Martian dust lifting, transport and associated processes

Submitted for the degree of Doctor of Philosophy

in Atmospheric, Oceanic and Planetary Physics

David Paul Mulholland

St. Anne's College, University of Oxford

Trinity 2012

Martian dust lifting, transport and associated processes

David Mulholland St. Anne's College

Submitted for the degree of Doctor of Philosophy Trinity 2012

Abstract

The dust lifting capacity of the UK Mars General Circulation Model has been extended through the development of a new wind stress lifting parameterisation, and the simulation of a finite, variable surface dust layer. This second addition, which was represented by the use of lifting thresholds that were adjusted at each surface gridpoint in response to the removal or deposition of dust, led to enhanced variability in the timing and peak magnitude of major dust storms produced in the model. These dust storms were realistic in many respects, and the observed global dust storm frequency of occurrence of roughly one in every three years was approximately reproduced by the model, but an artificial threshold decrease rate was required to maintain dust lifting on a multiannual timescale — this was believed to be due to inaccuracies in the net cross-equatorial dust flux, which showed a strong bias towards the northern hemisphere. Significant changes were seen in model dust lifting rates when the influence of a heterogeneous surface roughness length was included in the wind stress scheme, and the need for more sophisticated sub-gridscale methods in future dust lifting schemes, to cope with this and other effects, was noted. The inclusion of radiatively active water clouds in model runs also affected dust lifting rates, particularly in the vicinity of the polar caps in autumn, winter and spring. The dynamics behind the formation of small, cap-edge dust storms during these periods were examined in detail, and it was found that a cessation in dust lifting activity that occurs around winter solstice does so due to a combination of the radiative effects of global dust loading and polar hood ice clouds, and zonal variations in midlatitude topography. The direct interaction between dust and ice, in the form of nucleation and scavenging, was investigated. It was found that scavenging by water ice, if it is suitably efficient, could significantly reduce the dust content of the winter polar regions. However, the dust and ice vertical profiles measured in the aphelion cloud belt by Mars Climate Sounder were not reproduced by the model with any of the possible scavenging efficiencies used. It appears that scavenging cannot provide an explanation for the existence of sharply defined, elevated dust layers at low latitudes.

Acknowledgements

My greatest thanks goes to my supervisors, Peter Read and Stephen Lewis, for their advice and criticism on the work contained in this thesis; Peter in particular for his thoughtful suggestions, encouragement and helpful pointers on the trickier aspects of fluid dynamics, and Stephen for his expert assistance with the model and suggested directions of investigation. I thank Luca Montabone for providing the filtered Mars Climate Sounder data used in Chapter 6, and for our many enjoyable chats about Martian atmospheric science. I am also grateful to Peter Rogberg for providing assistance at the beginning of my project. Colleagues at the Laboratoire de Météorologie Dynamique in Paris — François Forget, Ehouarn Millour, Jean-Baptiste Madeleine, Aymeric Spiga and Constantino Listowski — provided advice and assistance through the provision of data and model routines, and I also thank the group for being so welcoming during my visit to Paris in 2010.

Of course I am very grateful to my parents, Geraldine and Paul, for their support over the whole of my academic career. I appreciated the advice of my friends Christopher, Connell and Michelle, who encouraged me to make the move to Oxford four years ago. Thanks to Chris, Jo and Richard for the lunchtime chat. I thank (among others) Laura Marling, Justin Vernon, Noel Gallagher, Christopher Owens and Lauren Laverne for providing the soundtrack to the writing of this thesis over the past nine months. Finally, I thank the Science & Technology Facilities Council for providing the funding for this work.

Contents

Chapter 1

Introduction

1.1 Mars' principal climate features

Mars, like Earth, can be classed as a rapidly rotating planet, and is in many ways the planet in our solar system with an atmosphere that is most similar to our own. This similarity, its close proximity to Earth and the possibility of it being habitable in the past, present or future have meant that it has been the most closely studied of the planets (other than the Earth) since the space age began. There are, however, some striking differences between Mars and Earth. Mars' atmosphere is thin — surface pressure is around 6 mb, more than a hundred times lower than that of Earth — and it is cold, with summer surface temperatures rarely exceeding 0°C (273 K). The atmosphere is predominantly carbon dioxide ($\sim 95\,\%$ by mass; other minor constituents include N_2, Ar, O and CO), which freezes out in winter when temperatures drop to ~ 140 K, forming polar ice caps that can extend as far as the midlatitudes.

Progress through the Martian year is usually described using the solar longitude, L_s, the angle the Mars-Sun vector makes with a reference vector, $L_s = 0°$, taken to occur at northern hemisphere (NH) spring equinox. In this way the year can be conveniently divided into seasons each spanning 90° of L_s, or 'months' each 30° in length. A Martian solar day, or *sol*, is very similar in length to the Earth day at 24.66 hours, but the year lasts for just less than 669 sols, making it almost twice as long as a year on Earth. Martian years are numbered according to the standard defined by Clancy et al. (2000) in which Martian

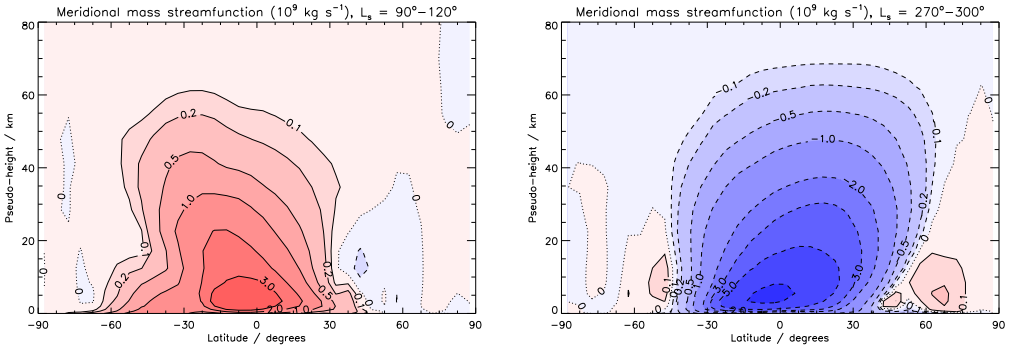

Figure 1.1: Zonal-mean circulation from the UK Mars General Circulation Model, in northern summer (left) and southern summer (right). Solid (dashed) contours denote anticlockwise (clockwise) motion.

year 1 (MY1) began on 11 April 1955, marking the beginning of modern-day, co-ordinated observations, coinciding with a Martian opposition and the detection of a major dust storm (see §1.2.2). The current Mars year (in 2012) is MY31, while the majority of observations referred to in this work are from MY24–29.

In a global-mean sense, the Martian climate can be divided into two distinct seasonal regimes, separated by northern autumn equinox. The northern spring and summer seasons, covering solar longitudes $L_s = 0$–$180°$, show relatively low atmospheric temperatures, low dust opacities, and high water vapour and ice abundances, while autumn and winter, $L_s = 180$–$360°$, are notably warmer, dustier and drier (Liu et al., 2003). This asymmetry is thought to derive both from the large meridional topographic gradient — most of the southern hemisphere (SH) sits several kilometres higher than the northern hemisphere (Zuber et al., 2000) — and from the eccentricity of the planet's orbit ($e \sim 0.09$, compared to $e \sim 0.02$ for Earth) coupled with the current timing of perihelion[1], which occurs at $L_s = 251°$ in late southern spring. The climate of the aphelion season is extremely repeatable from year to year — globally averaged atmospheric temperatures over several years compared by Liu et al. (2003) agreed to within $\sim 1\,\mathrm{K}$ — whereas the perihelion season displays significant variability.

The two rather different solsticial circulations are shown in Figure 1.1. Both periods are dominated, in a zonal-mean sense, by a cross-equatorial 'Hadley'-like cell, henceforth referred to as the principal meridional overturning cell (PMOC), extending to midlatitudes

[1]The closest approach of the planet to the Sun.

in both hemispheres with upwelling in the summer and downwelling in the winter hemisphere. The PMOC is stronger at $L_s = 270°$ than at $L_s = 90°$, however, which Richardson and Wilson (2002a) found to be primarily due to the topographic slope (see Figure 1.3), rather than the difference in insolation received between the two seasons. For the same reasons, zonal winds in the winter westerly jet are strongest during northern autumn and winter (McCleese et al., 2010).

1.1.1 Surface and atmospheric dust

The seasonal asymmetry manifests itself in a number of ways, most notably by providing winds strong enough to activate large-scale dust lifting away from the polar cap edges only during southern spring and summer. Dust storms initiated between $L_s = 180°$ and $\sim 330°$ have the potential to expand rapidly, taking advantage of strong flow in both the meridional and zonal directions, and with the added possibility of secondary lifting sites becoming activated, large portions of the atmosphere below ~ 50 km can be supplied with dust on a timescale of a few sols. A positive feedback is present here, as the increased dust loading warms the atmosphere and strengthens the PMOC, increasing the likelihood of further lifting (Haberle et al., 1982). Martian dust storms are described more fully in the next section.

1.1.2 Water vapour and ice clouds

The zonally averaged water cycle is shown in Figure 1.2, where the wet/dry annual pattern can clearly be seen. The Thermal Emission Spectrometer (TES) (Christensen et al., 1992) on board Mars Global Surveyor (MGS) is one of several modern instruments that has provided extended periods of continuous retrievals of vertically integrated water and ice abundances (Smith, 2002), supplementing earlier orbiter (Jakosky and Farmer, 1982) and ground-based measurements (Barker et al., 1970). There are some discrepancies between these various datasets concerning absolute abundance values, but all agree with the general picture seen here of a large vapour maximum at the north pole around summer solstice, and a smaller peak at the southern summer pole. In an annually averaged sense, the atmosphere of the northern hemisphere contains about 70% more water than the southern hemisphere (Smith, 2002). The prominent north polar vapour maximum forms as a residual water ice

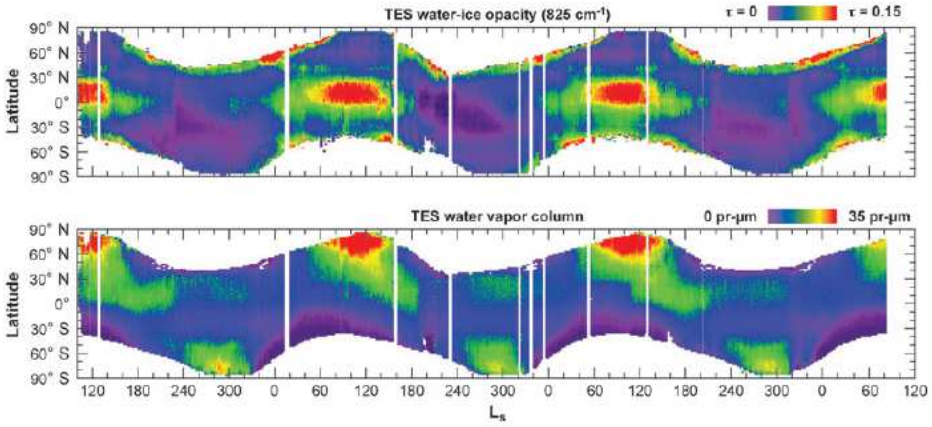

Figure 1.2: Zonally averaged infrared water ice optical depth (top) and water vapour column abundance (precipitable microns) for the duration of the TES observation period (from $L_s \sim 120^\circ$ in MY24 to $\sim 80^\circ$ in MY27), from Smith (2008).

cap is exposed to incoming solar radiation, following the disappearance of the seasonal CO_2 ice cap that covers it during the rest of the year, and sublimes, releasing vapour into the atmosphere. During spring and summer, some of the vapour in the northern hemisphere is advected into the rising branch of the PMOC, where it ascends to colder altitudes until the atmosphere becomes saturated, at which point (~ 20 km) an ice cloud layer forms, made of particles typically of radii 3–4 μm (Wolff and Clancy, 2003). This feature can be seen at 0–30°N in Figure 1.2, and is known as the aphelion cloud belt (ACB).

It has been pointed out (by Clancy et al. (1996)) that the level of saturation around aphelion prevents water vapour from spreading through the full height of the PMOC, thereby limiting the amount of water that is advected southwards. Around perihelion, on the other hand, the warmer atmosphere raises the level of condensation (to ~ 40 km (Smith, 2002)), so that vapour in the southern hemisphere can be freely transported northwards by the (stronger) PMOC and returned to the northern hemisphere. Evidence for this can be seen in the vapour 'tongue' formed at 0–30°N at $L_s = 180$–300°. This phenomenon is clearly dependent on the solar longitude of perihelion, and since this is known to precess with a period of 50 kyr (Ward, 1974), the so-called 'Clancy effect' has important implications for the long-term transfer of water between hemispheres (Montmessin et al., 2007). This involves an interplay with atmospheric dust, which helps to set saturation heights through the warming it causes in the atmosphere, but which could also be limited in its vertical

motion (and therefore also its meridional advection) by cloud formation, which occurs solely through heterogeneous nucleation around dust particles (Määttänen et al., 2005). The extent to which dust particles are *scavenged* — removed from the atmosphere or shielded from view — during ice nucleation is not known, but possible evidence for the effect was reported by Vincendon et al. (2008), who observed an anti-correlation between dust and ice abundances in parts of the south polar region during autumn.

The other significant cloud formation regions are the winter polar hoods, opacities for which were not retrieved by TES but have since become available through more recent observations (Benson et al., 2010, 2011). Peak IR opacities at each winter pole are comparable to the ACB opacity (~ 0.1) and again there are differences between the two winter seasons. Factors important for polar hood cloud formation appear to be temperature, availability of water vapour in the vicinity of the polar region, and cap-edge wave activity (Montmessin et al., 2004). Precipitating ice has recently been observed at northern high latitudes by the Phoenix lander's LIDAR instrument (Whiteway et al., 2009), which detected particularly large ice radii (up to $40\,\mu$m) near the surface.

The radiative impact of water ice clouds is complex and exhibits significant diurnal, and possibly seasonal, variation, but ice clouds are known to be responsible for temperature inversions in the lower atmosphere (Colaprete and Toon, 2000). Wilson et al. (2007) found modelled surface temperatures, when ice clouds were neglected, to be too cold at nighttime and too warm in daytime in cloudy parts of the Tharsis region, highlighting a nighttime warming effect from cloud IR emission, of $> 10\,$K in places. Further work estimated a warming of 5–10 K at and above the ACB, confirming that water ice clouds exert a significant impact on the large-scale temperature structure (Wilson et al., 2008). However, near-surface clouds at higher latitudes may induce a net cooling (Wilson et al., 2008), and further indirect temperature changes may be possible through vertical redistribution of dust.

The north polar residual cap covers an area of $10^6\,\mathrm{km}^2$ (from the pole to 80–85°N, with considerable zonal asymmetry) with a thickness of up to a few km (Zuber et al., 1998) — a vast reserve of water. What surface deposits may be present at the south pole are relatively unknown, as there is currently a CO_2 ice cap present there all year round, possibly due to the higher elevation of the south pole relative to the north, and/or because of a higher cap

albedo. This explains the significantly smaller vapour maximum seen at the south pole, as it forms from a combination of sublimation of ice deposited onto the seasonal cap during winter, transport from the northern hemisphere, and perhaps release from the regolith (Smith, 2002). Some of the interannual variability suggested in measurements of this peak (for example, compare MY24 and MY25 in Figure 1.2) is likely to be caused by variability in SH dust activity (Smith et al., 2009). The perennial CO_2 cap has been observed to be gradually shrinking (Malin et al., 2001), a potential example of the long-term cycling of condensable material between hemispheres that is believed to occur in response to changes in orbital forcing, such as the precession of the longitude of perihelion. Evidence of this transport has been left behind in the polar layered deposits, striped patterns in the polar regions believed to be a result of variation in dust-ice ratios during periods of deposition onto the poles (Cutts and Lewis, 1982).

1.2 Martian dust storms

As the simulation of Martian dust storms is a key component of the work presented here, it is important to review what is currently known about such storms, much of which comes from a multidecadal observational record, supported by theoretical models of increasing complexity.

1.2.1 Spatial variation

To facilitate discussion of regions of dust storm initiation and expansion, a labelled map of the spatial and topographic layout of Mars is shown in Figure 1.3. Martian surface altitude, referenced to the areoid[2], covers an extremely large range, from the volcanoes of the Tharsis range (including Olympus Mons, which at almost 22 km is the highest known peak in the solar system) to the deep craters of Hellas and Argyre. The meridional gradient mentioned above is evident from the figure. Of considerable importance to the circulation is the zonal wavenumber-three pattern in topography at $\sim 30°$N, which directs northern midlatitude disturbances into three channels (Acidalia-Chryse, Utopia-Isidis and Arcadia-Amazonis), each flanked by high topography to the west (Hollingsworth et al., 1996). Near the solstices, western boundary currents (WBC) form against these ridges as an intensification of the

[2]A geopotential surface analogous to Earth's sea level.

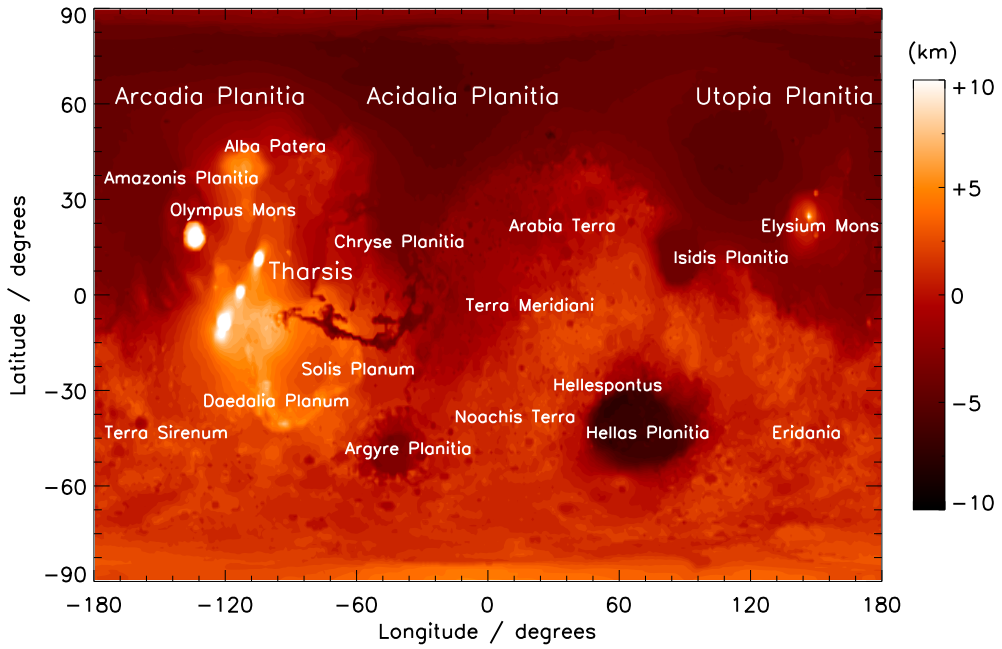

Figure 1.3: Map showing the areographic locations of regions mentioned during this work. Shading shows surface topography, as deviation from the areoid, in kilometres.

return-flow branch of the PMOC (Joshi et al., 1995); around northern winter solstice, this favours southward advection of disturbances — which can include frontal dust storms (of significant meridional extent) formed along the edge of the polar cap — along the three lowland channels.

Examples of how dust storms develop in relation to the topographic features shown above are provided in Figure 1.4. The plot shows dust optical depth output from TES data assimilated into a Mars General Circulation Model (GCM; alternatively Global Climate Model), focusing on the largest dust storm that occurred during southern spring in each of the three years assimilated (MY24–26). In MY24 and MY26, this storm was regional in scale; that is, it covered an area larger than $\sim 10^6\,\mathrm{km}^2$ at its greatest extent (following Cantor et al. (2001)) but did not significantly increase the atmospheric dust loading over the majority of the planet[3]. The sequence shown for MY24 actually features two storms that formed in quick succession: the first began in the northern Chryse region, moved

[3]The classification of storms is somewhat ambiguous; some storms classed as regional may in fact cover most or all longitudes, but in a limited latitude range. The term *planet-encircling* is often used in the present discussion as a conservative description of large, zonally extensive dust storms that may not qualify as being truly 'global'.

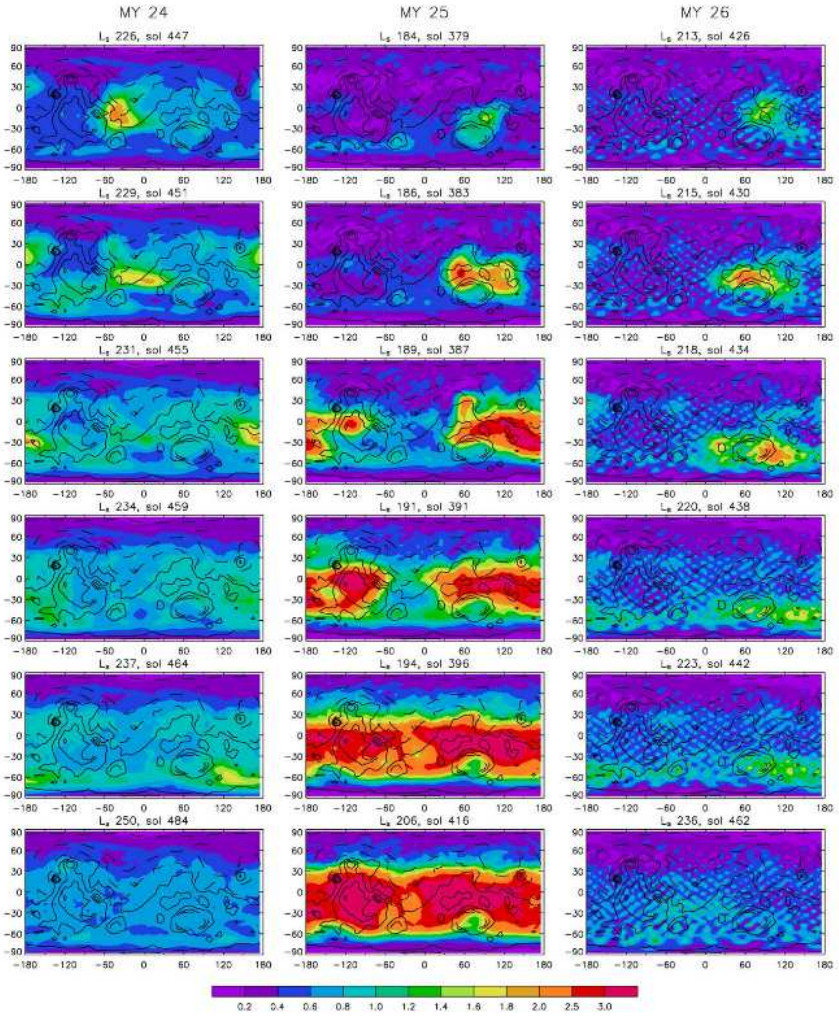

Figure 1.4: Progression of assimilated visible dust opacity at 610 Pa over several weeks in each of the three years seen by TES, showing regional and global storm development (from Montabone et al. (2005)).

southward and was dispersed zonally upon reaching 30°S (Liu et al., 2003), by which point a second peak in opacity had formed at Amazonis (180°E, 0–30°N). This storm also travelled southward and spread in the zonal direction, at 60°S, dying out in the process. Both of these events are examples of *flushing storms*, which form when frontal systems at the edge of the seasonal north polar cap raise dust and are subsequently advected southward by the mean flow, along one of the topographically steered WBC channels identified above. The path of a dust front is dependent on the local time of day, and the amplitude and dominant wavenumber of baroclinic eddies active in the region, as significant equatorward advection is achieved through successive instances of positive interference between tidal and travelling

waves (Wang et al., 2003). Flushing storms are observed only in early-to-mid autumn and mid-to-late winter, and are notably absent during the period $L_s \sim 240$–$300°$ (Wang et al., 2005).

In contrast to the other two years shown, MY25 contained a period in which visible opacity was raised significantly over virtually the entire planet: a global dust storm (GDS). The storm of MY25 was in fact the earliest GDS ever seen, beginning just after $L_s = 180°$. The initiation of this storm occurred with dust lifting on the northwestern rim of the Hellas basin (not shown in Figure 1.4; see e.g. Liu et al. (2003)), which spread eastwards around the basin. Through a combination of transport and further dust lifting away from the initial source, the storm cloud reached Daedalia Planum, where a secondary lifting site had already activated ($L_s \sim 190°$), and the storm soon became planet-encircling. Opacities peaked at $L_s = 195°$ but remained elevated globally past 240° (Cantor, 2007). Reasons suggested for the decay of large dust storms such as this have included an exhaustion of surface dust in source regions, and a decrease in peak near-surface windspeeds (i.e. weakened eddies) caused by an increased atmospheric static stability under conditions of high dust loading (Basu et al., 2006).

A comparison of MY25 and MY26 makes it clear that superficially similar storm initiation phases can lead to very different development. Both of these years saw opacity anomalies form in the Northern Hellas/Isidis region in early spring, yet after around 20 sols, one had grown into a global storm while the other had largely decayed after failing to expand away from Hellas. The factors affecting the growth or non-growth of major dust storms are currently the focus of much research and speculation, and are likely to include seasonal factors in the form of circulation patterns and CO_2 cap extent, transient wave activity (Basu et al., 2006) and surface dust cover (Szwast et al., 2006).

1.2.2 History of dust storm observations

The three years presented in Figure 1.4 alone give a sense of the very large interannual variability that characterises Martian dust storm activity. While the occurrence of some dust lifting in each year appears to be inevitable, global or planet-encircling storms are seen in certain years only, and not in others. To learn something about the likelihood of a GDS occurring in a given year, at least in a statistical sense, it is necessary to examine

the long-term trend in atmospheric opacity. As it is only in recent years that continuous spacecraft coverage of the Martian atmosphere has been possible, gathering a significant number of years of dust storm data requires drawing on the various observational methods that have been used to study Mars over the past century.

Early observations

Records of seasonal changes in the visible appearance of Mars began as early as the 19^{th} century, courtesy of astronomers such as Percival Lowell and E. M. Antoniadi. However, observations were limited in their clarity by the apparent size of Mars even at its closest approach (opposition), and by confusion between the emergence of dust clouds or hazes and albedo changes at the surface (Martin and Zurek, 1993); thus, the only observed pre-war storm believed with some confidence to have encircled the planet was in 1924, starting at $L_s = 310°$ (Phillips, 1924). A major dust storm that was certainly planet-encircling was widely observed in 1956, in the Mars year denoted MY1 (de Vaucouleurs, 1958). The storm began at $L_s = 254°$, featured active regions both at Hellespontus and near Argyre (Heintz, 1958), and lasted until $L_s \sim 304°$.

The spacecraft era

In 1971 (when viewing conditions once again became favourable) a huge global storm that spanned $L_s = 260$–$329°$ (MY9) was observed from Earth. It originated in the Hellespontus region and expanded to the west and south, forming a secondary core at Argyre (Martin, 1974b). Earlier in the year, it was noted, a regional storm had also formed at Hellespontus, but did not become global (Capen, 1974). A summary of the active regions during the 1971 global storm is shown in Figure 1.5. Very rapid expansion was noted by Martin (1974b) from about the fifth day of the storm, enabling it to encircle the globe within 15 sols. Derived cloud advancement speeds of as much as 700 km/h strongly indicated that the propagation of the storm involved significant additional lifting, rather than simply the dispersion of dust from the initial site.

The Mariner 9 spacecraft, the first to successfully orbit Mars, arrived midway through the 1971 storm, and was able to map temperature and dust opacity data during the decay period of the storm using its InfraRed Interferometer Spectrometer (IRIS) (Hanel et al.,

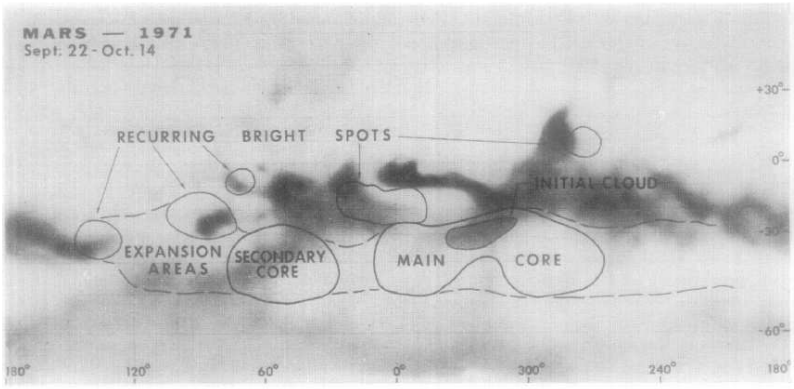

Figure 1.5: Latitude-longitude map showing the lifting centres of the 1971 global storm (taken from Martin (1974b)). Shading represents surface albedo.

1972). Work utilising this data (though now known to come from an abnormally dusty period) led to a modified-gamma distribution fit to dust particle sizes between 1 and 10 μm (Toon et al., 1977), and a reference vertical profile of mixing ratio that has been used by modellers ever since (Conrath, 1975). The following Mars year (MY10), after the termination of Mariner 9's mission, another planet-encircling storm was observed from Earth. Largely confined to the southern hemisphere, it apparently began (at $L_s \sim 300°$) with the almost-simultaneous formation of dust clouds at Solis Planum, Noachis and Hellas (Martin, 1974a).

The 1975 Viking mission was launched in two stages in quick succession, and in 1976 (MY12) a pair of orbiters, Viking 1 and 2, reached Mars. Each satellite also sent a lander to the surface; both were stationed in the northern hemisphere: Viking 1 (VL1) at Chryse Planitia ($\sim$ 22°N) and Viking 2 (VL2) at Utopia ($\sim$ 48°N). Between them, the orbiters returned enough data to cover (for the first time) an entire Mars year, over MY12–13; meanwhile, VL1 remained operational until November 1982, extending the mission into MY15. Two planet-encircling dust storms were recorded in 1977, through visible opacity measurements from cameras on the two landers (e.g. Colburn et al. (1989)) and from the Infrared Radiometers for Thermal Mapping (IRTM) on the two orbiters (Martin et al., 1979) (as well as visual observations by the orbiters). The first storm began at $L_s \sim 202°$, and the second, larger storm at $\sim 269°$, both at or to the west of Argyre (Thorpe, 1979). The second 1977 storm produced low-latitude 9 μm opacities as large as 2 (Martin et al., 1979). Although the orbiters had failed and visual data was no longer available from the

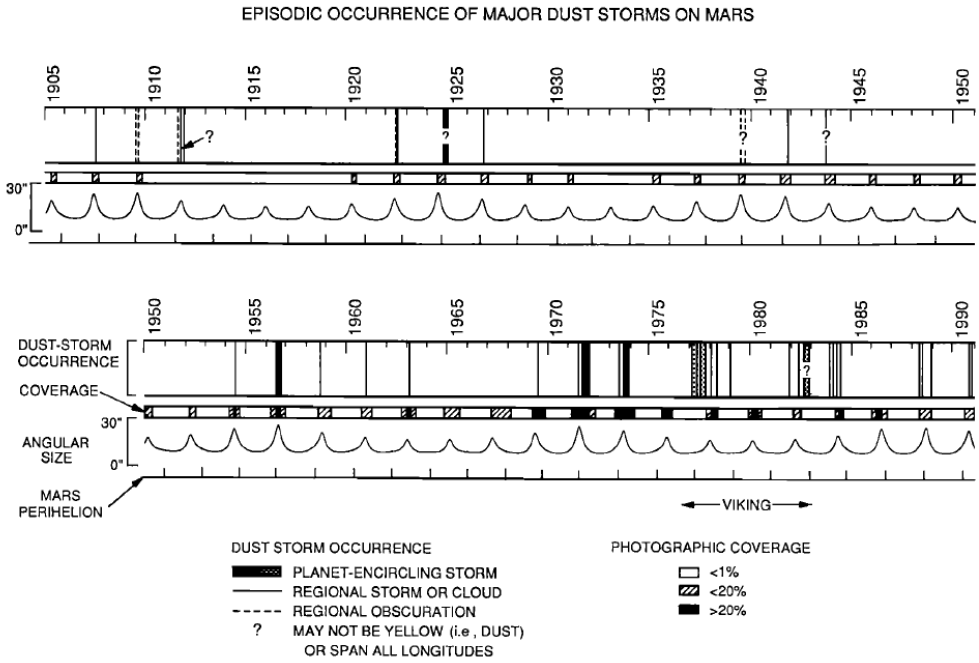

Figure 1.6: Timeline highlighting the major dust storms observed before 1990, and the changes in viewing conditions that occurred over this period (from Zurek and Martin (1993)).

landers by 1979/80, a further planet-encircling dust storm was detected in 1982 (MY15), occurring at a similar seasonal date to the first storm of 1977 (e.g. Tillman (1988)). The identification was done using the semidiurnal signal recorded by VL1's pressure sensor, which was shown to be a robust indicator of the onset of large dust storms (Zurek, 1981).

The observational record up to and including the late 1980s was neatly summarised by Zurek and Martin (1993) in the timeline shown in Figure 1.6. The group of very large storms seen during the Mariner and Viking missions in the 1970s led to the suggestion that such events were common and perhaps occurred every year (Zurek, 1982); on the other hand, Martin (1984) noticed that very few storms of comparable size had been seen from Earth over the preceding century, raising the possibility of a long-term climatic shift on Mars, from largely dust-free to highly dusty. However, Zurek and Martin (1993) highlighted the restriction on identifying changes in opacity that has been imposed historically by the cyclically varying angular size of Mars — in some years, it only briefly covered more than ~ 10 arcseconds of sky, making visual observation and photography from Earth difficult or impossible. In addition, oppositions during these 'distant' years are timed to occur during

the Martian aphelion season, away from the potential occurrence of any dust storms[4]. It can be seen from Figure 1.6 that those storms that have been detected from Earth, such as the one in 1924, occurred while favourable viewing conditions existed, i.e. near perihelic oppositions, which occur every 15–17 years[5]. The conclusion of Zurek and Martin (1993) was that it is not possible to comment on dust storm activity during many of the years in the Earth observational record. Taking this bias into account, the authors calculated that the probability of a planet-encircling dust storm occurring in a given Mars year is around one in three, and most likely between 18% and 55%. While some basic properties of Martian dust storms were known by this stage, it was clear that a much greater density of observational coverage over successive years would be needed to determine the true frequency of occurrence of global and regional storms.

Modern methods

In the past decade, this has become a reality. In addition to further opacity observations from ground-based and space-based telescopes (James and Lee, 1999) (though both are often still limited in their view of the perihelion season), since 1996 there has been a continuous orbiter presence around Mars, beginning with MGS and followed up by Mars Odyssey, Mars Express (MEx) and most recently Mars Reconnaissance Orbiter (MRO). Each of these has carried instruments to map aerosol (dust and/or water) optical depth, and the missions have been invaluable in providing a near-unbroken record of tropical and midlatitude dust opacities that now covers more than six full Mars years.

MGS reached Mars in September 1997 and began its full science mission in early 1999, which included high-resolution visible imaging of dust storms (including local events) from Mars Orbiter Camera (MOC) (Malin et al., 1998), supplementing TES maps of temperature and column-integrated dust opacity. In addition to the dust storms of MY24–26 described above, TES was able, during its aerobraking phase, to observe in detail a large regional storm in southern spring of MY23, beginning at $L_s = 225°$ in the Noachis region. This storm could reasonably be described to have been planet-encircling, as opacities above 0.5 were detected at all longitudes south of 40°S, though it was certainly not 'global'. A

[4]Note that, for this reason, the confinement of major dust storms to the period $L_s = 180$–$360°$ could only be confirmed as fact when the Viking mission observed several full Martian years in situ.

[5]As an example of the limitations of the ground-based telescopic method, note that none of the large Viking-era storms (in 1977 or 1982) could be detected from Earth.

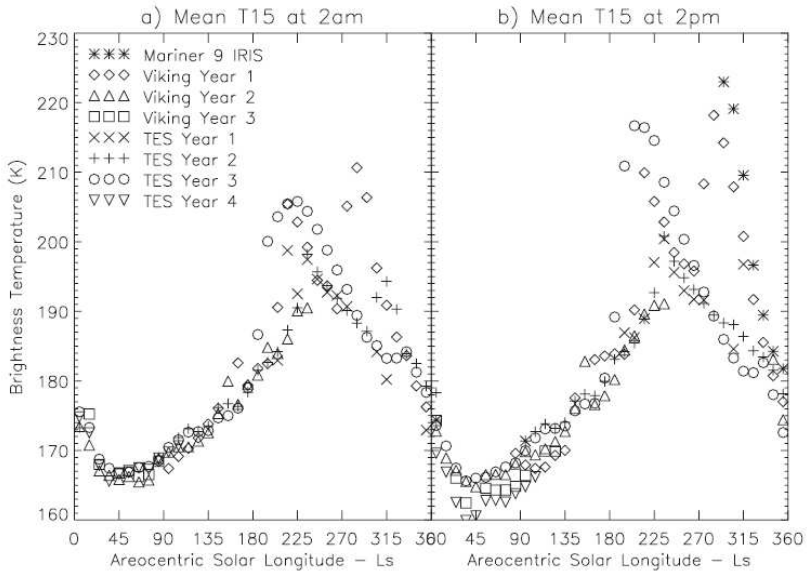

Figure 1.7: A summary of 60°N–60°S-mean 15 micron brightness temperature, representative of air temperatures in the lower/middle atmosphere, from multiple years of spacecraft observations (MY9–10, 12–14, 23–26) (taken from Liu et al. (2003)).

similar, more compact regional storm was seen later in the year, at $L_s = 309°$, forming northwest of Argyre (Smith et al., 2000). Thus, MY23, along with MY24 and MY26 (both of which also featured late-season regional storms), was a year without a GDS. The four TES years (MY23–26) are compared to the earlier periods observed by spacecraft in Figure 1.7. It can be seen that daytime atmospheric temperatures were greatly enhanced during the major storms of MY9, 12 and 25. Smaller responses to the various regional storms can also be detected (note that data from the second half of MY26 are not included in the plot). The figure also illustrates the high (low) degree of interannual variability that exists over $L_s = 180$–$360°$ (0–180°).

Concurrent with the end of the MGS lifetime, Mars Odyssey arrived in 2001 and began mapping dust opacity in late-MY25, using the THermal EMission Imaging System (THEMIS), a visible/IR multiband camera (Christensen et al., 2004). The overlap in viewing time of THEMIS and TES allowed comparison of the two sets of opacity measurements, between which close agreement was found (Smith et al., 2003). THEMIS has been observing for over five Mars years (though with more selective latitudinal coverage than TES) and has allowed monitoring of dust storm activity beyond the end of the TES observations in early MY27 (Figure 1.8). Like several preceding years, MY27 featured a regional storm

both before ($L_s = 225°$) and after ($L_s = 310°$) southern summer solstice. By contrast, a planet-encircling storm occurred in MY28 — it began at $L_s \sim 265°$ at Meridiani/Chryse and moved southward, enveloping most of the southern hemisphere before spreading northward. It was noted that regional storm activity around $L_s = 220°$ was less intense than in previous years, and there was an absence of significant dust lifting in the later part of MY28, while the planet-encircling storm was still decaying (Smith, 2009).

MEx dust opacity measurements have been made using the OMEGA imaging spectrometer, focusing on the south polar region (e.g. Vincendon et al. (2008)). MRO has recently provided important global-scale aerosol vertical structure through Mars Climate Sounder (MCS) (McCleese et al., 2007), covering multiple Mars years, allowing detailed study of seasonal variations in dust loading and the associated temperature response. MCS data are utilised in Chapter 6, where they are described further. Today, Martian dust activity is monitored more closely than ever before, by three orbiting satellites (Odyssey, MEx and MRO) and by the still-active Opportunity Rover (Lemmon et al., 2004).

1.2.3 Dust storm climatology

Through decades of observations of varying detail and coverage, a climatology of dust storm activity is in the first stages of being formulated. The basic features of this, as described throughout this section, are the following:

- dust storms of regional size or greater occur almost exclusively during the period $L_s =$ 180–360°, but show a wide range of variability in their timing and peak magnitude
- global dust storms (with some uncertainty over precise size classification) occur in approximately one of every three Mars years
- the largest dust storms tend to originate somewhere in the 30–60°S zone; Hellespontus appears to be the most popular initiation region
- significant dust storms regularly form near $L_s = 220°$ and 315°, at least in years without a GDS, and flushing storms are often at the heart of these events
- cap-edge dust lifting occurs throughout both autumn/winter seasons but shows distinct minima near the winter solstices

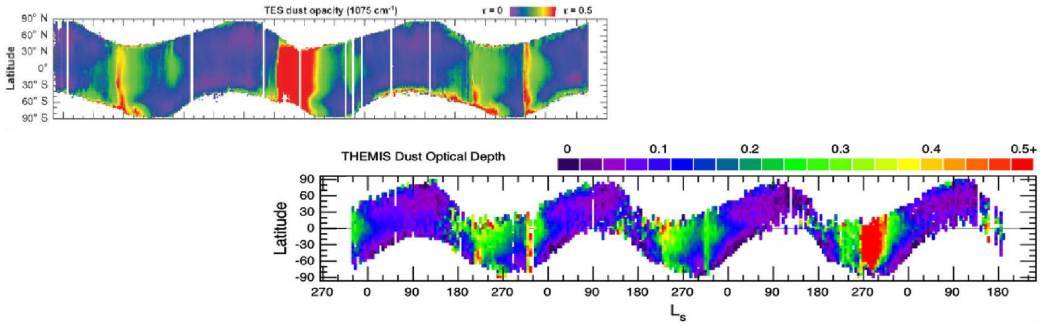

Figure 1.8: Zonally averaged visible (9 micron) dust optical depth from TES (Smith, 2008) and THEMIS (Smith, 2009), covering more than five Mars years, from $L_s = 120°$ of MY24 to 180° of MY29 (MY26 was observed by both instruments).

