX2, the second PDF-maximum; three are near the first maximum MAX1. The hypothesis that these maxima are related to the near-presence of relatively stable fixed points is thus supported. Note that the instability of the fixed points, although weak, makes it unlikely for them to be located exactly in the maxima of the PDF of the entire 231-dimensional phase space (after all, the system can never be exactly at an unstable fixed point). The fifth steady state lies rather eccentric to the PDF. We do not know why this is so.

5 Heteroclinic connections

The transitions between atmospheric flow regimes are usually thought of as stochastically induced phenomena. However, taking into consideration that in our model flow regimes can indeed be associated with steady states and that the model doesn't contain stochastic terms, it is, from a dynamical systems point of view, more natural to think of regime transitions as manifestations of heteroclinic behaviour. We hypothesize that the model is not very far (in parameter-space) from a situation in which there is one (or several) phase-space orbit(s) going exactly from a steady state associated with one regime to a steady state associated with another regime, and another orbit going back. Although such a situation requires the tuning of several parameters, and is thus unlikely to be exactly realised with the parameter setting that was chosen for the model, we expect that the influence of the connections is still felt by the system and manifests itself as transitions between regimes. In this section we will first describe in more mathematical detail what a heteroclinic connection is, and what the bifurcation scenario could be in which these connections are broken but still play a role. Then we both present evidence to support the hypothesis and show that the regime transitions in the model have a tendency to follow paths that are approximations of heteroclinic connections.

5.1 A bifurcation scenario

A heteroclinic connection, or heteroclinic orbit, is a solution of the model equations that connects one invariant set with another. The invariant sets can be steady states, periodic solutions, tori, etc. In what follows we will only consider heteroclinic connections between steady states.

Let us put these notions in more formal terms. Let $\dot{y} = G(y, \alpha)$ be a set of ordinary differential equations with parameter-set α. A heteroclinic connection $\tilde{y}(t)$ is a solution which asymptotically approaches in forward and backward time two different steady states y_- and y_+:

$$\lim_{t \to +\infty} \tilde{y}(t) = y_+, \qquad \lim_{t \to -\infty} \tilde{y}(t) = y_- \tag{5.1}$$

A set of heteroclinic connections forming a closed loop is called a heteroclinic cycle; a heteroclinic connection going from a steady state back to itself (i.e. $y_- = y_+$) is called a homoclinic orbit. For general background information on heteroclinic orbits, see Kuznetsov (1995). A detailed review of the theory of heteroclinic cycles is provided in Krupa (1997).

Our hypothesis is that regime behaviour, or vacillation, is eventually generated by heteroclinic connections going back and forth between invariant sets which are associated with regimes (i.e. a simple heteroclinic cycle). As previously mentioned, in this study we will assume the invariant sets to be steady states. The presence of this structure of two connections is not generic; in general, one or more parameters must be tuned in order to create the two connections. It is difficult to estimate a priori the number of parameters that must be tuned (the codimension). However, looking at the dimensions of the involved stable and unstable manifolds can give us some information. We choose the steady state

near MAX1 with 4 unstable eigenvalues as the first fixed point y_1 and the steady state near MAX2, which has 6 unstable eigenvalues, as the second fixed point y_2. Both fixed points are hyperbolic. $W^u(y_i)$ denotes the unstable manifold of y_i, $W^s(y_i)$ the stable manifold. A single heteroclinic connection from y_1 to y_2 must lie entirely in both $W^u(y_1)$ and $W^s(y_2)$. Since $\dim(W^u(y_1))= 4$ and $\dim(W^s(y_2))= 225$, and since in general 3 parameters must be tuned in order for a 4-dimensional (4-d) manifold and a 225-d manifold to have a 1-d intersection in a 231-d space , the creation of the 1-d intersection of $W^u(y_1)$ and $W^s(y_2)$ in the 231-d phase-space has codimension 3.

To formulate the argument in more precise terms, we need some definitions, taken mainly from the book by Guckenheimer and Holmes (1983). First of all, the *transversality theorem* tells us that the intersection of two manifolds of dimensions m_1 and m_2 in an n-dimensional space has, generically, dimension m_1+m_2-n. If $m_1+m_2 < n$ there is in general no intersection. Second, a *transversal intersection* of two manifolds in an n-dimensional space is an intersection in which the tangent spaces of the two manifolds in the intersection point(s) span the total n-dimensional space. A transversal intersection is persistent under small perturbations of the two manifolds. Finally, the codimension of a bifurcation can be defined as the smallest dimension of a parameter space which contains the bifurcation in a persistent (structurally stable) way. From these definitions follows that the creation of a 1-dimensional intersection of two manifolds M_1 and M_2 of dimensions m_1 and m_2 in an n-dimensional space with $m_1 + m_2 < n$ (which is a nontransversal intersection) has codimension $p = n + 1 - m_1 - m_2$, since in the space $\mathbb{R}^n \times \mathbb{R}^p$ (variables plus parameters) the two manifolds M_1 and M_2 (now of dimensions $m_1 + p$ and $m_2 + p$) have in general a transversal intersection of dimension $(m_1+p)+(m_2+p)-(n+p) = 1$. Applying this to the situation at hand shows that the 1-d intersection of a 4-d manifold with a 225-d manifold in a 231-d space happens in a persistent way in a parameter space of (at minimum) 3 dimensions. Thus, the codimension of the bifurcation in which this intersection is created is 3. For a more detailed treatment of these issues the reader is referred to Guckenheimer and Holmes (1983) and Kuznetsov (1995).

Unfortunately, a 1-d intersection of $W^u(y_1)$ and $W^s(y_2)$ doesn't necessarily take the shape of a heteroclinic connection from y_1 to y_2, as other singularities than the two fixed points may also play a role. For instance, the 1-d intersection could be a periodic orbit in phase-space, not connected to any of the two fixed points. As for the reverse connection, from y_2 to y_1: $W^u(y_2)$ and $W^s(y_1)$ generically have a 2-d (transversal) intersection, so this reverse heteroclinic connection at least seems to be more generic than the y_1 to y_2 connection. The dimensionality of $W^u(y_2)$ and $W^s(y_1)$ would even allow for a 2-d connecting surface rather than a single, 1-d heteroclinic connection. But again, other singularities can make the situation more complicated.

In other physical systems in which heteroclinic cycles play an important role in the dynamics, the existence of such cycles is usually connected to the presence of symmetries in the system. Symmetries induce the presence of invariant subspaces in phase-space, which can make the existence of a heteroclinic cycle generic, without needing to tune parameters (as is needed to create a heteroclinic cycle in a nonsymmetric system). Such a cycle will be robust for perturbations that preserve the symmetry, as these perturbations

leave the invariant subspace structure intact. See Krupa (1997) for a detailed account. Fluid-dynamical examples of heteroclinic cycles have been found in e.g. Rayleigh-Benard convection (Proctor and Jones, 1988) and the turbulent boundary layer (Holmes et. al. (1996, 1997) and references therein). At this stage we can only speculate about the (near-) presence of symmetries in the barotropic model, and their possible role in the formation of a heteroclinic cycle. In this context, it is interesting to note that the two PDF-maxima (and their corresponding steady states) are located roughly symmetric around the time mean state (where all PCs are zero). This can also be seen in the right panels of figure 4: the anomalies corresponding to MAX1 and MAX2 have roughly the same structure, but opposite sign. Another source of symmetry could be the dominant wavenumber 2 contribution to the orography. The question of symmetry, although impossible to answer at this moment, is of interest because a symmetry could be the common source of regime behaviour in different atmosphere models (e.g., if a certain (idealised) orography implies a symmetry that favors the creation of a heteroclinic cycle corresponding to regime transitions, it could explain the occurrence of regime behaviour in various atmosphere models with similar orography).

As mentioned previously, we hypothesize a point to be present in parameter space where a heteroclinic cycle exists between two fixed points corresponding to different regimes. Moving away from this bifurcation point (to more realistic physical conditions) one, or both, of the heteroclinic connections is almost certainly broken. If indeed a symmetry is related to the existence of the cycle, leaving the bifurcation point will probably perturb the symmetry (forced symmetry breaking). The phase-space structure that is left is expected to still bear traces of the former cycle, since it generates regime behaviour in the system. The structure should somehow account for the asymmetry of the regime behaviour, the fact that the zonal regime is visited more often than the blocked regime – the PDF is bimodal but not symmetric (in other studies, e.g. Itoh and Kimoto (1999), asymmetric PDFs have also been found). A possible explanation is that a homoclinic orbit connected to the zonal fixed point could be created when the cycle is broken. This would explain the findings in Crommelin (2002) where such a homoclinic orbit was detected. In that study, the possibility of heteroclinic behaviour was also mentioned. The behaviour of the current T21 model was very briefly studied there, and related to homoclinic dynamics; the heteroclinic influence was not detected due to the briefness of the inspection of the T21 model there.

It is known that the multiple equilibria associated with regime behaviour are the result of orographically induced saddle-node bifurcations that make the sheet of fixed points "fold" in parameter space, see Charney and DeVore (1979). Furthermore, Hopf-bifurcations due to barotropic instability occur naturally in barotropic models, see e.g. Legras and Ghil (1985) and De Swart (1988a, 1988b, 1989). Among the phenomena involved in the simultaneous occurrence of a Hopf- and a saddle-node bifurcation, a so-called fold-Hopf (or zero-Hopf) bifurcation, are heteroclinic connections, and homoclinic orbits resulting from perturbations, see Kuznetsov (1995). Thus, a scenario is possible in which orography and barotropic instability result in regime behaviour which originates from a heteroclinic cycle and which is deformed into more homoclinic-type behaviour, favoring the zonal regime, due

to forced symmetry breaking. We hope to report on this bifurcation scenario in a future paper. Let us furthermore mention that homoclinic bifurcations produce periodic orbits; in the cases of a Shilnikov-type or a bifocal-type homoclinic bifurcation an infinite number of (unstable) periodic orbits are generated. This could be related to low-frequency oscillations found by Plaut and Vautard (1994). However, much more study will be required to sort this out.

The regime behaviour in the barotropic model is clearly visible. Moreover, the regimes were found to correspond with nearby steady states. Having the above bifurcation scenario in mind, this means that the model cannot be too far from the point in parameter space where the heteroclinic cycle exists. The heteroclinic connections are most likely broken, but there will still be orbits that approximate the former connections and thus go from the vicinity of one steady state to the vicinity of the other. These approximations are what we will focus at in the next section. We will not attempt to prove the validity of the bifurcation scenario just sketched in the context of the barotropic model; such an attempt must start with the investigation of low-order models and of possible symmetries.