- elevated opacities following dust storms always decay fully by $L_s = 0°$, leaving low opacities (~ 0.1) away from the poles during $L_s = 0$–180°

Models of the Martian atmosphere must look to replicate as many as possible of these features, in order to be confident of giving a realistic representation of the important processes involved in dust storms.

1.3 Martian atmospheric modelling

Numerical simulation of the Martian atmosphere began with the fully nonlinear, two-layer model of Leovy and Mintz (1969), which produced quite realistic stationary waves and synoptic weather systems, and which for more than a decade was the only model of its kind. The model has today evolved into the NASA-Ames Mars General Circulation Model (MGCM) (e.g. (Pollack et al., 1990)). In the late 1980s and early 1990s several other models, which remain at the forefront of Mars atmospheric research today, were initiated, at the Geophysical Fluid Dynamics Laboratory at Princeton, USA (Wilson and Hamilton, 1996), the Laboratoire de Météorologie Dynamique in Paris, France (Hourdin et al., 1993), and at the University of Oxford, UK (Collins et al., 1996). Each of the models was converted from a terrestrial GCM: in the case of the Oxford model, that of Hoskins and Simmons (1975), which forms the dynamical core of the GCM (see §2.1).

Early modelling work focused on the effects of topography (Pollack et al., 1981) and atmospheric dust (using a prescribed spatial distribution, in a zonally symmetric model) (Haberle et al., 1982) on the large-scale circulation, with a view to explaining observations

such as those taken by the pressure sensors on the Viking landers. The first 3-D MGCM dust transport simulations came from Murphy et al. (1995), using the NASA-Ames MGCM, and included a particle size distribution and radiative feedbacks, but with a simple prescribed source of dust. Newman et al. (2002a) introduced a dust lifting and transport scheme to the Oxford MGCM, by parameterising the two main mechanisms proposed for dust lifting, namely dust devils (convective lifting) and near-surface wind stress. This approach has become the standard method used by other dust lifting MGCMs.

Typically, such dust-transporting runs now use radiatively active dust (Newman et al., 2002b) and some size dependence. Kahre et al. (2008), using a ten-bin particle size distribution in the NASA-Ames MGCM, showed that observed variations in atmospheric effective radius (namely, an increase during large storms) could be replicated without requiring any temporal variation in the size distribution lifted. The same model was used to consider a possible limit on the amount of dust available for lifting from the surface (Kahre et al., 2005). While such a constraint obviously does exist in reality, the use of finite dust amounts at each model gridpoint failed to produce a realistic multi-year dust cycle — areas important for wind stress lifting became depleted over several model years, without sufficient replenishment through deposition of dust carried from other regions. Therefore, dust storms ceased to occur after several model years. Basu et al. (2006), using the GFDL MGCM, 'tuned' several lifting parameters in order to best reproduce recent optical depth measurements, and found that a sufficiently high threshold for wind stress lifting can prompt the model to exhibit some interannual variability, with global dust storms produced in some but not all model years. However, these storms appeared, particularly in their decay phase, too closely linked to the annual insolation cycle, and there was limited spatial variation in initiation sites.

Modelling the water cycle has arguably been more successful in recent years, perhaps because the first-order features that must be (and have been) reproduced, the north polar vapour maximum and the ACB, show very little interannual variability. The approach has been to carry both water vapour and ice as tracers, with mass transferred from the former to the latter under conditions of supersaturation, either simply so as to keep the gridbox at saturation (Richardson and Wilson, 2002b) or utilising a microphysical scheme to calculate condensation rates (Montmessin et al., 2004). It has been found that allowing the ice

particle radius to vary, depending on supersaturation ratio and/or ice nuclei abundance, is necessary to accurately model cloud optical depths (Montmessin et al., 2004). The ice nuclei distribution has so far been prescribed to reproduce climatological dust opacities (Nelli et al., 2009). Most commonly, clouds have been radiatively inactive in model runs to date. Diffusion of water into and out of the subsurface regolith has been considered (Zent et al., 1993; Böttger et al., 2005), but is generally not included in current MGCMs.

1.4 Key questions

This project focuses on dust storm variability, both spatial and temporal, and the relationship between atmospheric dust and water ice clouds. Questions that it seeks to address can be summarised as:

1. How do dust cover and other surface properties affect major dust storm occurrence, and its interannual variability?

2. Why do cap-edge dust lifting rates vary as they do through autumn and winter, can models reproduce this variation, and is it related to the occurrence of larger dust storms?

3. How important is the interaction between atmospheric dust and water ice clouds, and what effect does this have on dust lifting and transport in the winter hemisphere?

The first of these is guided by the list of observed dust cycle characteristics presented in §1.2, several of which have not been satisfactorily simulated by free-running dust lifting GCMs thus far. In particular, no single model run has produced anything close to the range of storm sizes, timings and initiation regions seen within even a few Mars years. An important reason for this is likely to be the models' assumption that dust lifting is not constrained by the availability of surface dust at key locations, one which seems very dubious, for reasons discussed in Chapter 4. Another significant factor in determining the spatial distribution of common dust lifting sites could be a variation in surface properties, in particular the size distribution of dust and rocks present. The potential importance of such variation is explored later in Chapter 4, drawing on a substantial body of experimental and theoretical literature that exists concerning the modification of the near-surface flow

by nonerodible obstacles, and its effect on dust lifting. Preceding this, Chapters 2 and 3 introduce the model used in this work, review the previous dust lifting formulations and describe initial dust lifting work, highlighting some of the limitations of the 'infinite surface dust' approach mentioned above.

The occurrence of flushing dust storms during northern autumn and winter, and of circumpolar dust lifting in general, is strongly dependent on baroclinic wave activity along the boundaries of the seasonal polar caps. The weakening of these waves around winter solstice can have a significant impact on global dust loading, and therefore the reasons for the formation of this solsticial pause are the focus of Chapter 5. At this point, radiatively-active water ice clouds — mostly neglected until recently by models, but now appreciated to exert a significant impact on the winter polar atmosphere in particular — are introduced to the model, in a series of prescribed dust opacity simulations. Chapter 6 combines the water ice clouds with the dust transport model of earlier chapters, and investigates the effects on dust lifting patterns. The incorporation of atmospheric dust particles into nucleating ice crystals is also modelled, and the possible efficiency of the process is discussed. Mars Climate Sounder aerosol data is used in this section to assess the new dual-transport model's accuracy. Of particular interest is the ability of the model to reproduce the very clear winter poles seen by MCS. Finally, Chapter 7 draws together the results of the project and suggests some future directions for modelling of the Martian tracer cycles, particularly with respect to dust lifting.

Chapter 2

Dust lifting in a GCM

2.1 The UK Mars General Circulation Model

The model used in this project was the UK version of a Mars General Circulation Model shared between teams at LMD, Oxford, the Open University and the Instituto de Astrofísica de Andalucía in Spain, henceforth referred to as the UK Mars GCM, or UKMGCM. There are two versions, the other based at LMD, each with a distinct dynamical core. The physics package is the common element between the models, and includes contributions from all four groups. While the LMD model uses finite-difference methods for solving the hydrodynamical equations, by discretisation onto a longitude-latitude grid, the UKMGCM uses a pseudo-spectral dynamical solver, in which horizontal fields are represented by a series of spherical harmonic basis functions. The differing model cores provide a useful means of comparison to identify any weaknesses or idiosyncrasies inherent in either numerical method (Forget et al., 1999).

2.1.1 Dynamics

The dynamics of the UKMGCM involve solution of the hydrostatic primitive equations, discretised spectrally as mentioned above, for the variables temperature, ln(pressure), absolute vorticity and divergence, based on the model of Hoskins and Simmons (1975). Temperature, pressure and wind fields are then transformed onto a finite-difference physical grid, their increments are calculated for the physical timestep, and the tendencies (rates of change)

of each are transformed back into spectral space. An advantage of the spectral method is that derivatives and interpolations can be accurately calculated from the smoothly varying variable fields; the pseudo-spectral model also avoids the 'pole problem' suffered by finite difference models.

The model uses the terrain-following sigma system as the vertical coordinate (in finite difference form, with typically twenty-five levels, unevenly spaced so as to put more levels close together near the surface), where sigma is defined as $\sigma = \frac{p}{p_s}$, the pressure at a particular height divided by the surface pressure. Model geometric height then varies with surface elevation and atmospheric conditions, but typically extends to around 100 km. In the upper levels, a sponge layer is used to reduce unphysical reflections of vertically propagating waves from the top boundary. Horizontal resolution is determined by the number of basis functions retained in the calculation of the spectral fields, and with the triangular truncation method used here, resolutions are referred to by the notation Tn, where n is the total wavenumber. Resolutions used in this work range from T21 ($5° \times 5°$ dynamical gridpoint dimensions) to T63 ($1.875° \times 1.875°$). Physical grids are of slightly lower resolution than the dynamical grids.

2.1.2 Physics

There are two distinct builds of the UKMGCM used during the work described in this thesis, the difference between the two being in the physics package used. The older model version was used in Chapter 3 and most of Chapter 4, after which point the physics were upgraded, and the newer model adopted. The changes principally concerned the radiation scheme and the handling of tracers. The two model versions are referred to within the UK team as v3 (older physics) and v5 (newer physics), so these titles will be used in this thesis to make clear when each has been used.

Radiative transfer

The radiative transfer scheme in the UKMGCM calculates atmospheric absorption and emission due to carbon dioxide, dust and, in the latest version, water ice. The CO_2 radiative code uses an improved version of the wideband 15-μm scheme from Hourdin (1992). Near-infrared CO_2 absorption of incoming solar radiation, including non-local-thermodynamic-

equilibrium effects, becomes significant at altitudes greater than 50 km, and is calculated by the model using a parameterised form of the radiative transfer calculations of Lopez-Puertas and Lopez-Valverde (1995).

The effect of solar radiation on atmospheric dust is calculated following methods used in terrestrial GCMs, with upward and downward fluxes at each level found using the Delta-Eddington approximation. Two visible and two infrared (IR) bands are used, in addition to the 15 micron band, in which only dust absorption is considered (no scattering). In v3, dust single scattering properties from Ockert-Bell et al. (1997) and Forget (1998) are used in the visible and IR bands, respectively, and the two datasets are merged by fixing the ratio of visible to IR opacity, $\tau_{0.67\mu m}/\tau_{9\mu m}$, at a value of 2. Using these data, model atmospheric temperatures are too warm, so column opacities must be scaled accordingly. Subsequent measurements (Wolff et al., 2006) revealed atmospheric dust to be brighter than previously thought, and updated scattering parameters, including a larger single scattering albedo, are used in v5, removing the warm bias (Madeleine et al., 2011). In this version, scattering parameters corresponding to a gamma distribution of dust particle sizes, with effective radius $1.5\,\mu m$ and effective variance 0.3, may be used, or effective radius can be calculated at each gridpoint (when transporting dust) and scattering parameters retrieved from a look-up table. In this second case, then, $\tau_{0.67\mu m}/\tau_{9\mu m}$ also varies in space and time.

Water ice clouds can also be included in the radiative transfer scheme, using optical data from Warren (1984). Heating rates are again calculated in visible and IR bands, using scattering parameters calculated for the ice crystal effective radius in each gridbox via interpolation of a look-up table.

Boundary layer mixing

Since the horizontal resolution of GCMs is on the order of 100 km, only large-scale atmospheric motion is simulated directly, and localised phenomena, which can affect the general circulation, must be parameterised (Forget et al., 1999). Convection is accounted for using a simple adjustment scheme, which corrects unstable potential temperature profiles (where $\frac{\partial \theta}{\partial z} < 0$) back to the adiabatic lapse rate and estimates the vertical fluxes that would result from this instability. Turbulence within the boundary layer is handled using a turbulent closure scheme, allowing momentum, potential temperature and tracers to be mixed realis-

tically through the boundary layer. Newman (2001) found that vertical diffusion is crucial for raising a significant fraction of the dust lifted at the surface through the lowest model layers and into the bulk atmosphere, where it can be advected by the large-scale circulation.

Other physical parameterisations

Unlike terrestrial GCMs, models of Mars must include the periodic condensation and sublimation of a significant portion (up to a third) of the CO_2 atmosphere. Condensation occurs whenever a gridpoint falls below the CO_2 condensation temperature ($\sim 140\,K$ at Martian pressures), and latent heat is released to keep the gridbox at the condensation point. Under warming, sublimation occurs to keep the remaining CO_2 at the frost point. CO_2 snowfall is simulated, with descending solid particles allowed to resublime if they reach a sufficiently warmer lower gridbox, though CO_2 snow particles are not transported horizontally (their downward passage is assumed to be completed within a model physics timestep, which is at present an acceptable simplification given that the majority of CO_2 condensation occurs within the isolation of the winter polar vortices; however, the 3-dimensional transport of CO_2 ice particles is a likely addition to the model in the near future). It is assumed that snow particles are formed through heterogeneous nucleation and that there are abundant dust particles available to act as ice nuclei (the feedback on the dust particles is considered in Chapter 6 only). Surface albedo is increased where there is CO_2 ice cover, and emissivity is decreased during condensation to simulate freshly deposited snow and ice cloud effects (Forget et al., 1998).

Surface temperature is determined not only by incoming and outgoing radiation fluxes but also by thermal conduction in the soil beneath. The UKMGCM includes an 11- or 18-layer soil diffusion scheme, extending down to depths of order 10 m (depending on the thermal inertia of the surface gridpoint), so as to accommodate the penetration of the thermal wave on diurnal and annual timescales. Soil thermal inertia is vertically homogeneous in all the work presented here, and does not vary in response to the presence of surface ice. Gravity wave generation and topographic drag on the low-level flow, due to fine-scale topography within a surface gridpoint, are also parameterised, using the scheme of Lott and Miller (1997). Gravity waves propagate upwards and eventually break in the middle or upper atmosphere (according to the Lindzen (1981) saturation method that acts to keep

the Richardson number at critical value), exerting a force on the mean flow that is in the direction of the surface wind stress. Surface topography uses the data from Mars Orbiter Laser Altimeter (MOLA) (Zuber et al., 1992), and maps of surface albedo and thermal inertia (TI) come from TES (Mellon et al., 2000; Christensen et al., 2001).

2.1.3 Tracers

An atmospheric tracer is any component that moves with the circulation, possibly influencing it (active) or simply following the flow (passive). The UKMGCM contains a tracer transport scheme, and the associated physics necessary to model dust, water vapour, water ice, and a variety of chemical species. Tracers are advected using a semi-Lagrangian numerical method (Newman, 2001), including Priestley's method of mass conservation (Priestley, 1993).

Dust

The transport of radiatively active dust forms a major part of the work in this thesis. However, since atmospheric optical depth is such an important model forcing, a number of alternative methods of supplying this forcing are available to the model user, primarily to reduce the errors introduced into a simulation via the dust field by constraining it to observational data. Some amount of forcing by observations is also necessary if a comparison to data from a particular Mars year is desired, since dust opacity varies so dramatically from year to year.

The approach which has been used as standard in Martian GCMs for several years is to 'prescribe' an opacity field, either as a constant value (in this case, measured in the visible, at a reference pressure of 610 Pa) or as function of one or more of longitude, latitude and time. The function commonly used in this work replicates the observations of TES during one of three Martian years: MY24, MY25 or MY26. These fields were obtained from the results of the data assimilation project in which TES temperatures and column optical depths were inserted into the UKMGCM, using the technique of Analysis Correction (Lorenc et al., 1991; Lewis and Read, 1995; Lewis et al., 2007).

The 'MY24' opacity function, shown in Figure 2.1, actually includes values from MY25 during $L_s = 0$–$140°$, as TES data were not available for this period in MY24. Data were

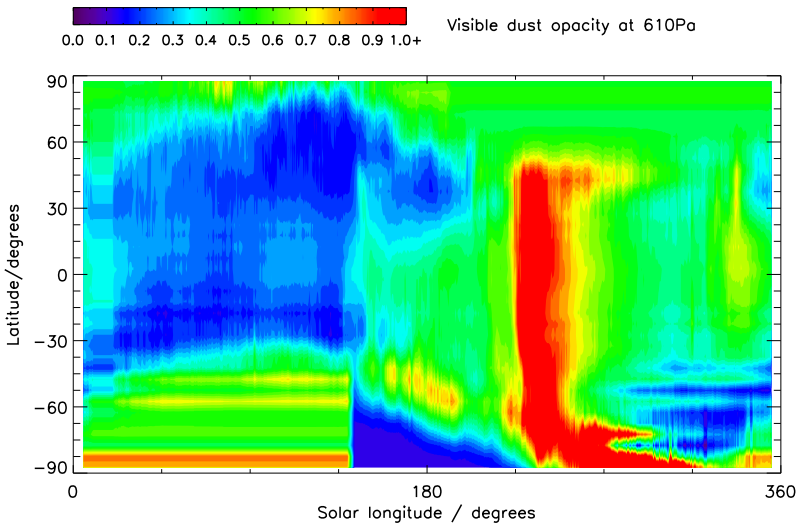

Figure 2.1: The 'MY24' dust scenario, zonally averaged: visible dust (extinction) optical depth at 610 Pa for a full Martian year, used as a model forcing in the absence of transported, radiatively active dust.

also not available at high latitudes during winter, due to operational constraints of the TES instrument: the approach used in this case was to keep the opacity at each gridpoint constant at the value it had when an observation was last available. This is likely to introduce significant error into the dataset, since some gridpoints necessarily retain the same opacity for almost half of the year; however the impact of this error on model circulation should be largely confined to the winter polar regions. Away from these areas, the MY24 dust scenario can be extremely useful, as results from model runs using the scenario may be compared directly to observations or assimilation results from the same year, and with the accurate column dust opacity forcing provided, any differences between model and observation point to other missing or incorrect physics.

When using this approach, the radiation code 'sees' a dust field made up of particles of radius 1.5 μm, whose mixing ratios q decay with decreasing pressure according to the standard 'Conrath function' (Conrath, 1975) (this is shown in Figure 6.8),

$$\frac{q}{q_0} = \exp^{0.007\left(1-\max\left[\left(\frac{p_0}{p}\right)^{\frac{70\,\mathrm{km}}{z_{max}}},1\right]\right)}, \tag{2.1}$$

where p is pressure, q_0 is the mass mixing ratio at a reference pressure p_0, and z_{max} defines the dust top height (in km), at which point dust mixing ratios have fallen to near-zero.

z_{max} is currently given by a simple analytic function that varies the height of the dust field in latitude and in time as a rough match to available observations, producing dust tops from 20 km in the clear season to around 70 km in the dusty season (Forget et al., 1999). Mixing ratios are then scaled so as to return the prescribed column optical depths. While the Conrath function was designed to best represent observations from Mariner 9, recent results from Mars Climate Sounder have suggested that it might not always be a good representation of Martian dust vertical profiles (Heavens et al., 2011).

A second, improved method utilises the TES dust scenarios but includes transport as well — dust is lifted at every surface gridpoint, at a constant rate, and transported both horizontally and vertically. However, at each timestep, before the dust field is passed to the radiative transfer scheme, mixing ratios are scaled at each gridpoint to fit column opacities to the TES observations. An advantage is that the vertical profile, both in mixing ratio and in particle size, is allowed to evolve naturally in the model, rather than being forced into the form described by equation (2.1); at the same time, the column-integrated dust optical depth is closely tied to observations.

The third and final option for model dust is full lifting and transport of radiatively active dust (Newman et al., 2002b). In this case, the intensity of surface dust emission is calculated using parameterised dust lifting equations, encompassing two different lifting processes (see next section). Once airborne, the dust is mixed upwards, advected, and settles under gravity. Heating and cooling rates are calculated at each physics timestep based on the mixing ratios of the freely evolving atmospheric dust field. This is the method that has been focused on in this work, and the following sections and chapters describe the development and use of lifting and other parameterisations related to dust transport, intended to improve both the stability and the variability of the modelled dust cycle.

The model may either transport one or more discrete particle sizes, or a particle size distribution, using a two-moment method in which a lognormal size distribution is assumed and dust mass mixing ratio and number mixing ratio are the two tracers used. The second method is newer and more efficient, though early runs in this work used the discrete-sizes method (see §3.1.3). Particles are mixed vertically by turbulent diffusion and convective adjustment, and sediment at a terminal velocity found by equating the Stokes drag force, scaled by the Cunningham slip-flow correction (Rossow, 1978), with the gravitational force

experienced by the particles:

$$w_{sed} = \frac{2g(\rho_d - \rho)}{9\mu}\left(\frac{D_p}{2}\right)^2\left(1 + \frac{8\lambda}{3D_p}\right), \tag{2.2}$$

where g is acceleration due to gravity ($= 3.72\,\mathrm{m\,s^{-2}}$ on Mars), ρ_d and ρ are particle density and atmospheric density respectively, μ is the gas molecular viscosity, D_p is particle diameter and λ is the atmospheric mean free path. Sedimentation is calculated using a 1-D, mass-conserving Van Leer I advection scheme (Hourdin and Armengaud, 1999).

Water

The model simulates the Martian water cycle (Böttger et al., 2005) using two tracers, water vapour and water ice mass mixing ratios, which undergo the same transport and (in the case of ice) settling processes as dust. The source of water for a model run is a prescribed residual polar ice cap occupying one-and-a-half of the northern-most latitude bands on the model grid, in a zonally asymmetric distribution. This surface ice sublimes in summertime to provide an upward vapour flux of

$$F = \rho C_d |u|(q - q_{sat}) \tag{2.3}$$

where ρ, u and q are the atmospheric density, wind velocity and vapour mass mixing ratio, all in the lowest model level, q_{sat} is the saturation vapour mixing ratio at the temperature of the surface and C_d is the drag coefficient. Vapour is then, like dust, vertically diffused up through the boundary layer, and transported by the advection scheme. If the vapour pressure of a particular gridbox is greater than the saturation vapour pressure, water vapour is converted into ice using a microphysical cloud scheme (Montmessin et al., 2002), which takes into account the number of ice nuclei available in the form of dust particles, whether as derived from the prescribed opacity field or as part of a transported dust field. Clouds may be radiatively active or inactive. Further details on this scheme are given in Chapter 6. Ice particles sediment under gravity and may resublime at a lower model level or reach the surface. Where surface water ice has formed to a depth greater than $5\,\mu$m, the surface albedo is adjusted to a value representative of water ice, unless CO_2 ice is already present.

Photochemistry

The model is also capable of simulating photochemical reactions, occurring primarily in the upper atmosphere, by carrying a range of chemical tracers, along with water vapour and ice (Lefèvre et al., 2004). The inclusion of CO_2 means that, if desired, the relative enrichment of non-condensable species at the winter poles (as a portion of the atmosphere is removed onto the surface causing the CO_2 mixing ratio to become significantly less than unity) can be simulated.

2.2 Previous dust lifting work

A dust lifting and transport scheme was first added to the UK model by Newman (2001). This was a thorough investigation into possible causes and mechanisms of emission of dust into the atmosphere, focusing on two lifting methods: convective lifting (dust devils) and near-surface wind stress.

2.2.1 Near-surface wind stress lifting

The theory behind the emission of dust particles from the surface due to the action of near-surface winds has been developed over the past seventy years, starting with the work of Bagnold (1954), with the primary goal of understanding and accurately modelling dust emission on Earth. A combination of experimental measurements of dust flux, from both field campaigns and wind-tunnel experiments (e.g. Gillette (1978)), and theoretical analysis of the physics involved in the process (e.g. Marticorena and Bergametti (1995)) has contributed to the formulation of semi-empirical expressions for the calculation of emitted dust flux, as a function of local surface wind velocity and a variety of other quantities.

Wind stress dust lifting has been observed to be a threshold-dependent process, occurring only when the drag velocity u_* exceeds a threshold value, u_*^t. Drag velocity is related to the atmospheric wind velocity near the ground — in the case of a global climate model, the velocity in the lowest model level, $u(z_1)$ — through

$$u_* = \frac{\kappa u(z_1)}{\ln\left(\frac{z_1}{z_0}\right)} \tag{2.4}$$

where κ is von Kármán's constant ($= 0.4$), z_1 is height of the lowest model level ($\sim 5\,\mathrm{m}$ in the UKMGCM) and z_0 is the surface roughness length (see §4.7). When the wind is strong enough, the process of *saltation*, in which larger sand particles flow across the soil bed, is initiated. During saltation, the particles collide with one another, and some are ejected upwards from the surface; however, since these particles are too massive to enter into suspension in the atmosphere, they return to the surface and, upon impact, give up their kinetic energy. Depending on the particle size distribution present at the surface and the properties of the soil, this can cause the ejection of micron-sized particles which are able to remain in suspension — this process is called *sandblasting*.

It is the combination of saltation and sandblasting that is assumed to be responsible for dust emission both on Earth and on Mars, because empirical and theoretical formulae for threshold drag velocity exhibit a minimum for sand-sized particles ($\sim 75\,\mu\mathrm{m}$ on Earth (Shao and Lu, 2000)), and surface windspeeds are generally too weak to raise micron-sized dust directly. Therefore, the vertical flux of micron-sized dust particles F may generally be expressed as

$$F(u_*, z_0) = \alpha_N(u_*)E(z_0)H(u_*, z_0)$$

where H is the horizontal (saltation) flux, α_N is the lifting efficiency and E is the fraction of the surface that is erodible (capable of emitting dust). The variables F, H and α_N technically also have a dependence on the size distribution of particles present on the surface; however, since no size-dependent parameterisation has yet been used for Mars (see §3.1.3), this has been omitted here for simplicity. Any size dependency can be thought of as being approximated by an implicit integration over particle size, with the assumption that the size distribution is homogeneous across the planet.

Newman (2001) assumed a uniform value of 1 cm for the roughness length z_0, and used a constant value of α_N, tuned to reproduce observed opacities. Therefore, the scheme of Newman (2001) can be further simplified to

$$F(u_*) = \alpha_N H(u_*).$$

The function chosen for H was that of White (1979):

$$H = \max\left[0, 2.61\frac{\rho}{g}(u_*)^3\left(1-\frac{u_*^t}{u_*}\right)\left(1+\frac{u_*^t}{u_*}\right)^2\right]. \tag{2.5}$$

Lifted flux F is therefore non-zero only when $u_* > u_*^t$. The threshold drag velocity u_*^t was calculated using the semi-empirical formula of Greeley and Iversen (1985):

$$u_*^t = A\sqrt{gD_p\frac{\rho_d-\rho}{\rho}} \tag{2.6}$$

where ρ_d is the density of Martian dust ($\sim 2500\,\mathrm{kg\,m^{-3}}$), D_p is the particle diameter and A is a factor with dependence on atmospheric density, particle density and diameter, and the interparticle cohesion parameter, I_p, which measures the tendency of small particles to stick together, as well as u_*^t itself. Calculation of A therefore required solving an equation (derived semi-empirically by Greeley and Iversen (1985)) using iterative methods. The result of this calculation is that high thresholds are predicted both for large and for very small particles, and lifting is predicted to be easiest for sand-sized particles with a diameter of order 10–100 μm. The size dependency (implied by the presence of D_p in the threshold expression) was removed by using the value of D_p for which u_*^t is lowest (i.e. $D_p \sim 100\,\mu$m).

2.2.2 Dust devil lifting

The second dust lifting mechanism considered was that of dust devils. These are convective vortices with low-pressure centers and strong updrafts capable of picking up dust and raising it as high as the top of the convective boundary layer (~ 1 km). Their horizontal extent can reach of order 100 metres. Martian dust devils have been observed by orbiting satellites (Fisher et al., 2005) and by surface landers (Greeley et al., 2006) (Figure 2.2).

Dust devils have been parameterised by modelling them as convective heat engines, following Rennó et al. (1998). Two different dust devil schemes were developed, one threshold-independent and the other threshold-dependent. In the first, vertical dust flux is given by

$$V_D = \alpha_D \eta F_s,$$

where α_D is a lifting efficiency (similar to α_N) η is the thermodynamic efficiency of the

Figure 2.2: A Martian dust devil in Gusev Crater in spring, as seen by the Spirit Rover (image courtesy of NASA/JPL/Texas A&M).

heat engine and F_s is the surface sensible heat flux. η depends on the pressure at the top of the convective boundary layer, and is largest when the boundary layer is deep. F_s is proportional to the temperature difference between the surface and the first atmospheric model layer.

The threshold-dependent method calculates the threshold vortex tangential velocity required to initiate lifting by using an empirical expression derived from laboratory experiments, and uses cyclostrophic balance to estimate the tangential velocity v_{tang} as ${v_{tang}}^2 \approx \frac{\Delta p}{\rho}$, where Δp is the difference in pressure between the centre of the vortex and the surrounding atmosphere. Δp can be calculated as a function of η, and it is possible to derive an expression for the lifted flux:

$$V_D = \alpha_D \frac{(\rho {v_{tang}}^2) - 15}{g}.$$

2.3 Successes and limitations of the existing scheme

The implementation of these schemes, using a single particle size and an unlimited surface dust supply, was described in Newman et al. (2002a,b). Various simulations were attempted, and some produced broadly realistic annual dust cycles, with low opacity during northern spring and summer and stronger lifting in southern spring and summer. The most successful setup used the non-threshold-dependent dust devil scheme to provide the aphelion background haze, and heavily threshold-dependent wind stress lifting to generate dust storms during the perihelion season. These included regional storms beginning at Hellas shortly after southern spring equinox (at a similar L_s and somewhat-similar location to the 2001 GDS) and flushing storms originating near Chryse (not dissimilar to the regional

storm of MY24, though occurring closer to summer solstice). Important wind stress lifting mechanisms were found to include the rising branch of the meridional overturning cell at $\sim 30°$S around southern summer solstice, western boundary currents in the major flushing storm channels at Chryse and Isidis Planitia, and slope winds at Tharsis and Hellas, including a thermal contrast contribution on the partially ice-covered slopes of the latter. Lifting thresholds (and specifically high thresholds) were necessary to achieve the rapid increases in opacity seen at the onset of dust storms. Similarly, dust storm decay through atmospheric stabilisation was found to be a viable shutdown mechanism for dust storms, provided dust lifting was sufficiently spatially limited by means of a high lifting threshold.

Dust devil lifting was found to follow the seasonal variation in the subsolar point due to the dependence of the parameterisation on the surface-air temperature difference, and annually averaged lifting rates showed peaks at 30°N and 30°S. The threshold-dependent version of the scheme resulted in more spatially confined lifting, as only gridpoints with a sufficiently deep convective boundary layer (high thermodynamic efficiency) were able to contribute. Attempts to recreate dust storm opacities in southern spring and summer using either of the dust devil mechanisms encountered the problem of generating overly large opacities in the aphelion season (using a constant lifting efficiency α_D), and so it was concluded that dust devils were unlikely to be responsible for storm generation. This has since been confirmed, as no correlation between dust devil activity and dust storm formation has been found in the observational record (Montabone et al., 2005; Cantor et al., 2006).

An important finding of the radiatively active simulations was the identification of feedback effects associated with each of the lifting mechanisms. For wind stress lifting, a strong positive feedback was noted — as dust is lifted, it warms the atmosphere and strengthens the PMOC and thermal tides, causing an increase in near-surface wind strengths, leading to further dust lifting (though slope winds and thermal contrast flows are weakened by a more isothermal, dusty near-surface atmosphere). Dust devil lifting, however, exhibits a negative feedback, since dust raised into the boundary layer causes the atmosphere near the surface to heat up, reducing the surface-atmosphere temperature gradient, and thereby reducing the heat input at the base of any convective vortices. The positive feedback that applies to wind stress lifting means it can initiate rapid increases in dust opacity during dust storms, while the negative feedback associated with dust devil lifting means it is a

self-limiting process, capable of providing a steady supply of dust to the atmosphere but not likely to be a major factor in dust storm generation and growth.

The work of Newman (2001) was a considerable step towards establishing an autonomous dust lifting and transport scheme within the UKMGCM, and the reproduction of several major features of the observed dust cycle is encouraging. However, the lack of interannual variability produced by multi-year model runs represents an important component of the Martian climate system that was missed by the model; as noted by Newman et al. (2002b), this could have been due to the neglected interaction with the water/ice and CO_2 cycles, the use of an unlimited surface dust reservoir, or unresolved local processes important for dust lifting. A greater catalogue of observational evidence is now available, and interannual variability can begin to be addressed through use of at least five continuous Mars years of column dust opacities. Also, useful lessons have been learned from other GCM dust lifting studies in recent years, concerning the use of particle size distributions and finite surface reservoirs. It was therefore timely to refine the existing dust lifting schemes, using the extra data available to constrain the model more closely and to evaluate its performance. This is done in Chapter 3, before new components are added to the lifting schemes in Chapter 4, and the interaction with water ice is implemented in Chapter 6.

Chapter 3

Initial dust lifting work

Since the work of Newman (2001), knowledge of the Martian dust cycle has increased, allowing greater constraints to be placed upon model parameterisations. Also, advances in computing power in the intervening years have made carrying more than one dust particle size a much more feasible option, and performing simulations at higher resolutions a possibility. A general evaluation of work done to date on the dust-transporting UKMGCM was carried out, resulting in the development of a new wind stress lifting scheme, which is described in the first section below. Other physical parameterisations remained the same as used by Newman et al. (2002a), but dust was now treated as a radiatively active tracer throughout, as in Newman et al. (2002b) and in subsequent work by other groups. The new parameterisation was tested through several multiannual simulations, each spun up from rest and each carrying a size distribution of dust particles, the results of which are summarised in §3.2. Horizontal spectral resolutions from T21 (as used by Newman et al. (2002b)) to T63 were used. A strong dependence of dust lifting rates on model resolution was discovered, and this is explored further in §3.3.

3.1 Developing the wind stress lifting scheme

It is now widely accepted that dust devils do not play an important role in the initiation of large dust storms (Basu et al., 2004; Cantor et al., 2006). This fact makes the setting of the dust devil lifting efficiency parameter considerably easier, since the most important contribution of dust devil lifting (in the model at least) is during northern spring and sum-

mer, L_s = 0–180°. The threshold-independent version of the dust devil parameterisation has been predominantly used during this work, but results with the threshold-dependent version were not greatly different — the role of dust devil lifting in the model dust cycle is to provide a fairly constant, low rate of lifting throughout the year, in order to maintain an appropriate atmospheric opacity during the aphelion season, when the model predicts little or no wind stress lifting.

Considerably more attention was given to the near-surface wind stress lifting parameterisation, since this is very likely to be the more important lifting mechanism for producing dust storms. The following parameterisation ignores the effect of surface roughness — this modification was added later (§4.7).

3.1.1 Calculation of lifting threshold

The estimation of the threshold wind velocity required for dust lifting and its application to the model are the most important considerations in a GCM dust lifting scheme. Ideally, the lifting threshold should be derived from physical principles, yet it must also be set at a level attainable by model winds with sufficient regularity to broadly reproduce the amount of dust lifting that is observed to occur.

As was explained in §2.2, it is assumed that saltation of sand particles is essential for the lifting of micron-sized dust particles; therefore, the minimum wind velocity required for dust lifting can be estimated by calculating the mobilisation threshold for particles with a diameter D_p of order 100 μm. In the work of Newman (2001), this threshold, in the form of a drag velocity, was given by equation (2.6). Several problems existed with this formulation. The interparticle cohesion parameter I_p has not been measured for Mars, but values of around $6 \times 10^{-7}\,\mathrm{N\,m^{-1/2}}$ are commonly used for both Earth and Mars (Greeley and Iversen, 1985). To allow any lifting at all it was necessary to use a value towards the lower end of allowed values; otherwise predicted thresholds were too high. Then, in subsequent years, it became apparent that the timing (within the model code) of calculation of the surface drag velocity needed to be corrected — it should be calculated after, rather than before, the vertical diffusion module is applied. This had the effect of lowering surface winds (due to upward turbulent mixing), making lifting more difficult. It was found that extremely small values of I_p were now necessary to allow lifting to take place.

A simpler method, which avoids the iterative calculation required previously, was found using the approach of Shao and Lu (2000). In their model of saltation, I_p is assumed to be a function of particle size, so does not explicitly appear in the formulae. It nonetheless produces a similar size dependence to the Greeley and Iversen (1985) method, with a minimum threshold for sand-sized particles in agreement with experimental studies (such as Fletcher (1976)). The threshold drag velocity for lifting is expressed as

$$u_*^t = \sqrt{A_N \left(\sigma_p g D_p + \frac{\gamma}{\rho D_p} \right)}, \tag{3.1}$$

where $\sigma_p = \dfrac{\rho_d}{\rho}$, $\gamma \approx 3{\times}10^{-4}\,\mathrm{kg\,s^{-2}}$ and $A_N \approx 0.0123$. The constants γ and A_N encompass particle arrangement, particle Reynolds number at threshold drag velocity and the inter-particle electrostatic force, and their values, which are not Earth-specific, were determined empirically by fitting to wind tunnel data. This formula can be differentiated with respect to particle diameter to find the minimum threshold velocity, for sand particles:

$$(u_*^t)_{min} = \sqrt{2A_N} \left(\frac{\gamma \rho_d g}{\rho^2} \right)^{\frac{1}{4}}, \tag{3.2}$$

which applies for particles of diameter

$$D_p = \sqrt{\frac{\gamma}{\rho_d g}},$$

or around $180\,\mu$m.

The dependence of the threshold velocity on near-surface atmospheric density ($\sim \rho^{-1/2}$) means that this is referred to as a 'constant threshold stress' approach, since surface stress is defined as $\zeta = \rho {u_*}^2$, so that threshold stress ζ^t is not a function of ρ. The previous method of calculating the threshold at each gridpoint had a more complex functional dependence on ρ, but actually produced relatively little deviation from a constant threshold assumption, so this simpler method sacrifices little accuracy. Indeed, Newman et al. (2005) reverted to such a method (though with a prescribed ζ^t), for experiments using altered orbital parameters, as did Kahre et al. (2006). In addition, using the Shao and Lu (2000) threshold requires no unrealistic assumptions of parameter values, as was previously the case with I_p.

However, thresholds calculated using equation (3.2) still proved to be too high for lifting

by model surface winds, except in very rare cases. This is not unexpected — two considerations, detailed below, can explain this discrepancy, and allow for reduction of these predicted thresholds to values more amenable to producing lifting in the GCM.

Sub-gridscale gustiness

General circulation models have typical horizontal resolutions of order 100 km, and physical timesteps of around 30 minutes, so miss the large degree of variability that exists on smaller scales, both spatial and temporal — surface wind is represented only by one (mean) value, so any local peaks or transient gusts in wind velocity at the surface are neglected. This is an important consideration for lifting, as it is a threshold-dependent process, so any lifting that does occur will depend only on the strongest winds that exist within a gridbox and during a timestep.

To account for this, an approach tried by Newman (2001) was to describe surface wind at each gridpoint using a Weibull probability distribution, applying some knowledge gained from Viking lander wind data. The total saltation (horizontal) flux was estimated by integrating over the probability density function, from threshold velocity to infinity. In this way, even if the mean surface wind was lower than the threshold, some fraction of the gridpoint still contributed a dust flux. The downside of this method was that lifting became ubiquitous, regardless of the surface wind speed, since there was always some contribution from the high-velocity tail of the distribution. This effect weakened the threshold dependence of the dust lifting parameterisation, and reduced interannual variability. The probability distribution approach also failed to reproduce the rapid increases in opacity observed during Martian dust storm initiation. For these reasons, it was removed from the wind stress lifting scheme (the issue of parameterising gustiness is returned to in §4.9).

Fluid and impact thresholds

Another factor requiring consideration is the difference in the effect of surface wind on stationary particles compared with its effect on particles in saltation. One may speak both of a *fluid threshold*, which is the drag velocity required to directly move/lift particles off the surface, and an *impact threshold*, which is the drag velocity needed to lift particles through collisions with other particles, or alternatively the drag velocity needed to sustain a flux of

saltating particles after they have been set in motion. The impact threshold is the lower of the two, since the contribution of the saltation flux enhances lifting, but on Earth the difference between the two thresholds is not thought to be very great (u_*^{it}/u_*^{ft} is around 0.8) (Kok, 2010).

Recent numerical studies (Almeida et al., 2008; Claudin and Andreotti, 2006; Kok, 2010) have found that the situation on Mars may be quite different, due to the combination of stronger windspeeds required for saltation and lower gravity allowing saltating grains to follow steeper trajectories and travel faster than on Earth. It has been estimated that the impact threshold may be less than 50% of the fluid threshold[1]. The Shao and Lu (2000) calculation applies to the fluid threshold, so using equation (3.2) as the criterion for dust lifting will overestimate the difficulty of lifting — in fact, surface wind need only exceed the fluid threshold briefly for saltation to begin, after which lifting can continue so long as the wind stays above the lower impact threshold. It seems, therefore, that a considerable hysteresis effect may be at work for Martian dust lifting, in which the occurrence of lifting at a particularly point in time is dependent upon the recent history of the wind stress at that point. This will apply particularly in areas where the temporal variability of surface windspeed is high, as even a relatively low time-mean windspeed may contain within it several gusts large enough for the drag velocity to exceed u_*^{ft} and allow dust lifting to continue for a significant period of time. However, this is a recent modelling development, and the implication of these numerical results is not yet clear.

Reduction to the calculated threshold

The lifting threshold formulation used at this point was therefore an approximation to both the issues mentioned above, applied in such a way as to remain simple whilst allowing some temporal variability in lifting. Both sub-gridscale gustiness and the effect of a low impact threshold should lead to more dust lifting than would result from their neglect, so the approach taken was to combine both these considerations into a single scaling factor that is applied to the fluid threshold (equation (3.2)). This is obviously a major simplification of what are undoubtedly important effects in Martian dust lifting, but at present it seems the most logical approach to take for a coarse-resolution GCM lifting scheme, until further

[1] $u_*^{it}/u_*^{ft} \approx 0.48$ according to Almeida et al. (2008); ~ 0.3 according to Claudin and Andreotti (2006); and ~ 0.1 according to Kok (2010).

information is available. In practice, the threshold was lowered until a reasonable amount of dust lifting became possible during a Martian year. The value of this scaling factor depends somewhat on resolution (§3.3), but at T31 a value of 0.7 was found to be sufficient to allow southern spring and summer dust storms to develop. This value is arguably close enough to unity to provide some confidence in the model mean drag velocities as being representative of the larger gusts within the gridbox that are actually responsible for dust lifting, through what may be a significantly localised (and time-dependent) saltation mechanism.

At T31, it was found that the range of scaling factor values that produced dust loadings that were at all realistic was fairly narrow. A value of $\gtrsim 0.75$ restricted dust lifting to only a few gridpoints for a brief period in southern spring, and did not permit the production of regional or global dust storms. Although storms could perhaps be produced by increasing α_N so as to allow the positive lifting feedback to take effect, the lack of wind stress lifting at other times in the year would remain, since the value of α_N only becomes relevant when model winds are strong enough to exceed the threshold. Using a threshold scaling factor of 0.6 or lower caused the production of dust storms during every year, as climatological model winds in southern spring midlatitudes exceeded the threshold for a sufficient length of time. Interannual variability in global dust opacity was small, because the internal variability of the surface wind field was unable to interact with the threshold value in such a way as to amplify this variability (through dust storm growth), as the threshold was set too low relative to the climatological wind stress. There was little benefit in attempting any further 'tuning' (attempting to fit resulting model dust loadings to observed climatology) around the compromise value of 0.7, given the simplifications mentioned above and the fact that, in reality, the threshold is unlikely to be spatially or temporally homogeneous. Instead, tuning was performed for a fixed threshold by varying α_N.

This end result is not dissimilar to the thresholds used in other models. Kahre et al. (2006) took a threshold stress value of 0.035 Pa from Greeley and Iversen (1985) and reduced it to 0.0225 Pa under similar considerations to those described above. In this work, the calculated threshold drag velocity equates to a threshold stress of ~ 0.041 Pa, which is then multiplied by $(0.7)^2$ to give 0.020 Pa. Basu et al. (2006) used a threshold of 0.055 Pa in order to maximise the interannual variability produced by the GFDL model — this is significantly higher than the other two figures, and apparently limited lifting in their model

(other than during dust storms, when the positive radiative feedback became active) to a very select number of gridpoints.

It is noted that an alternative theory on wind stress lifting, based on more recent visual observations (Sullivan et al., 2008), invokes direct lifting rather than saltation as the relevant process, by supposing that the clumping together of micron-sized dust particles on the surface produces larger, low-density aggregates that require lower wind stresses for lifting than individual dust particles (this is discussed further in §4.1). Using a density $\rho_d = 750\,\mathrm{kg\,m^{-3}}$ and particle size $D_p = 500\,\mu\mathrm{m}$ in equation (3.1) produces a threshold velocity $\sim 20\%$ lower than the value given by equation (3.2). This theory thus offers an alternative explanation for the need to reduce the threshold (3.2) for use in the model.

3.1.2 Saltation flux

When the surface wind stress at a particular gridpoint exceeds the threshold described previously (after it has been scaled appropriately), saltation and consequently dust lifting are assumed to occur. The commonly used relationship between sand flux H and drag velocity has been equation (2.5) (White, 1979). Experimental evidence (Iversen and Rasmussen, 1999) and more recent numerical modelling studies (Almeida et al., 2008) have suggested that the true dependence on wind shear may be slightly different, however. In particular, the numerical model of Kok and Renno (2008) was able to match a range of experimental results, and found a stronger saltation flux at wind speeds just above threshold than are produced by the White formula.

In addition, they considered the effects of electrostatic interactions between saltating particles and the surface, and found that this leads to a reduction in saltation intensity at higher near-surface wind velocities. During saltation, particles that are lifted, on average, acquire a negative charge, while the surface tends to become positively charged. The attraction between the saltating layer and the ground then acts to reduce the height of the saltating layer, reduces the saltating particles' velocities, and actually allows a greater concentration of particles near the surface, but overall reduces the intensity (measured as the height integrated mass flux) of the saltation, compared to the non-electrostatic case.

An approximate fit to the results of Kok and Renno (2008), including electrostatics, was

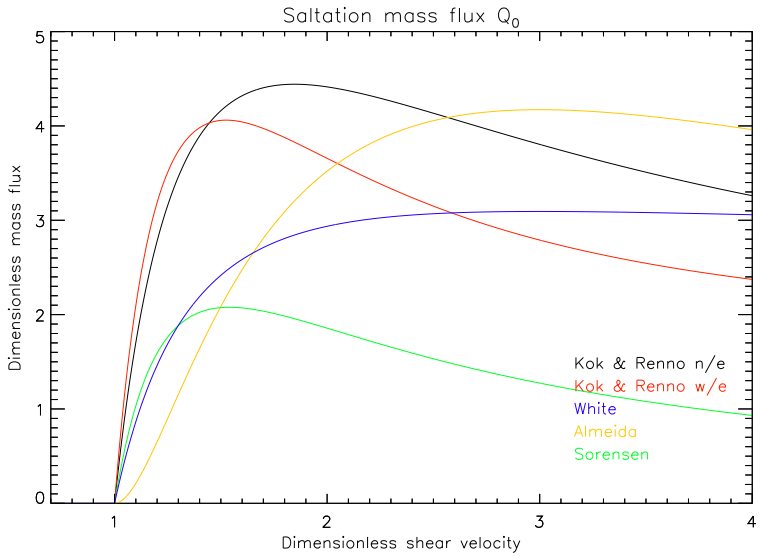

Figure 3.1: Various saltation flux functions, from White (1979); Sorensen (2004); Almeida et al. (2008) and Kok and Renno (2008) (with and without electrostatic effects), converted into a dimensionless flux $Q_0 = H(\frac{g}{\rho_d u_*^3})$ and plotted against dimensionless shear velocity u_*/u_*^t. The red curve is the one used in the UKMGCM lifting scheme throughout this work. (Note that the difference between the red and black curves at low u_*/u_*^t is not a 'real' effect and is merely an artefact of the analytic functions fitted to each set of results.)

found (by modifying the functional form of Sorensen (2004)) as

$$H = \max\left[0, 0.25\frac{\rho}{g}(u_*)^3\left(1-\left(\frac{u_*^t}{u_*}\right)^2\right)\left(7+50\left(\frac{u_*^t}{u_*}\right)^2\right)\right] \tag{3.3}$$

and inserted into the model dust lifting routine, to be activated whenever surface wind stress exceeds the threshold value.

A comparison of the functions (2.5) and (3.3), as well as several other forms suggested in the literature, is shown in Figure 3.1. Since the overall flux of dust lifted into the atmosphere is set proportional to the saltation flux using a tunable 'lifting efficiency' parameter, it is the shape, rather than the absolute magnitude, of the dimensionless flux curve shown that will determine the behaviour of the lifted dust flux. Therefore, it seems that the Kok & Renno function (the red curve) provides slightly more 'explosive' dust lifting than the White function when winds move just above threshold, while at higher velocities, perhaps $1.5u_*^t$ or greater, the effect of electrostatic interactions comes into play, limiting saltation and lifting somewhat. Model winds rarely reach this strength, so this seems unlikely to be an important effect for model (or actual) dust storm growth.

With a large degree of tuning involved in producing realistic model lifting, and a lack of detailed, in situ data on dust fluxes, asserting the validity of one saltation flux function over another is extremely difficult (even for dust emission on Earth) and seems not to be of primary importance in obtaining a first-order accuracy in the modelled dust cycle. Similar performance, including the possibility of interannual variability, was obtained using the White and the Kok & Renno functions (after adjusting α_N), indicating that the choice of function is not crucial at this stage, while other simplifications remain in the lifting parameterisation. The same could be said of the relationship between vertical (lifted) dust flux and horizontal (saltating) sand flux. Here, as in all other Martian lifting studies to date, the upward flux has been set proportional to saltation flux, with the constant of proportionality (the lifting efficiency, α_N) tuned to match resulting column dust opacities with those observed (seen in e.g. Figure 1.8). α_N may vary with drag velocity and/or location in reality (and 'effective' α_N must surely vary with time, as surface dust cover changes), but sufficient data on soil properties is not available with which to attempt to model these effects. The magnitude of α_N was largely deduced by tuning model wind stress lifting rates, given the threshold value explained previously. Again, this choice depends on model resolution as well as the particle size distribution being used (see below), and the possible use of a variable threshold (see §4.3), but values used for the work in this and following chapters were generally in the range $(1\text{–}50)\times10^{-5}\,\mathrm{m}^{-1}$ — at least of the same order as the terrestrial values of $\sim (2\text{–}5)\times10^{-5}\,\mathrm{m}^{-1}$ suggested by Gomes et al. (2003) or $2.75\times10^{-4}\,\mathrm{m}^{-1}$ measured by Gillette et al. (1997).

3.1.3 Size dependence

The work of Newman (2001) was carried out predominantly using a single particle size of $2\,\mu$m radius, but it included a preliminary attempt at carrying a more complete size distribution. This work has continued that approach. The majority of simulations described in this chapter and the next used six particle size bins, populated in one of two ways, as listed in Table 3.1. Here, using the older model physics of v3, the shortwave extinction coefficient, $Q_{ext}(0.67\,\mu\mathrm{m})$, was the only radiative parameter set to vary with particle size (as shown in the table). Sizes were chosen to fulfil a volume-ratio size distribution, i.e. for bin j, particle radius r_j satisfied $\dfrac{{r_j}^3}{{r_{j-1}}^3} = \text{constant}$.

Radius / μm	$Q_{ext}(0.67\,\mu\text{m})$	Lifted by
0.1	0.228	Both
0.251	2.27	Both
0.631	3.05	Both
1.58	2.39	Both
3.98	2.19	Both
10.0	2.10	Both

Radius / μm	$Q_{ext}(0.67\,\mu\text{m})$	Lifted by
0.251	2.27	Nsws
1.58	2.39	Nsws
10.0	2.10	Nsws
0.251	2.27	Ddev
1.58	2.39	Ddev
10.0	2.10	Ddev

Table 3.1: The two particle size distributions used when running with six size bins. The second was used to separate the effects of dust lifted by the two different mechanisms.

Martian atmospheric dust spectra have typically been fitted by observers (e.g. Tomasko et al. (1999)) by assuming either a modified-gamma or lognormal particle size distribution, with an effective radius $r_{eff} \sim 1.6\,\mu$m, fairly uniformly outside of large dust storms. An increase in average particle size, with $r_{eff} = 1.8$–$2.5\,\mu$m, has been noted during global storms such as that which occurred in 2001 (Wolff and Clancy, 2003). Assuming that the removal of dust particles from the atmosphere occurs primarily through their sedimentation under gravity, larger particles will be removed more quickly than smaller particles, since settling velocity is approximately proportional to particle size, for radii smaller than a few microns. This means that to produce the observed particle size for suspended dust, the distribution of dust as it is lifted from the surface must be shifted towards larger particle radii, in order to counteract the differential sedimentation that occurs once particles have been raised. Indeed, it was found during testing with the UKMGCM that a gamma-modified distribution with $r_{eff} \sim 2.5$–$3.0\,\mu$m for lifted dust could produce realistic atmospheric effective radii in both storm and non-storm conditions, similar to what was reported by Kahre et al. (2008).

The work of Chapter 6 used model v5, with a two-moment scheme for carrying a size distribution of dust particles. For this formulation, an effective radius and variance were specified for the lifted dust distribution (of $2.75\,\mu$m and 0.5 respectively), and a lognormal distribution was assumed (Hansen and Travis (1974) showed that lognormal and modified-gamma distributions give the same single scattering properties, provided the effective radii and variances are the same). The physics of v5 allow the use of size-dependent scattering properties, calculated online during simulation; however, the feature has not been used as standard in this work. Instead, when using the two-moment scheme, a single set of

scattering parameters were used for all gridpoints, corresponding to a lognormal distribution with effective radius 1.5 μm and effective variance 0.3. This was predominantly used in Chapter 6, where it did lead to a problem with atmospheric temperatures in the upper atmosphere — this is discussed in §6.2.3.

It is quite possible that the lifting mechanisms of dust devils and near-surface wind stress in fact tend to lift distinctly different particle sizes, but a lack of detailed observations of lifted dust flux leaves us without any knowledge of this relationship. Therefore, the same lifted size distribution was set for each of the lifting methods — a seasonal variation in atmospheric effective radius is still produced, as a consequence of the more concentrated nature of wind stress lifting during the dust storm season. The lifted size distribution included no dependence on drag velocity, due to the lack of information regarding regional or temporal variation of surface particle populations.

3.1.4 Summary of dust lifting scheme

Several aspects of parameterising dust lifting for an MGCM have been addressed. For estimating the threshold drag velocity necessary to initiate saltation, the formula of Shao and Lu (2000) was used, due to its simplicity in comparison with expressions previously proposed. Since model dust lifting using these threshold values was predicted to be very rare, the threshold velocities were reduced by around 30% to allow a suitable degree of lifting over the year. This was intended to account heuristically for the effects of sub-gridscale gustiness (since transient high-speed surface winds will tend to dominate dust emission) and the existence of a low impact threshold on Mars. The threshold choice and the overall scheme have a strong physical basis. The choice of saltation flux expression followed Kok and Renno (2008). While it is difficult at present to verify this flux dependence, particularly on Mars, it does slightly favour more explosive dust lifting in comparison with the previously used White (1979) formula, which would appear to be in qualitative agreement with what is observed in Martian GDS initiation. No size dependence has been used in the dust lifting calculations, due to a lack of knowledge of Martian soil properties, but the use of a fixed size distribution for lifted dust does replicate the general seasonal trend in atmospheric effective radius.

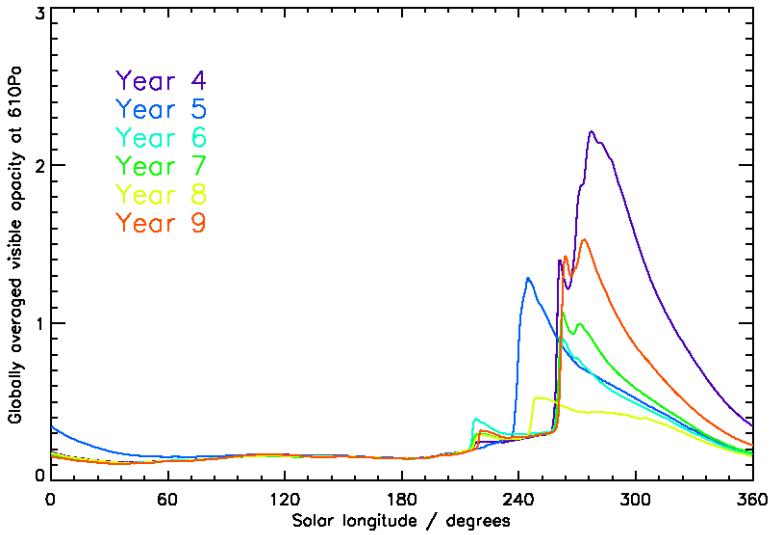

Figure 3.2: Six year series of globally averaged visible dust opacity, referenced to a 610 Pa pressure level, from a model run at T21, in which the threshold drag velocity was multiplied by 0.675, and lifting efficiencies $\alpha_N = 1.2 \times 10^{-4}\,\mathrm{m}^{-1}$ and $\alpha_D = 9.1 \times 10^{-9}\,\mathrm{kg\,J}^{-1}$ were used.

3.2 Simulations with unlimited surface dust

3.2.1 Choice of standard model resolution

Early dust-transporting model runs were performed at T21 ($5° \times 5°$ dynamical grid) resolution. The precise parameter values for the lifting efficiency and lifting threshold were obtained via a tuning process which aimed to

(a) sufficiently limit dust lifting in the aphelion season while allowing an increase in dust lifting and atmospheric opacity after $L_s = 180°$

(b) produce dust storms with an initiation phase displaying a rapid increase in opacity, and

(c) generate variability in storm magnitude and timing from one model year to the next.

In practice, this placed on the lifting threshold a lower limit, such that no dust storms were created north of the Tharsis range during northern spring/summer, and an upper limit due to the need for a minimal amount of lifting to be possible during the year. The lifting efficiency was constrained by demanding rapid increases in opacity upon the initiation of storms as well as limiting peak dust opacities to realistic levels. This last condition proved to be particularly vital, and model runs frequently crashed if α_N was set too high, as there

was not a 'shutdown' mechanism strong enough to overcome the intense lifting, which is further boosted by the positive radiative feedback.

With the 'correct' tuning, the model was able to produce reasonable globally averaged opacities throughout the Martian year, as well as some variability between years, as displayed in Figure 3.2. However, at T21 resolution, all the storms that were generated began in the Daedalia/Solis Planum region (around 120°W, 30°S), whereas the observational record over the past decades suggests multiple sites in the southern hemisphere are capable of initiating large dust storms. In the model, around perihelion, surface winds tended to become particularly strong at a very limited number of gridpoints at Daedalia, due to the crude representation of the complex topography present (Figure 3.3(a)). A larger peak in surface stress appeared in SW Hellas, and its effect can be seen in the regular peak in dust opacity around $L_s = 210°$ in Figure 3.2; however, unlike lifting at Daedalia, this did not grow into a larger dust storm, perhaps due in part to its occurrence earlier in the year.

It was found that an increase in resolution to T31 (3.75°) significantly reduced this wind anomaly relative to other SH peak winds such as those at Hellas, Argyre and Noachis (Figure 3.3(b)). This meant that a larger α_N could be used, which allowed other regions in the southern hemisphere to become potential dust storm sources. In unlimited surface dust runs such as these, optimum settings for the lifting parameters are effectively determined by that part of the southern hemisphere midlatitudes which sees the largest surface stress, since that is where storms will most readily form; therefore, the approach is vulnerable to gridpoints like those at Daedalia in the T21 model, where due to a lack of resolution winds are probably unrealistically high (at least compared to other gridpoints at that resolution).

Another consequence of the increased resolution was a significant increase in peak wind stresses in the northern mid-/high latitudes, extending southward along the topographic channels at Acidalia-Chryse ($\sim 50°$W) and Utopia Planitia ($\sim 100°$E), causing an increase in cap-edge dust lifting in autumn/winter and in the frequency of flushing storms. This suggests that the step from T21 to T31 leads to a significant improvement in the resolution of transient baroclinic waves along the cap edge, which provides the means to lift dust at these latitudes. At 45°N, gridbox widths on the model's dynamical grid are ~ 400 km and 300 km at T21 and T31 respectively, which can be compared to the Rossby radius of deformation, $L_d \sim 1000$ km, which controls the size of synoptic weather systems. The

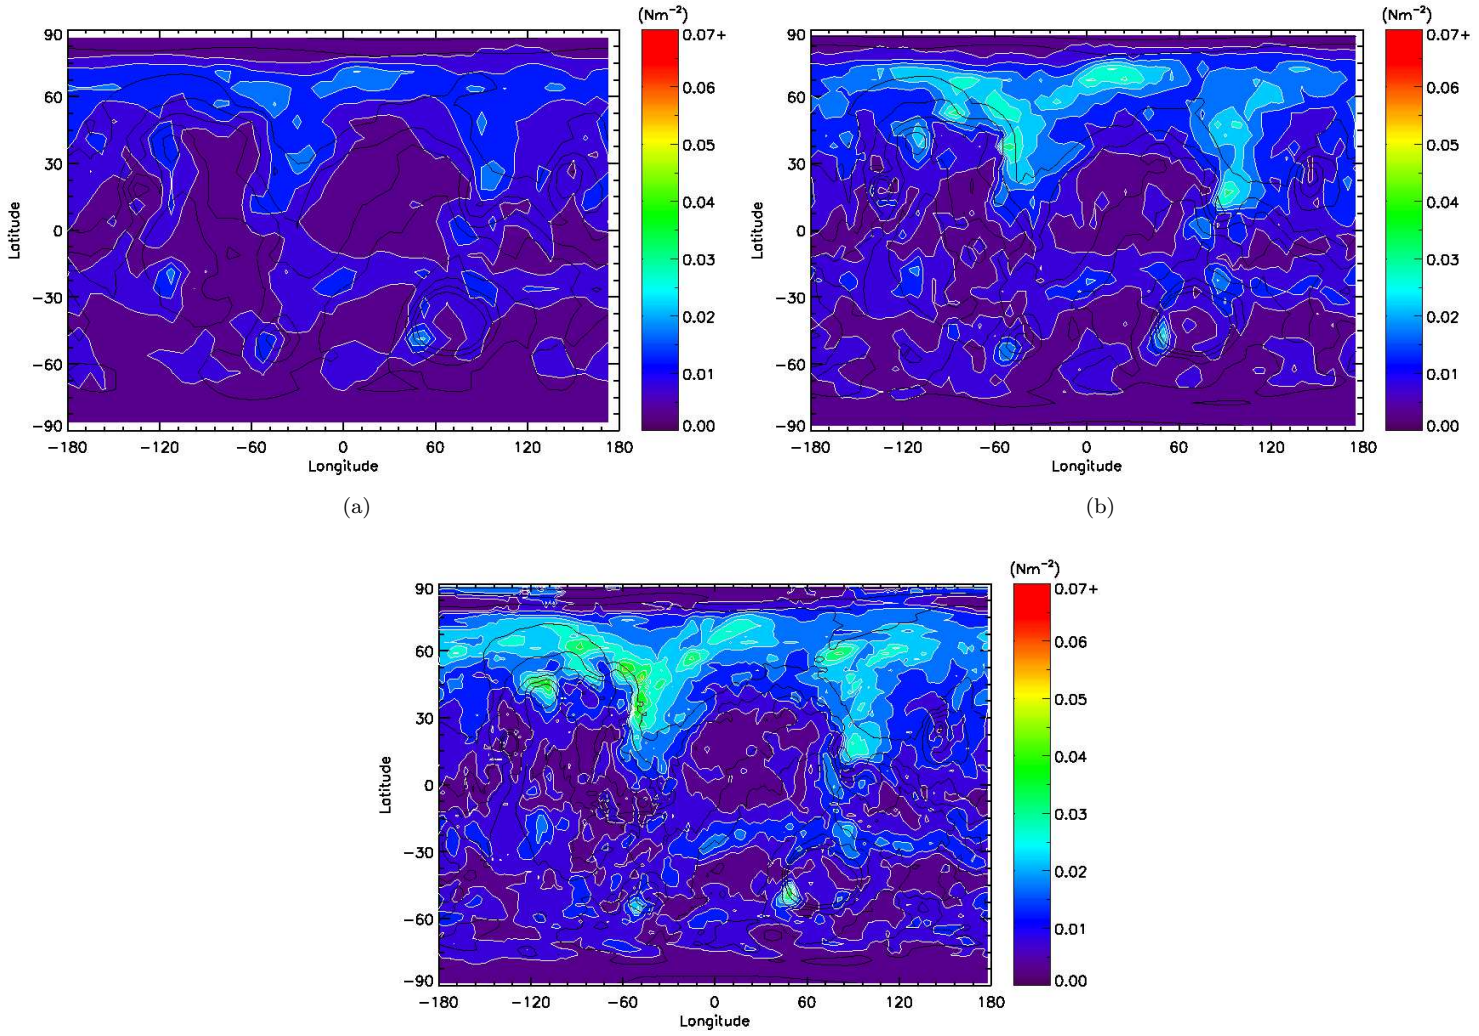

Figure 3.3: Maximum surface stress (in $\mathrm{N\,m^{-2}} \equiv \mathrm{Pa}$) from a year without dust lifting using the MY24 dust opacity scenario, at (a) T21, (b) T31 and (c) T63. Surface topography is overlaid in black contours.

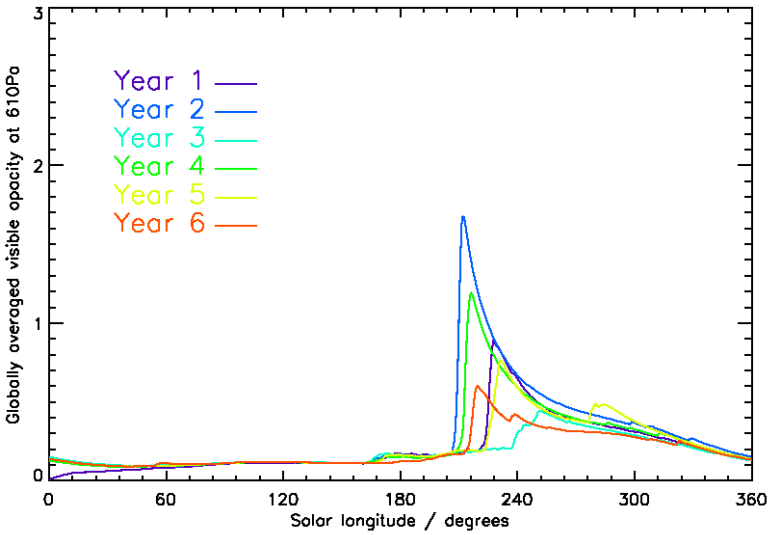

Figure 3.4: Globally averaged 610 Pa visible dust opacity from six years at T31, with threshold drag velocity multiplied by 0.7, and $\alpha_N = 1.4 \times 10^{-4}\,\mathrm{m}^{-1}$ and $\alpha_D = 9.1 \times 10^{-9}\,\mathrm{kg\,J}^{-1}$.

improved resolution is then perhaps due to the increase in model gridpoints per deformation radius from ~ 2 to ~ 3 when moving from T21 to T31. Further increases in peak wind stress in northern midlatitudes were seen when resolution was increased to T42 and T63 (Figure 3.3(c)), but for computational speed, T31 was adopted as the base resolution at which to perform the majority of the simulations described in this work.

3.2.2 T31 results

Similar globally averaged opacities to those at T21 were obtained for the higher T31 resolution (Figure 3.4). Again, some interannual variability was observed, but storms were mainly limited to the pre-solstice period, around $L_s = 210$–$250°$, and the prolonged, very large opacities typical of major global storms (such as the 2001 storm) were missing. There was also an absence of any significant storms in late summer ($L_s \sim 310°$), a time of year at which several regional storms have been observed, although one of the years shown here (year 5) did include two regional storms in the same year (at $L_s = 230°$ and $280°$). The preferred region for initiation of large storms shifted from Daedalia Planum, and showed some spatial variability. The majority of simulated storms occurred when a flushing event began at Utopia Planitia, travelled southward through Isidis, and reached the northern flank of Hellas, where it sparked lifting that could quickly escalate into a planet-encircling

storm. This pattern of storm development is shown in Figure 3.5 — in this example, lifting occurred as the flushing storm travelled southward, continued at Hellas then spread to the east at sols $\sim$ 467–469, reaching Daedalia as the storm became planet-encircling. The storm shown was relatively small, and lifting had ceased by $\sim$ sol 473 (not shown). In some other model years, a similar process occurred in the Chryse-Noachis regions.

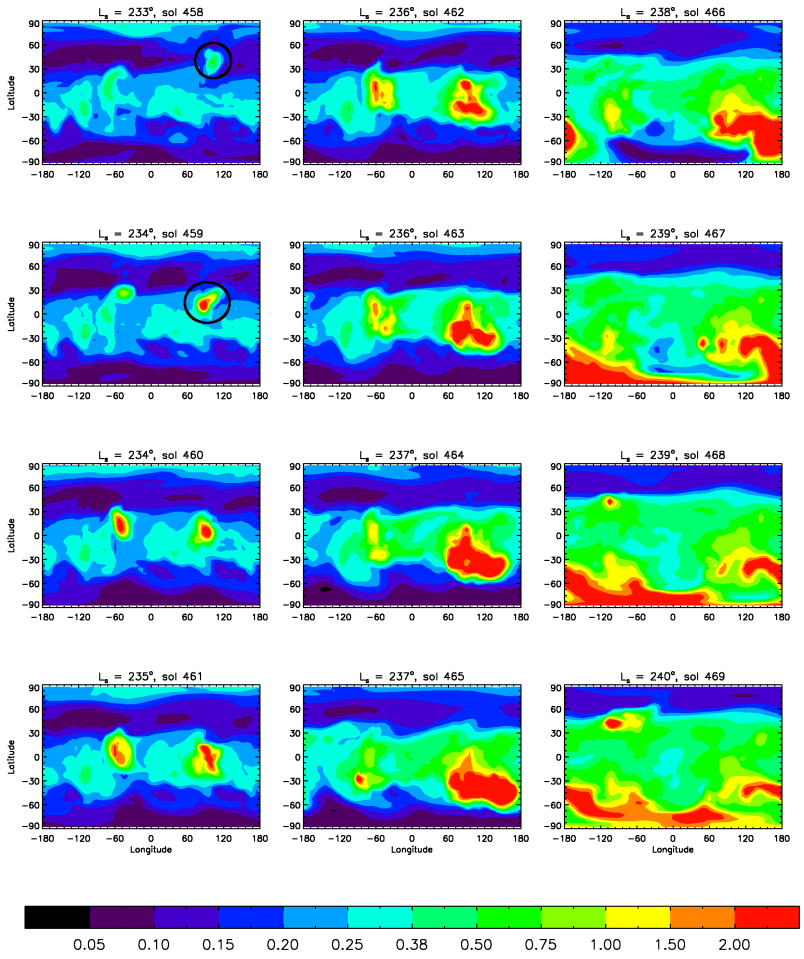

Figure 3.5: Progression in 610 Pa visible dust opacity for a typical Utopia-Hellas-type storm (circled in the first two panels), beginning at $L_s = 233^\circ$. (In this example, another flushing storm occurred simultaneously at Acidalia, but had little obvious effect on the development of the Hellas storm — single flushing storms are more typically seen in the model.)

This kind of storm initiation has been observed on a slightly smaller scale, and may be the dominant mechanism for the development of regional dust storms. A good example is

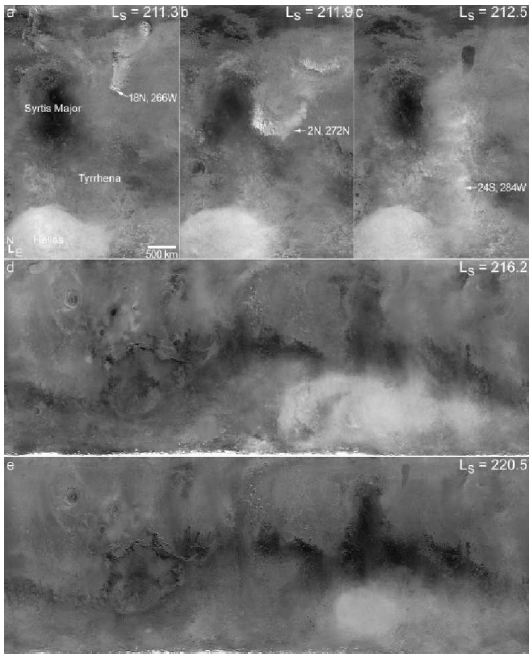

Figure 3.6: Wide angle red filter MOC images of the flushing storm in early southern hemisphere spring of MY26, from Cantor (2007). The storm originated in northern Isidis at $L_s = 211^\circ$ and dissipated by 220°, leaving behind a planet-encircling dust haze in the SH. Panels *d* and *e* cover a latitude range of 40°N–60°S.

given in Figure 3.6. These MOC images from SH spring of MY26 show a flushing storm that moved southward along the Isidis channel and spread zonally upon reaching northern Hellas, initiating further cap-edge dust lifting that reached 'as far west as 0°W at $L_s = 215.6^\circ$', and created a dust haze that encircled 'the southern hemisphere between 30.0°–65.0°S by $L_s = 220.5^\circ$' (Cantor, 2007). So, the notion of a SH response to the arrival of a flushing storm has some observational backing; however, in the MY26 example the resulting dust haze was much more diffuse than most of the cases seen in the model. In fact, as described in Chapter 1, observations strongly suggest that larger dust storms, of the magnitudes seen in several of the years in Figure 3.4, usually begin in the southern hemisphere, rather than as a consequence of a cross-equatorial flushing event.

Basu et al. (2006), using a finite-difference grid horizontal resolution of $5^\circ \times 6^\circ$, reported that most dust storms in their model formed in the Hellas region, with variability in timing affecting the size to which the storm could grow. Early storms developed in SW Hellas but were limited in size, similar to what was seen above for the UKMGCM. Storms developing after $L_s \sim 240^\circ$ tended to begin in northern Hellas and grew to global size; during these

events, lifting was dominated by an area on the northern rim of the Hellas basin, suggesting a similar behaviour to that seen at Daedalia in the UKMGCM T21 runs. Strong wind stresses persisted in this location until after $L_s = 300°$, which prevented the GDS from decaying and tied global opacities to the annual insolation cycle. By contrast, the T31 dust storms shown in Figure 3.4 shut down much more rapidly. In part this was because global opacities did not grow as large as in Basu et al. (2006), (though they could be made to do so by slightly increasing α_N). It was also possibly due to the different initiation mechanisms of storms in the two models — the UKMGCM did not produce a wind stress peak in northern Hellas of similar magnitude and duration, so that large storms were only sparked whenever a disturbance entered the Isidis/Hellas region from the north, and subsequently these storms were not sustained so effectively by lifting in northern Hellas. The GFDL model also produced storms at Chryse, but again these were larger and lasted longer than any events observed in the region. Chryse storms in the UKMGCM had lower peak opacities and were shorter in duration, though sometimes led to the development of larger storms in the southern hemisphere (an example of a Chryse flushing storm can be seen in Figure 3.5).

Simulations with the NASA-Ames model by Kahre et al. (2006), also on a $5° \times 6°$ finite-difference grid, did not show similar interannual variability, and rather were split into two types, 'baseline' and 'dusty', by using a different lifting efficiency for each. In both cases the annual dust opacity cycle featured a broad primary peak at $L_s \sim 300°$ and a secondary peak at $L_s \sim 210°$, and southern hemispheric lifting was concentrated quite strongly on the north and southwest rims of the Hellas basin.

3.3 Dependence of dust lifting on resolution

To explore the sensitivity of dust lifting to model horizontal resolution, several month-long comparative simulations were carried out, using either radiatively passive or radiatively active lifted dust. Each set of simulations included runs at T31 (3.75°), T42 (2.5°) and T63 (1.875°) resolutions, and model variables were output every 6 hours. The 'passive' runs used the MY24 dust opacity scenario and included wind stress lifting only, and were used to compare basic wind strengths and wave amplitudes away from the influence of altered dust loadings; conversely, the active runs used both lifting mechanisms, with dust opacities

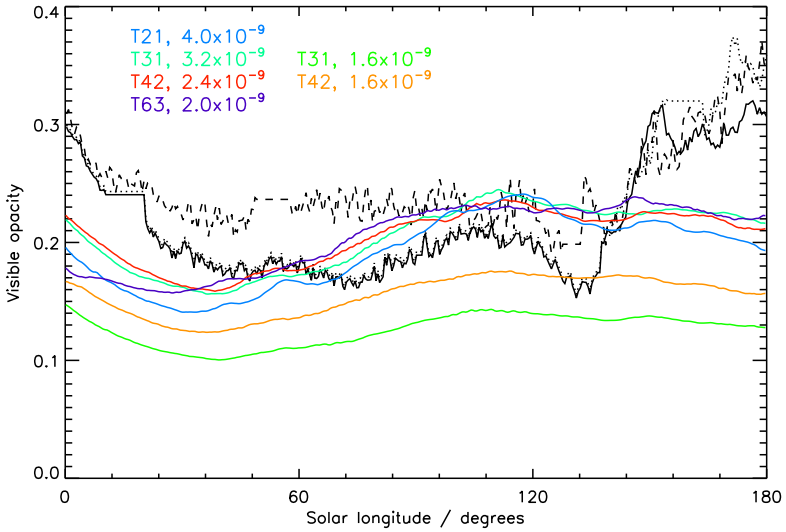

Figure 3.7: 610 Pa visible opacity, averaged between 45°N and 45°S, from runs using only dust devil lifting with values of α_D as shown (in units of kg J^{-1}). The three black lines show opacities (averaged over the same region) from the TES reanalysis of MY24 (solid), MY25 (dotted) and MY26 (dashed) (though 'MY24' actually uses MY25 values until $L_s = 140°$).

computed from the transported field, and provided information on dust lifting feedbacks.

3.3.1 Dust devil lifting

The resolution dependence of the threshold-independent dust devil scheme was assessed by tuning the dust devil lifting efficiency at each resolution in order to best match dust opacity over $L_s = 0$–$180°$ — when dust devil lifting is the dominant contributor to total opacity in the model (though perhaps not in reality) — to the three years in the TES reanalysis dataset, which show relatively little interannual variation during this part of the year (Montabone et al., 2005). As can be seen from Figure 3.7, dust devil lifting strength increased with increasing resolution — perhaps due to better resolution of gridpoints particularly favourable for convective lifting — so a progressively smaller α_D was needed. The negative feedback inherent to the scheme is also evident; for example, at T31, a doubling of α_D produced only a $\sim 70\%$ rise in opacity. The discrepancies between model and assimilation at $L_s = 0$–$30°$ and 150–180° were mainly due to local dust storms at the edge of the polar caps produced by wind stress lifting, which was not activated in these runs. The scaling of α_D identified here was applied to all subsequent simulations (including the active runs below), so that model dust devil opacity became approximately independent of

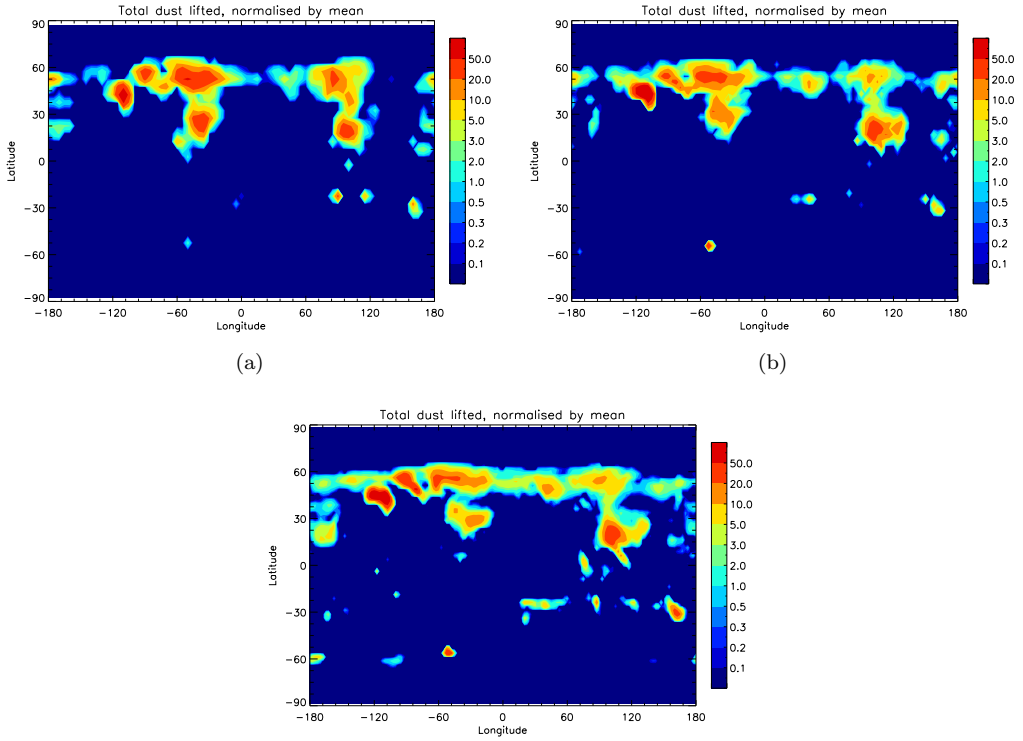

Figure 3.8: Total dust mass lifted at each gridpoint over the length of the passive 210° simulations, divided by the global mean in each case, at resolution (a) T31, (b) T42 and (c) T63.

horizontal resolution.

3.3.2 Wind stress lifting: northern autumn

The first experiment with wind stress lifting included was a passive run over $L_s = 210$–$240°$, a season in which dust lifting can typically occur both in southern midlatitudes and along the edges of both polar caps. The spatial distributions of lifting activity — total lifted mass at each gridpoint, normalised by the global-mean total — are shown in Figure 3.8, where it is apparent that as the resolution was increased, a greater fraction of gridpoints became able to take part in lifting. This resulted in more than twice as much total lifted dust mass at T63 than at T31. The reason for this was an increase in the frequency of occurrence of surface winds above the lifting threshold, through a combination of increased mean u_* and increased variability, as measured by the standard deviation $\sigma(u_*)$.