5.2 Numerical approximation

It is known that a heteroclinic connection from one steady state A to another B should start in the unstable manifold W^u(A) of state A and end in the stable manifold W^s(B) of state B. Furthermore, it takes infinitely long to follow the entire connection, from its starting point A to its end point B. In practice a heteroclinic connection is therefore approximated as a (finite time) connection starting in the linear approximation E^u(A) of W^u(A) at small but finite distance ε_{A} from point A and ending at distance ε_{B} from point B in the linear approximation E^s(B) of W^s(B) (with $0 < \varepsilon_{A,B} \ll 1$). A more detailed account is given in e.g. Doedel and Friedman (1989).

As said, it is unlikely that we are exactly in the bifurcation point and hence unlikely that exact heteroclinic connections exist; however, we can expect that there are still phase space orbits running from the immediate vicinity of one fixed point to the immediate vicinity of the other. To verify this, we took the most stable fixed point ($n_u = 4$) near the main PDF maximum (MAX1) as point A and the (only) fixed point near the submaximum (MAX2) as point B, and performed 2 series of 500 integrations of 100 days each. The integrations of the first series all started in E^u(A); the initial states were arbitrarily chosen using a random number generator, with the restriction that the distance of each initial state to A was fixed: if ψ^i denotes an initial state, the restriction was $||\psi^i - A|| = c\,||A - B||$, with $||.|| = <.,.>_1^{1/2}$, the square root of the kinetic energy inner product, and $c = 0.15$. This moderately small choice for c ensures that the model doesn't stay in the linear regime of A for too long, thereby keeping the total computation time within reasonable limits. During each integration the distance to B was monitored and the minimum distance to B was noted. These minimum distances showed large variances: using again the square root of the kinetic energy distance between A and B as unit, the minimum distances varied between 1 and 0.59. In the panels on the left of figure 7 the eight orbits that came closest to B (minimum distances between 0.59 and 0.67) are plotted in several projections onto

the leading EOFs. They clearly show a preference for a specific region of phase space.

The 500 integrations of the second series were started from arbitrarily chosen initial states in E^u(B) at fixed distance to steady state B. The distances to A were monitored and the minimum distances to A varied between 1 and 0.43. The eight orbits coming closest to A (minimum distances between 0.43 and 0.5) are plotted in the right panels of figure 7. Again, they follow roughly similar trajectories. Comparing the left and right panels of figure 7 shows marked differences between orbits going from A to B on the one hand and those from B to A on the other hand. In particular PC3 shows a separation: it is mainly negative during the transition from A to B and positive when switching back from B to A.

Although the projections of figure 7 suggest that the plotted orbits end up quite close to one fixed point or the other, the numbers show otherwise: the end distances are still of the order of half the distance between A and B. In order to find orbits that end up close to the desired fixed point, we have developed a new algorithm which makes uses of adjoint techniques, stemming from 4D-VAR data assimilation in weather forecasting.

The algorithm is an iterative method. From a given orbit a correction step for the starting point of the orbit is calculated. Applying this correction to the initial state should result in a new orbit ending up closer to the fixed point than the previous one. Suppose we want to find an orbit starting close to A and ending up close to B, and suppose we have a "first guess" orbit $\bar{\psi}(t)$ with $t \in [0, T]$. Let

$$\psi'_f = \mathrm{B} - \bar{\psi}(T) \tag{5.2}$$

and

$$G = < \psi'_f, \psi'_f >_1 , \tag{5.3}$$

with $< ., . >_1$ the kinetic energy inner product. The smaller G is, the closer the orbit ends near B, so we want to minimize G. Writing the model equation (2.1) in terms of streamfunction ψ and linearising it around states from the orbit $\bar{\psi}(t)$ yields the tangent linear model, which describes the linear evolution of perturbations ψ^p of $\bar{\psi}(t)$:

$$\dot{\psi^p} = L(\bar{\psi}(t))\, \psi^p \tag{5.4}$$

A second orbit $\tilde{\psi}(t)$ of the full nonlinear system (2.1), starting in $\tilde{\psi}(0) = \bar{\psi}(0) + \psi'(0)$, will end in $\tilde{\psi}(T) = \bar{\psi}(T) + \psi'(T)$. If $\psi'(0)$ is small enough, $\psi'(T)$ will approximately be the result of the tangent linear model (5.4) integrated over time T with initial state $\psi'(0)$.

The formal solution of the tangent linear model can be written using the resolvent $R(t_1, t_2)$ and reads

$$\psi^p(t_2) = R(t_1, t_2)\, \psi^p(t_1) \tag{5.5}$$

If we consider ψ'_f to be a perturbation of the orbit $\bar{\psi}(t)$ at its end-point we can relate it to a perturbation ψ'_i at its starting point by substituting $\psi^p(t_1) = \psi'_i$ and $\psi^p(t_2) = \psi'_f$ in the above equation, giving $\psi'_f = R(0, T)\, \psi'_i$. Writing shortly R for $R(0, T)$ and substituting in equation (5.3) yields

$$G = < \psi'_f, R\, \psi'_i >_1 = < R^*\, \psi'_f, \psi'_i >_1 = \, - < \Delta R^*\, \psi'_f, \psi'_i >_0 \tag{5.6}$$

in which R^* is the adjoint of R with respect to $< ., . >_1$ (note that the adjoint depends on the inner product) and $< ., . >_0$ is the squared norm inner product (see appendix A). For more information on the adjoint, in particular on the adjoint of the tangent linear barotropic model with respect to the kinetic energy inner product, see Talagrand and Courtier (1987) and Barkmeijer (1992).

We can now relate a change in the initial state of the orbit $\bar{\psi}$ to a change in the distance between the end point of $\bar{\psi}$ and the fixed point B:

$$\frac{\delta G}{\delta \psi_i'} = -\Delta R^* \, \psi_f' \tag{5.7}$$

This gives us the recipe to improve on the orbit $\bar{\psi}$: calculate the difference vector ψ_f', integrate this vector using the adjoint of the tangent linear model and let the Laplace operator Δ act on the resulting vector $R^* \, \psi_f'$. Shift the initial state $\bar{\psi}(0)$ in the direction of $-\Delta R^* \, \psi_f'$. The orbit $\tilde{\psi}$, that results from integrating the nonlinear system (2.1) with this new initial condition, should end up closer to the desired fixed point. See figure 8 for a schematic representation of the procedure. The method is applied iteratively under the constraint that the initial point of the orbit at each iteration step stays both in the linear approximation of the unstable manifold and at a fixed, small distance to the steady state.

The time-length T of the orbit $\bar{\psi}$ hasn't been considered yet. As was mentioned previously, a genuine heteroclinic orbit stretches over an infinite amount of time; however, the approximation of the connection, from E^u(A) to E^s(B), takes a finite amount of time. This time-length is not known a priori, so the algorithm just described, which tries to find the approximation, must allow for the possibility to adapt the interval T. This is taken care of by starting with a relatively small T and extending the length if at the previous iteration step the algorithm could only marginally improve on the orbit (that is, when the relative decrease of $G^{1/2}$ drops below, say, 1%).

To initialize the algorithm the 2 sets of eight "optimal" initial states, described previously, were taken from the 2 sets of 500 integrations. The "initial guess orbits" were set to a length of 25 days, starting from the two times eight initial states. The magnitude of the correction steps made to the initial states was initially set to a length of 0.1 times the distance between the initial point and the nearby steady state; this magnitude was decreased when lengthening the time interval T.

From both sets we present the orbits showing the best results (that is, the orbits coming closest to the fixed points while the algorithm couldn't improve on them anymore): 2 of the A to B (zonal to blocked regime) orbits have end distances to fixed point B of 0.35 (length: 167 days) and 0.37 (183 days). Compared to the earlier results with minimum distances to B between 0.59 and 0.67, this is a considerable improvement. It must be mentioned here that in the 200 year dataset there are only 9 events in which the distance to B is less than or equal to 0.37; the absolute minimum in the whole dataset is 0.30. For the reverse case (B to A) we have 3 orbits, with end distances to A of 0.15 (173 days), 0.13 (121 days) and 0.11 (217 days). The 200 year dataset has an absolute minimum of 0.11 and 2 events below 0.15. It must be stressed that all 5 orbits are solutions of the full nonlinear model (2.1), integrated with the same integration routine as was used to create the 200 year dataset.

The projections of the 5 orbits onto the PC1,PC3 plane are shown in figure 9. The characteristics already apparent in figure 7 show up again for these 5 orbits. The numbers show that the B to A orbits come much closer to fixed point A than the A to B orbits come to B. Moreover, the latter two are very similar to each other whereas the former three have clearly distinct trajectories (see figure 9). This supports the argument, given earlier on in this section, that the B to A connection is probably more generic (and could even be a 2-dim surface rather than a single curve) than the reverse connection.

5.3 Data interpretation

The dynamical structure that emerges by accepting the hypothesis that regime behaviour is linked to heteroclinic dynamics has its consequences for the behaviour of the barotropic model. We have seen that the approximations of the (nearby) heteroclinic connections travel through certain regions of phase space during transitions from one regime to the other. If the full barotropic system, as described by equation (2.1), really is influenced by these heteroclinic connections, this should be visible in the data as well. More specifically, the system should show in its data a tendency to have a negative PC3 anomaly when going from the zonal regime MAX1 to the blocked regime MAX2, and a positive PC3 anomaly when returning from the blocked to the zonal regime, since this is how the heteroclinic connections travel. Furthermore, during transitions from MAX1 to MAX2 the found orbits stay closer together than during transitions in the other direction. One would therefore expect the data of the MAX2 to MAX1 transitions to be more scattered than the data for the reverse transitions. Finally, the region with strongly positive PC2 and weakly negative PC1 anomalies is visited during MAX1 to MAX2 transitions but not during MAX2 to MAX1 transitions.

These assumptions can be checked using the 200 year dataset discussed earlier. Let us define the two regimes in a very coarse way, thereby including probably too many states in these regimes: all states with PC1 $>$ 0.004 count as states in the zonal regime, all states with PC1 $<$ -0.008 are said to be in the blocked regime. All transitions, i.e. all the segments of the data for which -0.008 $<$ PC1 $<$ 0.004, are divided in four categories: transitions from the zonal to the blocked regime (z $\rightarrow$ b), from blocked to zonal (b $\rightarrow$ z), from zonal back to itself (z $\rightarrow$ z) and from blocked back to itself (b $\rightarrow$ b). The beginning and ends of the transition segments are determined by PC1 crossing the value of 0.004 or -0.008. The duration of the transitions is not taken into account here. From the four different datasets that result PDFs are made. The number of datapoints ending up in each dataset is as follows: z $\rightarrow$ b: 10548; b $\rightarrow$ z: 6648; z $\rightarrow$ z: 26398; b $\rightarrow$ b: 2894. The projections of the PDFs of the b $\rightarrow$ z and z $\rightarrow$ b segments onto the planes of the leading three EOFs is shown in figure 10. The EOFs that are used are those of the full 200yr dataset. The PDFs show that the data indeed possesses the transition-characteristics that could be predicted from the trajectories of the near-heteroclinic orbits in figures 7 and 9. PC3 is negative during z $\rightarrow$ b and positive during b $\rightarrow$ z transitions. The region with weakly negative PC1, strongly postive PC2 is often visited during z $\rightarrow$ b but hardly during b $\rightarrow$ z transitions. The PDFs of the z $\rightarrow$ b transitions are more confined than those of

the b $\rightarrow$ z transitions (more specifically: the maxima of the z $\rightarrow$ b PDFs are much more pronounced, showing that these transitions are more confined than the b $\rightarrow$ z ones).