Intercomparison was more difficult in the active case at $L_s = 210$–$240°$, as (in this

particular year) dust storms of different sizes were produced after $L_s = 230°$ at T31 and T63, while at T42, no storm was produced at all, even though in the passive example, more lifting was produced before $L_s = 230°$ than was the case at T31. This illustrates the unpredictable manner in which storms can develop in the model — the system is highly nonlinear, and although a strong positive feedback exists for the wind stress lifting mechanism more generally (as a large dust loading causes deep atmospheric heating), lifting can also be sensitive to changes in near-surface lateral temperature gradients, such as can be produced on a more rapid timescale following injections of dust into the boundary layer. In this way, stronger lifting in one area can actually inhibit further lifting and reduce the total lifted mass. Nevertheless, considering the no-storm period $L_s = 210$–$230°$, the same trend was seen as in the passive case: increasing lifted dust mass with increasing resolution. Figure 3.9 shows that, when resolution was increased from T31 to T63, mean u_* increased over much of the planet, but particularly at the region of complex topography in the Tharsis range at 0–30°S, 120°W, and near other large topographic slopes such as to the north of Alba Patera, around Elysium Mons and on the northern and southern edges of Hellas. Similar results were obtained from the passive runs, except for the increase at 0–60°E, 30°S, which was unique to the active runs and can be attributed to a strengthening of the PMOC, and the summer subtropical westerly jet, under greater dust loading. Increases in u_* were seen across the 60°S band, which marks the edge of the seasonal CO_2 cap in this season. Mesoscale-resolution simulations by Toigo et al. (2002) produced dust lifting in this region, and it was attributed mainly to the thermal contrast flow set up by the strong temperature gradient at the cap edge (similar to the terrestrial 'sea breeze' circulation), which was evidently best resolved at T63 in the present case. At the southern rims of the Argyre and Hellas basins, this thermal contrast flow can combine with downslope winds to produce particularly favourable dust lifting conditions, which may be the reason that several large storms have been observed to originate in these regions. While increases in $\sigma(u_*)$ were seen around NE Argyre and NW Hellas, a strong decrease was seen to the NE of Hellas (100°E, 30°S) (similar changes were seen in the passive case, meaning that they were not a result of increased dust opacities). This might suggest that the Utopia-Isidis-Hellas storm initiation mechanism (see §3.2.2, Figure 3.5), seen regularly at T31 but less so in observations, will be less effective at higher model resolutions.

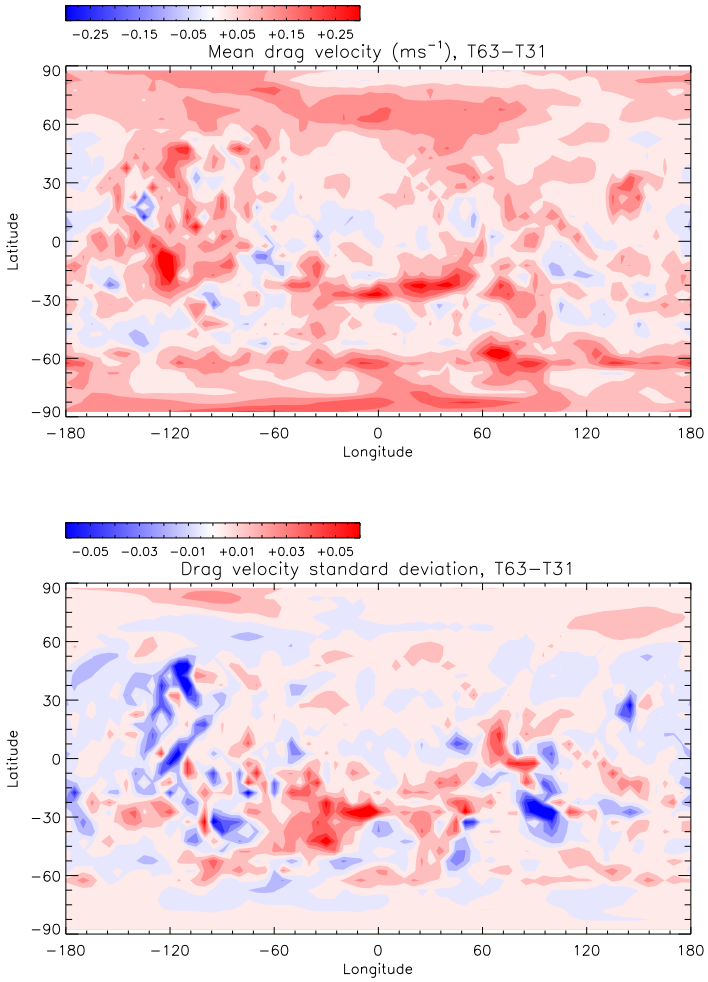

Figure 3.9: From the active 210° simulations, the change (T63 − T31) in average u_* (top) and average standard deviation $\sigma(u_*)$ (bottom) (both in $\mathrm{m\,s^{-1}}$).

3.3.3 Wind stress lifting: northern summer

A second set of model runs were performed at L_s = 150–180°, to further examine the potential for dust storm production along the edge of the receding south polar cap, as observed at this time of year by Cantor et al. (2001). Again, u_* increased with resolution, but at all resolutions lifting was only possible at one southern hemisphere site, on the southwestern rim of Hellas. In the northern hemisphere, travelling wave strength increased with increasing resolution (as it did in the previous set of runs). This is shown in Figure 3.10 in root-mean-square (RMS) variance of meridional velocity v at 2.5 km above the ground, filtered to include only waves with periods of between 1.5 and 10 sols. A different travelling

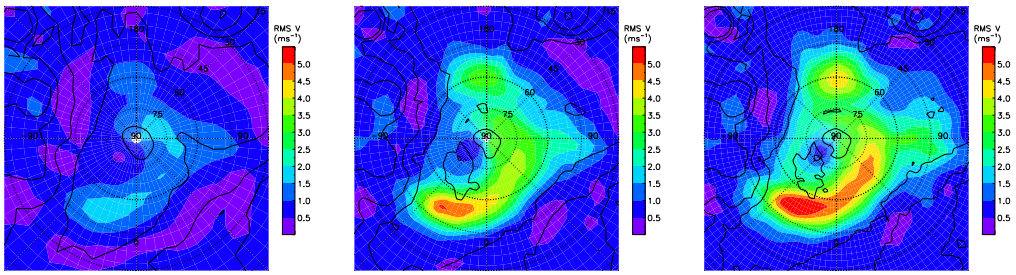

Figure 3.10: RMS meridional velocity, filtered to include only waves with period 1.5–10 sols, from the passive 150° runs, at 2.5 km height in the north polar region. Resolutions are T31 (left), T42 (middle) and T63 (right). Topography is shown in black contours.

wave behaviour was seen in the southern hemisphere (top row of Figure 3.11): waves were weaker at T63, relative to T42. This appears to have been due to the stronger numerical dissipation that was used at T63 in order to retain model stability, as the peak near-surface Eady growth rate (see §5.3.3) in the southern hemisphere (which occurred near the south pole) increased slightly at T63 relative to T42, suggesting that eddy activity should have increased, as it did in the northern hemisphere. However, it is then surprising that the dissipation did not have the same effect in the northern hemisphere. The increases in peak wave amplitude in the NH (both at $L_s = 150°$ and at $L_s = 210°$) suggest that springtime midlatitude dust lifting will be enhanced at higher model resolutions, but the same cannot necessarily be said for the SH, based on these results.

In the active case, lifting in the SH remained spatially confined, though it was activated at the edge of the Argyre basin at T42 and at further gridpoints in north and west Hellas at T63. Transient waves in meridional velocity, as shown in the bottom row of Figure 3.11, were rather different than in the passive results. Wave amplitude was larger and peaked in midlatitudes, particularly on the western rims of the basins, where the waves contributed to maxima in u_*. The differences here highlight the errors that could be introduced by the use of the prescribed opacity scenario, since TES observations were not available at the south pole at this time of year, requiring assumptions to made in the creation of the scenario. Again, transient waves weakened from T42 to T63.

It has been seen here that, even at T63, the model did not predict cap-edge dust storms to be as widespread as those observed by Cantor et al. (2001), who noted a particular concentration of storms in the Noachis region, where no dust lifting occurred in these runs. Model lifting at this time was mainly restricted to SW Hellas. These results suggest that

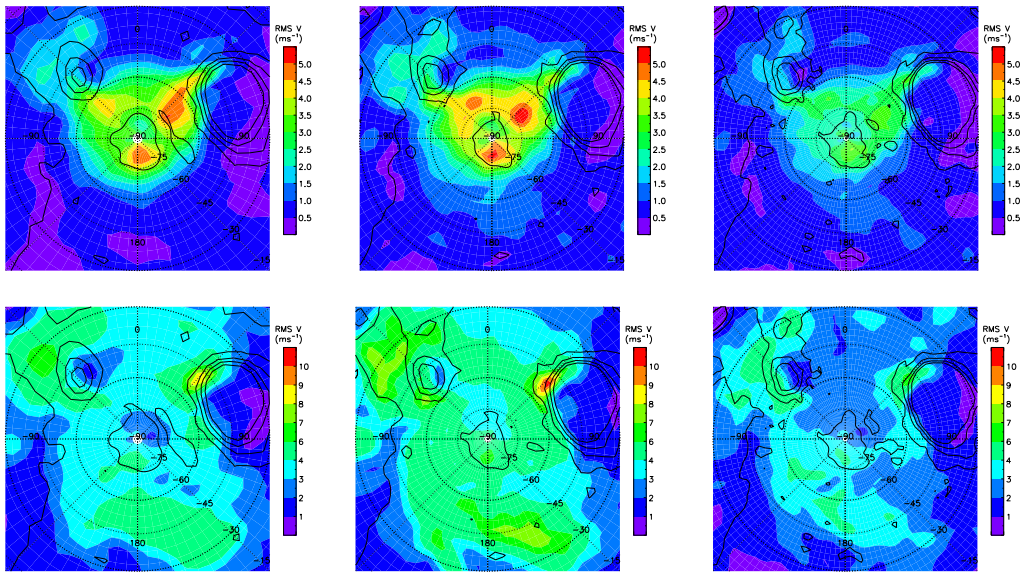

Figure 3.11: Top row: transient RMS V (as in Figure 3.10) from the passive 150° runs at T31 (left), T42 (middle) and T63 (right). Bottom row: the same from the active 150° runs. Note the different colour scales for each row. Topography is shown in black contours.

to properly model cap-edge flow in the southern hemisphere, at least near equinox when the seasonal cap still extends to midlatitudes, requires a resolution close to that of Toigo et al. (2002) (i.e. mesoscale, $\lesssim 1°$); the differences between hemispheres seen here may be due to a particular feature of the surface topography. However, the issue of SH cap-edge dust lifting is returned to in Chapter 6, when water ice clouds are added.

3.3.4 Resolution dependence summary

The results of the resolution sensitivity experiments are summarised in Figure 3.12. Total mass lifted increased with resolution, but did so less steeply when dust was radiatively active. This suggests that although dust lifting exhibits a positive radiative feedback during the initiation of regional-scale storms, when lifting is more localised (note that the storms after $L_s = 230°$ were excluded from these calculations) the feedback tends to be negative, most likely due to an increase in static stability as dust induces a warming throughout the boundary layer. The decrease in temporal variability of the surface wind at T63 in the active 150° run can be attributed to a large increase in dust opacity over T42.

Theoretically, the scaling factor applied to the threshold drag velocity should be marginally increased as resolution increases, since the model is resolving a greater amount (spatially)

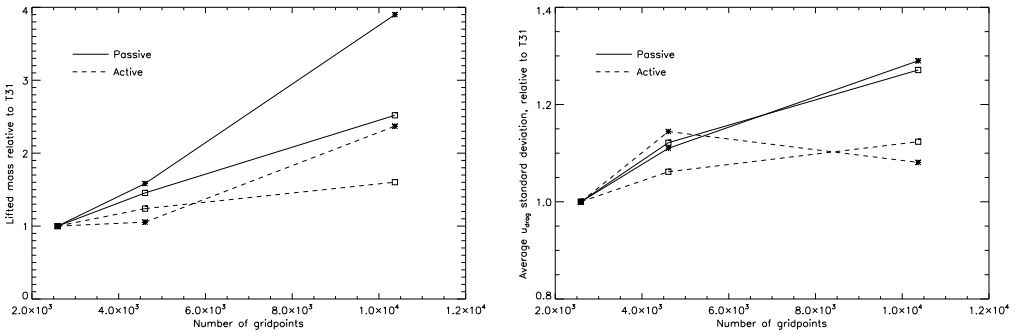

Figure 3.12: Total dust mass lifted (in kg, left) and average $\sigma(u_*)$ (in $\mathrm{m\,s^{-1}}$, right) from the four simulations described — open squares are from the 210° runs and stars are from the 150° runs. The active 210° values are averages covering 210–230° only.

of surface wind gustiness. However, the differences between the passive and active trends in $\sigma(u_*)$ make this very difficult to quantify. Calculating the scaling factors needed to match total lifted masses in the active simulations above suggested a possible increase to ~ 0.8 by T63 (from 0.7 at T31). On the other hand, having a greater fraction of gridpoints possess the potential for dust lifting is probably a more realistic result, in which case an alternative method to keep lifted mass constant would be to reduce α_N; the magnitude of this reduction can be estimated from the active runs above, and a value of $\sim 0.6\alpha_N(\mathrm{T31})$ was produced for T63.

3.4 Limitations of the unlimited surface dust method

Due to the wide range (with regard to timing and location) of dust storms that have been observed on Mars, it is not easy to say how realistic storms of the kind described in §3.2.2 may be; what is certain, though, is that this variation in date of occurrence and areographic origin was not fully captured by the model. The very large, sustained dust loadings of (for example) the 2001 GDS were also missing from the multi-year sequence shown in Figure 3.4, which displayed a more moderate variation in peak opacity from year to year.

Another drawback of the current formulation is that a very precise level of tuning of the lifting parameters (efficiency and threshold) is required in order to produce dust lifting rates that are in any way realistic. While possible at T31 resolution, increasing the model resolution or adding any new physics of significance would require a retuning of the parameters, which would soon become a prohibitively lengthy process. Whilst there clearly

exists a 'true' lifting efficiency, one would not expect that over- or under-estimating α_N by even a few percent should produce such a dramatic change in model output (huge GDSs every year, or no GDSs at all, respectively). Other, similar, climate model parameters — for example, cloud nucleation efficiency — do not in general produce this behaviour, and the key difference is that a negative feedback usually exists to act against the error in the parameter (if cloud nucleation efficiency is too high, the supply of atmospheric water vapour rapidly becomes depleted, and nucleation rates decrease). In the case of dust lifting, the negative feedback could be the supply of dust on the ground, and so the inclusion of this feedback should allow greater flexibility in the α_N setting.

Therefore, both to improve the variability produced by the model and to facilitate its use at different resolutions, work was begun to link the surface state of the model to the lifting processes, to include an explicit negative feedback on dust lifting and to represent the spatial and temporal heterogeneity that exists in the ability of the Martian surface to supply dust to the atmosphere.

Chapter 4

Surface improvements

4.1 Motivation for using a time-varying lifting threshold

As discussed at the end of the previous chapter, the assumption of an infinite supply of surface dust may lead to significant errors in the predictions of dust lifting schemes. In reality, surface densities of dust on the planet's surface are finite[1], and probably small enough in certain locations to limit the total amount of dust that is available to be lifted during a large storm. The qualitative spatial distribution of surface dust (particles smaller than $\sim 100\,\mu$m) was derived from Thermal Emission Spectrometer (TES) albedo measurements by Ruff and Christensen (2002), who found the southern hemisphere to be, in general, more sparsely covered by dust grains than the northern hemisphere. Periodic changes in surface albedo, notably in the wake of major dust storms, were detected by Szwast et al. (2006), indicating the redistribution of significant fractions of this surface dust (depth changes on the order of a few microns). Furthermore, Strausberg et al. (2005) noted that during the 2001 global dust storm (GDS), dust lifting on the slopes of Hellas, the site of the initiation of the storm, appeared to cease after $\sim 10^\circ$ of L_s, while the storm was very much still in progress. Local surface windspeeds remained large at this time (Montabone et al., 2005), raising the possibility that the Hellas region was exhausted of surface dust, preventing further emis-

[1]The rate of erosion is not well known but the phenomenon may be assumed to be negligible on inter-annual timescales.

sion from taking place. These changes in surface dust may also play an integral role in driving the strong interannual variability that has been observed in dust storm occurrence and intensity, by introducing into the climate system a long-term memory (as surface dust distributions can persist indefinitely, in the absence of significant wind stress), something that the Martian atmosphere, with a radiative timescale of a few sols, does not possess.

The idea that a time-varying threshold for wind-stress dust lifting might be appropriate for Mars was first put forward by Pankine and Ingersoll (2004), who found that adding such a parameter prompted hitherto-unseen interannual variability in their (very) simplified global Mars model. Specifically, they used a threshold which increased following a large dust storm, and decreased gradually in the absence of any dust lifting. Suitable tuning of the increase and decrease rates allowed global dust storm frequency to be adjusted to approximately match that currently observed. They suggested that this 'self-organised criticality' (in which a system regulates itself so as to remain close to a threshold point) could be an explanation as to why Mars appears to be finely tuned to produce GDSs only under certain favourable conditions. However, with the use of a low-order model (a zonally symmetric model with simplified dynamics), with effectively only two surface gridpoints (one in each hemisphere), they were unable to examine storm development or the redistribution process, or ascertain whether or not the approach would work equally well in a fully-fledged GCM. A first attempt at adapting this idea for a GCM was made by Wilson and Kahre (2010), and did produce increased interannual variability in the magnitude of dust storms produced in their model, though not in the initiation dates of the storms. Also, model GDSs were frequently too large, and occurred in groups every twenty or so years, rather than with a frequency of one every three years, as was estimated by Zurek and Martin (1993) and as has continued to hold true over MY23–28 (§1.2.2).

The physical rationale offered for the Pankine and Ingersoll parameterisation is that the local threshold lifting stress in a particular location should vary according to the ratio of 'erodible' dust particles (sitting out in the open) to nonerodible elements (namely rocks or pebbles too large to be moved by wind, but which absorb a fraction of the wind momentum). It may be anticipated that after a burst of emission from a certain site, the dust particles that remain will be those that are sheltered by nonerodible elements (leading to the formation of so-called 'desert pavements'), and that moving these will require

larger wind stresses. In this way, areas may not become totally depleted of surface dust, while nonetheless ceasing to be viable regions for intense lifting: it has been observed that the rate of dust removal from an area containing a mixture of dust/sand and nonerodibles decreases over time as the nonerodibles are uncovered and begin to shelter the remaining erodible material (Greeley and Iversen, 1985). To reverse this trend, either dust is returned to the area from elsewhere and deposited on or around the nonerodibles, or small-scale processes (dust devils, perhaps) release the dust from under the roughness features, leading to a reduction in threshold stress and improved conditions for dust lifting.

An alternative yet — in the context of the lifting scheme being used — complementary interpretation comes from more recent suggestions that saltation does not occur particularly often on Mars. This is suspected to be the case because of a lack of long-term movement seen in sand dunes (Zimbelman, 2000; Malin and Edgett, 2001). Instead, dust may be lifted directly from the surface. While the wind stress required for the lofting of micron-sized particles is much higher than is usually experienced on Mars, if the dust on the surface exists rather as larger, loosely-packed agglomerates, with a mass density lower than that of the dust particles themselves, the threshold for lifting of these objects would be significantly lower, perhaps making direct detachment a viable mechanism on Mars (Sullivan et al., 2008). Once airborne, it is proposed that the dust aggregates separate into the observed micron-sized components. Support for the use of a time-varying threshold in this context comes from the experimental work of Merrison et al. (2007), using a Martian dust analogue material in a wind tunnel, under Martian conditons. The authors claimed that the aggregation of surface dust grains could require the presence of a surface dust layer of some minimum depth. Therefore, as dust is introduced to a certain location, particles lying on the surface could more easily form mobilisable aggregates, creating an inverse relationship between dust surface density and gridbox-average lifting threshold.

The low-density aggregate hypothesis for dust lifting on Mars is interesting, although it has not been investigated in detail in this work. It is not clear that applying it to the UKMGCM would alter results much, since the reduction in lifting threshold caused by the use of a lower dust density has, in the present case, been replicated by a prescribed reduction in the model lifting threshold (§3.1.1). Without any knowledge of how quickly the sand-sized dust aggregates form on the surface or break up once in the atmosphere, their

invocation in the model would merely provide a different physical justification for using a low threshold stress, one within reach of wind stresses generated in model simulations. The variable threshold approach taken in this chapter can be justified for either of the 'saltation' or 'aggregates' models of dust lifting.

4.2 Implementation of the variable threshold scheme

4.2.1 Threshold response to dust lifting and deposition

The wind stress threshold for dust lifting, $\zeta^t = \rho(u_*^t)^2$, was therefore programmed to vary in response to the change in dust surface density at a particular gridpoint, at a rate ${\zeta^t}_{inc}$ (with units of $\mathrm{Pa}\,(\mathrm{kg\,m^{-2}})^{-1}$) linearly proportional to the amount of dust removed or deposited in a physical timestep; that is, the threshold increases as dust is removed from the gridpoint, and decreases whenever dust is deposited onto the gridpoint (the 'inc' suffix is used because, as will be seen later, the dominant effect imparted by this term in key dust source regions in the southern hemisphere is a threshold increase). The true behaviour of stress threshold with changing surface density is unknown, so the value of ${\zeta^t}_{inc}$ was chosen simply so that thresholds were raised enough in the course of a large storm to cut off further lifting from at least some gridpoints; similarly, the linear rate of increase was used purely for simplicity. The lifting efficiency α_N (now set larger than before, to counteract the higher thresholds which result from the use of this method) and ${\zeta^t}_{inc}$ could to some extent be varied in tandem, as more efficient lifting required faster-growing thresholds to prevent runaway lifting. Various combinations of these two parameters were tested.

To prevent the possibility of lifting outside of the storm season, through year-on-year threshold reduction in dust sink regions, a minimum threshold stress was specified, equal to the original threshold value of $\rho(0.7u_*^t)^2 = 0.020\,\mathrm{Pa}$, with u_*^t again given by (3.2). Thus, the new parameterisation acted only to increase surface lifting thresholds from the previously used value. Following the Pankine and Ingersoll argument, this minimum represents a saturation of dust cover in the unsheltered areas of a gridbox, such that adding more dust does not decrease the effective lifting threshold any further. In terms of the aggregate model, it corresponds to a saturation surface dust depth above which lifting does not become any easier, probably due to a limit on the size and density of dust aggregates that can form.

4.2.2 Need for an artificial 'resupply rate'

When the scheme as described so far was applied in a multi-year simulation (with thresholds at each gridpoint initialised at the minimum value of 0.020 Pa), the results were very similar to those of Kahre et al. (2005) (see their Fig. 3), which used a 'straight' finite surface dust method (i.e. without threshold changes). That is, for several model years, major dust storms were produced in a manner largely similar to that seen in previous infinite surface dust work, but dust storms then ceased to occur, and the model became 'locked' into a state that did not feature a significant amount of wind stress lifting in the southern hemisphere. This shut-down of lifting was a result of the threshold increases produced in parts of the southern hemisphere following the removal of surface dust during the storms that occurred in the initial years of the simulation, as dust deposition in these regions evidently did not occur at a rate sufficient to counteract this loss. The time taken to reach this shut-down point decreased as ${\zeta^t}_{inc}$ was increased, since this parameter controls the extent to which thresholds are affected by the removal of a given dust amount (a similar result was shown by Kahre et al. (2005), whereby the decline in global dust opacity could be delayed for a number of years by initialising gridpoints with a deep dust reservoir).

This net movement of dust away from key source regions is a symptom of a more general interhemispheric transport bias, in which dust is transferred from southern to northern hemisphere, that exists in the model, as discussed later (§4.4.2). To correct for this bias, so as to allow to continuation of dust storms over multiple model years, another parameter, ${\zeta^t}_{dec}$ (units of $\mathrm{Pa\,s^{-1}}$), was added, applying a constant rate of decrease to surface thresholds at every gridpoint, in addition to any decreases that occur 'naturally' in the model through deposition of atmospheric dust. This may be interpreted as the continual addition of extra dust to each gridpoint that is above the minimum threshold, so is clearly a crude way in which to counteract the interhemispheric bias. In fact, there is reason to believe that such a term could be justified even in a model without this transport bias, by virtue of redistribution of surface *within a gridbox*. (see §4.4.2). The value for this parameter was set by requiring that some dust lifting remain possible each year and that global dust storms occur approximately every three years. As would be expected, increasing the rate

of resupply increases the frequency of large dust storms in the model — approximately,

$$\frac{\zeta^t{}_{inc}}{\zeta^t{}_{dec}} = \frac{\text{gap between successive GDSs}}{\text{dust surface density reduction at initiation sites during a GDS}}.$$

Thus, through these parameters the model could be tuned to try to reproduce the long-term dust storm behaviour that is observed in reality.

4.2.3 Treatment of dust devils

Szwast et al. (2006) determined, from MGS surface albedo measurements, that dust devil lifting probably plays little or no role in the redistribution of surface dust. Model parameterisation of dust devil lifting, however, does show a net annual dust movement. Furthermore, among the regions that are depleted in this way are those that are important in storm generation — surface dust removal by dust devils, though relatively weak, is particularly noticeable along the 30°S latitude band, and on the Acidalia-Chryse and Utopia-Isidis storm tracks. With this in mind it seemed sensible to separate the two lifting methods, allowing dust lifted by convective processes to provide the background opacity while ignoring its effect on the surface dust distribution. In part this was motivated by a suspicion that dust devil lifting in the model could be overestimated, in order to make up for unresolved wind stress lifting events, particularly at $L_s = 0$–$180°$.

This is one area where the different versions of the model used required a slightly different approach. In the work described in this chapter, using v3, dust transport involved six tracers ('three plus three', with each size bin duplicated) as listed in Table 3.1, and so the effects of the two lifting methods could easily be kept apart (lifting thresholds were affected only by the removal/deposition of one set of three tracers). In v5, with the two-moment transport scheme available, the transported dust field represents the combination of the two types of lifting, and so the desired situation of having only dust lifted by wind stress affect the thresholds was not possible to achieve. Attempts were made to accommodate dust devil lifting into this formulation, either by allowing dust devil lifting to raise thresholds and combined sedimentation to lower them, or by only allowing sedimentation to affect thresholds. Both of these methods were problematic, due to the changes imparted to the threshold field by the convective lifting — in the first case, regions around the 30°S zone, potentially important for lifting, lost significant amounts of dust year-on-year, raising

thresholds much too high, and in the second case thresholds were decreased too rapidly through the sedimentation of dust lifted by dust devils. It was evident that significant quantities of dust raised by dust devils tended not to sediment back to the location from which it was lifted.

Ideally, dust devil lifting should have been subjected to the effects of the changing surface dust distribution as well, since dust availability clearly affects lifting by dust devils as much as wind stress lifting, but it was not at all obvious how this should be implemented, particularly since the commonly used dust devil scheme does not feature any thresholds. Including dust devil lifting in the variable threshold scheme would have required adding a further unconstrained parameter, which was undesirable with respect to the simplicity and illustrative ability of the scheme. Moreover, the background dust haze provided in the model by this mechanism does not need to vary interannually in order to match observations, and it is unknown how much convective lifting varies between years in reality. In the end, the approach used within v5 was to create a second two-moment tracer denomination to carry dust lifted by dust devils. Thus, the two dust mechanisms were separated as before, and surface thresholds were affected only by (and affected only) dust lifted by wind stress. In reality, it may be that dust devils play an important role in moving dust out from underneath nonerodibles, so a future form of this scheme could try to vary the size of the resupply term depending on local dust devil activity.

The results presented in §4.3–4.6 come from the earlier implementation (v3); the transfer of the scheme to v5 is covered briefly in §4.9 and in Chapter 6.

4.3 Effects of use of the scheme

4.3.1 Improved variability

The variable threshold scheme was employed in a multidecadal model integration, using a value for ${\zeta^t}_{inc}$ large enough that the increase in lifting thresholds at a GDS source region was sufficient to shut off further dust lifting at some of these gridpoints during the course of the storm, a value for ${\zeta^t}_{dec}$ chosen to provide the desired GDS frequency of around one every three model years, and using a higher efficiency α_N than before to counteract the larger thresholds (recall that thresholds were allowed only to increase relative to their

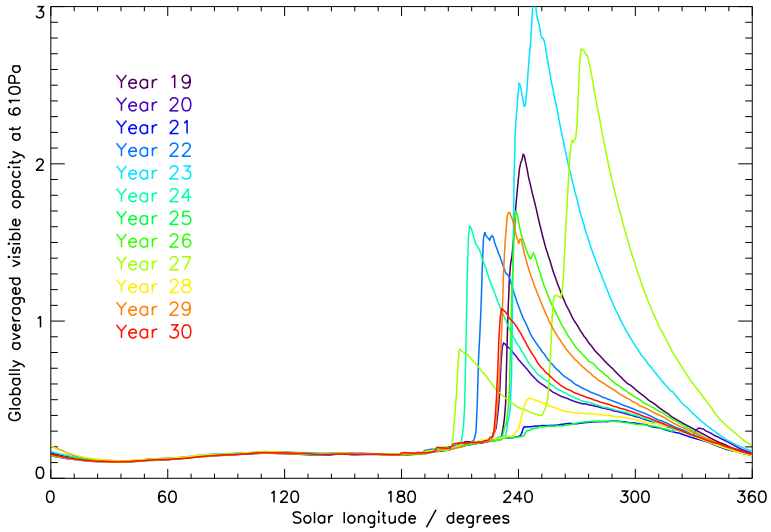

Figure 4.1: Globally averaged optical depth over several years at T31 ($\alpha_N = 4 \times 10^{-4}\,\mathrm{m}^{-1}$, ${\zeta^t}_{inc} = 0.5\,\mathrm{Pa}\,(\mathrm{kg\,m}^{-2})^{-1}$, ${\zeta^t}_{dec} = 1 \times 10^{-10}\,\mathrm{Pa\,s}^{-1}$), showing the variation in timing of storm initiation. Year numbering began when the model surface threshold stress field was properly spun-up (Figure 4.3); only the last twelve years of the run are shown.

original value). The stability of the model was greatly improved by the negative feedback on lifting provided by ${\zeta^t}_{inc}$; however, runaway dust lifting, causing a model crash, was still an occasional occurrence, due to very rapid regional opacity increases during the initiation of major storms in certain cases. In order to successfully carry out the multidecadal simulation, a condition was added to the wind stress lifting scheme to shut off dust lifting at any gridpoint where the column visible opacity exceeded 10, as a last resort to prevent a crash in these select instances. This condition therefore slightly limited the rate of expansion of GDSs in a few years of this long run, but is not expected to have affected the peak global opacity attained or the decay phase of the storms, since such large local opacities are only possible during an explosive lifting phase at the beginning of a storm.

Figure 4.1 shows a 12-year sample of the longer, multidecadal simulation, which was begun after a number of spin-up years during which the threshold field reached a quasi-equilibrated state (see Figures 4.3(a) & 4.4). Interannual variability in dust storm magnitude and initiation date increased dramatically, in comparison to the unlimited surface dust simulation seen in Figure 3.4. Over the full duration of the run, large dust storms (those producing a globally averaged visible opacity of greater than ~ 1) began at dates in the range $L_s = 205°$–$280°$. This represents a significant portion of the observed window for

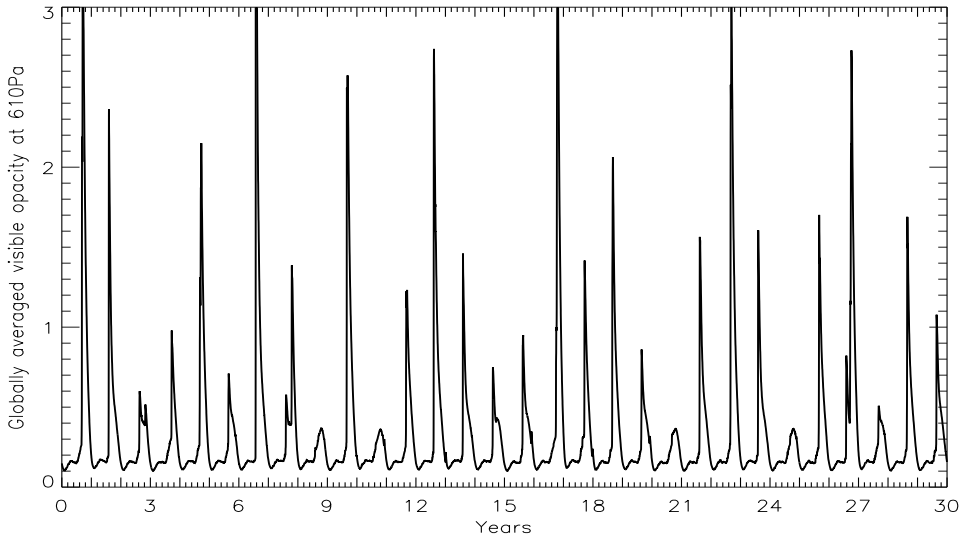

Figure 4.2: Globally averaged optical depth over a period of thirty model years, from the same run as in Figure 4.1.

global storm occurrence: the only GDS known to have initiated later than $L_s = 280°$ was that of 1924 (Phillips, 1924). Even the largest model storms decayed sufficiently rapidly — through a combination of lifting being made more difficult by increasing threshold stresses and a reduction in surface winds due to increased atmospheric static stability under high dust loadings — so as to return global opacity levels to their standard 'background' state of ~ 0.2 by $L_s = 0°$, in agreement with observations. Also, two sizeable storms could occur in the same year, such as at $L_s = 205°$ and $250°$ in year 27, which bears a resemblance to MY12, in which storms began at around $L_s = (202°, 269°)$.

Figure 4.2 shows the chronological progression of globally averaged opacity over thirty model years. The peak opacity attained in a given year clearly depends strongly on the size of the storm that was produced in the previous year, as the largest storms tended not to occur in successive years. A rough periodicity of 4–7 years can be discerned from the plot. If global dust storms are defined as those for which global opacity reaches ~ 2, then the '1 in 3' observed frequency of occurrence was roughly replicated by the model in this case. However, also present were a significant number of model years in which opacity exceeded 1 (19 of the 30 years shown), resulting from dust loadings which were too large to be said to be typical of the regional storms that usually occur in non-GDS Martian years (see Fig. 4, Montabone et al. (2005)). Observations seem to indicate a separation in size between

global and regional dust storm peak opacities (though not decisively so, as yet), which the model did not simulate.

4.3.2 Surface dust redistribution

Within a few model years, the threshold stress field evolved, following dust lifting and deposition, towards a 'climatological' state, as shown in Figure 4.3(a) (in which is plotted the threshold field averaged over the final six years of the simulation), about which anomalies occurred from year to year, affecting the ability of various source regions to supply dust to the atmosphere. This mean state featured significant threshold increases over 15–45°S (the location of the summer subtropical jet), and along the topographic channels at Chryse, Isidis and Amazonis, denoting persistent dust lifting. Conversely, the low-TI 'continents' at Amazonis/Tharsis, Arabia and Elysium were net accumulators of dust (gridpoints that remain at the minimum allowed threshold can be inferred to be sink regions), particularly on their eastern edges. Differences between this spatial pattern and the peak wind stress distribution shown in Figure 3.3(b), taken from a non-storm, passive dust run, highlight the effect of lifting feedbacks, principally an intensification of the southern subtropical jet during spring/summer dust storms.

The source regions seen here correlate well with those implied by the Dust Cover Index (DCI) (Figure 4.3(b)), as derived from TES data by Ruff and Christensen (2002). Though only a qualitative measure of surface dust depth, it indicates that the low-TI continents are very dusty, while most of the southern hemisphere and the Chryse and Utopia flushing channels are relatively depleted of dust. The similarities between the model surface dust field and the observed DCI over most of the planet suggest that the current surface dust field can be explained in terms of present-day wind regimes, rather than requiring the invocation of a past climate state to redistribute dust so as to produce the present DCI. However, possible transport due to dust devil lifting is not considered in this reasoning.

The contribution of various locations on the planet to storm generation and the variation of threshold stresses about their climatologies are highlighted in Figure 4.4, which shows the threshold stress variation in response to dust lifting from several key regions. Large storms still generally began when a flushing storm from the northern hemisphere (originating in Acidalia or Utopia) crossed the equator, whereupon lifting was initiated in the southern

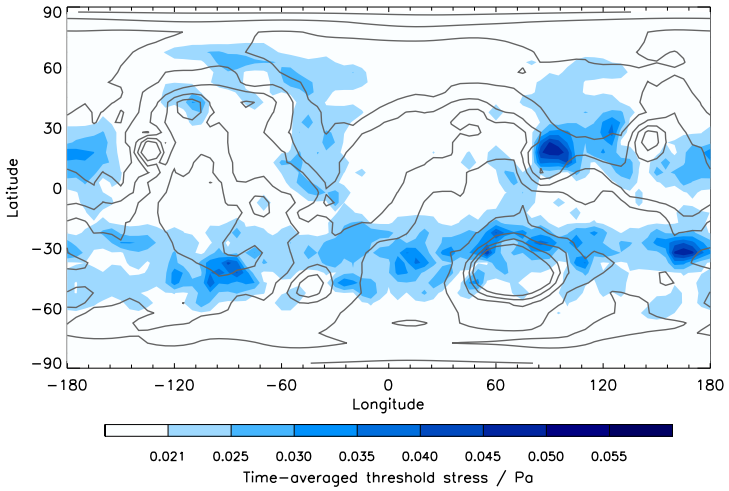

(a)

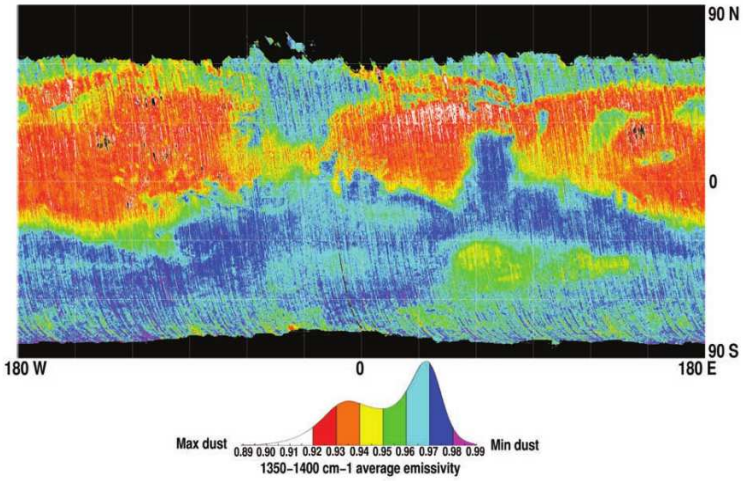

(b)

Figure 4.3: (a) The climatological threshold stress field, estimated using the model field averaged over 6 years. Thresholds in sink regions are always close to the minimum allowed value of ~ 0.02 Pa, while source region thresholds are permanently elevated. (b) The 'Dust Cover Index' for Mars, from Ruff and Christensen (2002) (blue areas have low dust surface density, suggesting that they act as source regions).

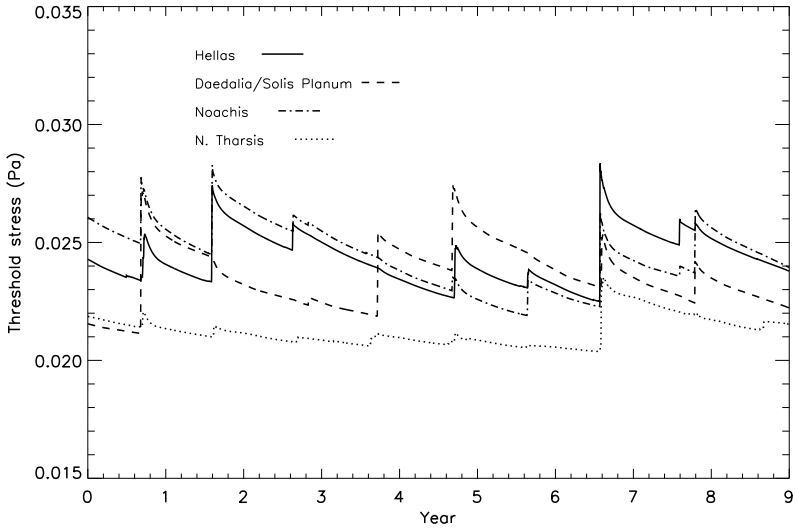

Figure 4.4: Behaviour of effective threshold stresses averaged at key SH lifting regions and at Tharsis in the NH, over a portion of the run (the first nine years shown in Figure 4.2).

midlatitudes, usually either near Daedalia Planum or northeast of the Hellas basin. Often lifting was subsequently activated over large parts of the southern midlatitudes, as seen in the coincident rising of thresholds at the three southern hemisphere sites shown, during certain years. During other years (e.g. in the fourth year plotted), smaller events began at one of the sites (again, usually sparked by a flushing storm, though not always), but failed to become planet-encircling, at least in part due to the thresholds at the other sites being too high to allow the sites to become involved in dust lifting.