The b $\rightarrow$ b and z $\rightarrow$ z transitions have not been discussed yet. The former has not enough realisations to be really interesting, but the latter does have many realisations. Although these transitions don't involve transitions from one regime to another but rather from the zonal regime back to itself, the concept of global bifurcations can still shed some light on these transitions. If the hypothesis is correct that the z $\rightarrow$ b $\rightarrow$ z heteroclinic cycle gets broken and results in a homoclinic z $\rightarrow$ z connection, one would still expect this homoclinic connection, and hence the data from the z $\rightarrow$ z transitions, to show traces of the original cycle. The preference for certain phase space regions during the first half (leaving the zonal regime) and second half (returning to the zonal regime) of the z $\rightarrow$ z transition may still be somewhat similar to the preferences of the z $\rightarrow$ b resp. the b $\rightarrow$ z transition. In order to check this, all segments of the z $\rightarrow$ z transitions were cut in two. All datapoints of the first half of each z $\rightarrow$ z segment were collected in one dataset, the points of the second halves in another. The PDFs of these two datasets are shown in figure 11. The projection on the PC1,PC2-plane doesn't show the characteristics anymore, but the other two projections do: the PC3 anomalies still show the characteristics described earlier: negative when leaving the zonal regime, positive when returning.

The very coarse way in which the two regimes were defined above, and in which the transitions were counted, has the consequence that the "heteroclinic characteristics" that are found back in the data are not very detailed. We could improve on this by selecting the data more precisely, e.g. by defining the regimes more strictly and by imposing a maximum time limit in which the transition from one regime to the other should be completed. However, such stricter selection would obscure the very general influence of the heteroclinic behaviour. The coarse definition and selection of regime- and transition states shows that the heteroclinic characteristics are visible not only for those transitions that run precisely between narrowly defined regimes, but also for a much broader category of states, namely the large majority of transitions between only crudely defined regimes. From the 46488 states that were counted as transitional states, 43549 (93.7%) fell in either the z $\rightarrow$ b, b $\rightarrow$ z or z $\rightarrow$ z cases, which were shown to possess on average the characteristics of heteroclinic transitions.

6 Discussion

Looking at the regime behaviour, displayed by the barotropic model, from a dynamical systems point of view gave a new perspective on the dynamics of regime transitions. The thought that these transitions should be related to heteroclinic connections, since the model is entirely deterministic, guided us to the notion that the transitions must be structured in phase space. The transitions from the zonal to the blocked regime will, on average, follow paths through different regions of phase space than the transitions from the blocked to the zonal regime. The 200 yr model dataset indeed shows clear differences between the two transitions, even with a coarse filtering; these differences match well with the

approximations of (the ruins of) the hypothesized heteroclinic connections. A few remarks must be made to conclude this paper:

• The new algorithm, presented in this study, to approximate the ruins of the heteroclinic connections was able to improve on the optimal transition orbits that resulted from choosing many arbitrary initial states in (the linear approximation of) the unstable manifolds of the two steady states. Nevertheless, the algorithm still needs improvement in order to make it work more easily and probably with better results, in different model contexts. In particular, the handling of the time-interval and the choice of the length of the correction step can be more sophisticated. However, such improvement is beyond the scope of this paper.

• The difference between the two regimes mainly projects onto EOF1, which resembles the Arctic Oscillation pattern. The difference between the z→b and b→z transitions is most clearly visible in EOF3, which projects strongly onto the North Atlantic Oscillation pattern (with positive PC3 corresponding to the negative NAO-phase and vice versa). Combining these findings in one picture results in figure 12, a highly simplified, schematic picture of the dynamics of regime behaviour, NAO and AO in the barotropic model. Starting in the zonal regime, which shows a polar vortex that is weaker than average, the model evolves via a positive NAO phase to the blocked regime. The blocked regime is characterised by a blocking over Europe and, on a larger scale, a strong polar vortex. Leaving the blocked regime, the block persists for a while but shifts to the west, to the North-Atlantic. This transition phase thereby gets the features of a negative NAO phase. Finally, the model atmosphere is back in the zonal regime.

• The great (phase space) variety of the b → z transition orbits compared to their z → b counterparts (which showed up both in the data and in the approximations to the heteroclinic connections, and which can be explained partly by looking at the dimensions of the stable and unstable manifolds of the involved fixed points) raises interesting questions about the formation and destruction of blockings in our model. There seems to be less freedom in the ways blockings can be created than there is in the ways blockings are destroyed. The breaking up of the zonal regime should follow a rather specific scenario for the system to evolve to a blocking situation within a certain amount of time (say two months).

• This conclusion cannot end without pointing out that the model used in this study, a T21 barotropic model with realistic forcing, obviously has its limits: the spatial resolution is rather coarse, it is limited to northern hemisphere flow, and baroclinic dynamics is absent (but note the implicit presence of the time-averaged impact of neglected (e.g. baroclinic) processes in the forcing). Nevertheless, the bimodality and regimes observed in this model have several realistic aspects. Blocked flows, variations of the polar vortex and the NAO are all known features of the atmospheric flow. The details may be inaccurate but in essence these phenomena of the barotropic model are realistic. The regimes MAX1 and MAX2 resemble the AO in its negative and positive phase, respectively (Thompson and Wallace, 1998). They combine a weakening resp. strengthening of the wintertime polar vortex with anomalies of the same sign (but opposite to the polar anomaly) over Europe and the North-Pacific (see figure 4). In the real atmosphere, the anomaly over

Europe stretches towards North-America. This stretching is absent in the regimes of the barotropic model, but somewhat present in our EOF1. Our regimes can also be compared to the clusters found by Corti et. al. (1999). Their clusters D and B resemble our MAX1 and MAX2. Furthermore, the pattern of their cluster A is comparable to a negative PC3 anomaly of our model: positive NAO-phase, negative anomaly over the North-Pacific, positive anomaly over North-America. These similarities beg the question whether transitions from cluster D to B have a tendency to go via cluster A, which would be the same order as we found in the barotropic model. In our model data no regime was found that corresponds to cluster A. However, the preferred transition routes have a small "loop" at the midpoint of the MAX1 to MAX2 transition (figure 9). This extra curvature may result in a local PDF maximum in time-averaged data (note that Corti et. al. use monthly mean data, whereas we have used daily data). The PDFs for the z $\rightarrow$ b transitional states (figure 10) have indeed a maximum at the location of the small loop. This observation suggests that there may be two types of regimes: one type associated with (marginally unstable) equilibria of the large-scale flow, the other type induced by local curvatures of the preferred transition orbits. Finally, the resemblance of cluster C and our positive PC3 anomaly is less strong (although not completely absent), so the return transition from B to D via C may be less clear. This could be related to our findings that in the barotropic model the MAX2 to MAX1 transitions are less confined than the reverse ones.

We have argued that the regime behaviour of the barotropic model is quite realistic in its qualitative aspects. The concept of heteroclinic connections gave a new perspective on the regime behaviour in the model. It was shown that these connections have important implications for the details of the regime transitions, an aspect which has been largely ignored in the literature. This perspective can very well be used for further, more detailed studies of transitions in both models and observations.

Acknowledgments. F.M. Selten and W.T.M. Verkley are kindly thanked for supplying me with the code of their barotropic model. W.T.M. Verkley is furthermore acknowledged for his valuable suggestion that 4D-VAR data assimilation techniques might be useful for detecting global bifurcations. I am grateful to J.D. Opsteegh and F. Verhulst for many useful discussions and comments. The comments of three anonymous reviewers helped to improve the paper. Financial support for this investigation came from the Netherlands Organization for Scientific Research (NWO).

A Appendix: The inner product

The inner product used throughout this paper is the kinetic energy inner product, defined by

$$< \psi_1, \psi_2 >_1 = \frac{1}{2} \int (\nabla \psi_1) \cdot (\nabla \psi_2) \, \mathrm{d}\Omega \tag{A.1}$$

The integral is taken over the entire sphere; $\mathrm{d}\Omega$ denotes a volume element on the sphere. ψ_1 and ψ_2 are streamfunction fields on the sphere. For the barotropic model, the kinetic energy of a state with streamfunction ψ is given by $< \psi, \psi >_1$. The kinetic energy inner product belongs to a class of inner products

$$< \psi_1, \psi_2 >_k = \frac{1}{2} \int (\nabla^k \psi_1) \cdot (\nabla^k \psi_2) \, \mathrm{d}\Omega, \quad k \in \mathbb{N} \tag{A.2}$$

It is easy to show that $< \psi_1, \psi_2 >_1 = - < \xi_1, \psi_2 >_0$, where $\xi_1 = \Delta\psi_1$, the vorticity field associated with ψ_1.

B Appendix: Gradient and Jacobian

The scalar function F to be minimized in section 4 is the discretized version of the functional $\mathcal{F}$, which is defined as the squared norm of the vorticity tendency integrated over the sphere:

$$\mathcal{F} = \int (\dot{\xi})^2 \, \mathrm{d}\Omega. \tag{B.1}$$

All integrals in this appendix are assumed to be taken over the sphere; $\mathrm{d}\Omega$ denotes again a volume element on the sphere. Using the definition $\delta\mathcal{F}(\xi) = \mathcal{F}(\xi + \delta\xi) - \mathcal{F}(\xi)$ and substituting equation (2.1) gives:

$$\begin{aligned} \delta\mathcal{F} &= 2 \int \dot{\xi} \, (\delta\dot{\xi}) \, \mathrm{d}\Omega \\ &= 2 \int \dot{\xi} \, (J(\delta\xi, \psi) - J(\delta\psi, \xi + f + h) + k_1 \Delta\delta\psi + k_2 \Delta^4 \delta\psi) \, \mathrm{d}\Omega \\ &= 2 \int \delta\psi \, (J(\dot{\xi}, \xi + f + h) + \Delta J(\psi, \dot{\xi}) + k_1 \Delta\dot{\xi} + k_2 \Delta^4 \dot{\xi}) \, \mathrm{d}\Omega \end{aligned} \tag{B.2}$$

The above derivation uses $\delta\xi = \Delta\delta\psi$, as well as the following identities:

$$\int \eta \, (\Delta\zeta) \, \mathrm{d}\Omega = \int \zeta \, (\Delta\eta) \, \mathrm{d}\Omega \tag{B.3}$$

$$\int \eta \, J(\zeta, \nu) \, \mathrm{d}\Omega = - \int \nu \, J(\zeta, \eta) \, \mathrm{d}\Omega \tag{B.4}$$

These identities stem from the identity $\int \nabla \cdot \mathbf{A} \, \mathrm{d}\Omega = 0$, which is valid for any smooth vector field $\mathbf{A}$ since the integration runs over the sphere.