The value used for ${\zeta^t}_{inc}$ was $0.5\,\mathrm{Pa}\,(\mathrm{kg\,m^{-2}})^{-1}$, which led to rises of up to $0.005\,\mathrm{Pa}$ in the effective (i.e. following the reduction by a factor of $(0.7)^2$) threshold stresses at the SH lifting sites (area-averaged; individual points could increase by $0.015\,\mathrm{Pa}$ during an intense burst of lifting). Threshold stresses fluctuated about a mean value of $\sim 0.025\,\mathrm{Pa}$, as seen in Figure 4.4. Thus, variations of no more than $\pm 15\%$ about this mean value were sufficient to affect drastically the resulting dust storm behaviour. For one thing, this illustrates the sensitivity of the lifting scheme to the precise choice of threshold. The value for ${\zeta^t}_{dec}$ was then chosen to counteract the threshold increases caused by ${\zeta^t}_{inc}$, so as to allow GDS to form with the desired frequency. This was found to require ${\zeta^t}_{dec} \sim 1 \times 10^{-10}\,\mathrm{Pa\,s^{-1}}$, which equates to a dust resupply rate of $4.75\,\mu\mathrm{m\,yr^{-1}}$. Compared to estimated settling rates of $\sim 12\,\mu\mathrm{m\,yr^{-1}}$ (Rover Team, 1997), this is a high rate of dust 'creation' in the model, much

more than is needed to balance loss to the polar regions, or the forced loss of dust through the use of the minimum threshold, which was estimated as being only $\sim 0.1\,\mu\mathrm{m\,yr}^{-1}$. This term is discussed further in §4.4.2.

4.4 Analysis of selected model years

4.4.1 Dust storm characteristics

The interannual variability produced can be understood better by examining the dust lifting activity of individual model years. Figure 4.5 shows the progression of a dust storm in each of two successive years (years 25 & 26 above). In both cases, the storm began in the pre-solstice period ($L_s = 230$–$245°$) as a flushing event in the Chryse region, but the subsequent developments were very different. In the first year shown, a burst of dust lifting occurred at 55°N, 45°W at $L_s = 243°$, coincident with a local increase in transient wave activity, producing a local visible opacity of > 2, and the resulting storm moved southward from Chryse towards Solis Planum. It reached $\sim 60°$S, south of Daedalia, and dissipated after 10 sols. A second lifting event was seen soon after, again at Chryse, but failed to develop into a flushing storm.

In the second year, the lifting that led to storm formation occurred earlier ($L_s = 231°$) but in the same location as in the previous year. The peak in opacity was larger in magnitude and spatial extent than before, but flushed southward along a similar path. However, on this occasion the storm grew in strength as it crossed the equator, and upon reaching Solis Planum it sparked intense dust lifting at a number of gridpoints (sol 459). The storm then grew rapidly, spreading westward by activating further lifting sites around Daedalia and Sirenum ($\sim 45°$S, 150°W), and also moving eastward into northern Noachis. The initial SH lifting at Solis Planum only lasted for a few sols before dropping out, by which time neighbouring regions had taken over as dust sources. The storm became planet-encircling at around $L_s = 238°$, 10 sols after initiation, and covered most of the globe at its maximum extent. Global opacity began to decline as early as $L_s = 240°$, when lifting had largely stopped at the Solis/Daedalia sites, but did not return to the 'background' level until $L_s = 330°$.

The effect of these dust storms on the threshold surface stress field is shown in Figure

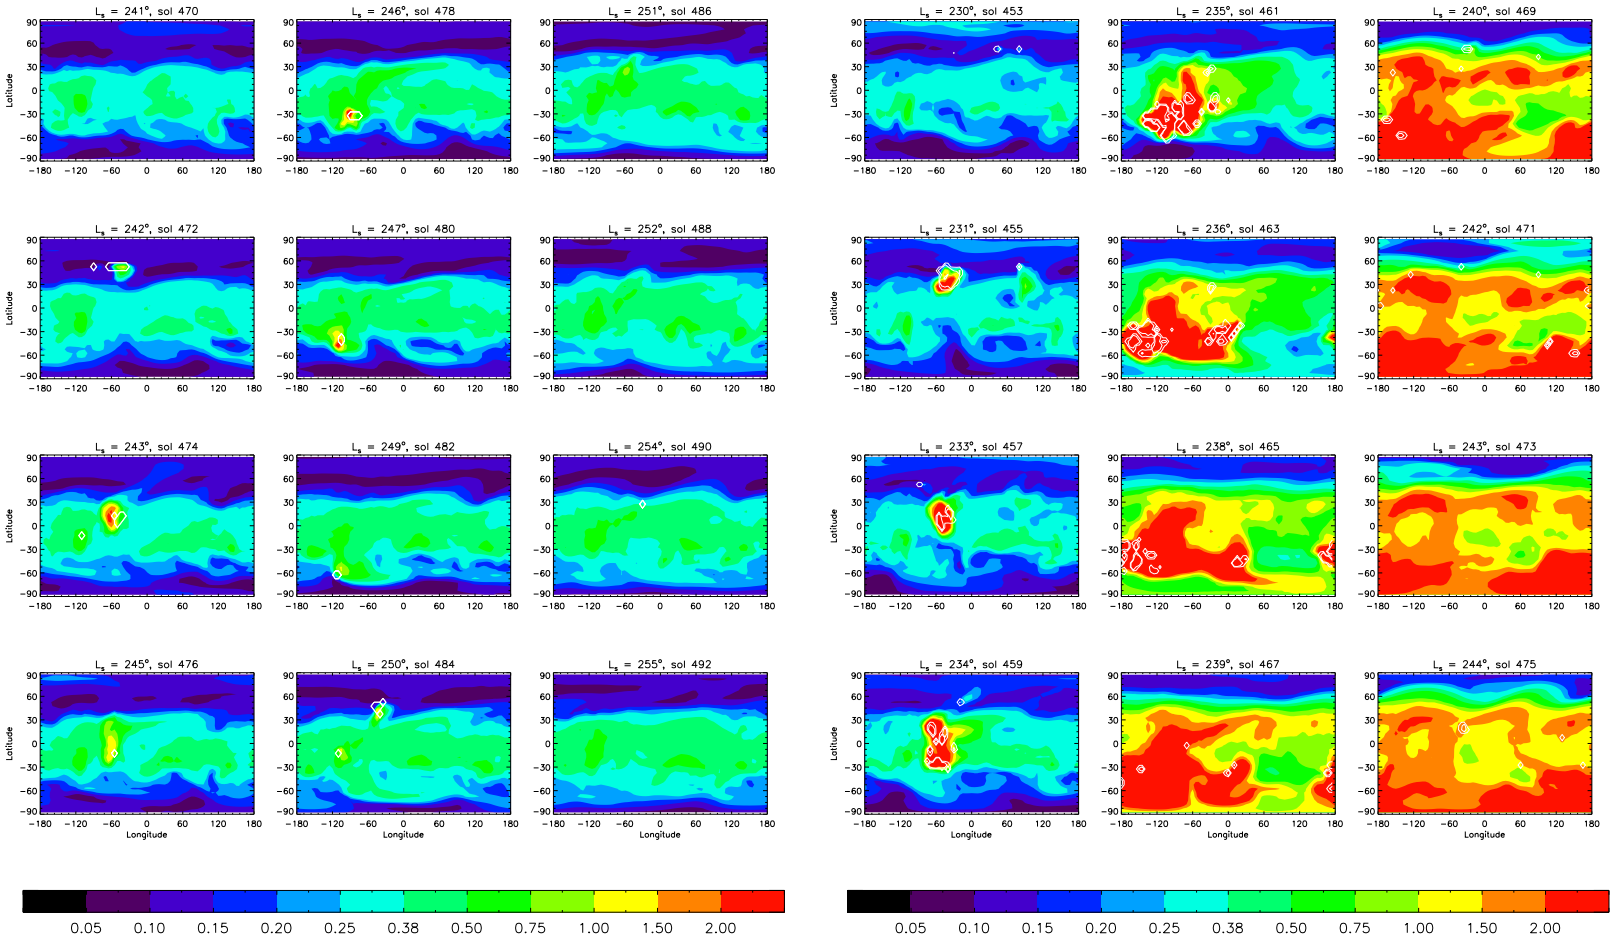

(a) (b)

Figure 4.5: Progression in 610 Pa opacity for sections of years 25 and 26 of the longer series: (a) $L_s = 241$–255° of year 25, and (b) $L_s = 230$–244° of year 26. White contours mark where wind stress dust lifting is occurring — contour levels are $(10^{-8}, 10^{-7})\,\mathrm{kg\,m^{-2}\,s^{-1}}$.

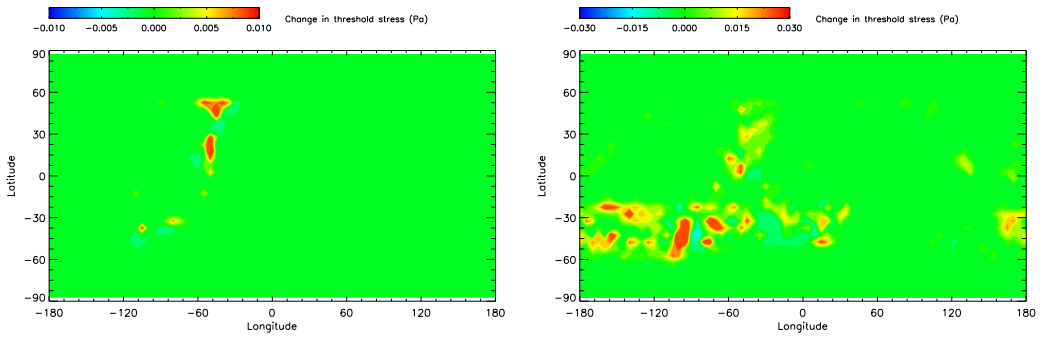

Figure 4.6: Change in the threshold stress field (in Pa) between $L_s = 240°$ and 260° in year 25 (left) and between $L_s = 230°$ and 260° in year 26 (right). Note the different colour scales used.

4.6 (note that different scales are used for each of the plots). As expected, surface dust removal in the first of the two years occurred primarily in the northern hemisphere, while in the second year much of the western hemisphere between 20°S and 60°S lost dust. It was the large increases in threshold stress at Solis Planum, representing a significant depletion of surface dust, that caused the cessation of lifting seen there a few sols into the storm.

Figure 4.7 shows the spatial distribution of mean and transient RMS surface wind stresses for the two years during the pre-storm periods, at $L_s = 230$–$240°$ in year 25 and $L_s = 220$–$230°$ in year 26. Both measures of surface stress were very similar between the two years. If anything, mean stresses were slightly stronger in the first year at 30°S, due to the more developed subtropical jet that had formed by the later date shown. Travelling waves were also stronger at $\sim 70°$N in the first year than in the second. Despite this, it was the second year which produced a planet-encircling storm following the period shown, whereas the flushing storm that occurred in the first year dissipated upon reaching the southern hemisphere.

The effect of increased dust opacity on transient waves, a topic that is addressed in more detail in Chapter 5, is demonstrated in Figure 4.8. In the first year, the main increase to dust loading around perihelion occurred due to the seasonal increase in dust devil lifting, and surface transient waves remained relatively strong at 60–80°N, particularly at Acidalia. In the second year, the global dust storm was underway by $L_s = 250°$, greatly weakening transient waves at high latitudes in both hemispheres. Note that waves were suppressed at 60–80°N even though increases in dust opacity were more limited at these latitudes than over the rest of the planet (Figure 4.5(b)).

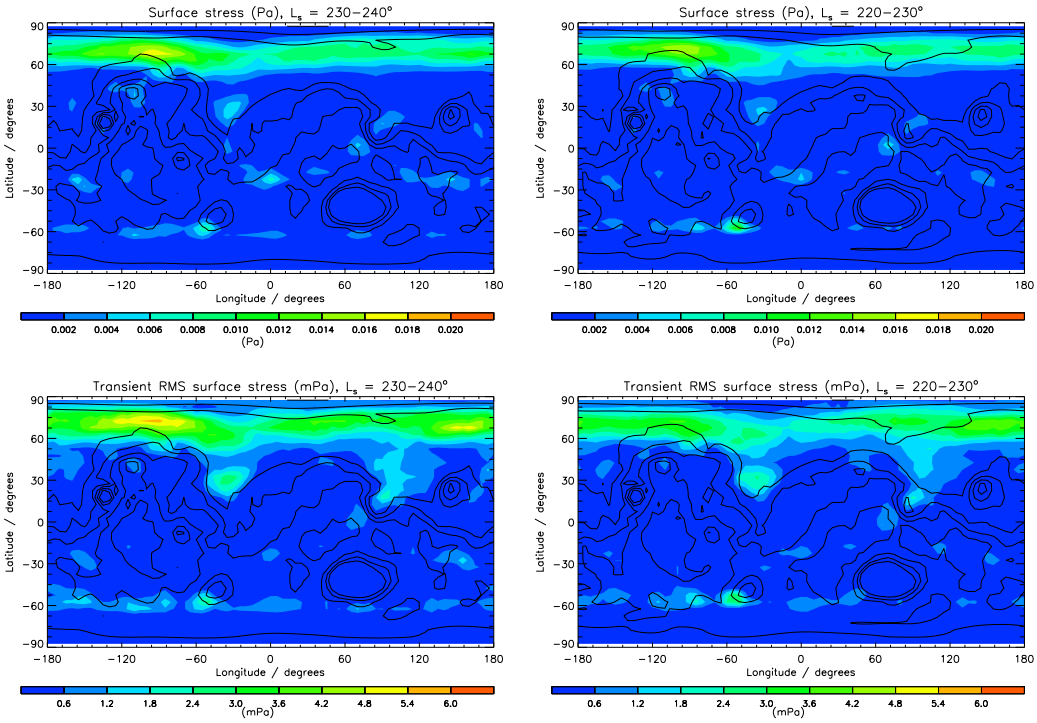

Figure 4.7: Mean (top row) and transient (bottom row) surface stress, averaged over 10° of L_s prior to the storms that occurred in year 25 (left column) and year 26 (right column).

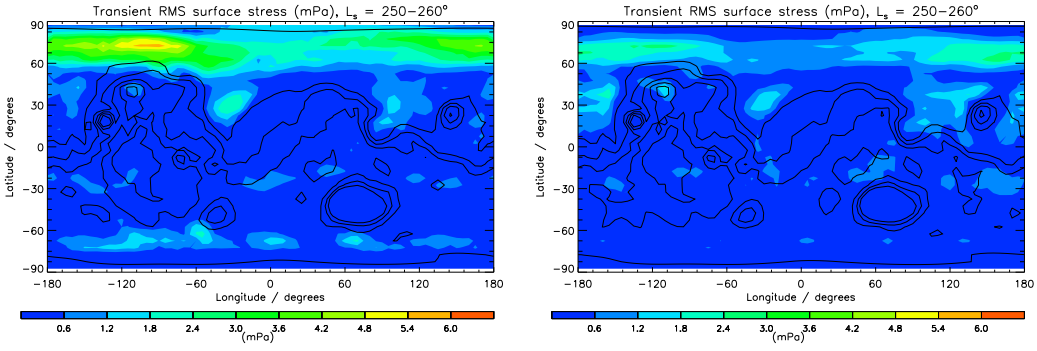

Figure 4.8: Transient RMS surface stress (in mPa), averaged over $L_s = 250$–$260°$ in years 25 (left) and 26 (right), showing the weakening of eddies that occurs during a GDS (year 26).

The evolution of both mean and wave components of surface stress during a GDS is shown in more detail in Figure 4.9, which covers the same period as Figure 4.5(b). Mean stress (and surface wind magnitude) increased as the storm grew and dust opacity increased, first west of Daedalia, then in northern Noachis and Eridania. By sol 475, stresses were raised across the 30°S band and along the flushing storm channels, signifying a global strengthening of the PMOC. Transient stresses responded in a similar way, but increases around Daedalia were more widely spread in area than those of the mean stress, which were limited to narrow bands at 30°S and 60°S. By sol 475 a decrease in eddy activity due to high dust loading had begun; it evidently proceeded rapidly after this point, to give way, within $\sim 10°$ of L_s, to the field shown in Figure 4.8(b).

Rather than distinct wind regimes, in this case the key difference between the pre-solstice periods of the two years was found in the threshold stress field, which declined slowly over time due to the resupply term used. Since increases to threshold stresses due to lifting were small in year 25 (Figure 4.6), by the time the storm season arrived in year 26 thresholds were generally lower than they were before the storm of the previous year. Figure 4.10 shows that much of the 30°S band saw thresholds fall by more than 0.005 Pa, which represents a significant change in the context of the long-term evolution of southern hemisphere thresholds (c.f. Figure 4.4). This change made lifting at Daedalia much more likely, and also facilitated the growth of the dust storm, as this initially occurred along the 30–60°S band through the activation of secondary dust lifting sites. The importance of the threshold stress to the development of large dust storms was tested by rerunning year 26 after replacing the threshold stress field at the start of the year with the (larger) values from the start of year 25 (i.e. negating the action of the resupply term and any dust deposition over the course of the first year). The resulting storm again began near Chryse, but was much more limited than the original year 26 storm, with a peak globally averaged opacity of less than 1. The magnitude of this storm lay between the magnitudes of the storms in the two years above, indicating that dust lifting is controlled by a combination of interannual variability in wind stresses and variation in the threshold stress.

In Figure 4.9(b), there is the hint of an increase in eddy stresses in the western hemisphere at 20–50°S slightly *before* the arrival of the dust cloud (panels at 231° and 233°). This could be explained if the radiative heating anomaly at the location of the dust cloud

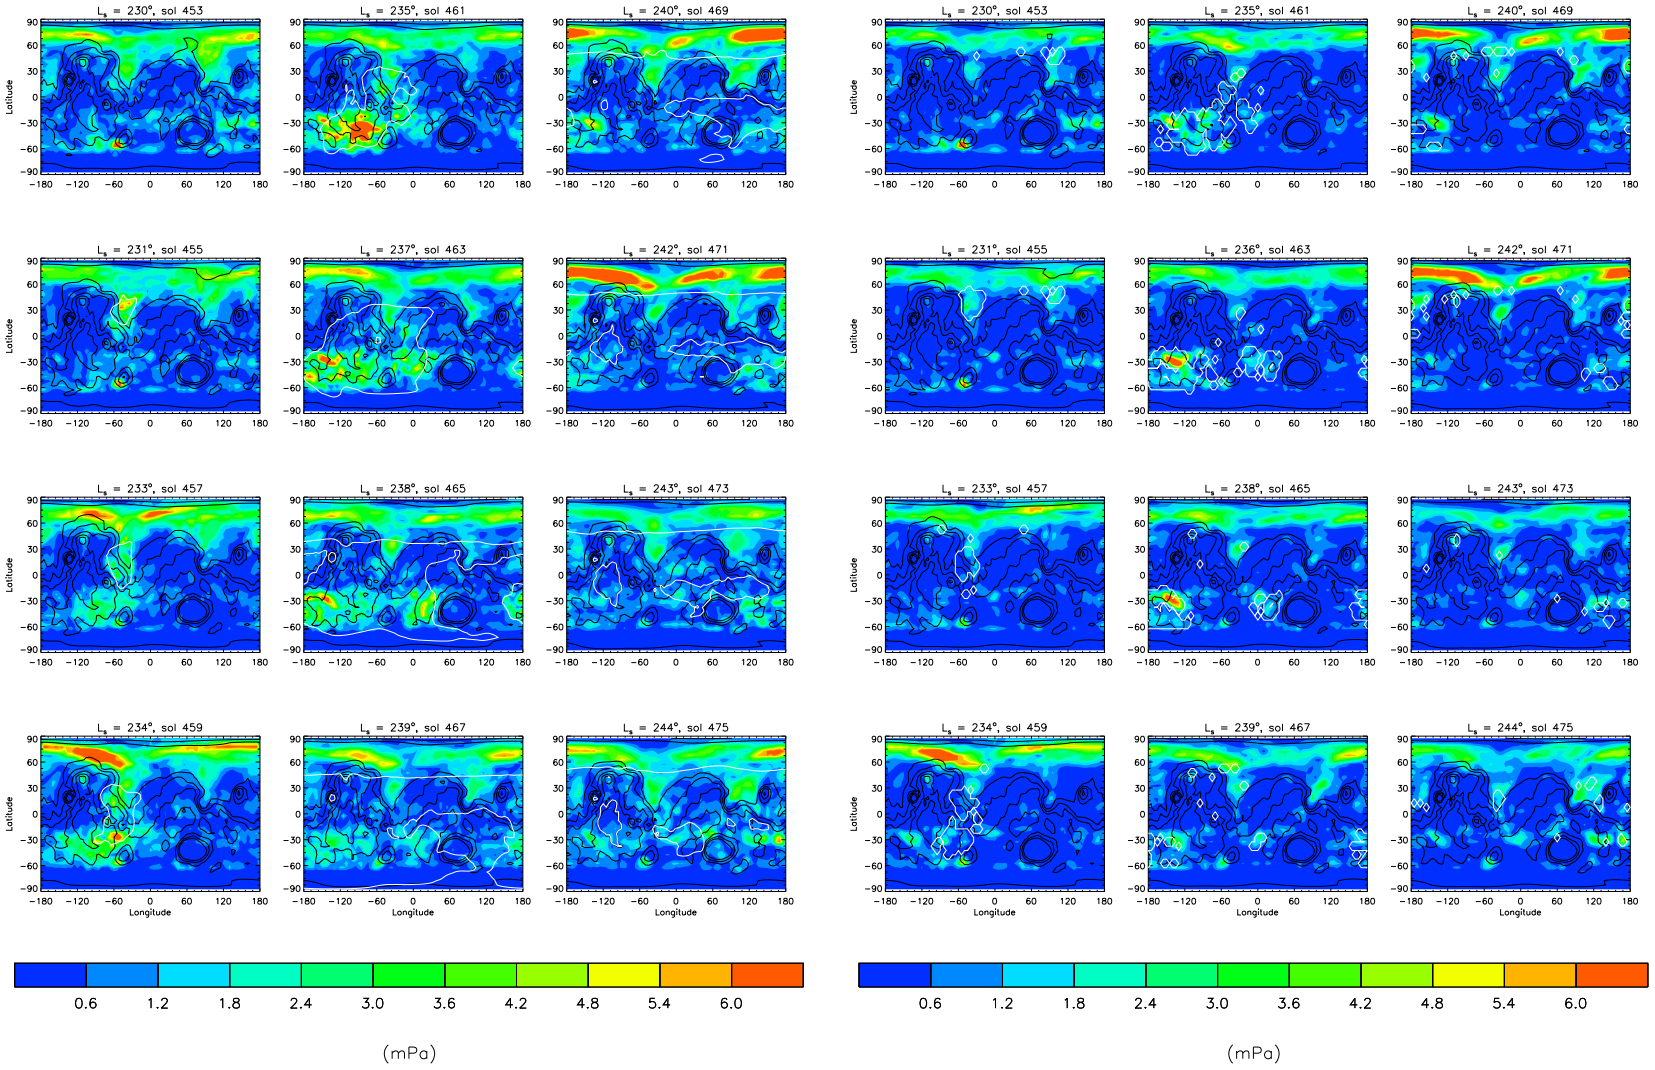

Figure 4.9: Time series of mean (left, in Pa) and transient (right, in mPa) surface stress during the development of the dust storm in year 26. The timing coincides with that shown in Figure 4.5(b). In the right plot, the white contour marks 610 Pa optical depth values of 1, to show the progression of the dust cloud.

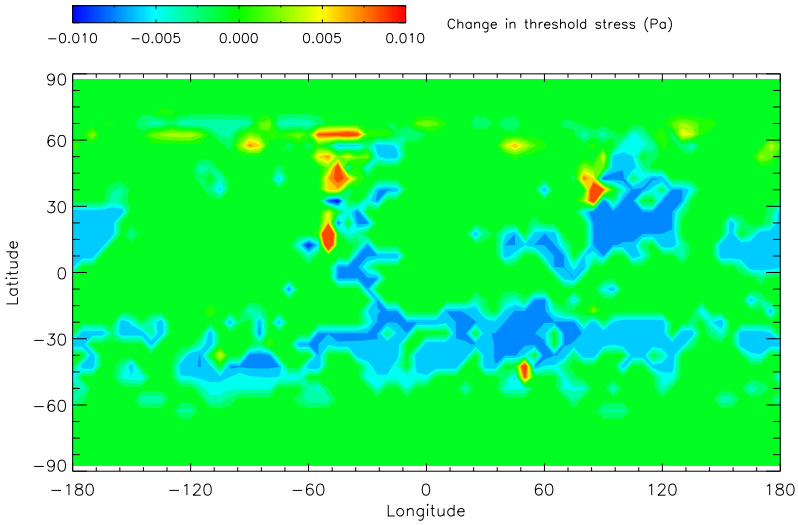

Figure 4.10: Difference in threshold stress (Pa) between $L_s = 230°$ in year 26 and 240° in year 25, showing the enhanced potential for lifting in the SH at the beginning the second storm season, in comparison to the first.

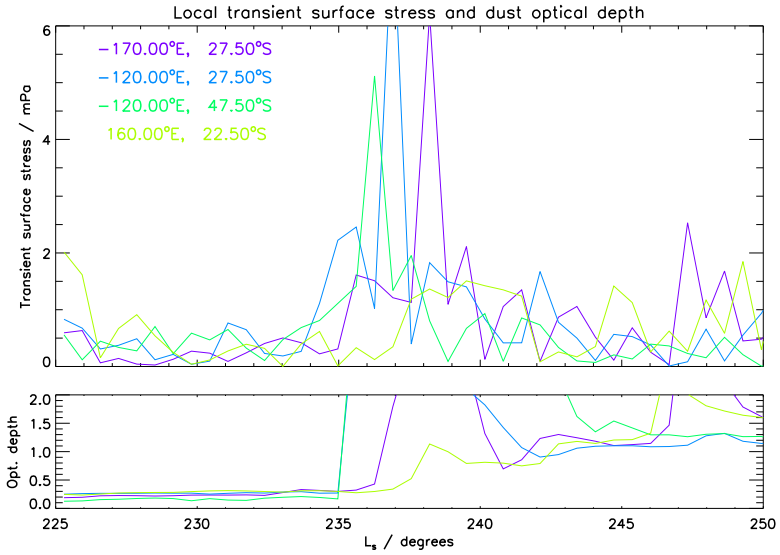

Figure 4.11: Daily ($\sim 0.5°$ of L_s) averaged local transient surface stress and 610 Pa dust opacity values at several gridpoints in the Daedalia and Eridania regions, from year 26. A flushing storm began in Chryse at $L_s = 231°$, and reached 30°S at 234°.

when it first formed, at 40°N, were to excite planetary waves, with wavetrains extending meridionally as well as zonally. The situation is shown in more detail for several gridpoints in Daedalia and Eridania, in Figure 4.11. The separation in time between increases in eddy wind stress and in opacity is small — at most 2 sols — but an advance wave response may indeed be present, and may then contribute to the rapid growth of the dust storm when it does arrive in the southern midlatitudes. The effect is clearest in longer-period (> 5 sols) waves, and can also be detected in the TES reanalysis dataset; for example, in Noachis at $L_s = 223 - 224°$ ahead of the arrival of the MY24 flushing storm. This may be evidence of a teleconnection pattern, like those that have been extensively analysed on Earth (e.g. Trenberth et al., 1998), but with a wave response that is excited on a much shorter timescale ($\lesssim 5$ sols), which is plausible given the differences in radiative timescale between the two planets. The importance of the phenomenon in this model example is uncertain, however, as it was swamped by the more substantial increases in eddy activity that occurred upon the arrival of the flushing storm in the southern hemisphere (at $L_s = 234°$).

4.4.2 Dust mass transport and deposition

Cross-equatorial transport

It is instructive to examine the net cross-equatorial dust transport simulated by the model during the above years, both to potentially justify the use of the artificial resupply rate and to investigate the role of dust transport in maintaining the GDS frequency of around one in three years. Several authors (e.g. Haberle (1986), Szwast et al. (2006)) have sought to explain the observed variation in dust storm intensity through a periodic redistribution of dust between the two hemispheres: it has been assumed that a global storm, originating in the southern hemisphere at a time when the upper branch of the PMOC is directed northward, causes a net transfer of dust into the northern hemisphere, which must be returned in subsequent years in order to allow initiation of another GDS. Dust may be transported southward by the upper PMOC branch during northern spring/summer, and potentially via the flushing storm channels, along the western boundaries of which flows the lower branch of the PMOC, later in the year. It is not known exactly on what timescale this replenishment takes place, but Christensen (1988) observed a brightening in the northern hemisphere following the dust storms of 1977, and Putzig and Mellon (2007) showed that

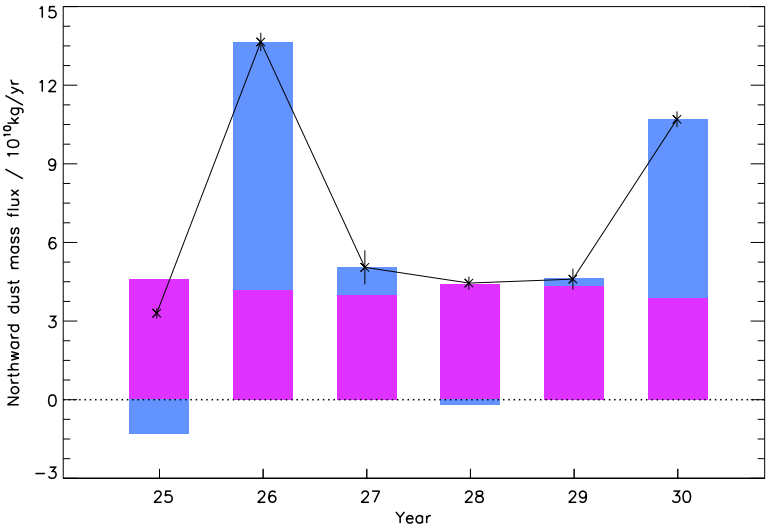

Figure 4.12: The net rate of transfer of dust from the southern hemisphere to the northern hemisphere (in $10^{10}\,\mathrm{kg\,yr^{-1}}$), in each of six model years (solid line), and the contributions to this from dust devil lifting (pink bars) and wind stress lifting (blue bars). The error bars attached to the overall values represent the level of accuracy in the calculation allowed by the less-than-perfect mass conservation in the model (cross-equatorial fluxes are of order 1% of the annual global dust lifting and deposition rates). The individual contributions are slightly more approximate as they were calculated neglecting a possible change in atmospheric dust mass between the beginning and the end of the year, except in years 28 and 29, where the fluxes due to each mechanism were recalculated to include the atmospheric mass change of each dust 'type' (lifted by dust devils or wind stress), to ensure that the direction of the wind stress flux was robust.

significant changes in surface albedo occurred in both hemispheres between MY24 and MY26.

The annual cross-equatorial dust flux was calculated for the final six years of the sequence shown previously, by comparing the total masses of dust lifted and deposited in each hemisphere during the year, and adding a correction to account for any change in atmospheric dust loading between the beginning and end of the year. The results are displayed in Figure 4.12. In each of the years, the net mass transfer when considering both lifting mechanisms was northward, at an average rate of $7 \times 10^{10}\,\mathrm{kg\,yr^{-1}}$. Since the period chosen contained both GDS and non-GDS years, this sustained loss of dust from the southern hemisphere is at odds with the multiannual dust redistribution process outlined above. However, as mentioned earlier, transport of convectively lifted dust — with a northward flux of $\sim 4 \times 10^{10}\,\mathrm{kg\,yr^{-1}}$ — appears to be more consequential in the model than observations would suggest (and is resolution-dependent: see §4.5). Such a regular northward dust flux seems hard to justify, unless wind stress lifting is vastly underestimated at present, as

it would make it very difficult to keep the SH supplied with dust. In any case, dust lifted convectively did not take part in the redistribution simulated by the model, as the dust devil scheme was not constrained by a finite dust surface density.

For dust lifted by near-surface wind stress only, the net annual transport was dependent on the dust storm activity of the year in question, but was northward overall, at a rate of $\sim 3 \times 10^{10}\,\mathrm{kg\,yr^{-1}}$. In only two of the six years (years 1 and 4), neither of which featured a major dust storm, was the net transport southward, and it was certainly not sufficient to balance the loss to the NH during the other (dust storm) years. Most of this lost dust ended up in the low-TI regions or in northern high latitudes (see Figure 4.3(a)). This explains why the resupply term was needed: southern hemisphere dust cover, vital for the formation of regional and global storms, was not being replenished rapidly enough during non-storm years to lower thresholds after a period of lifting in the SH. The southward flux of the first year (the largest storm during which was the flushing storm shown in Figure 4.5(a)) does, however, provide an indication that replenishment of SH dust sources during years without any major storms could be possible. This issue is discussed further in §4.5, in light of results at higher horizontal resolution.

Seasonally, it was found that the direction of cross-equatorial flux (for dust lifted by both mechanisms) followed the direction of the upper branch of PMOC: accordingly, it reversed at $L_s \sim 180°$. During northern spring/summer, some of the dust lifted by dust devils (0.4–0.8% of the total) was advected across the equator from the northern to the southern hemisphere, by a combination of stationary waves (found to be the largest contributor to cross-equatorial flux during this period) and the mean flow; cross-equatorial flux due to wind stress lifting was negligible during this period. During southern spring/summer, transport was in the opposite direction, was dominated by the PMOC and was generally much stronger. Over this period, both forms of lifting were active in each hemisphere, but both were stronger in the SH than in the NH. The stronger PMOC at this time of year led to more efficient cross-equatorial transport: 2–5% of the total lifted was transferred between hemispheres, depending on the year. The differences in transport between the four 'dust storm' years were surprisingly large, and net interhemispheric mass transport did not simply depend on peak storm magnitude. What can be said is that the most efficient south-north transport occurred during $L_s \sim 210$–$240°$, and flux tended to be higher when the dust storm

began in Noachis or Daedalia than when it began at other longitudes (Hellas/Eridania).

These results are quite different to those reported by Haberle et al. (2006) for the Ames model — their runs, which contained a dust storm around $L_s = 270°$, produced a northward flux of dust lifted by wind stress, but this was approximately cancelled by a southward flux of dust lifted convectively (of $\sim 10^{11}\,\mathrm{kg\,yr^{-1}}$), leading to an annual cross-equatorial transport that was beneath the level of mass conservation in their model[2]. It is surprising that the Ames model produced a dust devil flux so different to that in the UKMGCM, since both used effectively the same parameterisation for dust devil lifting (Kahre et al., 2006). One reason may be that the regular southern solsticial dust storm in their model tended to reduce dust devil lifting when it would normally be at its strongest in the SH, thus limiting the potential for northward transport (dust devil dust flux over $L_s = 0$–$180°$ in the UKMGCM was indeed directed north-south). Another could be that dust is less efficiently mixed up through the boundary layer after being lifted, which would affect its ability to enter the upper branch of the PMOC rather than the lower branch (see §6.6.1). If nothing else, this discrepancy indicates that the global-scale effects of dust devil lifting and transport are not well constrained at present, and therefore justifies the decision to exclude its influence from much of the discussion in this chapter.

The resupply term

It is appropriate to say a little more about possible justification for the use of the resupply term, ${\zeta^t}_{dec}$, by interpreting its physical meaning. As stated previously, the term was introduced for pragmatic reasons, to prevent the model from becoming locked in an unrealistic, low-dust state during a multiannual simulation. The reason for its need has been identified as the long-term model mass transport bias towards the northern hemisphere, which must inevitably lead to the cessation of dust storms in the southern hemisphere if a finite surface dust reservoir is used.

The proper way to correct this model transport bias, if it is indeed a bias, is through model improvement. However, the need for a parameterised reduction in threshold (equivalent to an increase in *available* dust) can extend beyond the crude correction of transport bias. The introduction of the scheme by Pankine and Ingersoll (2004) suggested that a

[2]The level of mass conservation in the UKMGCM has been estimated at ∼0.2% of the total dust lifted, so the results discussed here should be robust.

decrease in threshold over time could be caused by a redistribution of the surface dust layer on a microscale level, as dust is shifted out from sheltered locations (near nonerodible elements) into areas from which it can more easily be mobilised by winds. In a GCM, such a process would occur within a single gridbox, and must therefore be parameterised. A similar argument can be made if dust lifting in fact occurs by the direct detachment of sand-sized aggregates (§4.1), as the availability of dust for lifting from a particular gridpoint would then be contingent upon the formation of these aggregates, a process which again must be parameterised, occurs on some characteristic timescale, and is in this case represented by ${\zeta^t}_{dec}$. In either case, then, there exists a plausible mechanism by which the 'effective dust cover' of a gridpoint may increase over time without the input of any extra dust through deposition.

Polar dust deposition

A further area of interest is the rate of dust deposition in the polar regions, due to its effect on ice cap albedo (though not yet in the model) and its significant role in the formation of the polar layered deposits. Annually, deposition in the north polar region was greater by around a factor of two than at the south pole. North polar deposition peaked after $L_s = 180°$ each year as cap-edge lifting began, but as no such storms were produced in the model around the south pole, the only supply of dust to the south pole was from dust lifted at more northerly latitudes and transported via the global circulation. South polar dust deposition peaked during large storms, due to expansion of the PMOC and of a secondary reverse cell present at 40–70°S, as well as the increase in atmospheric dust content. It is not clear how efficiently dust is exported to the south pole during dust storms in reality (Wang, 2007).

Deposition rates for 60–90° latitude, averaged over each of the three years, were 0.011–$0.016\,\mathrm{kg\,m^{-2}\,yr^{-1}}$ (or 4.4–$6.4\,\mu\mathrm{m\,yr^{-1}}$) in the NH, and 0.005–$0.012\,\mathrm{kg\,m^{-2}\,yr^{-1}}$ (or 2.0–$4.8\,\mu\mathrm{m\,yr^{-1}}$) in the SH. Assuming that some of this dust became isolated from the rest of the planet (since very little lifting occurs poleward of $\sim 75°$), these figures provide only a small part of the justification for the resupply rate (of $4.75\,\mu\mathrm{m\,yr^{-1}}$), as the polar region accounts for a small fraction of the overall surface area. Pollack et al. (1979) estimated a settling rate for the north polar region of $\sim 0.2\,\mathrm{kg\,m^{-2}\,yr^{-1}}$ from Viking data, but this

is likely to be an overestimate, as discussed in Chapter 6 — their global average rate of $0.02\,\mathrm{kg\,m^{-2}\,yr^{-1}}$ may be closer to the true polar deposition rate, and is certainly closer to the values estimated here for the model. Another settling rate estimate was made by Drube et al. (2010) for the Phoenix lander, situated at 68°N — the value is rather uncertain, and comes from a limited period during summertime, but they estimated $0.05 \pm 0.04\,\mu$m per sol, which puts the annual rate at anything between 0.005–$0.15\,\mathrm{kg\,m^{-2}\,yr^{-1}}$, depending on the density of dust assumed.

A comparison of the surface deposition rates of dust and CO_2 ice during periods of CO_2 condensation implied a dust content in the southern seasonal cap of $\sim 10^{-6}\,\mathrm{kg/kg}$, and a higher dust content in the northern cap of $\sim 10^{-5}\,\mathrm{kg/kg}$. The ice caps were dustier during the initial stages of their formation, as condensation rates were lower. Deposition onto the north polar cap was enhanced by up to an order of magnitude during major dust storms, producing temporary dust contents of $\sim 10^{-4}$ by mass (the south polar cap was subliming and receding at this time of year). Observations of the dust content of the seasonal caps are few, but indicate substantial spatio-temporal variability. Langevin et al. (2007) estimated no more than $\sim 10\,\mathrm{ppm}$ by volume ($\sim 10^{-5}$–$10^{-4}\,\mathrm{kg/kg}$) of dust in the upper layer of bright regions of the southern seasonal cap, but suggested that this could increase towards the end of the recession of the cap — such an increase in dust content was simulated by the model, at $L_s = 150$–180°. Langevin et al. (2006) identified a (temporarily) very dusty upper layer in the low-albedo 'cryptic region' of the south pole around $L_s = 220°$. They attributed this to inflow from nearby local dust storms, which are severely under-represented in the model, so it is not surprising that the model failed to simulate any peaks in dust content of this magnitude during southern spring.