From equation (B.2) follows

$$\frac{\delta\mathcal{F}}{\delta\psi} = 2 \left(J(\dot{\xi}, \xi + f + h) + \Delta J(\psi, \dot{\xi}) + k_1 \Delta\dot{\xi} + k_2 \Delta^4 \dot{\xi} \right) \tag{B.5}$$

Discretizing the equations, i.e. expressing all fields in terms of a finite number of spherical harmonics, yields the approximation F of $\mathcal{F}$; ψ is replaced by $(\psi_1, \ldots, \psi_N)^{\mathrm{T}}$. The functional derivative (B.5) becomes the vector $\partial F / \partial \psi_i$, $i = 1, \ldots, N$: the gradient of F with

respect to $(\psi_1, \ldots, \psi_N)^{\mathrm{T}}$. This is the gradient needed for the minimization algorithm.

An expression for the Jacobian, the matrix that describes the model linearised around some state, can be derived in a similar way as the expression of the gradient above. The Jacobian (denoted by L) is defined as

$$L_{ij} = \frac{\partial \dot{\psi}_i}{\partial \psi_j} \tag{B.6}$$

which may be thought of as a series of gradients of scalars $\dot{\psi}_i$. Define

$$D_i = \dot{\psi}_i = \int Y_i^* \, \dot{\psi} \, \mathrm{d}\Omega \tag{B.7}$$

in which Y_i^* is the complex conjugate of the spherical harmonic with (multi-)index i; the gradient can be shown to be

$$\frac{\delta D_i}{\delta \psi} = J(\Delta^{-1} Y_i^*, \xi + f + h) + \Delta J(\psi, \Delta^{-1} Y_i^*) + k_1 Y_i^* + k_2 \Delta^3 Y_i^* \tag{B.8}$$

Discretizing in spherical harmonics gives the Jacobian:

$$L_{ij} = \left(\frac{\delta D_i}{\delta \psi} \right)_j \tag{B.9}$$

References

Achatz, U. and G. Branstator, 1999: A two-layer model with empirical linear corrections and reduced order for studies of internal climate variability. *J. Atmos. Sci.*, **56,** 3140–3160.

Armbruster, D., J. Guckenheimer, and P. Holmes, 1988: Heteroclinic cycles and modulated travelling waves in systems with O(2) symmetry. *Physica D*, **29,** 257–282.

Aubry, N., P. Holmes, J. L. Lumley, and E. Stone, 1988: The dynamics of coherent structures in the wall region of a turbulent boundary layer. *J. Fluid Mech.*, **192,** 115–173.

Barkmeijer, J., 1992: Local error growth in a barotropic model. *Tellus*, **44A,** 314–323.

Branstator, G. and J.D. Opsteegh, 1989: Free solutions of the barotropic vorticity equation. *J. Atmos. Sci.*, **46,** 1799–1814.

Charney, J. G. and J. G. DeVore, 1979: Multiple flow equilibria in the atmosphere and blocking. *J. Atmos. Sci.*, **36,** 1205–1216.

Corti, S., F. Molteni, and T. N. Palmer, 1999: Signature of recent climate change in frequencies of natural atmospheric circulation regimes. *Nature*, **398,** 799–802.

Crommelin, D. T., 2002: Homoclinic dynamics: a scenario for atmospheric ultra-low-frequency variability. *J. Atmos. Sci.*, **59,** 1533–1549.

De Swart, H. E., 1988a: Low-order spectral models of the atmospheric circulation: a survey. *Acta Appl. Math.*, **11,** 49–96.

De Swart, H. E., 1988b: *Vacillation and Predictability Properties of Low-Order Atmospheric Spectral Models.* PhD thesis, Rijksuniversiteit Utrecht.

De Swart, H. E., 1989: Analysis of a six-component atmospheric spectral model: Chaos, predictability and vacillation. *Physica D*, **36,** 222–234.

De Swart, H. E. and J. Grasman, 1987: Effect of stochastic perturbations on a low-order spectral model of the atmospheric circulation. *Tellus*, **39A,** 10–24.

Doedel, E., 1986. AUTO: Software for continuation and bifurcation problems in ordinary differential equations. California Institute of Technology.

Doedel, E. J. and M. J. Friedman, 1989: Numerical computation of heteroclinic orbits. *J. Comput. Appl. Math.*, **26,** 155–170.

Guckenheimer, J. and P. Holmes, 1983: *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields.* Springer-Verlag, 459 pp.

Holmes, P., J. L. Lumley, and G. Berkooz, 1996: *Turbulence, Coherent Structures, Dynamical Systems and Symmetry.* Cambridge University Press, 420 pp.

Holmes, P.J., J.L. Lumley, G. Berkooz, J.C. Mattingly, and R.W. Wittenberg, 1997: Low-dimensional models of coherent structures in turbulence. *Phys. Rep.*, **287,** 337–384.

Holton, J. R., 1992: *An Introduction to Dynamic Meteorology, 3rd edition.* Academic Press, 511 pp.

Itoh, H. and M. Kimoto, 1996: Multiple attractors and chaotic itinerancy in a quasi-geostrophic model with realistic topography: Implications for weather regimes and low-frequency variability. *J. Atmos. Sci.*, **53,** 2217–2231.

Itoh, H. and M. Kimoto, 1997: Chaotic itinerancy with preferred transition routes appearing in an atmospheric model. *Phys. D*, **109,** 274–292.

Itoh, H. and M. Kimoto, 1999: Weather regimes, low-frequency oscillations, and principal patterns of variability: A perspective of extratropical low-frequency variability. *J. Atmos. Sci.*, **56,** 2684–2705.

James, P.M., K. Fraedrich, and I.N. James, 1994: Wave-zonal-flow interaction and ultra-low-frequency variability in a simplified global circulation model. *Quart. J. R. Met. Soc.*, **120,** 1045–1067.

Kimoto, M. and M. Ghil, 1993a: Multiple flow regimes in the northern hemisphere winter. part I: Methodology and hemispheric regimes. *J. Atmos. Sci.*, **50,** 2625–2643.

Kimoto, M. and M. Ghil, 1993b: Multiple flow regimes in the northern hemisphere winter. part II: Sectorial regimes and preferred transitions. *J. Atmos. Sci.*, **50,** 2645–2673.

Knobloch, E. and J. Moehlis, 2000. Burst mechanisms in hydrodynamics. In *Nonlinear Instability, Chaos and Turbulence, Volume II*, ed. L. Debnath and D. Riahi, Computational Mechanics Publications, Southampton, 237-287.

Krupa, M., 1997: Robust heteroclinic cycles. *J. Nonlinear Sci.*, **7,** 129–176.

Kuznetsov, Yu. A., 1995: *Elements of Applied Bifurcation Theory.* Springer-Verlag, 515 pp.

Legras, B. and M. Ghil, 1985: Persistent anomalies, blocking and variations in atmospheric predictability. *J. Atmos. Sci.*, **42,** 433–471.

Mo, K. C. and M. Ghil, 1988: Cluster analysis of multiple planetary flow regimes. *J. Geophys. Res.*, **93,** 10927–10952.

Palmer, T. N., 1999: A nonlinear dynamical perspective on climate prediction. *J. Climate*, **12,** 575–591.

Plaut, G. and R. Vautard, 1994: Spells of low-frequency oscillations and weather regimes in the northern hemisphere. *J. Atmos. Sci.*, **51,** 210–236.

Proctor, M. R. E. and C. A. Jones, 1988: The interaction of two spatially resonant patterns in thermal convection. Part 1. Exact 1:2 resonance. *J. Fluid Mech.*, **188,** 301–335.

Reinhold, B. B. and R. T. Pierrehumbert, 1982: Dynamics of weather regimes: Quasi-stationary waves and blocking. *Mon. Wea. Rev.*, **110,** 1105–1145.

Roads, J. O., 1987: Predictability in the extended range. *J. Atmos. Sci.*, **44,** 3495–3527.

Selten, F. M., 1995: An efficient description of the dynamics of barotropic flow. *J. Atmos. Sci.*, **52,** 915–936.

Talagrand, O. and P. Courtier, 1987: Variational assimilation of meteorological observations with the adjoint vorticity equation. Part 1: Theory. *Q. J. R. Meteorol. Soc.*, **113,** 1311–1328.

Thompson, D. W. J. and J. M. Wallace, 1998: The Arctic Oscillation signature in the wintertime geopotential height and temperature fields. *Geophys. Res. Lett.*, **25,** 1297–1300.