4.5 At higher resolution

Three model years were run at T63 resolution (a 1.875° dynamical grid), with atmospheric and surface variables (including threshold stress) initialised by interpolating the fields output at the end of the T31 simulation above. The same lifting and threshold change parameters as above were used, though for stability, dust lifting was shut off at any gridpoints with a visible opacity greater than 5 (a higher limit of 10 was used for the third year). Some caution is required with the interpretation of wind stress lifting patterns in these results,

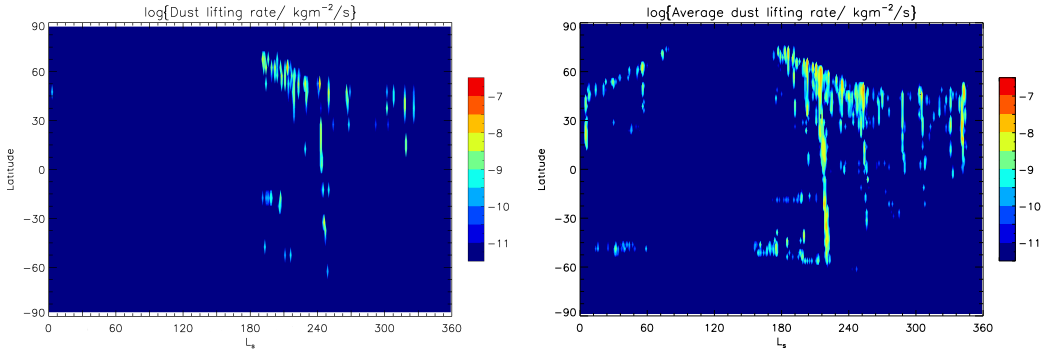

Figure 4.13: Logarithm of the zonally-averaged dust lifting rate, in $\mathrm{kg\,m^{-2}\,s^{-1}}$, for a year without a major dust storm at T31 (left) and T63 (right), using the same lifting parameters.

as the simulation was not continued for long enough to allow the threshold field to adjust to the wind stress lifting patterns of the higher resolution model (i.e. there exists at T63 a different mean threshold state to the one shown in Figure 4.3(a)). Nonetheless, some significant changes to both types of lifting were apparent. As expected from the testing in Chapter 3, dust lifting occurred more readily than at the lower resolution, not least in the NH baroclinic zone, as seen in Figure 4.13, which shows the first of the years run. At least five flushing storms were produced during this year. While this number may be unrealistically high (Wang et al., 2005) (suggesting NH thresholds were a little too low in this first year), it is clear that a better resolution of frontal systems at the northern cap edge was achieved at this resolution. In particular, late summer storms appear to have been produced more readily than previously. Figure 4.13 also illustrates the absence of a solsticial pause (see Chapter 5) in cap-edge dust lifting in the current model, at least whenever global opacities have not been raised by a pre-solstice dust storm: in the T63 run, flushing storms occurred at $L_s = 252°$ and 288°, when no such storms are usually observed in reality. Cap-edge lifting in the southern hemisphere was increased somewhat by the increase in resolution in the first year, though the lifting at $L_s = 0$–$60°$ was not repeated in the next two years, due to the increase in thresholds caused (as the model adjusted to its new climatology). In the second year (not shown), two dust storms occurred during NH spring and summer, confirming that higher minimum lifting thresholds should be used at this higher resolution.

Obtaining reliable cross-equatorial transport figures was difficult at this resolution, due to the prohibitive computational cost of running for more than a few Martian years. How-

ever, several important results were noted. In the first year, which featured no major dust storm, the transport was strongly southward, even when considering both lifting mechanisms ($5 \times 10^{10}\,\mathrm{kg\,yr^{-1}}$). In addition to the southward transport occurring due to the series of flushing storms, the situation was helped by a weakening, by around 50%, of the net northward transport due to dust devil lifting, which occurred primarily through stronger southward transport over $L_s = 0$–$180°$, produced by a combination of a decrease in SH dust lifting relative to the NH and a slightly expanded PMOC.

The wind stress lifting fluxes from both these and the T31 results of the previous section allow an evaluation of the plausibility of transport during non-GDS years (primarily via flushing storms) being rapid enough to offset GDS loss to the northern hemisphere within a few Martian years. Given the typical observed ratio of two non-GDS years to every one GDS year, the southward flux during the former would need to be around half as large as the northward flux in the latter, in order to support continued GDSs in the model without the need for the resupply term. Taking the average of the four dust storm years at T31 (each featured a peak visible opacity of at least 1) as a rough estimate of the northward wind stress transport during a GDS year gives a value of $\sim 4.5 \times 10^{10}\,\mathrm{kg\,yr^{-1}}$, while the two non-storm years provided a transport of only $\sim 0.6 \times 10^{10}\,\mathrm{kg\,yr^{-1}}$ in the opposite direction. Therefore, an increase of a factor of 3–4 in flushing storm transport would be required if the GDS frequency were to be explained in this way.

However, cap-edge lifting and flushing storm frequency were almost certainly underestimated at T31 resolution (1–3 flushing storms were produced per year), and a much larger southward wind stress flux of $\sim 7 \times 10^{10}\,\mathrm{kg\,yr^{-1}}$ was recorded in the non-GDS year at T63 resolution. The occurrence of five flushing storms in a single year is probably unrealistic (Wang et al., 2005), so a more reasonable estimate of the north-south mass transfer rate due to flushing storms might be $(3$–$4)\times 10^{10}\,\mathrm{kg\,yr^{-1}}$. By contrast, the GDS transport at T63 resolution, estimated from two available years of dust storms, was $\sim 6 \times 10^{10}\,\mathrm{kg\,yr^{-1}}$ northward — only a slight increase over the T31 values. Therefore, had the ratio of dust storm to non-storm years in the model matched the observed value of one in three (it was closer to two in three), these figures suggest that the net flux over multiple years would have been close to zero (in a very rough sense), and the resupply term would not have been needed.

As an explanation of the observed GDS frequency, a potential flaw in this argument is the neglect of cross-equatorial transport due to dust devil lifting. However, the significant change in the rate of this transport caused by an increase in resolution highlights the current lack of certainty regarding it. In particular, dust devil transport during $L_s = 0$–$180°$ should depend very much on the vertical distribution of dust mass in the tropics, and in this regard the cross-equatorial transport was most likely biased towards the northern hemisphere, since the model did not adequately reproduce the dust mass mixing ratio profiles seen around aphelion by Mars Climate Sounder (Heavens et al., 2011) (this is explained further in §6.6.2). In reality, it may be that a significant amount of southward dust transport occurs over $L_s = 0$–$180°$, further assisting the return to the southern hemisphere of dust lost during a GDS. Considering wind stress transport only, there is reason to believe that the northward dust flux resulting from a typical GDS (with some uncertainty as to exactly how large this may be) could be balanced by transport in the reverse direction during non-GDS years, over a period of around three Martian years.

4.6 Summary and comparison with dust storm climatology

Interannual variability in dust storm magnitude and timing of initiation was greatly increased by the addition of temporally varying dust lifting thresholds, and a range of dust loadings during southern spring and summer were produced during a single multi-year simulation. The increase in variability was attributed to the oscillation of surface thresholds at several important sites in the southern hemisphere, in response to emission and deposition of dust mass. In fact, the initiation of a major dust storm at one of these sites in a given year was primarily determined by its surface dust supply, implied through local threshold magnitudes. It is not known, in reality, how strongly the distribution of surface dust controls the generation of large dust storms in the SH, but this does illustrate the great sensitivity of modelled dust lifting to the precise choice of threshold, and, therefore, the importance of estimating thresholds accurately.

After several years of simulation, the dust surface density field showed good qualitative agreement with the observed pattern of Martian surface dust cover. The model produced

a net northward cross-equatorial dust mass flux, which necessitated the addition of an artificial threshold decrease rate in order to allow the continued generation of dust storms over the course of a multi-year simulation. At standard model resolution, for the southward mass flux due to cross-equatorial flushing storms to offset the northward flux due to GDSs on a timescale of ~ 3 years would have required an increase in the former by a factor of 3–4. However, results at higher model resolution and uncertainties in dust vertical profiles provided evidence that this could be a plausible explanation for a GDS frequency of one every three years.

The variable threshold model generated storms with a range of initiation dates, but its behaviour can be broken into several broad categories, and evaluated against the observed dust storm climatology, as identified in §1.2.3. Firstly, no storms were produced in the model as early as the GDS of 2001, during $L_s = 180$–$205°$. The model regularly produced a peak in surface stress in SW Hellas around $L_s = 180°$, but with the threshold and lifting efficiency chosen, this was never sufficient to initiate anything larger than a local storm. As suggested by Basu et al. (2006), this is likely to be partly a result of unresolved mesoscale circulation patterns that are important around Hellas and other regions: it should be recalled from §3.3 that the model appears to miss out significant quantities of lifting along the SH cap edge, perhaps due to the importance of local slope and/or thermal contrast flows in this process. However, the absence of a 2001-like storm in the model is neither surprising nor of great concern, as the evidence so far suggests it was an exceptional event. The 2001 GDS occurred about 20° earlier than any other planet-encircling storm on record, and expanded rapidly even though the southern subtropical jet should have been fairly weak at that time of year. It is likely that thermal contrast and/or topographically influenced flows are particularly important for dust storms beginning in early spring, as the extension of the south polar cap puts the baroclinic zone close to the basins near 60°S over $L_s \sim 180$–$220°$.

The model often produced either a regional or global storm during $L_s = 210$–$240°$, which is known to be a particularly dusty time of year (Liu et al., 2003). However, modelled global storms after 240° were fairly rare, while several of those observed (MY9, MY12, MY28) have begun around $L_s = 260$–$270°$. The most obvious explanation for this deficiency is that the model had usually produced a large storm earlier in the year, after which the decrease in surface winds in the dusty, isothermal atmosphere made it difficult to produce any more

dust storms until the original event had cleared. None of the three years mentioned above are known to have featured a large dust storm beginning as late as $L_s \sim 230°$, as generally occurred in the model; furthermore, it is interesting that the one example of a post-240° GDS in Figure 4.1 (in year 27) was preceded by a storm that began unusually early, at $L_s \sim 205°$, and had almost fully decayed by 250°. It is also notable that no recorded planet-encircling storm has begun during 210–240°. This period generally sees high flushing storm activity in the northern hemisphere, which may actually increase stability in the southern hemisphere and prevent GDS initiation there until flushing storms have ceased (see Chapter 5). This suggests that many of the storms shown during $L_s = 210$–240° in Figure 4.1 tended to grow too large (which contributed to the net northward dust flux), and highlights the fact that their genesis, from flushing events that cross the equator, was unrealistic.

There was also a lack of late-summer (post-300°) storms in the simulation described above. In some years this can again be explained by elevated dust opacities following on from global storms earlier in the year; in other years though, opacities had returned to close to the background level, and the lack of further dust lifting is harder to justify. It is quite possible that the post-solstice rise in eddy activity, which is responsible for the resumption of NH flushing storms, was too weak in the model, and dust lifting may also have been hindered by increases to the thresholds in the NH, caused by the removal of dust earlier in the year, that had not yet been reversed by dust replenishment. The T63 simulation showed enhanced late-summer dust lifting, presumably as fronts were better resolved along the north polar cap edge. An improvement was also seen in test runs using the newer model physics (v5), including the two-moment dust transport scheme and updated scattering properties: several flushing storms were produced around $L_s = 330°$. These cases featured slightly lower dust opacities in late summer than those seen above — as the two-moment scheme resulted in faster mass sedimentation rates than the size bin technique, allowing major storms to decay more completely by $\sim 330°$ — which led to stronger waves in the NH. The development of late-summer regional storms may therefore be a quite subtle effect, but it is one that is close to being captured by the standard resolution model at present.

The most serious deficiency in the simulation presented above was the rarity of dust storms that originated in the southern hemisphere, independent of a NH flushing storm, as SH storm initiation has been observed in multiple instances in reality. As noted already,

poor resolution of topographic flows was probably one reason for this, since several of those storms observed began in the vicinity of the Hellas or Argyre basins. SH storm initiation was seen occasionally in the model, in Noachis or Daedalia, but these storms never grew particularly large. To investigate the importance of the threshold level in allowing storm generation in the southern hemisphere, another multi-year run was performed using a lower minimum threshold of 0.015 Pa. This did cause the occurrence of three SH storms in nine model years (each beginning at Daedalia Planum), but it also led to the unrealistic development of storms at L_s = 60–120° during several of the years simulated. These formed due to persistent wind stresses in the northern Tharsis region, south of Alba Patera, which allowed lifting to take place when the local thresholds returned (mainly via ${\zeta^t}_{dec}$) to values near their minimum. Previously, with the higher minimum threshold, this lifting was prohibited. The lower threshold run also produced a large storm at 210° (usually beginning as a NH flushing storm) with increased regularity, meaning that interannual variability was reduced. The choice of threshold (or in this case, minimum threshold) clearly has a significant effect on dust lifting magnitude and variability, but the use of a lower minimum alone does not appear to be a viable approach by which to prompt the model to generate large storms spontaneously in the southern hemisphere.

4.7 Heterogeneous surface roughness

4.7.1 Background

For many years, the only measure of the aerodynamic roughness length, z_0, of the Martian surface came from the Viking landers, for which a value in the range 0.1–1 cm was estimated (Sutton et al., 1978); as a result, values within this range have typically been assumed globally in MGCMs thus far. Other surface landers have made similar estimates of z_0, including 3 cm at the Mars Pathfinder landing site (Sullivan et al., 2000), and a range of values z_0 = 0.07–0.60 cm along the path of the Spirit Rover (Greeley et al., 2008), indicating significant variation in roughness length over the planet's surface.

Recently, a collaboration between workers at LMD, the Laboratoire ATmosphères, Milieux, Observations Spatiales (LATMOS), and the Laboratoire Interuniversitaire des Systèmes Atmosphériques (LISA) (all in Paris) has yielded a global z_0 map covering the

full Martian surface, at a resolution ($\frac{1}{8}^\circ \times \frac{1}{8}^\circ$) well in excess of typical MGCM resolutions. The dataset was derived from rock abundance maps obtained from TES thermal inertia data (Nowicki and Christensen, 2007), as explained in Hèbrard et al. (2012).

The completed z_0 map is displayed in Figure 4.14. At full resolution, values of z_0 range from ~ 0.001 cm to 2 cm. The southern hemisphere, particularly in the key lifting band at 30–60°S, is characterised by large roughness lengths, > 0.5 cm, with a large number of pixels around 1 cm. By contrast, the northern hemisphere is largely smooth, with values of $z_0 \lesssim 0.1$ cm over much of the low-TI plains at 15°S–30°N. The flushing storm channels of Acidalia-Chryse and Utopia-Isidis, however, exhibit fairly large roughness lengths. The polar regions, when smoothed to a lower (model) resolution, feature mainly moderate roughness values of $0.1 < z_0 < 0.5$ cm, along with some pixels with high values of thermal inertia, assumed to be ice-covered, at which z_0 was set to 0.01 cm. This setting was also used poleward of 87°, where no data were available. The number of pixels with z_0 around 1 cm means that this value was a good choice for use in the UKMGCM until now, but the z_0 field is certainly not homogeneous, and exhibits a wide range of values (around three orders of magnitude) over the planet's surface. The map is consistent with the original Viking estimate of 0.1–1 cm, as it includes values of 0.3 and 0.9 cm at the two Viking landing sites. There are greater discrepancies between map and in situ z_0 at the Pathfinder site and at some of the Spirit locations, but these may be due to microscale surface variations below the resolution of the map ($\sim$ km).

The distribution of roughness on the Martian surface has important implications for dust lifting, as described in the following subsection.

4.7.2 Lifting theory with variable surface roughness

The roughness length of the surface is of importance in the MGCM as it is needed to define the drag coefficient used to calculate surface turbulent fluxes, and to determine how much drag is exerted on the atmosphere by the surface, and vice versa. With regard to the near-surface wind stress dust lifting scheme, it is necessary to know how much of the atmospheric drag is exerted on the underlying erodible surface itself, and how much of it is 'lost' on the roughness elements. The raising of dust is therefore easiest from a smooth surface, where the wind flow is unimpeded by nonerodible elements. For rough surfaces,

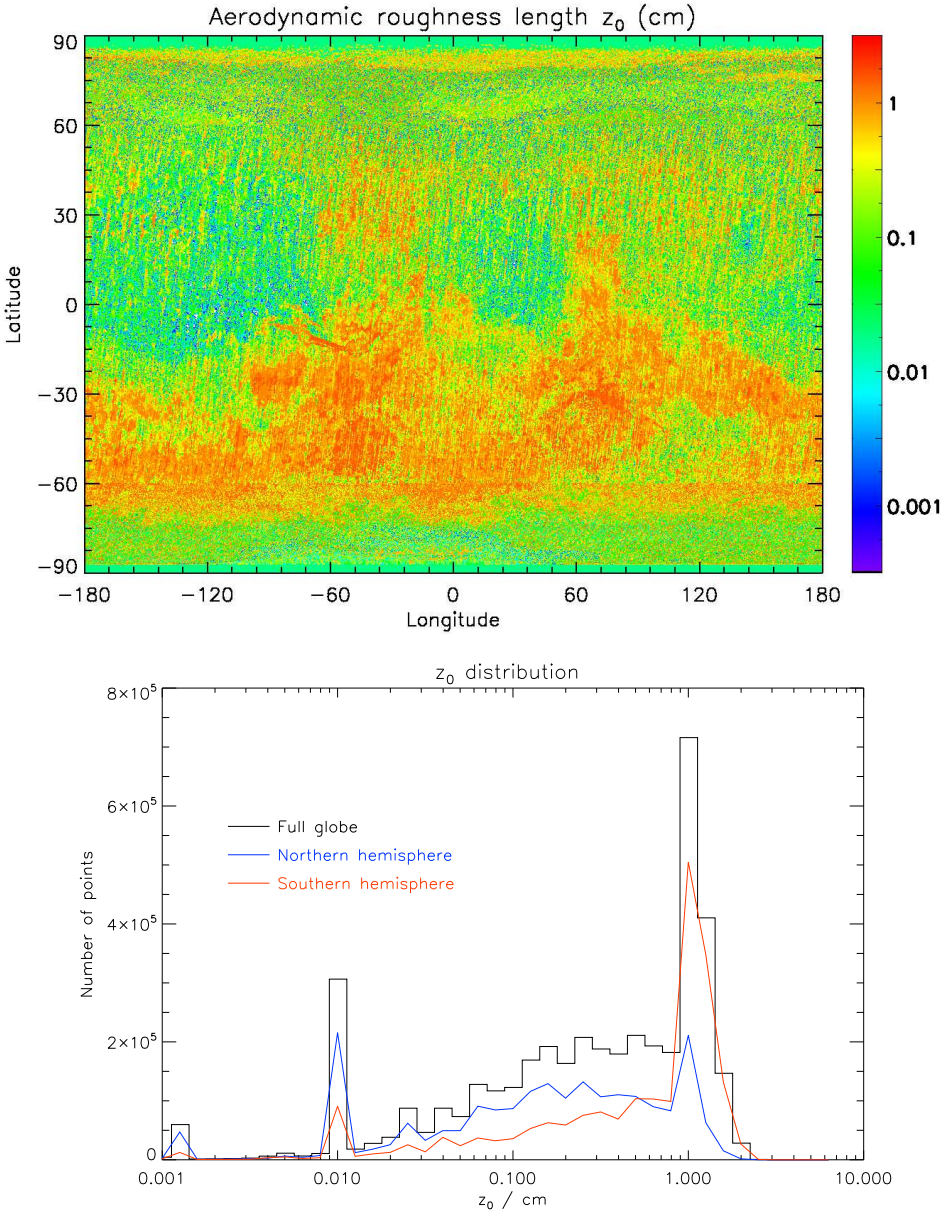

Figure 4.14: Martian surface roughness length (in cm) at $\frac{1}{8}^\circ$ resolution (top), and the distribution of pixel values, binned in equal widths of log-z_0 (bottom).

a larger drag velocity is required to achieve the same saltation flux as that which would occur if the surface were smooth (free from nonerodible elements); thus, the threshold drag velocity u_*^t increases as surface roughness increases. This is expressed through the drag partition function, f_{eff}, defined as the square root of the fraction of the total surface stress that is imparted to the erodible surface, and which is also equal to the ratio of the threshold drag velocity for a smooth surface to the threshold drag velocity for the rough surface:

$$f_{eff} = \left(\frac{\zeta_s}{\zeta}\right)^{\frac{1}{2}} = \frac{u_{*s}^t}{u_*^t}$$

The threshold drag velocity is therefore given by

$$u_*^t = \frac{u_{*s}^t}{f_{eff}} = {f_{eff}}^{-1}\sqrt{A_N\left(\sigma_p g D_p + \frac{\gamma}{\rho D_p}\right)} \tag{4.1}$$

where now it is the threshold for a smooth surface, u_{*s}^t, that is given by the formula of Shao and Lu (2000), equation (3.2). f_{eff} is always less than one, so the threshold velocity for a rough surface is increased from its calculated value for a smooth surface. Before now, f_{eff} had been implicitly set equal to one; therefore the lifting threshold at all gridpoints was (quite significantly) lower than it should have been. Of course, since thresholds were scaled to allow dust lifting to occur in the model (§3.1.1), this was not of absolute importance; however, the inclusion of a heterogeneous z_0 field implies that thresholds should not be treated uniformly. While a uniform scaling (accounting for sub-gridscale gustiness, etc.) of thresholds may still be used as a first guess, the initial threshold values, calculated from equation (4.1), vary greatly over the planet.

Details on the calculation of f_{eff} as a function of z_0 are provided in Appendix B. Its use in Equation 4.1 results in substantial increases to the calculated lifting thresholds for medium-to-large roughness lengths. At the same time, the drag velocity u_* itself depends on z_0: rough surfaces feel a stronger drag force from a given surface wind than smooth surfaces do. With regard to where on the planet wind stress dust lifting will be easiest, these two effects are in opposition, as shown in Figure 4.15(a). Compared to the previous scheme (which assumed $z_0 = 1\,\mathrm{cm}$ everywhere), for a given model surface wind $u(5\mathrm{m})$, gridpoints with $z_0 < 1\,\mathrm{cm}$ will feel a weaker drag velocity, while rougher areas will feel a stronger u_*. However, the relative increase to the threshold drag velocity ($\propto {f_{eff}}^{-1}$) is always greater

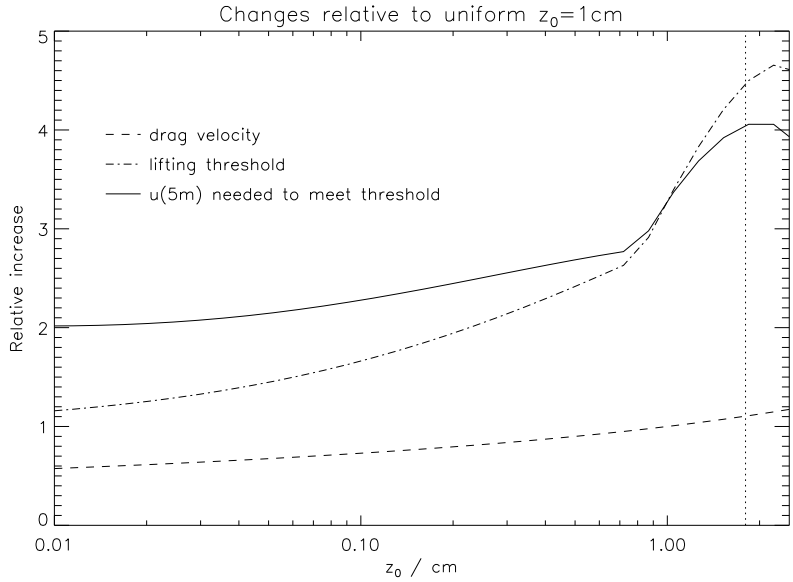

(a)

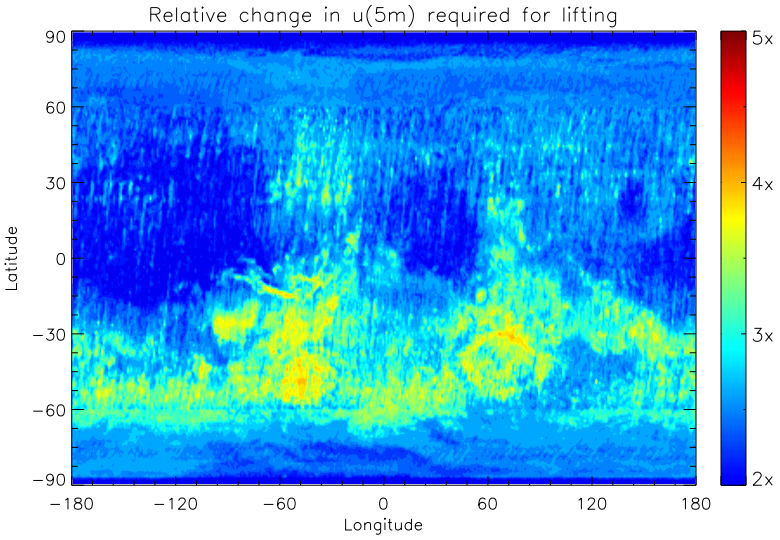

(b)

Figure 4.15: (a) Relative increases resulting from the use of a heterogeneous z_0 field instead of a uniform value of 1 cm when calculating the drag velocity (dashed), and from the inclusion of z_0 in the calculation of the threshold drag velocity (dash-dotted), and the relative increase in windspeed in the lowest model level required to meet the new lifting threshold. The dotted vertical line marks $z_0 = 1.8$ cm, the largest value present in the roughness field at 1° resolution. (b) The relative increase in windspeed in the lowest model level required to meet the lifting threshold, as a function of longitude and latitude. The minimum increase needed is a doubling of u(5m).

than one, since this effect was neglected previously. The increase in threshold is by far the dominant effect for high z_0, meaning that lifting at these gridpoints requires the largest increase in model wind relative to the old model setup. All gridpoints now require *at least* a doubling in model wind to reach the threshold for dust lifting (a further consideration, neglected from the curve in Figure 4.15(a), is the change that may be caused to the model wind u(5m) itself, as described below in §4.8.1).

The areographic distribution of this increase in required surface wind is shown in Figure 4.15(b), where it can be seen that, in general, lifting should now become much more difficult in the southern hemisphere, while large parts of the northern hemisphere experience comparatively little change. It may be noted that this pattern is quite similar to the dust cover field shown in Figure 4.3(b), implying, perhaps surprisingly, that dust lifting should be most difficult in the areas (predominantly at 30–60°S) that are low on surface dust. Therefore, it must be true that Martian surface winds do become strong enough, on occasion, to remove dust from these high-z_0 regions; conversely, in the smooth northern hemisphere plains, a high dust cover has persisted despite lifting being favoured by the low-z_0 conditions. This suggests that the distribution of surface dust (as a result of the lifting that occurs) is determined primarily by surface wind strengths rather than the roughness of the surface.

Figure 4.15(b) also shows similarities to the typical surface stress field resulting from runs using the variable threshold scheme as described earlier (e.g. Figure 4.3(a)). Therefore, it would seem that the increased difficulty caused by z_0 in lifting dust from 30–60°S has already been partially captured by the model's variable threshold scheme. However, the correlation is not perfect, and the ability of various regions to provide dust for lifting must depend on a combination of the roughness of the surface and the amount of surface dust present.

A further effect of increased surface roughness was applied following Laurent et al. (2008), who used an erodible fraction, E, which measures the areal fraction of each surface gridpoint that can contribute to dust emission (with the rest of the total area being occupied by nonerodible roughness elements). Lifted dust fluxes for both mechanisms are multiplied

by this fraction, defined empirically by Laurent et al. (2008) as

$$E = 0.7304 - 0.0804 \log(z_0) \text{ if } z_0 \geq 0.003\,\text{cm}, \tag{4.2}$$

and $E = 1$ otherwise. Its effect is to further inhibit dust lifting from rough areas, though the smallest value of E (for the largest z_0) is ~ 0.7, so the reduction to the emitted flux is not dramatic.

4.8 Implementation of the roughness length map

4.8.1 Climatic changes

The roughness length map of Figure 4.14 was inserted into the UKMGCM, smoothed to the model resolution in the same way as done for the other surface datasets, replacing the global z_0 value of 1 cm used previously. During simulations, a 'smooth' setting of $z_0 = 0.01$ cm was applied at any gridpoint with a thick (> 10 cm) covering of CO_2 ice. Dust lifting is prohibited at these gridpoints already, so the choice of roughness length here affected only the drag exerted on the atmosphere by the surface. The newer model version, v5, was used from this point onwards; however, in contrast to the runs of Chapter 6, the size-bin approach to dust transport was retained for the remainder of the work in this chapter (i.e. the two-moment tracer scheme of v5 was not utilised). A year-long run was performed using prescribed dust (from MY24) to assess the impact of the change in surface roughness on the climate in general, aside from its effect on dust lifting. Since the original z_0 value of 1 cm is near the upper end of the range of the map, most gridpoints obtained smoother surfaces than were previously assumed, with the northern hemisphere plains seeing the most dramatic change (an order of magnitude decrease in z_0, at this resolution).

Figure 4.16 shows the change in daytime and nighttime temperatures in the lowest model layer (a height of ~ 5 m) and in the sixth model layer (~ 300 m), averaged annually, when using the z_0 map compared to when using the uniform value of 1 cm. As expected, the most noticeable changes were found where z_0 is lowest. The decrease in near-surface temperatures during the day, and increase at night, occurred due to a reduction in upward mechanical mixing: turbulent diffusion, in the model parameterisation, is less efficient over

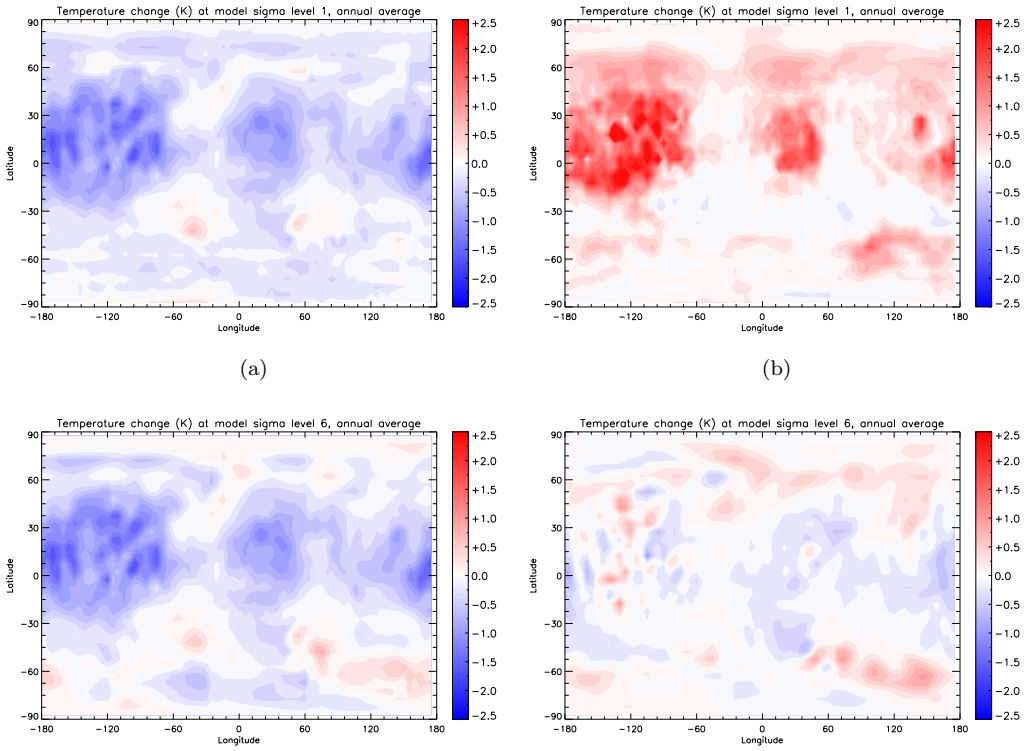

Figure 4.16: Change in temperature (K), annually averaged, caused by use of the roughness map, at the first model sigma level (~ 5 m above the ground) (a) during the day and (b) at night, and at the sixth model sigma level (~ 300 m above ground) (c) during the day and (d) at night.

smooth surfaces, so heat and momentum (as well as any tracers present in the boundary layer) tend to remain trapped near the surface. In daytime, when the surface is warmer than the atmosphere, a lower z_0 means that the near-surface atmosphere receives less heat through turbulent mixing, and is therefore cooler, by a few K over the smoothest gridpoints. Conversely, at nighttime, the colder surface means that a reduction to upward mixing warms the near-surface atmosphere, again by several K in places. This warming was seen only in the lowest few model layers, since the nighttime boundary layer is extremely shallow; during the day, the cooling applied over several km.

The overall effect of a lower surface roughness was, then, a reduced near-surface diurnal temperature range, as shown in the scatter plot of Figure 4.17. While the smoothest gridpoints saw a reduction in range of around 3 K, for an order of magnitude decrease in z_0, sensitivity studies showed that similar results were obtained if surface thermal inertia was increased by around 10–20%. Therefore, z_0 has rightly been assumed to be of secondary

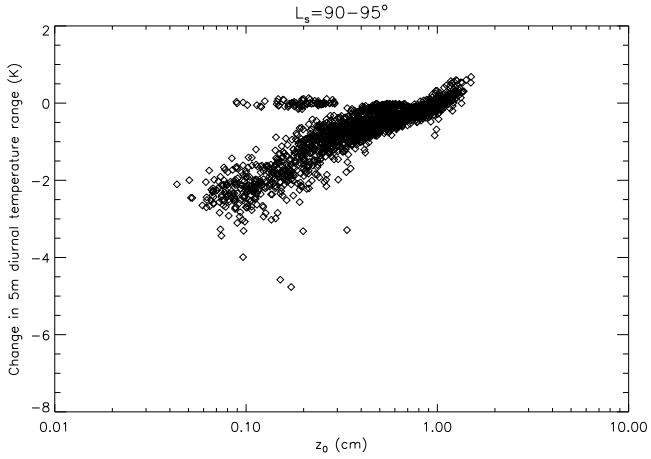

Figure 4.17: The change to the diurnal temperature range (K) at the lowest model level when moving from a uniform roughness length of 1 cm, around northern summer solstice.

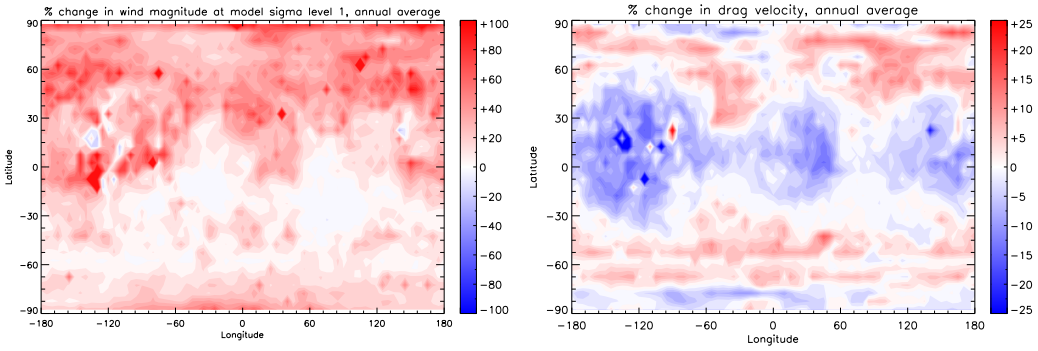

Figure 4.18: Annually averaged change (%) in (a) wind magnitude at the lowest model level (i.e. u(5m)) and (b) drag velocity (note the different scales). Gridpoints containing CO_2 ice have been excluded from the average; at such points, surface wind was strongly increased due to the low value of z_0 (0.01 cm) applied.

importance in simulating the behaviour of the Martian boundary layer, in comparison to other surface fields. The diurnally averaged effect of the temperature changes was a ~ 1 K warming near the ground where z_0 was lowest.

Vertical mixing of momentum was also reduced over smooth surfaces, leading to an increase in wind magnitudes in the lowest model level (Figure 4.18(a)). Averaged over the year, excluding any ice-covered gridpoints, near-surface winds increased by as much as 40–50% over the low-z_0 northern plains. Near-surface winds were strongly enhanced over the polar caps, where the surface is assumed to be very smooth (not shown in Figure 4.18(a)). The surface drag velocity, u_*, which is used to calculate the vertical flux of dust due to wind stress lifting, is obtained from the lowest-level wind through equation (2.4), so a low

z_0 implies that a lower fraction of the near-surface momentum goes into exerting a drag force on particles on the surface. Figure 4.18(b) shows that in the northern plains this counteracted the increase in near-surface wind to produce a lower drag velocity, though reductions were generally less than 10%. Therefore, the dust lifting ability of rough areas (such as the band at 30–60°S) was enhanced, in the sense that high-z_0 gridpoints now felt stronger drag velocities than they did previously. However, as shown above (§4.7.2), the increase to the lifting threshold is the more important roughness-induced change, and dust lifting on the whole becomes much more difficult at high-z_0 gridpoints than at low-z_0 gridpoints.

As a consequence of the weakened vertical mixing, daytime boundary layer height decreased by up to 1 km in the low-z_0 regions, and turbulent kinetic energy decreased throughout the boundary layer. Little change was seen in atmospheric static stability, suggesting that convective adjustment, the other mechanism for vertical mixing in the model, is not strongly affected by surface roughness.

4.8.2 Effect on dust devil lifting; comparison of lifting rates to observations

The decrease in the daytime temperature gradient between the surface and the lowest model layer over smooth surfaces, combined with the slight lowering of the boundary layer top, reduced the efficiency of the heat engine used in the model's dust devil parameterisation, lowering global lifting rates (Figure 4.19). The parameterisation was most strongly affected over areas of low z_0; however, the negative feedback associated with model dust devil lifting meant that resulting optical depths were not reduced by more than $\sim 30\%$. The addition of the erodible fraction condition (4.2) reduced rates by a further 15%, fairly uniformly over the course of the year.

Observations of dust devils agree with the model prediction that they are most frequent in the southern hemisphere, during local spring and summer — perhaps an order of magnitude more so than in the northern hemisphere (Whelley and Greeley, 2006). Dust devil counting is usually done using the darkened tracks they leave behind as they remove surface dust as a proxy, but not all dust devils leave tracks — it is thought that surface dust cover needs to be low for the dust devil's path to be seen (Fisher et al., 2005). Despite this

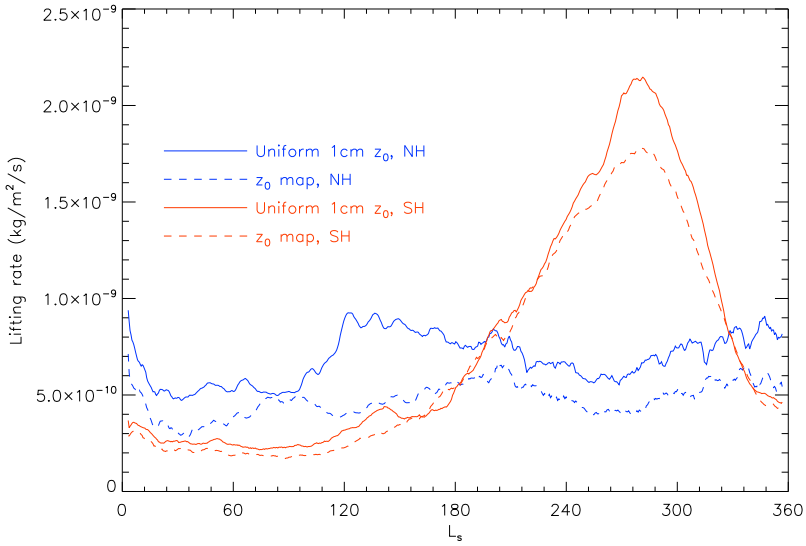

Figure 4.19: Hemispheric averages of dust devil lifting rate for runs with uniform (solid) and heterogeneous (dashed) surface roughness, using only dust devil lifting. In the latter, the effect of the erodible fraction, E, was not included, so the differences between the solid and dashed lines result only from changes to the surface-atmosphere temperature gradient and boundary layer height forced by the use of the z_0 map.

possible bias towards the southern hemisphere, where dust cover is lower, the latitudinal variation in dust devil activity is believed to exist, driven by the eccentricity of Mars' orbit and the asymmetry in the strength of the solsticial overturning cells (Whelley and Greeley, 2006; Stanzel et al., 2008). In each hemisphere, the dust devil season begins in mid-spring and ends in early autumn, and dust devils are most commonly found at $\sim$ 40–70° in latitude (Whelley and Greeley, 2008). The model reproduces the preference for dust devil lifting in local spring and summer, and the lifting peak in the southern hemisphere is larger (by around a factor of two) than that in the northern hemisphere. The potentially problematic aspect of model dust devil lifting is that northern hemisphere lifting remains significant during local autumn and winter, in contrast to some observations, particularly those of Cantor et al. (2006). On the other hand, Stanzel et al. (2008) found 'frequent' NH dust devils during this period. Also, in comparison to most observations, lifting in the model appears to peak too close to the equator throughout the year.

Neakrase and Greeley (2010) suggested, from laboratory simulations of dust devils, that tangential wind velocity and lifted sediment flux may decrease considerably with increasing roughness, and that the threshold tangential velocity may increase. However, their study

involved only four different roughness densities, and did not produce a functional dependence for any of these quantities on z_0. Therefore, this effect could not be incorporated into the dust devil lifting parameterisation at present (in fact, the only explicit dependence of dust devil upward flux on surface roughness comes through the use of the erodible fraction). If this relationship does hold true on Mars, dust devil lifting rates would be expected to be most strongly reduced in the southern hemisphere, so the fact that dust devils are in fact most commonly observed in the SH is perhaps surprising. As the best observational estimate of the lifting ratio between the southern and northern hemispheres is that it is probably $\lesssim 10$ (Whelley and Greeley, 2006), it is not possible to use the model to quantitatively gauge the strength of the influence that z_0 (and/or surface dust cover) must exert on dust devil lifting rates.

Nevertheless, dust devil lifting in the second half of the year may be overestimated by the model. Although global dust opacity does increase during this period in reality, it does not appear to do so in the repeatable, solstice-centric manner produced by the dust devil parameterisation (seen e.g. in year 28 in Figure 4.1); rather, increases in dust optical depth appear to be associated with (wind-stress-initiated) regional dust storms, and show some preference for occurrence either side of solstice (see e.g. Figure 2.1). Thus far, the dust devil lifting scheme has been used to provide a realistic background dust loading during northern spring and summer, when little or no wind stress lifting is produced by the model. However, it has not been firmly established how much of the dust in the Martian atmosphere is lifted by dust devils, and as a result of this, the dust devil lifting efficiency α_D had to be tuned to reproduce the resulting global dust opacity, rather than the dust lifting rates themselves. The problem with this method is that the dust opacity observed during northern spring/summer is not solely the result of dust devil lifting — Cantor et al. (2006) observed that a significant number of dust storms occur during these seasons, primarily along the southern polar cap edge at $L_s = 0$–60° and $L_s = 120$–180°. It has been seen already that the model (at least at standard resolution) currently underestimates the amount of cap-edge dust lifting, particularly in the southern hemisphere, which means that dust opacity during $L_s = 0$–180° is very low unless dust devils are included.

It is possible to make a tentative comparison of the model's global annual dust devil lifting rate with observational estimates that have been made. These vary significantly, due to

the difficulty in reliably estimating dust devil frequency, in measuring the mass lifted by an individual dust devil and in extending such a figure globally to cover the range of sizes of dust devils that can form. Estimates at present are spread over an order of magnitude: Whelley and Greeley (2008) calculated a global lifting rate of $0.0016\,\mathrm{kg\,m^{-2}\,yr^{-1}}$, Cantor et al. (2006) quoted a minimum bound of $0.004\,\mathrm{kg\,m^{-2}\,yr^{-1}}$, and Greeley et al. (2006) extrapolated local observations from the Spirit Rover to produce a global figure of $0.013\,\mathrm{kg\,m^{-2}\,yr^{-1}}$. The total dust devil flux in the years examined in §4.4 was $\sim 0.02\,\mathrm{kg\,m^{-2}\,yr^{-1}}$, which is larger than these observational estimates. This is perhaps not surprising, given that small-scale wind stress lifting events are underestimated by the model. Another way of judging the accuracy of model dust devil lifting is by comparing the contribution of each of the two types of lifting to the annual global dust load. Whelley and Greeley (2008) combined their value for dust devil flux with the estimate made by Cantor et al. (2001) of the annual lifting rate due to local and regional dust storms ($0.0035\,\mathrm{kg\,m^{-2}\,yr^{-1}}$) to conclude that dust devil lifting contributes approximately half as much dust as wind stress lifting, or about a third of the total annual flux. The average wind stress lifting rate from the last year of the series in §4.4, which featured a pre-solstice regional dust storm, was $0.005\,\mathrm{kg\,m^{-2}\,yr^{-1}}$, meaning that dust devils provided 80% of the total lifted dust in this year. For years featuring global dust storms, the dust devil percentage was as low as $\sim 50\%$ — this is still higher than the Whelley and Greeley (2008) estimate, but it should be noted that their estimate of annual dust devil flux was low in comparison to other observations, and contained large error bars. Kahre et al. (2006) similarly found that dust devils must provide around half of the total lifted flux in order to best reproduce the observed trend in opacity in the Ames model.

4.9 Effect on wind stress lifting

It was immediately apparent that the significant threshold increases caused at high-z_0 gridpoints meant that lifting would only be possible with a much bigger reduction to the threshold of equation (3.2). Indeed, with a factor of 0.425 applied to the calculated u_*^t, only one gridpoint produced any lifting over a year. With the scaling factor lowered to 0.2, lifting became more widespread, but (in a run with fixed threshold stresses) was dominated by several gridpoints north of Tharsis following equinox at $L_s = 180°$. This region was particularly favoured for lifting due to the combination of strong wind stresses, aided by

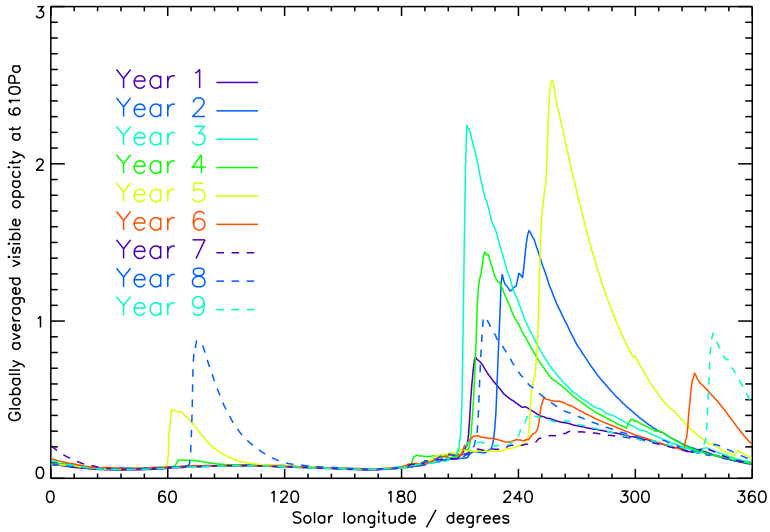

Figure 4.20: Globally averaged dust optical depth from a multiannual simulation using heterogeneous surface roughness and the variable threshold scheme, with model v5 (but still using the size bin approach to dust transport). The minimum used was the calculated threshold velocity multiplied by 0.2, and lifting parameters were $\alpha_N = 2.2 \times 10^{-4}\,\mathrm{m}^{-1}$, $\alpha_D = 1.3 \times 10^{-8}\,\mathrm{kg\,J}^{-1}$ (using the threshold-dependent dust devil scheme), ${\zeta^t}_{inc} = 2\,\mathrm{Pa}\,(\mathrm{kg\,m}^{-2})^{-1}$, and ${\zeta^t}_{dec} = 1 \times 10^{-10}\,\mathrm{Pa\,s}^{-1}$. Year numbering began when the model was properly spun-up with respect to the threshold stress field.

the topographic slopes present, and low thresholds calculated for the relatively smooth surface.

To avoid this unrealistic concentration of dust lifting at one northern hemisphere site, the variable threshold scheme developed earlier in this chapter was employed. In this formulation, a minimum threshold was used as before by choosing a base value of the scaling factor for u_*^t, though now the minimum threshold stress varied spatially through the additional influence of the z_0 field. A scaling factor of 0.2 was used. The resupply term acted on the smooth lifting threshold, u_{*s}^t, rather than actual threshold u_*^t. As before, a multiannual run was performed, with initial years used to spin up the surface threshold field. Results from the later years of the run are shown in Figure 4.20. It can be seen that similar results to those of Figure 4.1 were obtained over the period L_s = 180–300°, with a range of storm sizes and initiation times produced over the years shown. A notable difference was in the formation of regional storms in late summer (L_s = 320–340°) in two of the nine years, due to one of or both a change to the spatial distribution of lifting caused by the use of z_0, and the slightly lower opacities produced at this time by the model v5

than seen in Figure 4.1. Both of these years also featured a pre-solstice regional event, and were thus reminiscent of the no-storm TES years of MY24 or MY26.

The other notable feature of Figure 4.20 is the development of dust storms in late northern spring in three of the years in the sequence, as was found earlier when using a lower minimum threshold. These again formed in northern Tharsis whenever thresholds had settled back down towards the minimum value. This highlights the problem with using a single threshold scaling factor, which must be set very low to allow lifting to occur in the southern hemisphere. The relation between calculated and effective threshold may actually vary significantly from place to place, depending on various factors, as is discussed later in this section.

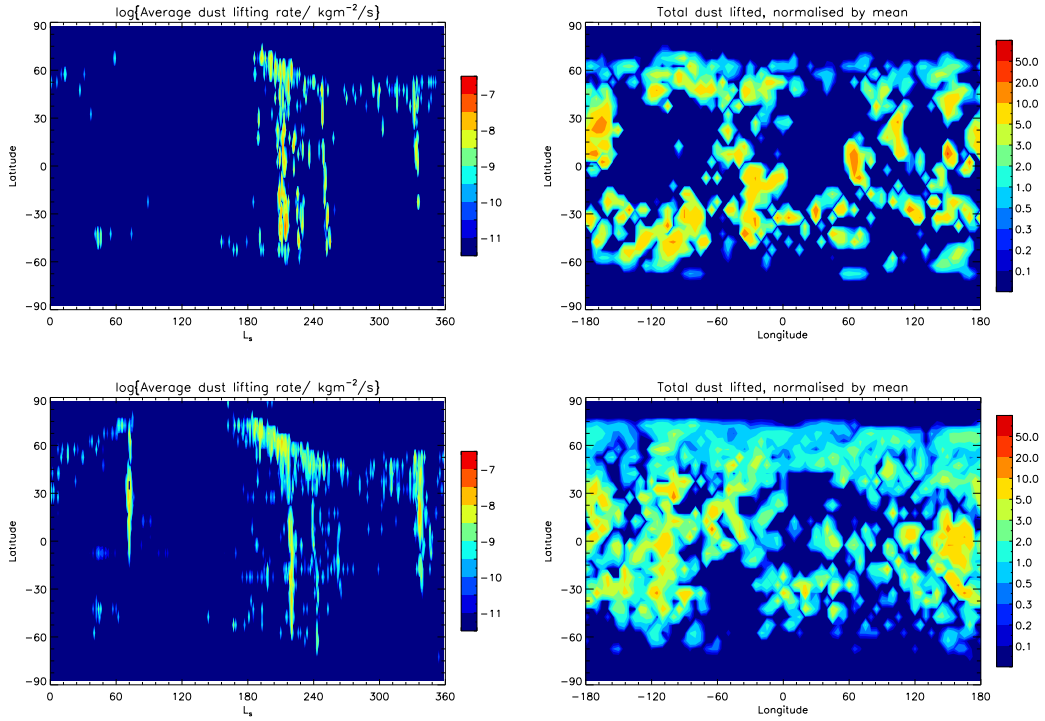

Figure 4.21: Results, averaged over three model years, from simulations using model v5, without (top row) and with (bottom row) the roughness length map. Appropriate scalings were applied to the lifting threshold in each case. Plots in the left column show the logarithm of the zonally averaged dust lifting rate, in $\mathrm{kg\,m^{-2}\,s^{-1}}$, and plots in the right column show the total dust mass lifted at each gridpoint, normalised by the global mean.

The z_0 field acts in combination with the wind regime to produce a spatial distribution of preferred lifting areas that is different to that seen previously. Figure 4.21 compares dust lifting rates between two three-year sequences, both using the variable threshold scheme

and the model v5: the first neglected z_0 when calculating the lifting threshold, and reduced the resulting thresholds by a factor of 0.7, while the second used the z_0 map as described above, with a scaling factor of 0.2 applied to the calculated thresholds. While the model's interannual variability means individual years cannot be compared directly, a three-year series is long enough to illustrate the general changes to lifting patterns that result from using the z_0 map. In the left-hand plots it can be seen that dust storms were still generally produced during L_s = 210–250°, but there was more lifting at 180–210°, both in the southern hemisphere and at the NH polar cap edge, when the roughness map was used. Late-summer lifting appeared to be stronger when using the roughness map (as was suggested above), and this lifting was better sustained along the receding edge of the NH cap in spring. The right-hand plots show that lifting was more evenly distributed across the globe in the z_0 run, with particular increases at Amazonis/Tharsis and at northern high latitudes, due to the smooth surface found in those regions. Conversely, lifting was weaker across the rough 30°S band.

4.9.1 Sub-gridscale gustiness revisited

Such a large reduction to the calculated lifting threshold suggested that the approach used so far to handle sub-gridscale effects (the use of a single scaling factor) is not suitable when roughness lengths are included in the threshold drag velocity calculation. Gustiness is a likely candidate for playing the dominant role in allowing lifting to take place when model-resolved winds would suggest that no lifting should be possible over most of the planet. Therefore, a fresh attempt at parameterising gustiness was needed, if the z_0 map was to be used in dust lifting simulations.

As noted in §3.1.1, Newman (2001) used a Weibull probability distribution — since shown by Fenton and Michaels (2010) to provide a good fit to microscale wind variation in large eddy simulation (LES) model runs — to model the sub-gridscale variation in surface windspeed, but found that this significantly weakened the threshold dependence of the wind stress lifting scheme and therefore the interannual variability in global dust opacity. With a Weibull distribution, the fraction of the gridbox for which drag velocity is above the

threshold u_*^t can be expressed as

$$\exp\left[-\left(\frac{u_*^t}{u_*}\right)^\kappa\right],$$

where u_* is the drag velocity value for the gridbox and κ is the shape parameter. Therefore, a larger κ leads to reduced lifting when the wind velocity is below the threshold, by focusing the distribution more firmly on the mean wind and reducing the contribution of the high-u_* tail. Newman used $\kappa = 1.5$, taken from Viking Lander data (Lorenz, 1996); however, this data recorded variations on an hourly scale, whereas the model, with a timestep of 30 min, requires parameterised wind variation on a timescale of minutes, for which a smaller spread of values would be expected. A higher value for the shape parameter, $\kappa = 2.5$–3, is more appropriate (Fenton and Michaels, 2010). As an example of the effect of this on lifted fluxes, with $u_* = 0.5\,\mathrm{m\,s^{-1}}$ and $u_*^t = 2.5\,\mathrm{m\,s^{-1}}$, integrating for the area under the probability distribution curve results in an equivalent fraction of 10^{-5} of the gridbox producing dust emission when $\kappa = 1.5$, but only 10^{-25} when $\kappa = 2.5$ is used (though the difference between the two choices of κ is greatly reduced when the mean wind is closer to the threshold value). The use of a larger shape parameter therefore returns a threshold dependence to the lifting parameterisation, by weakening the lifting contribution of gridpoints with low mean wind velocities.

Even with parameterised gustiness, thresholds at the roughest gridpoints remained out of reach of model winds, so a reduction to the calculated threshold drag velocities was still required. The factor used was now larger: ~ 0.4–0.6 allowed a reasonable fraction of surface gridpoints to reach the lifting threshold within a year, when using $\kappa = 2.5$. The use of the Weibull distribution allowed storm-magnitude dust lifting to take place at higher values of the scaling factor, but the value of 0.7 used previously at T31 remained too large whenever z_0 was included in the threshold calculation. A larger κ value of 3 required a factor smaller than 0.4, as it represented lower amounts of gustiness.

The parameterised gustiness was spatially uniform, so the preferred lifting sites seen above remained dominant. In reality, there are likely to be some regions that typically experience more gustiness than others, implying that a variable κ value should be used when computing dust fluxes. In fact, one of the principal causes of variability in gustiness could be the surface roughness, since it was observed earlier (§4.8.1) that turbulent KE decreased

over low-z_0 gridpoints. Cakmur et al. (2004) parameterised gustiness in a terrestrial atmospheric GCM by modelling components u and v with normal distributions, whose standard deviation was calculated online as a sum of contributions from three sources: PBL turbulence, and dry and moist convection. They found that dry convection tended to dominate in desert regions with strong surface heating, a situation that may be comparable to the Martian environment. An outcome of modelling this gustiness was that higher thresholds could be used in matching simulated and observed dust emissions, which is the aim in the present case. It should be possible to estimate wind variability from the turbulent diffusion scheme in the UKMGCM (as a function of TKE), and including the contribution from (dry) convection may be possible in the near future, using a new flux-limited convection scheme that is currently under development at LMD.

4.9.2 Sub-gridscale z_0 variation

The results above offered some explanation, through wind velocity variations on small spatial and temporal scales, for the low thresholds required to generate dust storms in the model, but a significant reduction to the calculated threshold was still needed: at least 40% to the threshold drag velocity, equivalent to a $> 60\%$ reduction to the threshold stress. Another sub-gridscale effect which could, if considered, enhance lifting in the model concerns the z_0 map itself, which is smoothed to model resolution using an arithmetic mean method that will tend to bias the gridpoint z_0 value towards the higher pixel values contained with the gridpoint. Where dust lifting is concerned, any low-z_0 pixels — perhaps as much as an order of magnitude below the mean — within the gridpoint area would be preferred locations in which for dust lifting to occur, due to their lower thresholds. This effect is part of the explanation given by Gillette (1999) to explain the existence of dust emission 'hot spots' on Earth.

Okin (2005) investigated the effect of sub-gridscale surface heterogeneity on dust emission parameterisations by using a stochastic approach to calculating the threshold drag velocity. He used a normal distribution to sample parameters in the threshold velocity expression controlling the soil and roughness characteristics (equivalent to u^t_{*s} in (4.1) and β, σ in (B.2)), and found that emission was dominated by small areas with randomly low lifting thresholds. Allowing for a spread in threshold values meant that more gridpoints

were able to participate in emission, leading to a greater overall lifted flux. Okin also found that including the variance of the parameters was roughly equivalent to lowering the single-value threshold by a certain amount, dependent on the size of the spread in parameter values assumed.

This result undoubtedly applies equally to dust lifting in the UKMGCM, since the same threshold formulation and a similar flux expression are used. A further test of the influence of sub-gridscale surface variations on estimated dust emission was carried out, focusing on the variance in z_0 itself, which directly affects both the threshold drag velocity and the conversion from model wind velocity to drag velocity. The error in lifting flux estimates at T31 resolution was calculated explicitly by applying a range of 5m-height windspeed values to the full-resolution z_0 map and to the 5° (the resolution of the physical grid at T31) regridded version, and in each case calculating the saltation flux in the same way as is done in the model. A near-surface atmospheric density of $0.02\,\mathrm{kg\,m^{-3}}$ was assumed. It was found that the error was a function of windspeed, as total flux was underestimated at T31 at relatively low $u(5\mathrm{m})$ values, but became an overestimate once the wind was increased beyond $\sim 18\,\mathrm{m\,s^{-1}}$, due to the contribution of low-z_0 gridpoints (Figure 4.22). The behaviour at high u arose because once the windspeed at a T31 gridpoint exceeds the threshold, the whole gridpoint takes part in dust lifting, while at full resolution there remain portions of the area in which the threshold is not exceeded. The low-u situation is the opposite, as T31 windspeed remains below the gridpoint threshold while saltation has begun in parts of the gridpoint area at full resolution. (It was found that the two curves are in closer agreement if the low-resolution map is derived using a geometric mean, rather than an arithmetic mean, as was used for the calculations in this and the following figure, because the geometric mean reduces the weighting of any large-z_0 values within a T31 gridpoint. For fields such as surface roughness, where variation over more than an order of magnitude within a gridpoint can typically occur (at least in this particular map), the geometric mean is probably more appropriate than the arithmetic mean; however, the problems caused by smoothing the field to a lower resolution remain qualitatively the same.)

Saltation fluxes at T31 resolution were then recalculated using different threshold values in order to find the threshold needed to match the T31 flux to the flux as calculated from the full-resolution map; that is, the effective threshold drag velocity produced by considering

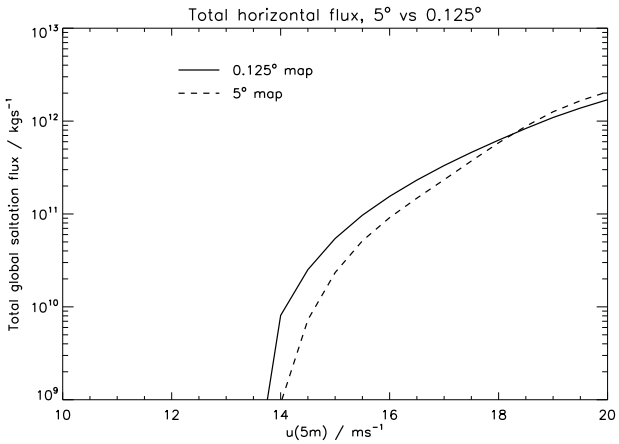

Figure 4.22: Global saltation flux as a function of the 5 m windspeed used to calculate drag velocities and threshold drag velocities, for the 5° map (dashed), calculated using an arithmetic mean, and the full-resolution, 0.125° map (solid).

sub-gridscale variations in z_0. Fluxes were first matched on a global basis, for which it was found that reductions of only a few percent were required to increase the T31 flux enough to match the full-resolution value (when $u(5\text{m})$ was below $\sim 18\,\text{m}\,\text{s}^{-1}$; at greater windspeeds, an increase to the threshold of up to 10% was needed). However, dust lifting occurs far from uniformly across the planet, so a more relevant calculation is to match the fluxes at each T31 gridpoint, by regridding the full-resolution fluxes accordingly for comparison. This was done for a range of $u(5\text{m})$ values, and a sample of the results are shown in Figure 4.23. As expected, larger threshold reductions were required in the high-z_0, high-$\sigma(z_0)$ areas — as much as 40–50% at times. The low-z_0 northern plains generally required increases to the calculated T31 thresholds, as they produced too great a saltating flux once the windspeed became moderately large.

It is clear that using the gridpoint-mean value for z_0 to compute the saltation flux skews dust lifting rates towards low-z_0 areas, which are predominantly found in the northern hemisphere. Significant reductions to the calculated lifting thresholds can be justified on the basis of sub-gridscale z_0 variation, but may not be applied uniformly across the surface; in fact, in some areas the threshold should generally be increased, to counteract the overestimation of lifted flux. However, the threshold corrections required are also a function of windspeed, due to the non-linear, threshold-dependent nature of the saltation flux parameterisation.

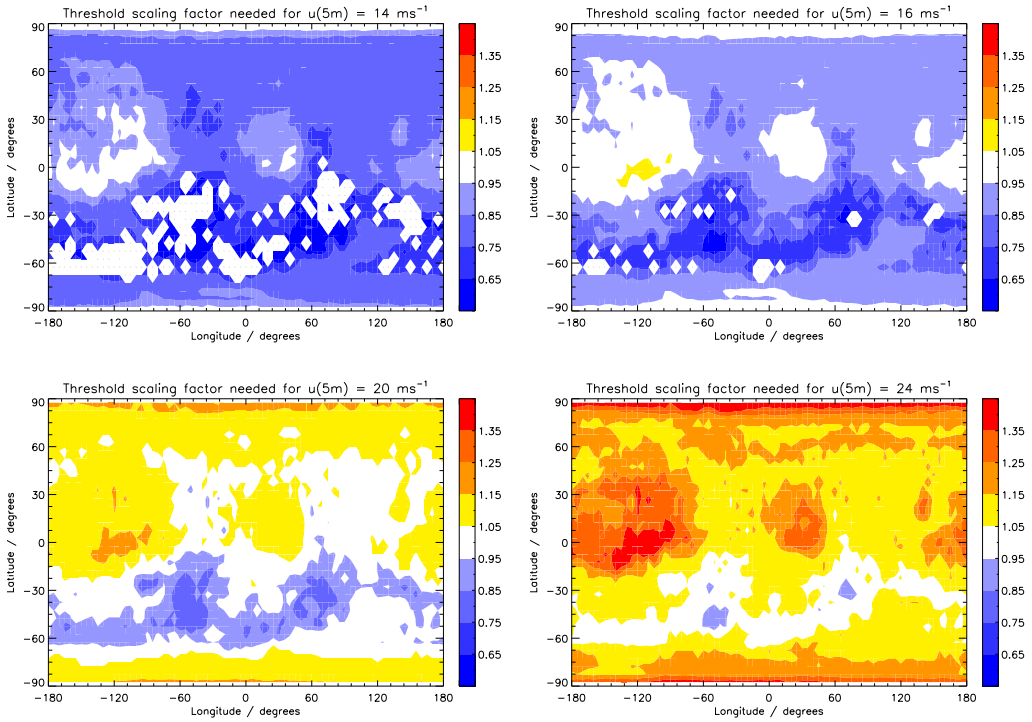

Figure 4.23: The scaling factor required for the threshold drag velocity at each T31 gridpoint to replicate the true (0.125°) saltation flux, for 5 m windspeeds of 14, 16, 20 and 24 m s^{-1} (which equate to drag velocities of $\sim 0.9 - 1.3\,\mathrm{m\,s^{-1}}$). Gridpoints marked with a dot are those at which no saltation occurred, even at full resolution.

4.9.3 Conclusions on dust lifting using realistic z_0

It is clear from the work in this second half of this chapter that the proper inclusion of surface roughness length in the calculation of dust lifting thresholds has a huge effect on modelled dust emission rates. Threshold stresses in the southern hemisphere in particular are much larger than values previously used in dust lifting MGCMs, meaning that, without proper handling of sub-gridscale effects, lifting becomes prohibitively difficult in southern hemispheric areas known to be important sources of dust. At first, this necessitated the use of significantly reduced thresholds: u_*^t had to be lowered by more than 70% (which equates to a 90% reduction in threshold stresses) to generate a storm season peak in dust opacity. However, it was estimated that wind gustiness could account for a 10–30% threshold velocity reduction, and that mesoscale (and potentially microscale) heterogeneity in z_0 could be responsible for a further 20–40% reduction, if it is assumed that in large dust storms most dust is lifted in the southern hemisphere when winds are not far above threshold. Therefore,

the calculated thresholds *could* be justifiably lowered in the UKMGCM so as to allow the generation of realistic dust opacities during southern spring and summer. However, such scaling factors would need to vary substantially in both space and time, and should ideally be estimated online depending on the ambient conditions at each gridpoint. This would be a complicated parameterisation, and as such is beyond the scope of this work, but the findings of this chapter suggest that a correct treatment of the z_0 field would bring the spatial distribution of dust lifting back towards the form it took when the influence of z_0 was ignored.

Due to uncertainty in the use of the roughness length z_0 in the wind stress lifting parameterisation, this aspect of the dust lifting scheme was not carried forward into Chapter 6, in which radiatively active water ice clouds are added to the model. The variable threshold scheme was retained, as it was found that the presence of ice clouds quite strongly affected dust lifting preferences, leading to a dominance of certain regions/periods that could not be avoided with an unlimited surface dust approach. First, though, water ice clouds were added to the model while using prescribed dust opacities instead of lifting and transport, in the context of a dynamical analysis of winter polar and midlatitude baroclinic activity, known to be vital in the formation of flushing dust storms.

Chapter 5

The Martian solsticial pause

Outside of regional and global dust storms, there is increasing evidence that the majority of local storms, a result of wind stress lifting, occur along the edges of the polar caps that are present in each hemisphere from autumn to spring (Cantor, 2007). The presence of the sharp meridional temperature gradient at the ground makes the region baroclinically unstable, and travelling waves, with zonal wavenumber 1–4, are commonly seen, extending several scale heights above the surface (Banfield et al., 2004). However, the amplitude of the eddies varies significantly in time. Near the surface (the lowest scale height or so), transient eddy activity is strongest in early autumn and in late winter, and exhibits a minimum at winter solstice, a phenomenon known as the 'solsticial pause', first noticed by Barnes (1980) in Viking surface pressure data. Conversely, transient temperature perturbations at higher levels have been observed to reach a maximum around solstice (Wang et al., 2005).

A similar phenomenon is seen on Earth (though not in all years), in the form of a suppression of the northern Pacific storm track around midwinter, at a time when linear baroclinic theory predicts a maximum in storminess (Nakamura, 1992). Several important differences exist between the Earth and Mars cases, however: while transient perturbations exhibit minima around the full latitude circle on Mars, the midwinter suppression on Earth is limited to the Pacific Ocean (no such effect occurs over the Atlantic); also, the weakening of eddy activity occurs only near the Martian surface, but does so throughout the troposphere on Earth.

On Mars, this suppression exerts, through a reduction in surface windspeed variance and

magnitude (Wang et al., 2005), a strong control over dust lifting along the seasonal polar cap edge. The seasonal progression of frontal dust storm frequency in northern midlatitudes has been seen to display a double-peaked structure, with a minimum in dust lifting activity observed around solstice (Cantor, 2007). It was because of the important connection to dust lifting that this phenomenon was investigated in more detail. It had been observed during the work of the previous chapter that there was no pronounced weakening of travelling waves around northern winter solstice seen in dust-lifting UKMGCM simulations, and the same could be said of simulations using prescribed dust. However, a significant improvement was seen upon the addition of a new radiatively active water ice cloud parameterisation, so by examining the contribution of the ice clouds to the midlatitude and polar climate, it was possible to investigate the mechanism(s) behind the formation of the solsticial pause. The suppression of waves around winter solstice had previously only been analysed for the idealised case of a dust loading appropriate for a GDS, either prescribed (Hourdin et al., 1995; Kuroda et al., 2007) or in a transport simulation (Basu et al., 2006), though the phenomenon is clearly not restricted to such instances (Wang et al., 2005).

5.1 Observational evidence and characteristics

Direct TES observations of winter eddy activity were presented by Wang et al. (2005), but it was the availability of the 3-year TES reanalysis dataset, obtained by assimilating temperature and dust opacity observations into the UKMGCM (Lewis et al., 2007) that allowed a detailed examination of the nature of the phenomenon, to understand what aspects the free-running model could and could not reproduce. Root-mean-square (RMS) variance of temperature at 2.5 km above the surface is presented, for the full dataset, in Figure 5.1. The data here and in many subsequent plots have been bandpass filtered to remove the diurnal tide and any long-period, quasi-stationary waves, retaining only signals with a period within the range 1.5–10 sols. For all plots of RMS, an averaging window of length 20 sols was used. It can be seen that transient eddies were fairly weak (1–3 K) in the midlatitudes of each hemisphere at its winter solstice ($L_s = 90°$ for the southern hemisphere and $L_s = 270°$ for the northern hemisphere), but showed peaks in activity either side of solstice, in late autumn and in early spring (L_s = (0–60°, 120–180°) for the southern hemisphere, L_s = (180–240°, 300–360°) for the northern hemisphere). Transient wave

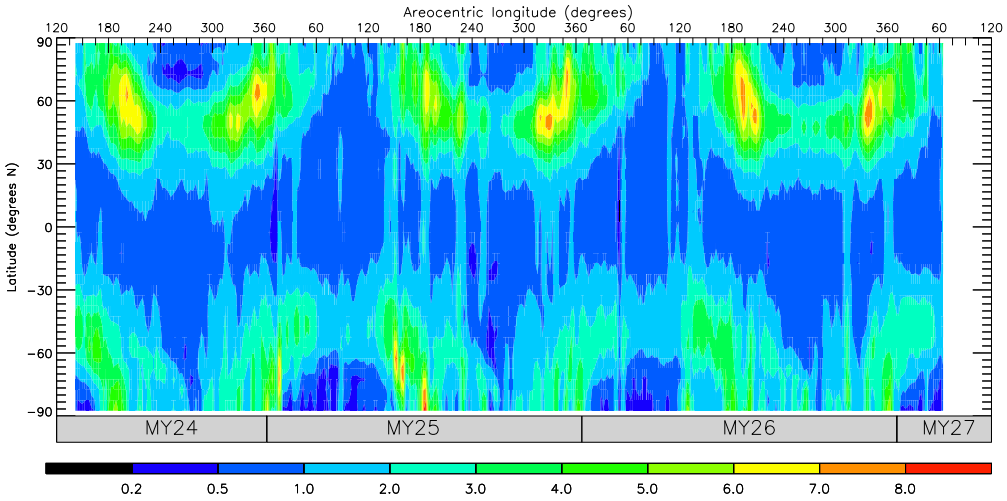

Figure 5.1: Zonal-mean of root-mean-square transient temperature (in K) at 2.5 km above the surface, from the TES reanalysis dataset.

activity was generally stronger in the northern hemisphere than the southern hemisphere; peak RMS values in the northern hemisphere were 7–8 K in each of the three years covered by the assimilation.

None of the three TES years featured a large dust storm around winter solstice — MY25 did feature a global storm, but it had almost completely dissipated by $L_s = 270°$, while the other two years saw only regional dust storms, none occurring around winter solstice. None of these years would have experienced the very strong enhancement to the solsticial meridional circulation that would be caused by a Viking-like midwinter storm as used in previous modelling studies, though there would certainly have been some strengthening of the circulation due to the increase in background dust loading that occurs globally during northern winter every year. Despite this, the northern hemisphere solsticial pause was seen clearly, and covered at least 70° of L_s, in each of the years.

Figure 5.2 illustrates the vertically limited nature of the solsticial pause by showing the variation of eddy heat flux with latitude and pressure, for the case of MY24, again from the assimilation. Positive values of $\overline{v'T'}$ imply poleward (northward) eddy heat transport, indicating that baroclinic instability was present in the northern hemisphere throughout the autumn/winter seasons. While the solsticial pause can be easily identified in the $\overline{v'T'}$ minimum below ~ 300 Pa, above 20 Pa, a maximum developed around solstice. This feature is more easily understood, and occurred due to the large meridional temperature gradients

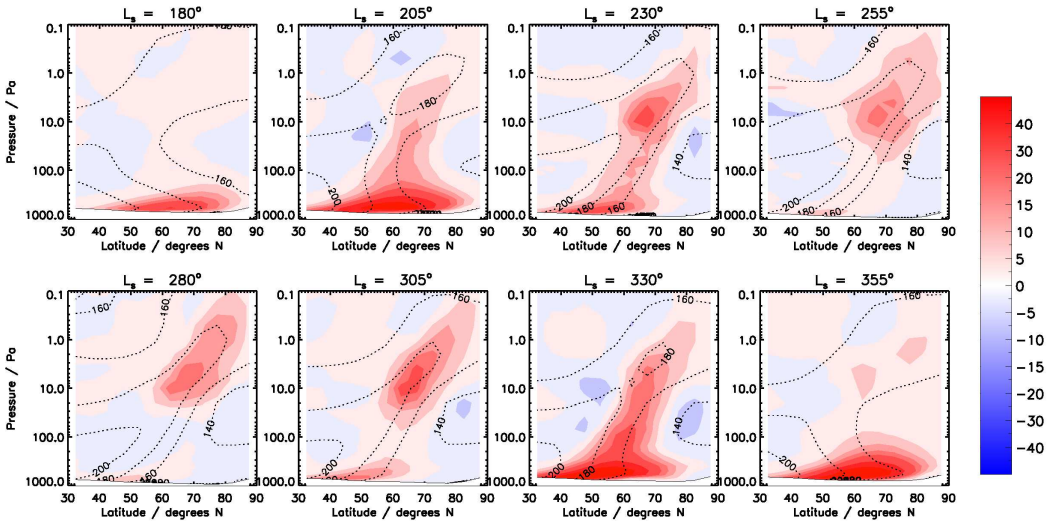

Figure 5.2: Zonal-mean eddy heat flux $\overline{v'T'}$ (colours, in K m s^{-1}) and temperature (dotted contours, in K) from the reanalysis of MY24.

that existed in this part of the atmosphere around solstice, as a consequence of the middle atmosphere polar warming (Deming et al., 1986; Kuroda et al., 2007). The fact that the Martian solsticial pause is so clearly limited to ~ 300 Pa and below suggests that arguments pertaining to the structure of the jet in the middle atmosphere that have been invoked to explain the terrestrial suppression (e.g. Zhang and Held, 1999) may not be particularly applicable to the Martian case. It is also worth noting that the solsticial pause developed despite the presence of strong meridional temperature gradients near the ground, at 40–60°N.

The basic state of the atmosphere in the MY24 assimilated record is shown further in Figure 5.3, during the pre-solstice (surface) eddy maximum, the solsticial minimum and the post-solstice maximum. An extension of high temperatures into the polar middle atmosphere was seen throughout this period, but was at its strongest at solstice. The westerly jet exhibited a tilt, upwards towards the pole, during this period. Around solstice, the core of the jet at 10 Pa was strengthened and shifted $\sim 10°$ to the north, resulting in a greater tilt (away from the vertical) in the lower part of the jet. By late winter, the body of the jet had moved back towards the equator, its core had weakened and it had regained a homogeneous tilt across its full height.

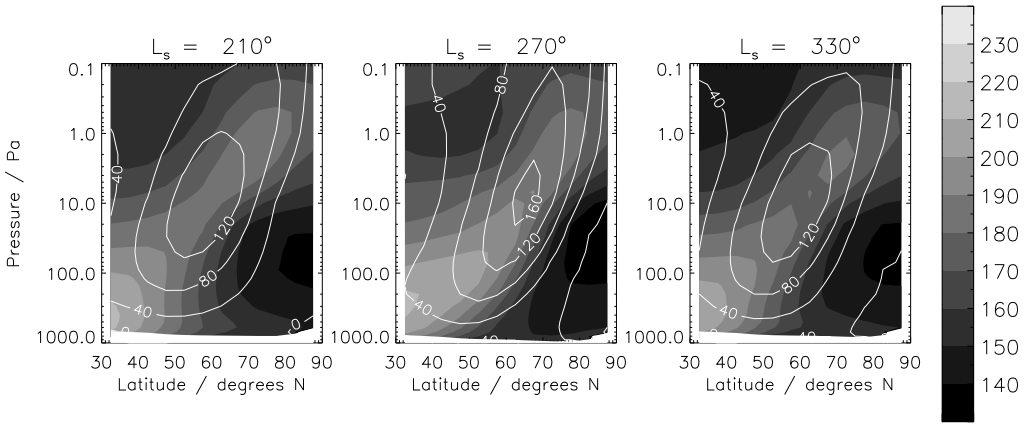

Figure 5.3: Zonal-mean zonal wind (white contours, $\mathrm{m\,s^{-1}}$) and temperature (shading, K) from the reanalysis of MY24.

5.2 Model simulations

To test whether or not the model reproduces the observations presented in the previous section, and to investigate the conditions necessary for the formation of a solsticial pause, a series of simulations were performed using a variety of model formulations, all with v5. Prescribed dust opacities were used in this chapter (unless stated) in lieu of a transported dust field, to enable more direct comparisons with observations and to fix the dust opacity field so as to study the sensitivity of the solsticial pause to several other effects. The simulations are named and summarised in Table 5.1. Two used the 'MY24' dust scenario described earlier, one including the radiative effects of water ice clouds (τ^*_{MY24}) and the other neglecting clouds (τ_{MY24}). A similar pair of runs with (τ^*_{low}) and without (τ_{low}) clouds used instead a constant visible dust opacity of 0.2, referenced to the 610 Pa pressure surface. Finally, τ_{high} simulated a large, perennial dust storm by using a constant 610 Pa dust opacity of 2 (and neglected water ice clouds). All model runs were carried out at T31 resolution, as was used for the data assimilation. The cloud scheme added to the model at this point was a radiatively active version of the one presented in Montmessin et al. (2004), which predicts ice particle sizes and growth rates, taking into account temperature, humidity and local density of dust nuclei. It is described in more detail in the next chapter, where its typical output is shown, but it is sufficient for the moment to say that the scheme

Run name	Dust opacity	Clouds?	$\frac{(T_{RMS})_{min}}{(T_{RMS})_{max}}$
τ_{MY24}	MY24	No	0.65
τ^*_{MY24}	MY24	Yes	0.35
τ_{low}	0.2	No	0.84
τ^*_{low}	0.2	Yes	0.74
τ_{high}	2.0	No	0.43
MY24 assimilation			0.31

Table 5.1: Details of the settings used in each of the simulations performed.

was able to represent the key features of the observed water cycle to an acceptable degree of accuracy.

Table 5.1 includes a quantity $\frac{(T_{RMS})_{min}}{(T_{RMS})_{max}}$ which measures, in a crude way, the extent to which a solsticial pause was produced in the northern hemisphere for a particular run. This quantity was calculated by finding the maximum value of the RMS temperature at a height of 2.5 km within the region 30–90°N at each point in time, and taking the ratio of the minimum of this quantity during $L_s = 240$–$300°$ and its maximum during $L_s = (180$–$240°$, 300–360°). Low values therefore signify a deep minimum in transient eddy activity. In summary, four of the five runs, the exception being τ_{low}, produced a discernible solsticial minimum in near-surface eddy activity. τ^*_{MY24} came closest to replicating the results of the assimilation, while τ_{MY24} produced a weaker minimum, suggesting that both the presence of water ice clouds and an increase in global dust loading close to solstice can contribute to the development of a solsticial pause. A large dust storm (τ_{high}), with a greater dust loading than the peak MY24 value, delivered a deeper solsticial minimum than τ_{MY24}; however, the minimum was shorter in duration than τ^*_{MY24} and the assimilation of MY24.