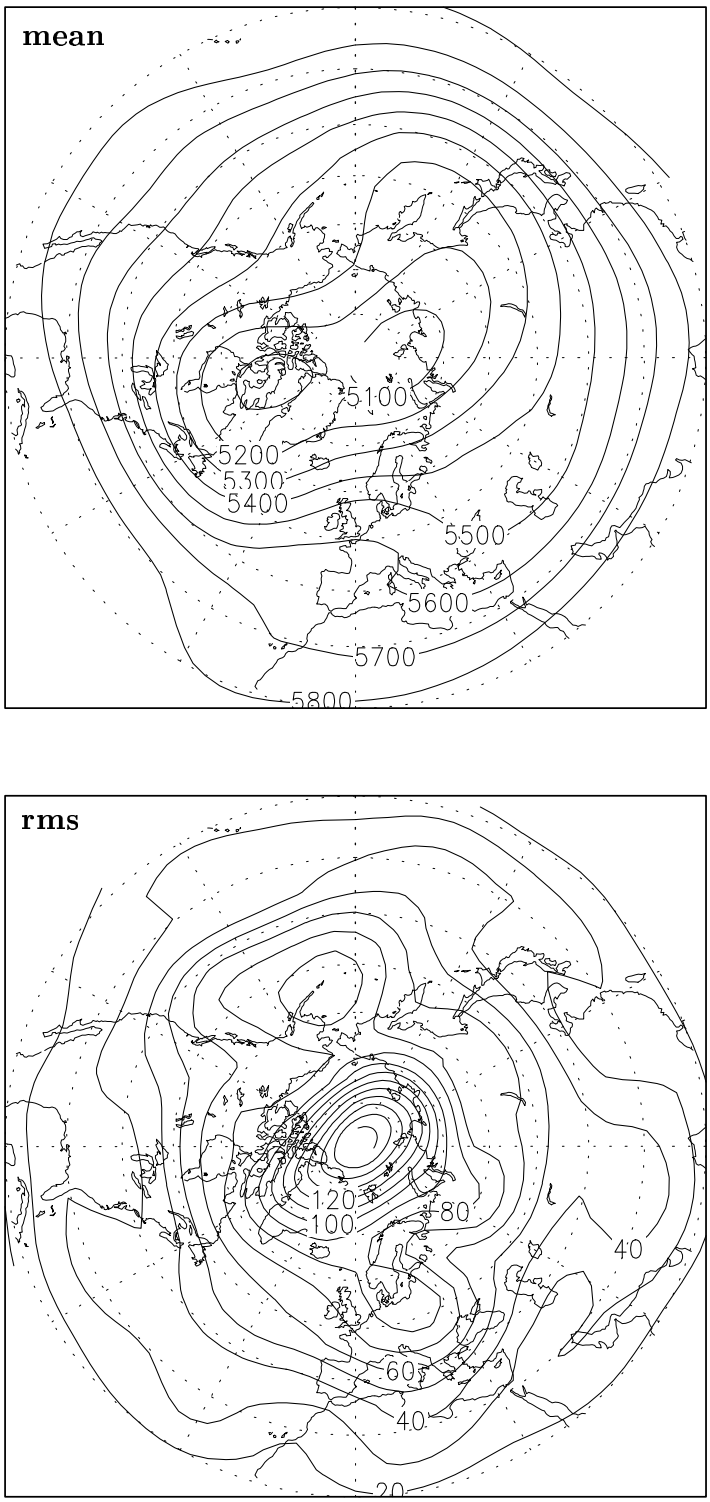

Figure 1: Top: time mean of the 500 hPa geopotential height Z_{500} as produced by the T21 barotropic model. Contours are drawn every 100 meter. Bottom: Z_{500} rms of the T21 model. Contour interval is 10 meter.

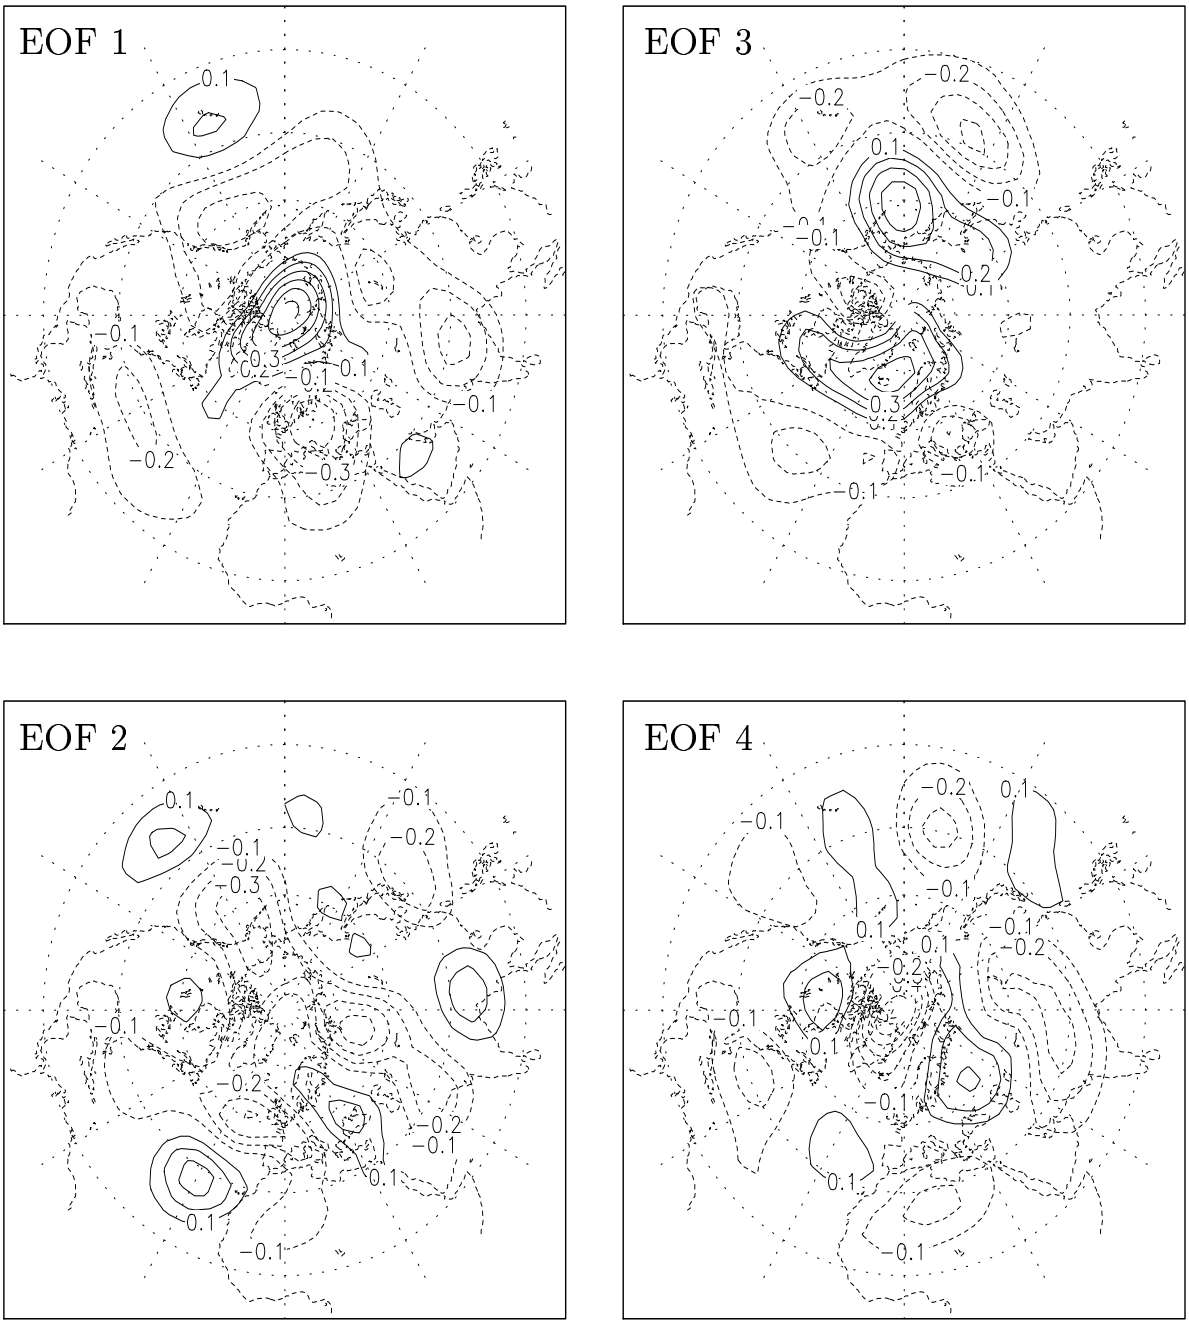

Figure 2: EOFs 1-4 calculated from the 200 year dataset produced by the T21 barotropic model. The kinetic energy norm was used for the calculation.

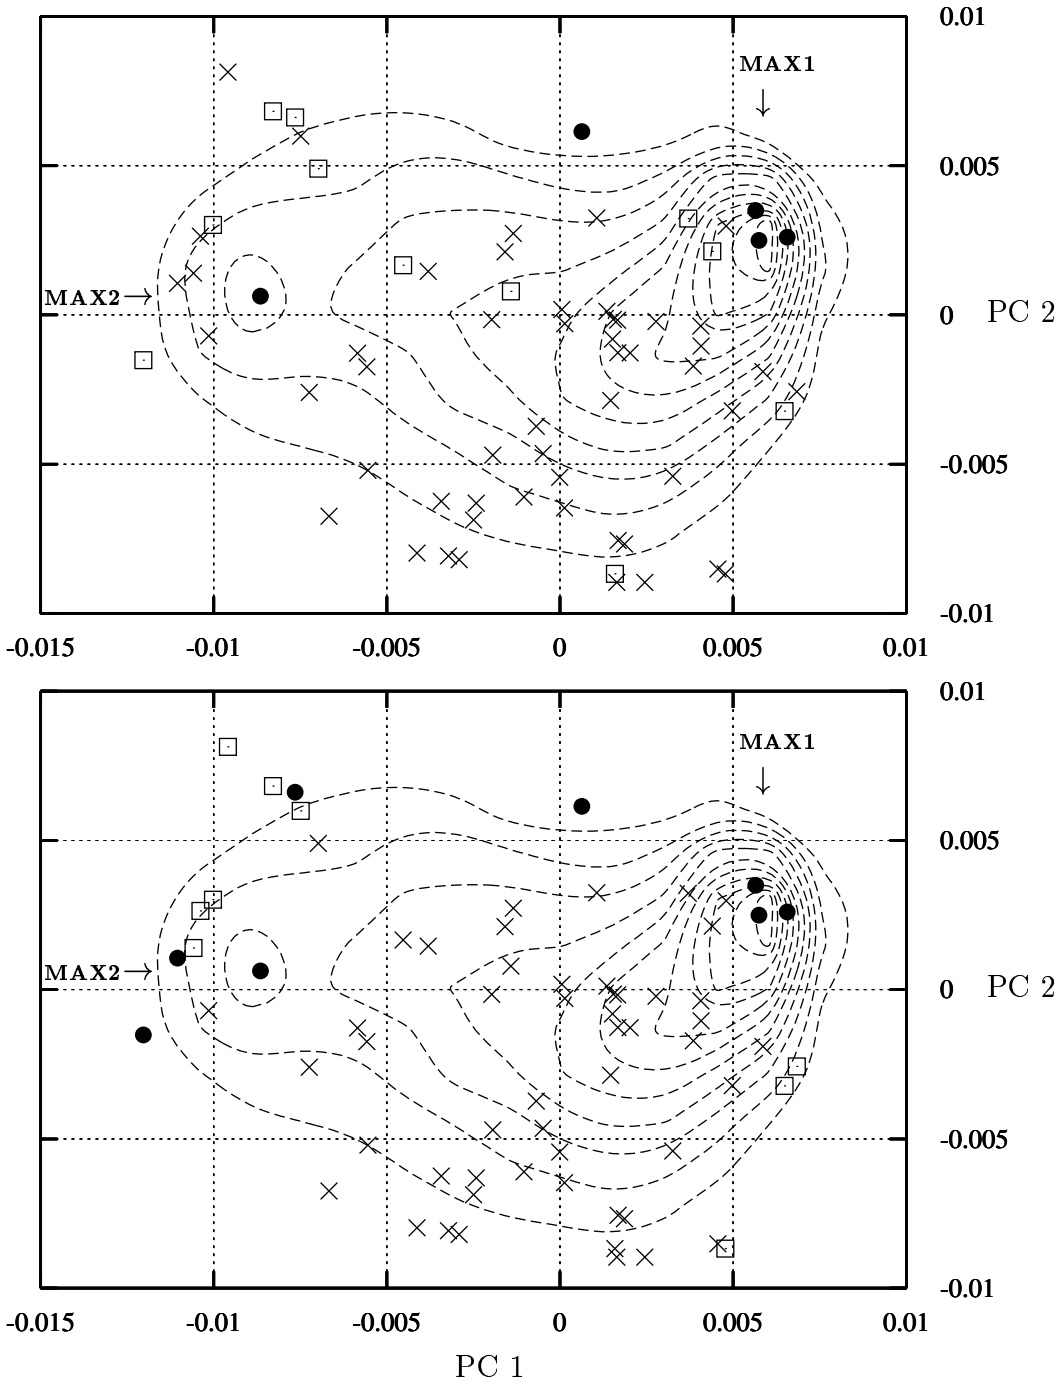

Figure 3: PDF and steady states in PC1,PC2–plane. The dashed lines are the contours of the PDF (contour distance: 200 datapoints per bin, number of bins: 15 × 15). Top: steady states with number of unstable eigenvalues $n_u \leq 6$ (•); $7 \leq n_u \leq 10$ (⊡); $n_u > 10$(×). Bottom: steady states with ratio of projection onto the leading 2 EOFs $R_2 > 0.5$ (•); $0.35 < R_2 < 0.5$ (⊡); $R_2 < 0.35$(×).

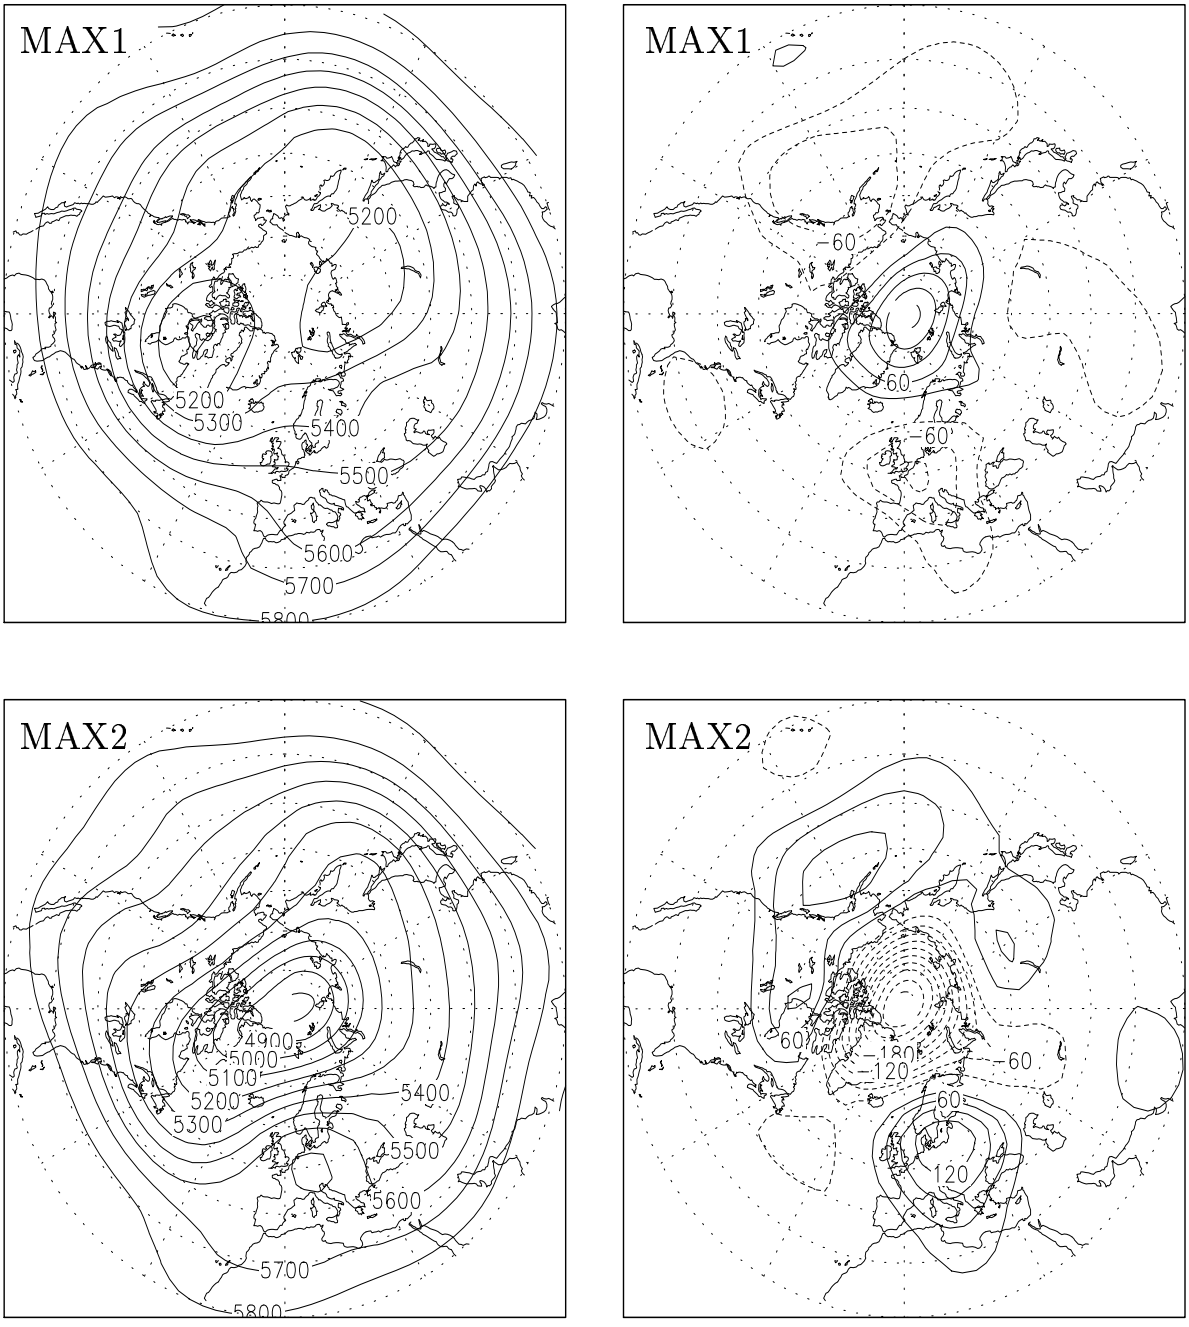

Figure 4: Patterns corresponding to the maxima of the PDF. Top: MAX1, bottom: MAX2 (see text). Shown left are patterns of the 500 hPa geopotential height field Z_{500} (contour interval 100 m). On the right are the corresponding Z_{500} anomalies (contour interval 30 m).

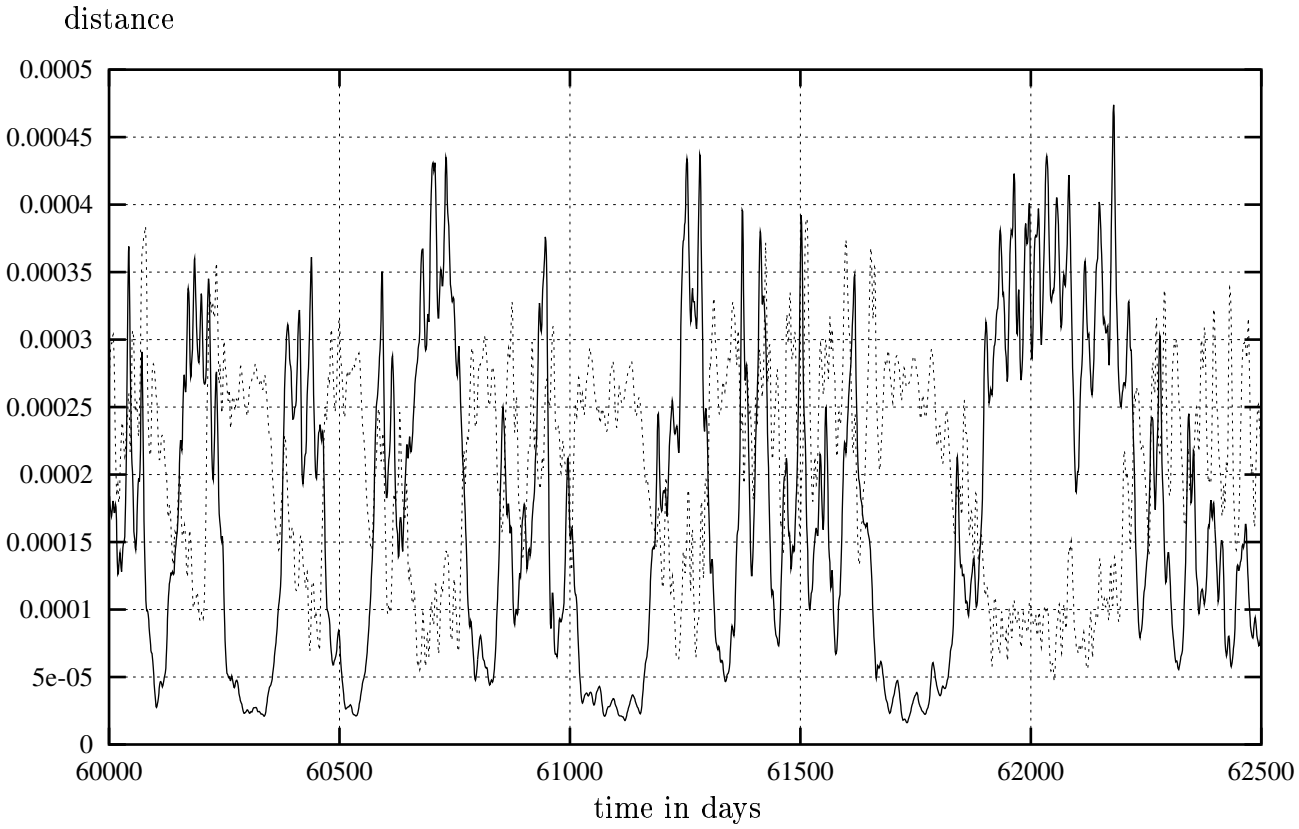

Figure 5: Phase space distances to two maxima MAX1 (solid) and MAX2 (dotted) during a data segment of 2500 days. Distances are given as differences in turbulent kinetic energy; horizontal axis denotes time in days.

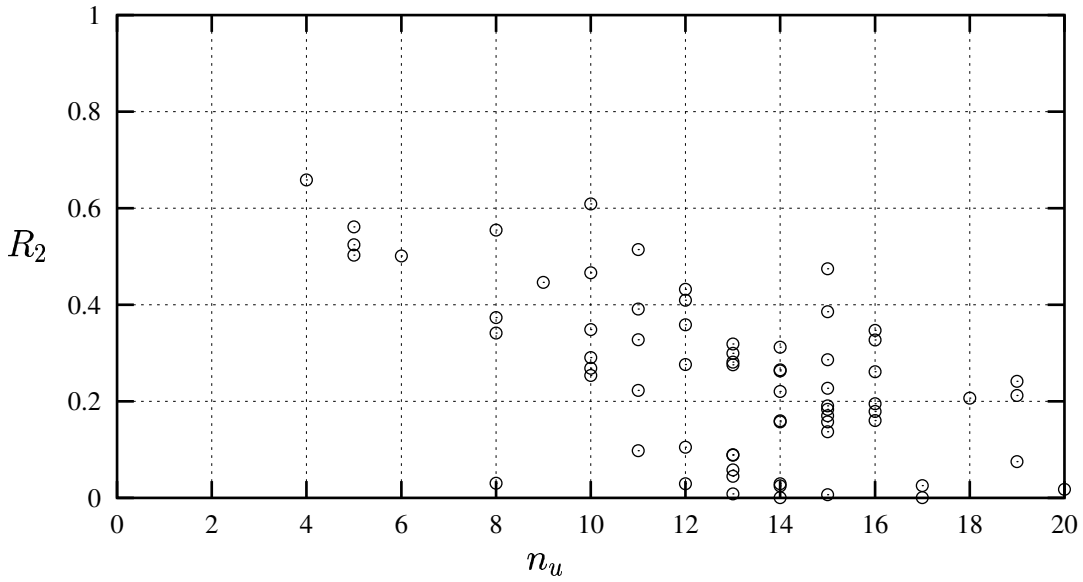

Figure 6: Number of unstable eigenvalues n_u versus ratio R_2 of projection onto the leading 2 EOFs (see text) for each steady state of the barotropic model.

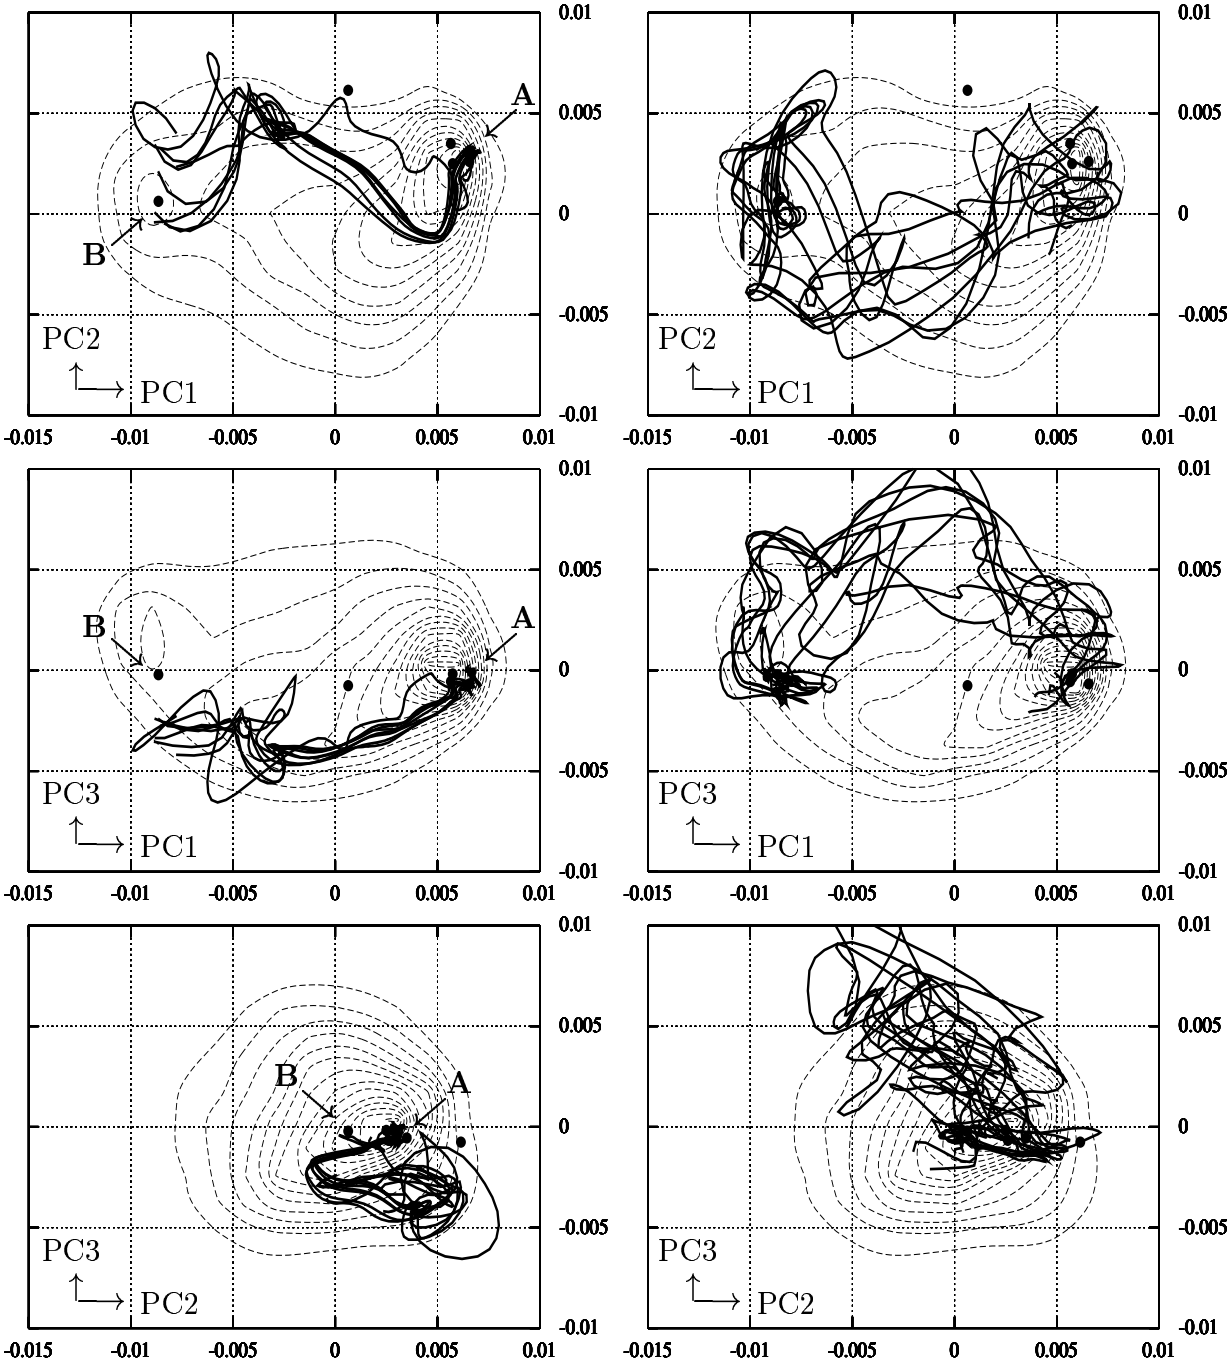

Figure 7: Projection of transition orbits onto planes of PC1,PC2 (top), PC1,PC3 (middle) and PC2,PC3 (bottom). Left: orbits starting close to steady state A near MAX1. Right: orbits starting close to steady state B near MAX2. Also shown are the contourlines (dashed) of the PDF of the 200 year dataset and the locations of the 5 selected steady states (see text). The locations of A and B are indicated in the left panels.

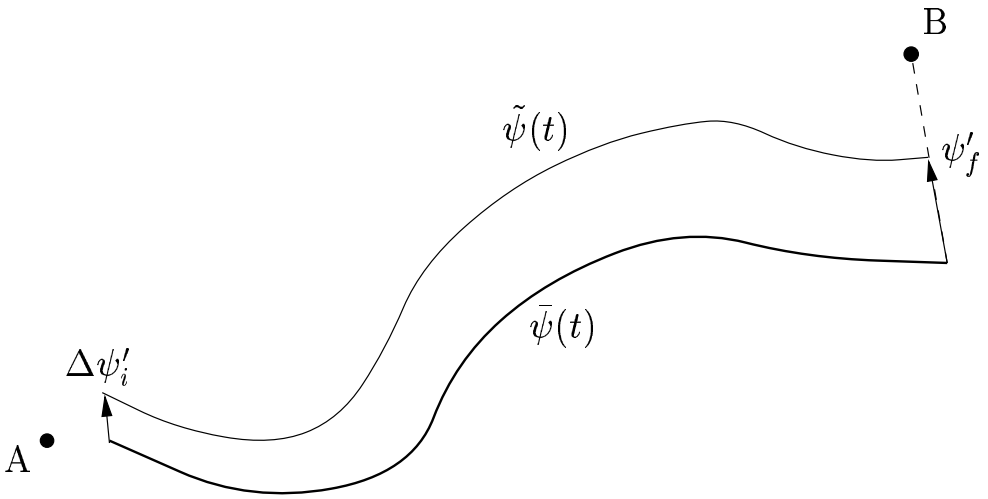

Figure 8: Schematic representation of the adjoint method to find approximate heteroclinic connections. The vector ψ'_f pointing from the end point of the orbit $\bar{\psi}(t)$ to the fixed point B is integrated with the adjoint of the tangent linear model, resulting in ψ'_i. By moving the initial state in the direction of $\Delta\psi'_i$ a new orbit (of the full nonlinear system) $\tilde{\psi}(t)$ is obtained ending closer to B.

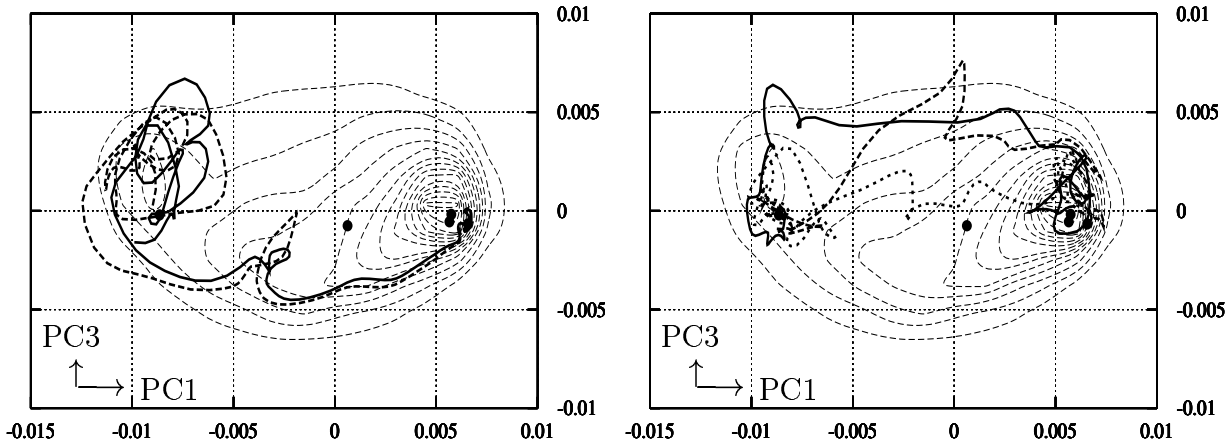

Figure 9: Approximations of nearby heteroclinic connections between regimes. Left: zonal to blocked, right: blocked to zonal. Shown is the projection of the orbits onto the PC1,PC3-plane.

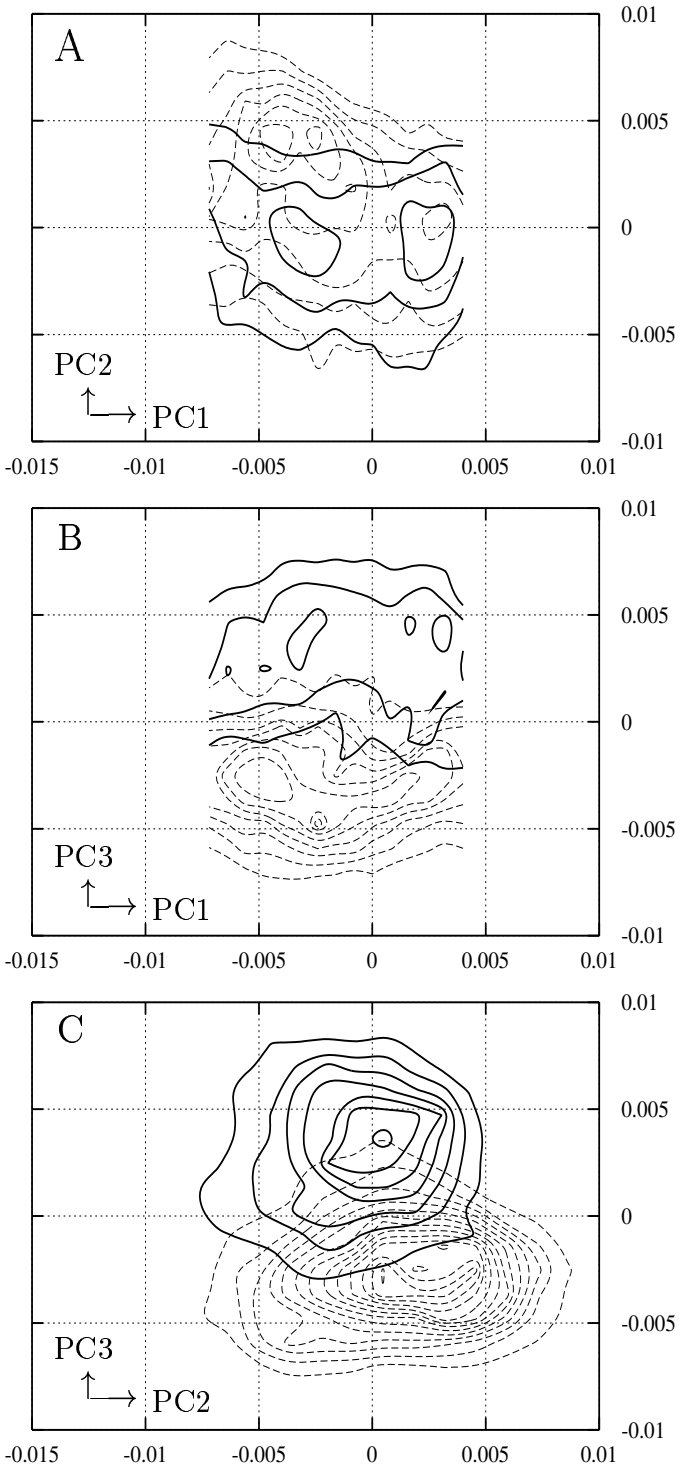

Figure 10: PDFs of transitional states in 200yr dataset. Dashed lines: zonal → blocked; solid thick lines: blocked → zonal. Contour interval is 25. Projections shown are PC1 vs. PC2 (A), PC1 vs. PC3 (B) and PC2 vs. PC3 (C)

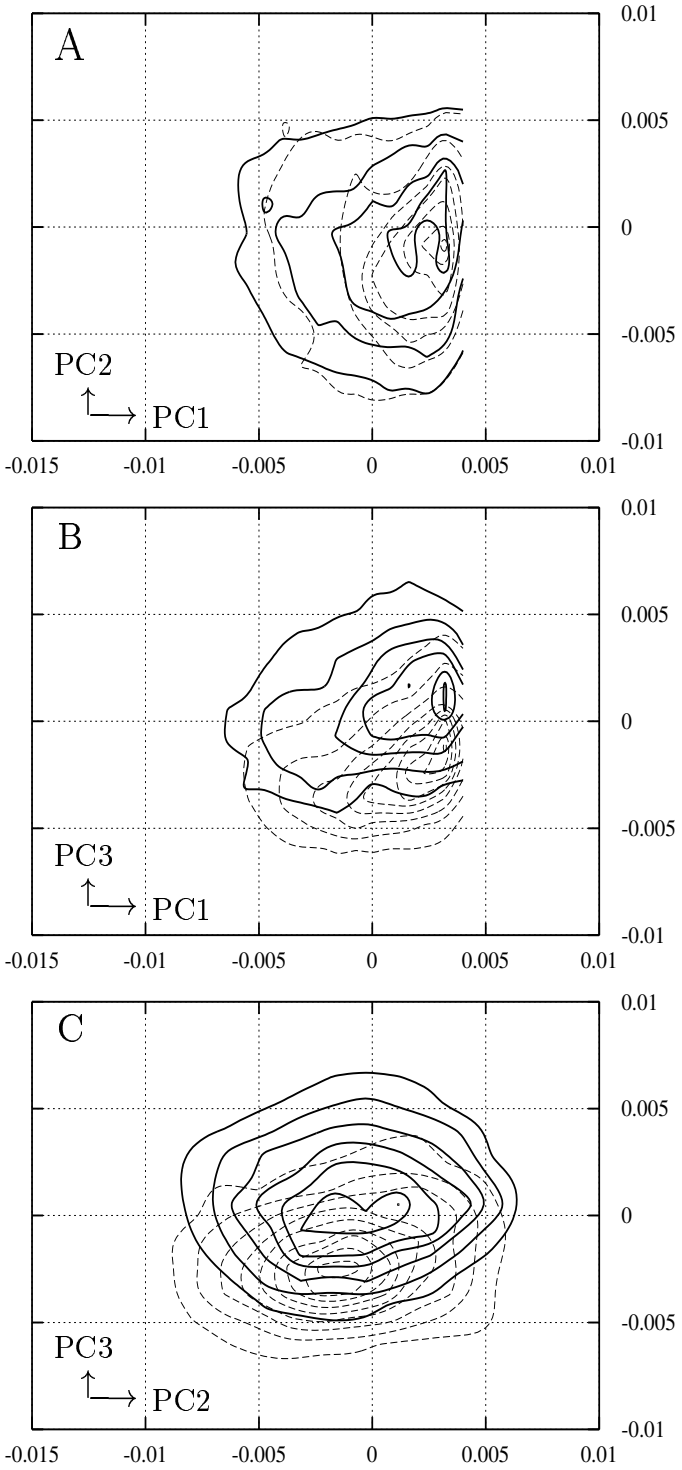

Figure 11: PDFs of zonal → zonal transitional states in 200yr dataset. Dashed lines: first half of each segment; solid thick lines: second half of each segment. Contour interval is 50. Projections shown are PC1 vs. PC2 (A), PC1 vs. PC3 (B) and PC2 vs. PC3 (C)

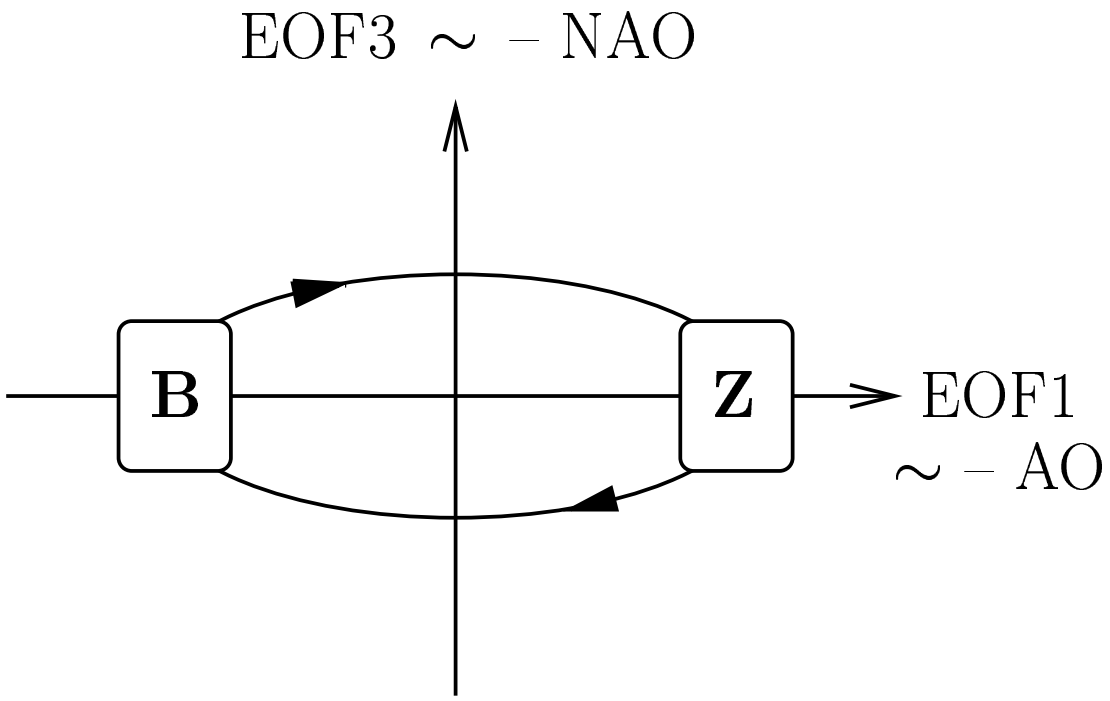

Figure 12: Schematic representation of the interrelationships of regime behaviour, NAO and AO, as emerging from the barotropic model. B denotes the blocked regime, Z the zonal regime. Negative PC1 corresponds to a stronger polar vortex (positive AO phase), negative PC3 corresponds to the positive NAO phase.