5.2.1 Transient eddy activity in the GCM

The top row of Figure 5.4 shows the zonal mean of RMS transient temperature at 2.5 km above the ground in the northern hemisphere from runs τ_{low} and τ_{MY24}, performed without including the radiative impact of water ice clouds. The constant opacity case showed no sign of a solsticial pause in the northern hemisphere, with a peak RMS of 7–8 K throughout

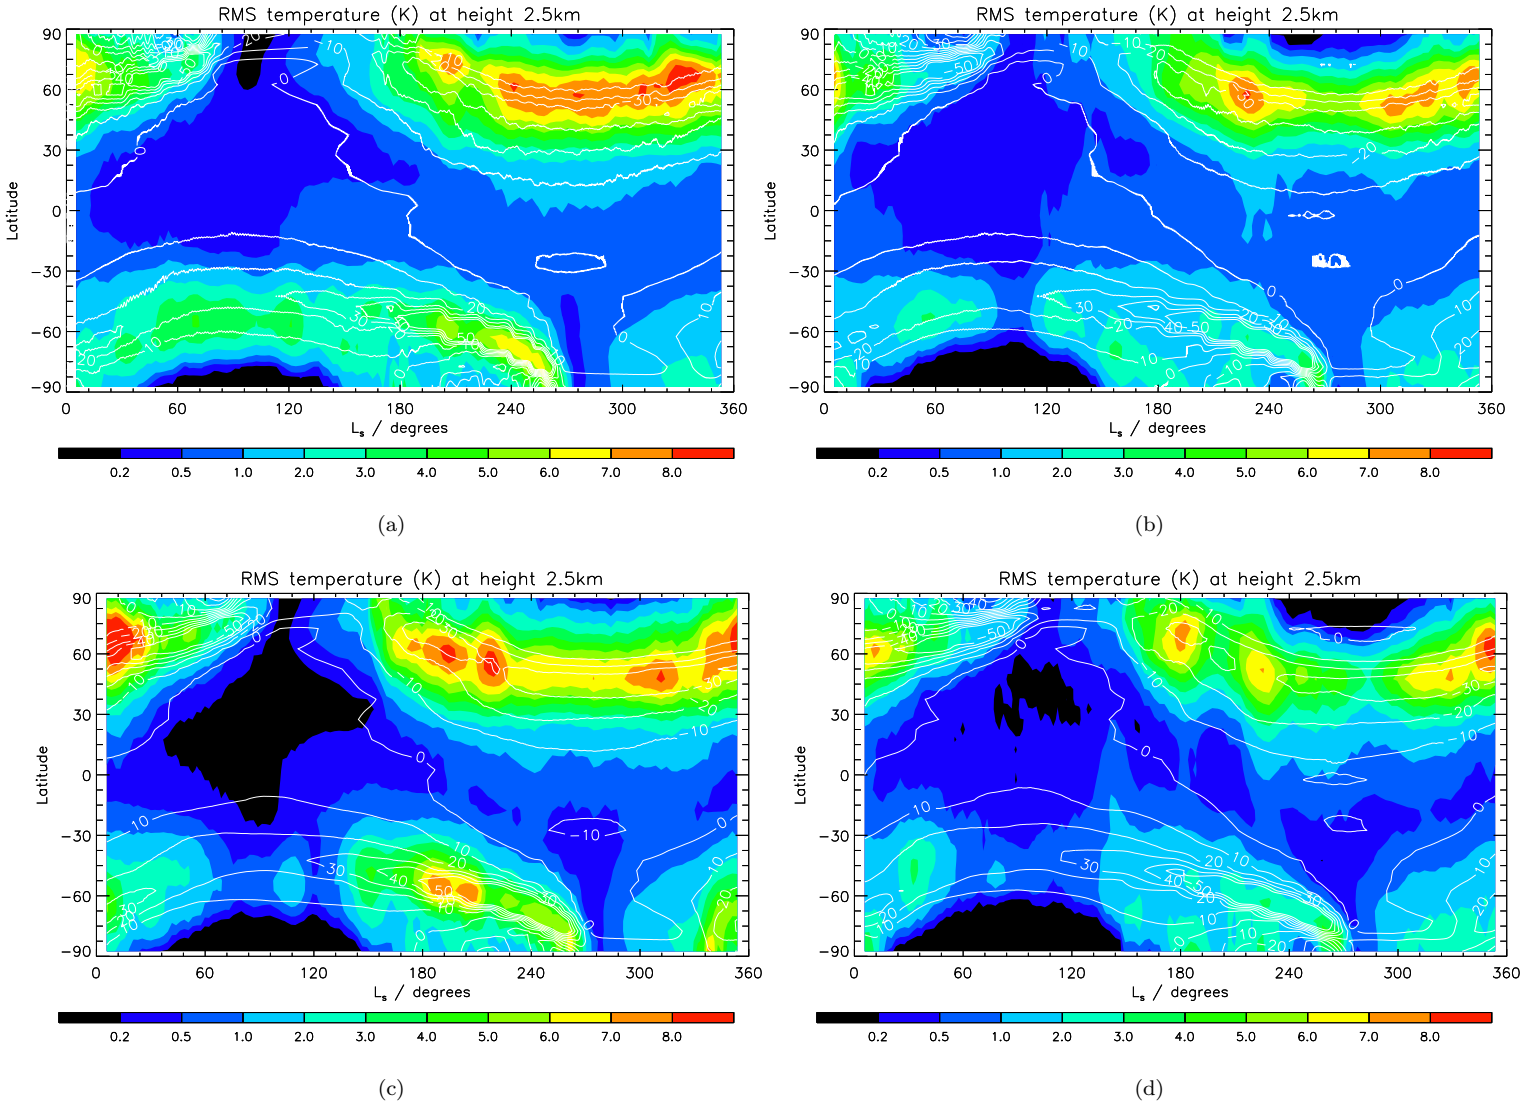

Figure 5.4: Zonal-mean RMS transient temperature (K) at a height of 2.5 km above the surface, output from the GCM from runs (a) τ_{low} (b) τ_{MY24} (c) τ^*_{low} (d) τ^*_{MY24}. Meridional temperature gradient $\frac{dT}{dy}$ at 200 m above the ground (10^{-6} K m^{-1}) is overlaid in white contours.

the period $L_s = 180$–$360°$. With the MY24 dust scenario, a minimum in RMS developed between $L_s = 260°$ and $L_s = 290°$, but it was much smaller in amplitude and shorter in duration than those seen in the assimilation. The overlaid meridional temperature gradients show that in both cases the peak in eddy activity lay to the north of the latitude where the temperature gradient at 200 m height was strongest, suggesting a strong meridional tilt in the location of this peak gradient over the lowest few kilometres. Evidently the MY24 dust scenario, with elevated opacities over most of the planet from $L_s \sim 210°$ onwards, goes some way towards forcing a solsticial pause, perhaps for a similar reason to that identified by Kuroda et al. (2007) for a large solsticial dust storm; however, this realistic dust loading fell short of reproducing the results from the assimilation for that year. The southern hemisphere showed similar results, with fairly constant eddy activity for τ_{low} and a short (though otherwise quite realistic) minimum for τ_{MY24}.

The bottom row of the same figure shows the results of a similar pair of runs to those above, but with radiatively active clouds now switched on (τ^*_{low} and τ^*_{MY24}). Focusing again on the northern hemisphere, with MY24 dust opacity the solsticial pause was now wider and deeper than before, and resembled the observed pause in the assimilated record. Post-solstice eddy activity reached its peak value later in winter, and autumn eddy activity began much earlier than it did in τ_{MY24}, with an additional peak forming at $L_s = 180°$. The minimum was, however, of smaller magnitude and slightly shorter duration than the observed solsticial pause, even with radiatively active ice clouds — RMS temperature did not drop off sharply enough before $L_s = 240°$. The southern hemisphere pause was also enhanced, and now appeared wider than those in the assimilation. τ^*_{low} also simulated solsticial pauses in both hemispheres. The northern hemisphere pause was shallower than that of τ^*_{MY24} and was centred around $L_s \sim 280°$. In both simulations the peak RMS was generally located at the latitude of strongest temperature gradient, meaning that in early northern autumn and late winter the baroclinic zone sat 10–15° further south than it did in the runs without water ice clouds.

For τ^*_{MY24}, zonal wavenumbers 1–3 were all reduced in strength at solstice, relative to autumn and spring (Figure 5.5). $s = 3$ eddies did not exhibit the sharp peaks in pre- and post-solstice activity seen in the TES data by Wang et al. (2005); instead the wavenumber 1–3 signals were rather coherent at 30–60°N, while northward of this, wavenumbers 1 and

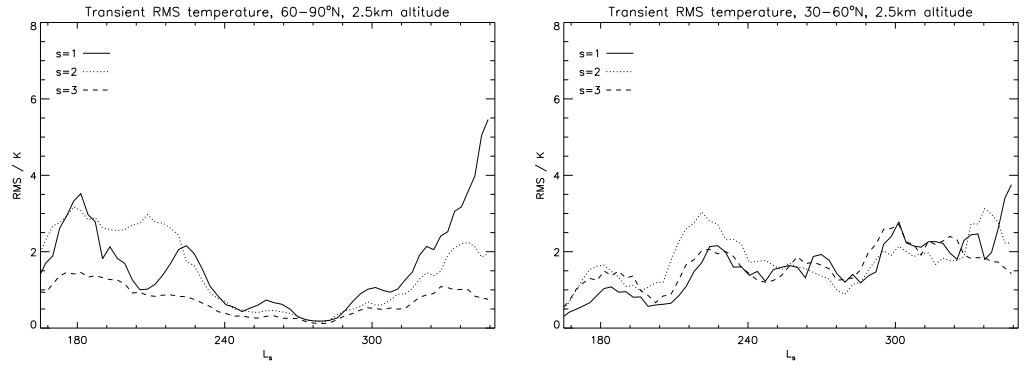

Figure 5.5: Zonal-mean transient RMS temperature (K) at 2.5 km from τ^*_{MY24}, averaged over 60–90°N (left) and 30–60°N (right) (compare with Figure 5.4(d)).

2 were dominant. The confinement of $s = 3$ perturbations equatorward of 60°N and the greater strength of $s = 1$ poleward of 50–60°N are in agreement with Wang et al. (2005).

5.2.2 Dust transport simulations

It should be noted that the dust opacities used to construct the MY24 scenario, while agreeing very well with direct observations over much of the year, are extremely uncertain at mid- and high latitudes in both winter hemispheres, as TES opacity data in these latitudes were not obtained for these periods. As a result of the approach taken in the data assimilation procedure, dust opacities remain constant over a period without available observations. The realism in the seasonal trend in eddy activity of τ^*_{MY24} compared to the assimilation provides some confidence that plausible polar dust opacities are being used; moreover, the above results show that a solsticial minimum formed, at least for a limited period, regardless of dust scenario used, provided that radiatively active water ice clouds were included in the simulation.

To ascertain that the polar dust distribution is not of key importance, a pair of dust lifting and transport simulations were performed, one with active water ice clouds and one without. These runs did not use the variable threshold scheme described in the previous chapter, but featured broadly realistic dust and ice opacities (see next chapter). In the dust lifting run without clouds, even with a regional storm occurring at a similar time to that of MY24, no strong minimum in northern hemisphere transient eddy activity was seen. With clouds included, a much more realistic solsticial pause, similar to that of τ_{MY24}, was

produced. This result strengthens the argument that clouds, in the UKMGCM at least, make an important contribution to the jet structure during northern autumn and winter, and one which encourages the suppression of surface baroclinic instability around solstice.

5.3 Changes to atmospheric baroclinicity

Having found that the inclusion or removal of water ice clouds strongly affects the ability of the GCM to simulate the solsticial pause, the model runs were analysed more closely to determine the manner in which the clouds exert their influence. The most obvious mechanism for this (with regard to the solsticial pause) is that the clouds alter the structure of the atmosphere so as to reduce baroclinicity at low levels around solstice.

5.3.1 Basic atmospheric state

Figure 5.6(a) shows the zonal-mean state of τ^*_{MY24}, the model run which produced the most realistic NH solsticial pause. The change to the zonal-mean zonal wind caused by the addition of water ice clouds is shown in Figure 5.6(b). Polar hood clouds were confined below 100 Pa, and were generally thickest near the seasonal cap edge (which, for reference, was located at approximately (60°, 50°, 50°)N at $L_s = (220°, 270°, 320°)$). Cloud optical depth in this model run reached a maximum around the time of a regional dust storm at $L_s = 240°$, and declined thereafter. It is also notable that cloud optical depth dropped to a minimum at $L_s = 200°$ (not shown), coincident with the gap between the two pre-solstice peaks seen for τ^*_{MY24} in Figure 5.4(d).

The addition of clouds slowed the zonal wind by 5–10 $\mathrm{m\,s^{-1}}$ near the ground at 60–80°N around solstice, and below 100 Pa the jet was shifted equatorward in such a way as to enhance the poleward tilt of this lower section. The change to the thermal structure near the ground around solstice encouraged this, as temperature contours at 40–70°N were roughly aligned with the vertical before and after solstice but sloped strongly around solstice as the surface at 40–70°N cooled. Since cloud formation is determined primarily by atmospheric temperature (Benson et al., 2011), the clouds adopted approximately the same sloped structure. A cooling of up to 5–10 K occurred in the vicinity of and below the clouds, the effect of which was to reduce the meridional temperature gradient on the clouds' poleward side,

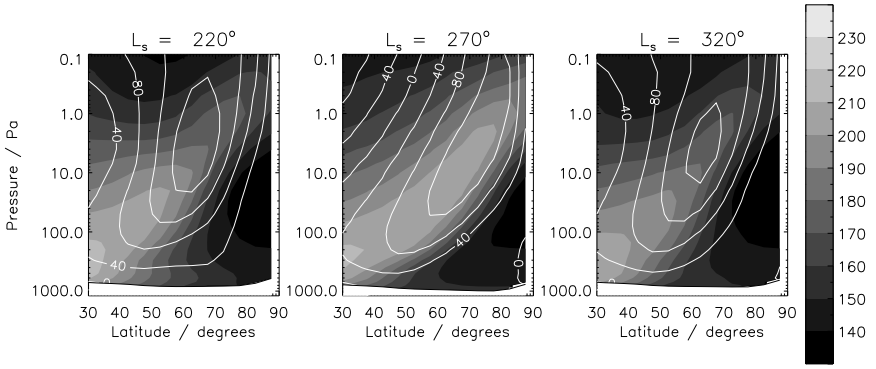

(a)

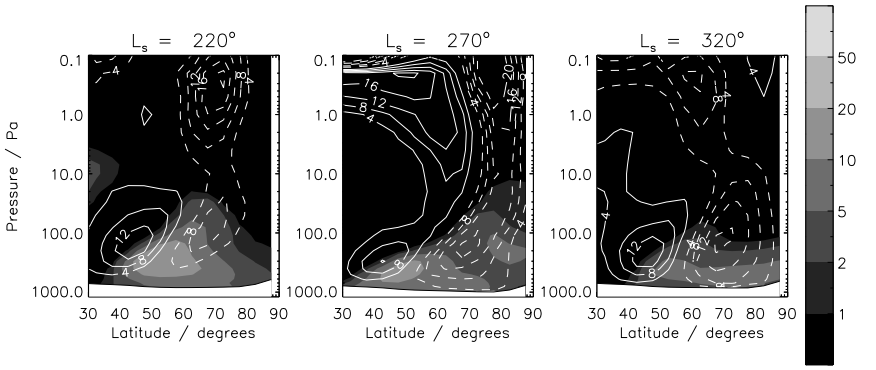

(b)

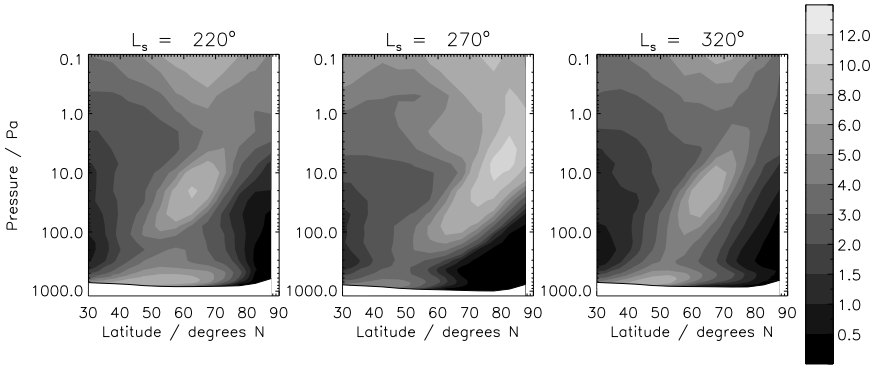

(c)

Figure 5.6: (a) Zonal-mean temperature $\bar{T}$ (K, shading) and zonal wind $\bar{u}$ ($\mathrm{m\,s^{-1}}$, contours) for τ^*_{MY24}; (b) change in $\bar{u}$, $(\tau^*_{MY24} - \tau_{MY24})$ ($\mathrm{m\,s^{-1}}$, contours), and water ice cloud mixing ratio for τ^*_{MY24} (10^{-6} kg/kg, shading); (c) Transient RMS temperature (K) for τ^*_{MY24}.

and strengthen it on the equatorward side. Above 10 Pa, the jet core was weakened by 10–20 $\mathrm{m\,s^{-1}}$ before and after solstice, and shifted slightly equatorward around 1 Pa at solstice. In the MY24 assimilation, the reduction in zonal wind at 60–80°N began earlier in the year than it did in τ^*_{MY24}.

5.3.2 Near-surface temperature gradients

A comparison of Figures 5.4(b) and 5.4(d) shows that the maximum (most negative) meridional temperature gradient very close to the ground in the northern hemisphere was weakened by the addition of clouds (seen particularly during the pre-solstice period), since the peak in clouds tended to lie to the south of the peak gradient, implying a reduction to the peak gradient as explained above. The latitude of the peak gradient was moved $\sim 5°$ to the north pre-solstice in τ^*_{MY24}, but reached $\sim$ 45–50°N at solstice in both τ_{MY24} and τ^*_{MY24}, in good agreement with the assimilation.

Eddy activity did not particularly follow the strength of the temperature gradient; for example, compare the very similar gradients seen in southern midlatitudes at $L_s = 120$–$210°$ in all four simulations shown in Figure 5.4, with the different temperature variances that resulted. The assimilation itself (MY24, not shown) featured a 200 m temperature gradient that was fairly constant over $L_s = 220$–$310°$, and was stronger at solstice than it was at the time of the pre-solstice maximum ($L_s = 200$–$220°$), though it did increase around the time of the post-solstice maximum. Therefore, near-surface baroclinic eddy activity cannot simply be explained in terms of the varying magnitude of the surface temperature gradients.

5.3.3 Linear growth rates

A better measure of the potential for baroclinic instability of the low atmosphere is provided by the Eady growth rate, σ (e.g. James and Gray (1986)):

$$\sigma = 0.31 \frac{f}{N} \frac{\partial \bar{u}}{\partial z} \tag{5.1}$$

where f is the Coriolis parameter, N is the atmospheric static stability and $\bar{u}$ is the zonal-mean zonal wind velocity. Since the baroclinic zone moves equatorward from autumn to winter with the growth of the seasonal CO_2 cap, the change in f (which is $\propto \sin$ (latitude))

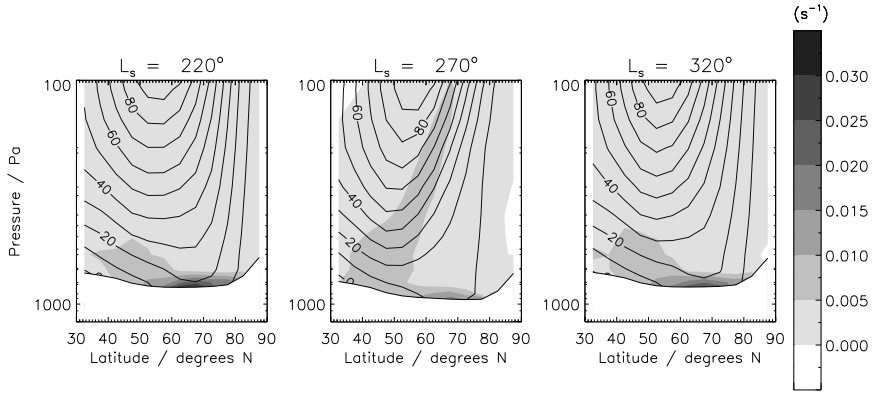

(a)

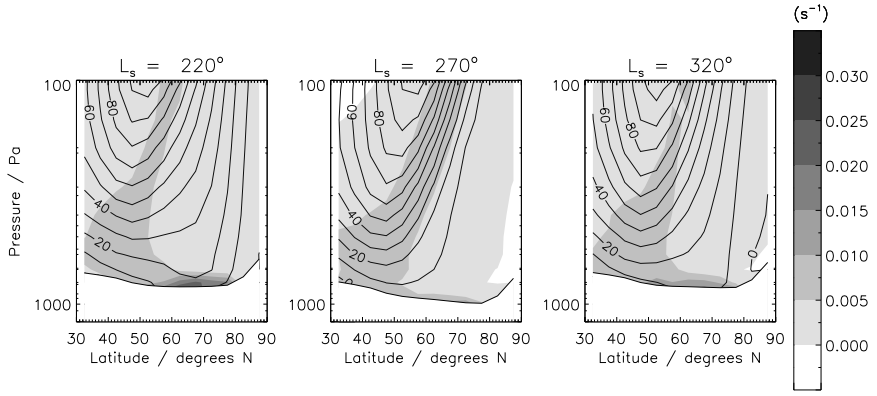

(b)

Figure 5.7: Zonal mean vertical wind shear $\dfrac{\partial \bar{u}}{\partial z}$ (shading, s^{-1}) and zonal wind $\bar{u}$, (contours, $\mathrm{m\,s}^{-1}$) near the surface, for runs (a) τ_{MY24} and (b) τ^*_{MY24}.

will cause σ to decrease by a factor of ~ 1.3 from 70° to 45° latitude. On the other hand, Wang et al. (2005) found that in the TES data, N^2, at a height of 2.5 km above the north polar cap edge, peaked at $L_s = 260°$, at around twice its value at 200°, implying an increase in N by a factor of ~ 1.4 from autumn to winter (and these results were fairly well reproduced by the model, both with and without ice clouds). Therefore, variations in f and N may be expected to approximately cancel each other out with regard to baroclinic growth rates, leaving the vertical wind shear as the most important term in the equation for the present discussion.

The growth rate at 20 Pa in τ^*_{MY24} (not shown) exhibited a maximum around solstice,

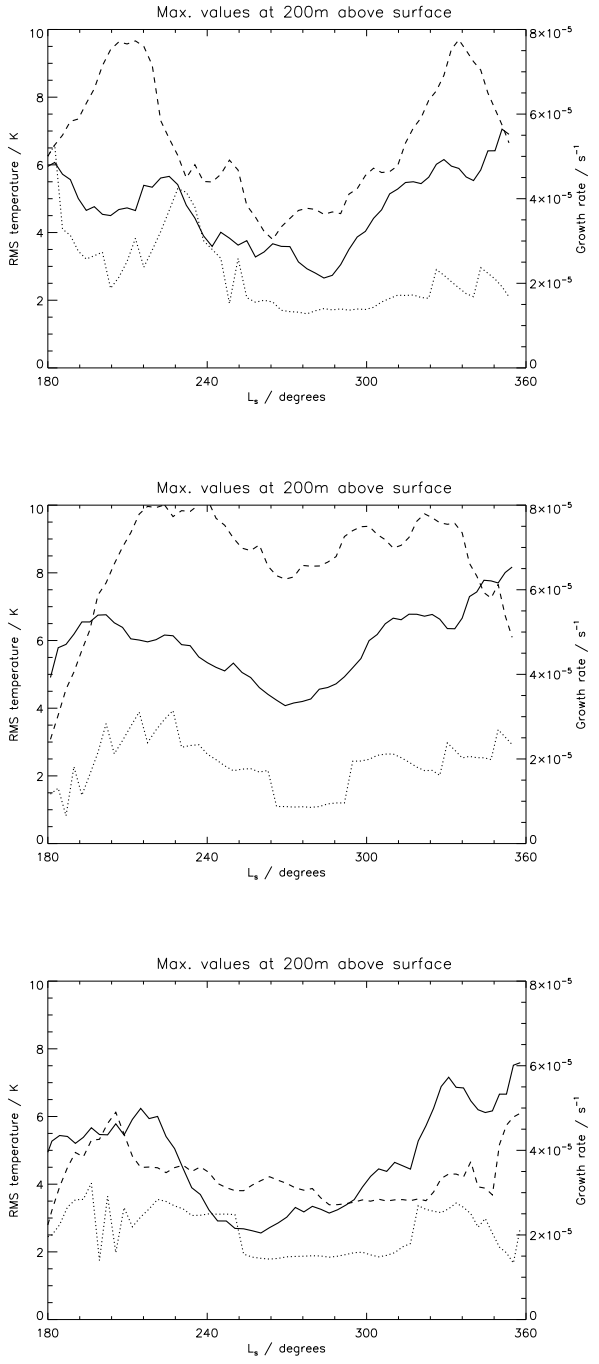

Figure 5.8: Top: Peak values of RMS transient eddy temperatures (solid) and Eady growth rates (dashed) within 30–90°N at 200 m above the ground, and growth rate at the latitude of the peak in RMS temperature (dotted), during northern autumn and winter, for τ^*_{MY24}. Middle: The same for τ_{MY24}. Bottom: The same for MY24 of the assimilation.

at $\sim 70°$N, in agreement with the simulated upper RMS temperature peak (Figure 5.6(c)). The addition of clouds strengthened this upper maximum in comparison with τ_{MY24}, by increasing the polar warming and the associated vertical wind shear, $\dfrac{\partial \bar{u}}{\partial z}$. σ also maximised at solstice at $\sim 50°$N at altitudes as low as 400 Pa. However, the growth rates that are important with respect to the solsticial pause are situated very near the ground, within the lowest few hundred metres, and varied (in this case) primarily due to changes in vertical wind shear. Without clouds (τ_{MY24}), a relatively large near-surface wind shear of 0.02–0.04 s^{-1} existed at 60–80°N throughout the period shown (Figure 5.7(a)). It was slightly reduced around solstice, compared to its values in late autumn and early spring. When clouds were used, this reduction in shear around solstice was more substantial (Figure 5.7(b)). The jet also appeared sharper and tilted further from the vertical around 100 Pa, with an equatorward shift closer to the ground, at $L_s = 270°$.

The correlation between Eady growth rate and RMS transient eddy temperature is seen in Figure 5.8, which shows the maximum values of the two quantities at 200 m above the surface within 30–90°N, along with the growth rate at the location of the (200 m) RMS temperature maximum. For τ^*_{MY24}, the peak temperature variance in the northern domain showed a correlation with the maximum Eady growth rate, and both fell to half their $L_s = 210°$ values by 270°. However, the locations of the two maxima differed somewhat after $L_s \sim 210°$. As can be inferred from Figure 5.7(b), the peak Eady growth rate was found at 60–80°N throughout the season, whereas the maximum in temperature variance was centred south of 60°N after about $L_s = 210°$ (Figure 5.6(c)). The post-solstice increase in RMS temperature and eddy activity ($L_s = 290$–$320°$) was not matched by a significant local rise in growth rate, but rather by an increase to the growth rate to the north of this location, at $\sim 70°$N. It is notable that the pre- and post-solstice maxima in peak Eady growth rate were very sharp with respect to L_s, consistent with the limited periods of strong NH eddy activity seen in Figure 5.4(d) (and Figure 5.1).

The same plot for τ_{MY24} shows that its more moderate solsticial RMS temperature reduction was accompanied by reductions in both local and domain-maximum Eady growth rates. The local growth rate at solstice again fell to less than half of its pre-solstice value, suggesting that the maximum RMS temperature, which had only decreased to $\sim \frac{2}{3}$ of its pre-solstice value, was to some extent being maintained by the presence of a larger

growth rate, situated to the north, which displayed only a slight solsticial minimum. This could be possible through equatorward propagation of disturbances from 60–70°N, since the meridional flow near the surface is relatively strong (along the topographic channels) in the autumn period, as evidenced by the propagation of flushing dust storms in the same direction.

For the assimilation, the situation appears more complicated, and RMS temperature peaks were less well predicted by the Eady growth rate. While the maximum RMS temperature fell to less than half of its peak value at solstice, the local growth rate decreased by a smaller fraction, and this decrease was delayed by $\sim 30°$ of L_s. The latitude of peak growth rate coincided more closely with peak eddy activity than it did in the model, as peak growth rates moved southward to ~ 60°N over $L_s = 220$–$330°$. The post-solstice rise in RMS temperature at $L_s = 315°$ was matched by a brief increase in local growth rate, however, as the CO_2 cap began to recede and the latitude of peak eddy activity moved slightly northward. It should be noted that the quasi-geostrophic assumption behind the Eady growth rate would be tested at such low altitudes, within the turbulent planetary boundary layer.

5.3.4 Similarities with 'dust storm' simulation

Similar results were observed in the dust storm simulation τ_{high}, which produced a significant northern hemisphere pause (see Table 5.1) despite not including water ice clouds. As shown in other studies, an increased dust loading enhances the mean meridional circulation and strongly increases the polar warming over northern high latitudes (as much as 60 K warming at the 10 Pa level, in comparison to τ_{low}) (Figure 5.9). This resulted in a stronger tilt of the entire jet, which lead to a reduction in vertical shear near the surface at 60°N, similar to that seen previously. It was at winter solstice that the principal meridional overturning cell (PMOC) was strongest and the tilt of the jet most pronounced, and thus the instability at the ground was suppressed. $s = 1$ waves were most strongly affected, and $s \geq 3$ less so, giving a decrease in the dominant wavelength, in agreement with Barnes et al. (1993) and Kuroda et al. (2007).

In this simulation the variation in peak RMS temperature was well-matched by the local Eady growth rate, as both dropped to around half their pre-solstice values at solstice (Figure

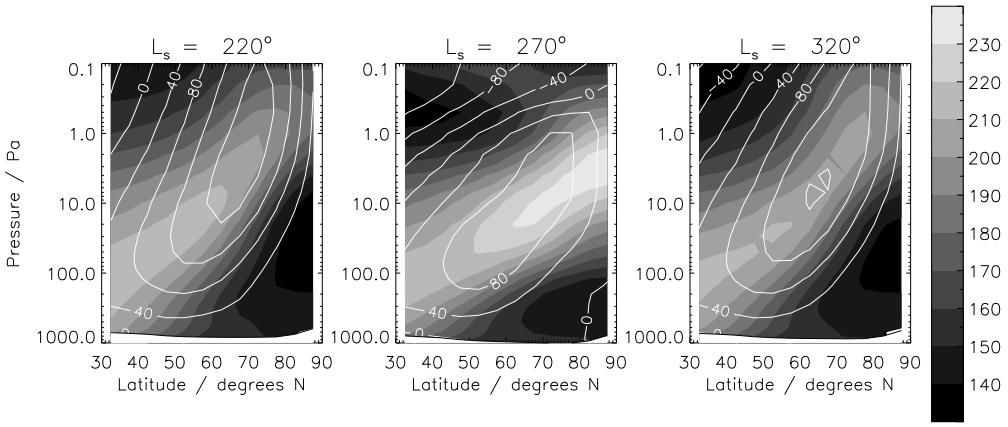

Figure 5.9: As Figure 5.6(a), but for τ_{high}.

5.10). The maximum growth rate within the domain reached a minimum, at $L_s = 280°$, of around $\frac{1}{3}$ of its pre-solstice value. By contrast, in τ_{low}, a similar $\sim 50\%$ solsticial reduction in local Eady growth rate did not lead to any reduction in peak temperature variance. Meanwhile, the maximum growth rate in the domain reached a maximum around solstice. The differences between τ_{low} and τ_{high} here lend support to the idea that consistently high growth rates at higher latitudes can contribute to the maintenance of eddy activity in the midlatitudes (along the polar cap edge, where the surface temperature gradient is strongest).

It appears therefore that in each of τ^*_{MY24}, τ_{MY24} and τ_{high} (as well as τ^*_{low}, not shown), the simulated solsticial pause formed due to a meridional tilting of the westerly jet — for the dust storm case this affected the whole jet and was due to the enhanced seasonal expansion of the PMOC, whereas for the more typical non-storm case, the equatorward shearing occurred primarily near the ground, due (in the model) to the influence of the polar hood clouds.

5.3.5 Southern hemisphere

The southern hemisphere winter jet, though weaker than in the north, also exhibits a poleward tilt over its full height, which is significant over $L_s \sim 30$–$150°$. A strong reduction in windspeed poleward of 60° near the ground was seen in the model at this time, and it

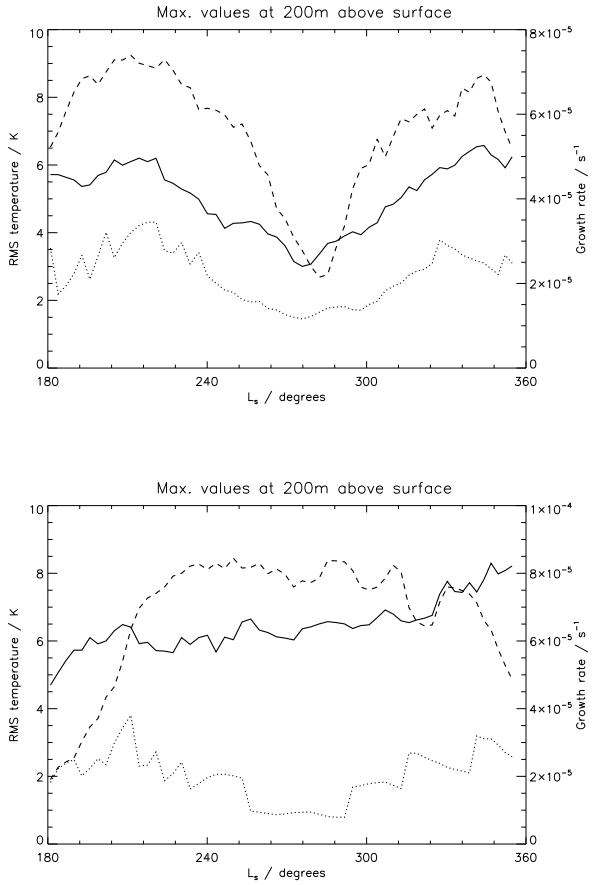

Figure 5.10: As in Figure 5.8, but for τ_{high} (top) and τ_{low} (bottom).

was slightly enhanced when clouds were included. Clouds also reduced $\bar{u}$ and $\dfrac{\partial \bar{u}}{\partial z}$ at the ground at 40–50°S, giving a deeper solsticial pause in τ^*_{MY24} than in τ_{MY24}. By contrast, the wind shear did not weaken around solstice for τ_{low}, and no pause was produced.

The southern hemisphere minimum in τ^*_{MY24} was seen despite consistently large Eady growth rates at ~ 70°S (Figure 5.11). Variations in RMS temperature were in agreement with the local predicted growth rates, as around solstice both fell to $\sim \frac{1}{3}$ of their pre-solstice values, though the post-solstice rise in RMS temperature was not accompanied by an increase in either the local or the peak growth rate. The fact that, unlike in the northern hemisphere, the presence of larger Eady growth rates at higher latitudes did not sustain eddy activity around solstice could be explained by the greater separation in latitude of the cap edge and the peak growth rate region that existed around solstice — during $L_s = 60$–140°, the peak temperature variance was located at 30–40°S (c.f. 40–50°N), while the peak in

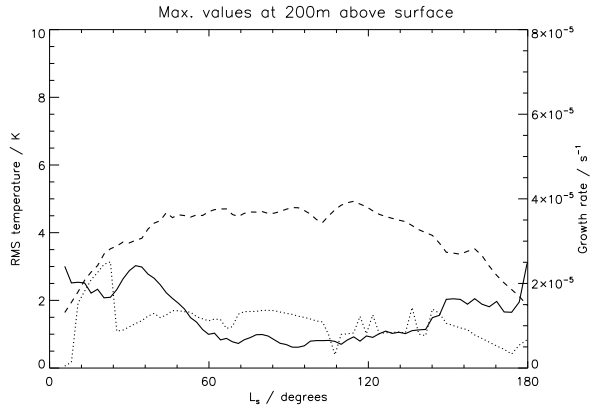

Figure 5.11: As Figure 5.8, but for 30–90°S during southern autumn/winter, in simulation τ^*_{MY24}.

growth rate remained at 60–70°S — or by the existence of less favourable conditions for the equatorward propagation of disturbances. The 60–70°S growth rates may have contributed to the larger post-solstice RMS temperature maximum, which occurred in closer proximity, over $\sim$ 40–60°S, as the south polar cap contracted. In the assimilation of southern autumn and winter of MY25 (not shown), the situation was quite different, as both measures of growth rate at 200 m minimised around solstice, along with temperature variance.

5.4 Other causes

In addition to the change to the basic state caused by the presence of water ice clouds, several other potential causes of the solsticial pause were investigated, in part taking inspiration from the mechanisms suggested for the terrestrial midwinter suppression. The first such mechanism considered (which is closely related to the ideas presented in the previous section) is an increase in the barotropic shear of the near-surface atmosphere, which may suppress baroclinic activity via the 'barotropic governor' (James and Gray, 1986). Another possibility is that heating/cooling caused by the presence of the clouds suppresses baroclinic waves directly, through diabatic damping (Chang, 2001). In contrast to Earth, this would not be through condensational heating, owing to the tenuous nature of ice clouds on Mars, but could occur rather through the radiative impact of the clouds. Finally, the solsticial pause could primarily be the result of a topographic effect, activated only when clouds are present by means of a shift in the location of the region of peak baroclinicity.

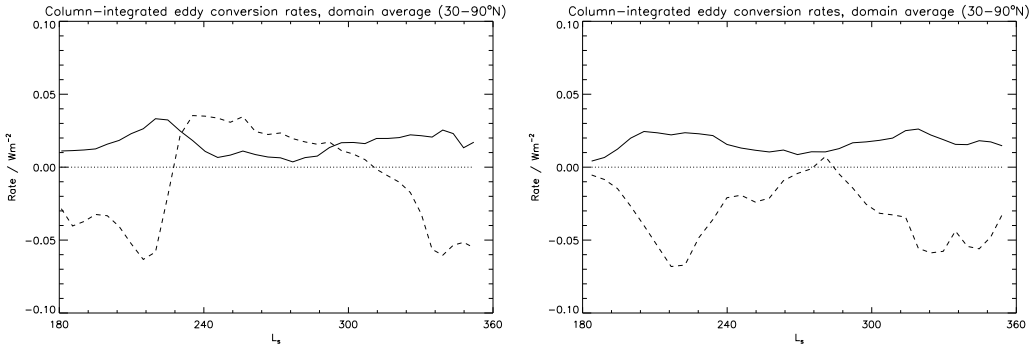

Figure 5.12: Left: Rate of generation of eddy kinetic energy (in $\mathrm{W\,m^{-2}}$), integrated over the model height and averaged over 30–90°N, through baroclinic conversion from eddy available potential energy (solid) and through barotropic conversion from zonal-mean kinetic energy (dashed), from τ^*_{MY24}. Right: the same from τ_{MY24}.

5.4.1 Barotropic shear

James and Gray (1986) noted that flow with a substantial barotropic component can become significantly stabilised with respect to baroclinic instability, due to deformation of any disturbances which should occur at a rate proportional to the horizontal wind shear, $\bar{u}_y$. By comparing this rate to the expected linear growth rate, they suggested that instability is significantly reduced when

$$\frac{N\bar{u}_y}{f\bar{u}_z} \gtrsim 0.31 \tag{5.2}$$

James and Gray envisaged the barotropic governor as converting eddy kinetic energy (EKE) into zonal mean kinetic energy, and additionally reducing the rate of baroclinic conversion from zonal available potential energy (APE) to EKE. The northern hemisphere at winter solstice actually generated significant transient eddy energy barotropically in the model (i.e. converting zonal mean KE into EKE), though not when clouds were neglected (Figure 5.12). For both τ^*_{MY24} and τ_{MY24}, total barotropic and baroclinic conversion rates over 30–90°N were anticorrelated. MY24 of the assimilation displayed a similar pattern to τ^*_{MY24} of barotropic generation around solstice and dissipation before and after, though there were differences in the timing of the switchover points, and pre- and post-solstice dissipation rates were lower for the assimilation. These energy conversion rates were calculated

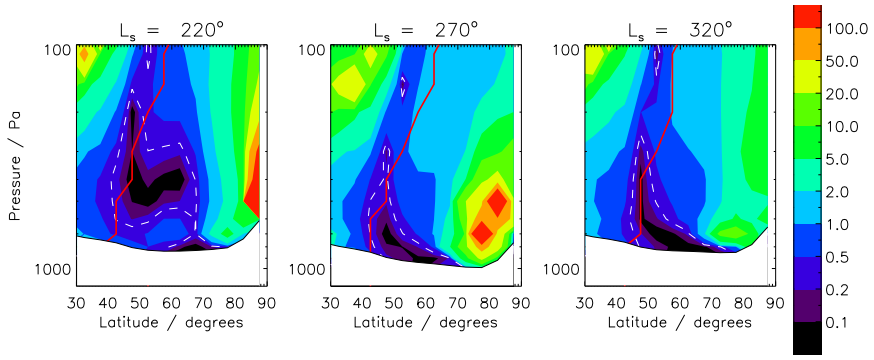

Figure 5.13: The dimensionless quantity $\frac{N\bar{u}_y}{f\bar{u}_z}$, calculated for simulation τ^*_{MY24}, during autumn/winter in the northern hemisphere. The dashed contour shows where this quantity is equal to 0.31, and the solid line marks the latitude of the strongest meridional temperature gradient at each model height level. The volume within the 0.31 contour should be 'protected' from the barotropic governor's effect.

using the formulae from Peixoto and Oort (1992):

$$C(P_E, K_E) = -\int \overline{\omega'\alpha'} dm \tag{5.3}$$

$$C(K_M, K_E) \approx -\int \overline{u'v'} \cos\phi \frac{\partial(\bar{u}/\cos\phi)}{R\partial\phi} dm \tag{5.4}$$

where ω is the vertical pressure velocity, α is the inverse of the specific density, ϕ is latitude, R is the planetary radius and m is the mass of the atmosphere per unit area; primed quantities denote deviations from the time-mean and an overbar denotes a zonal average[1]. The rates were integrated over the domain 30–90°N. Baroclinic conversion was primarily responsible for the pre- and post-solstice surface eddy maxima, and the solsticial pause can be identified as a minimum in the baroclinic conversion rate in Figure 5.12. Baroclinic contributions to this mass-weighted total came predominantly from the lowest atmospheric scale height. Barotropic conversion, however, was much less vertically limited, and the solsticial maximum seen for τ^*_{MY24} in Figure 5.12 contained energy conversion from throughout the atmosphere, including near the location of the upper RMS temperature maximum. Therefore, this increase in barotropicity represented changes in the atmosphere as a whole; however, this did include a near-surface increase over 30–70°N.

[1]Other terms in the Peixoto and Oort expression for $C(K_M, K_E)$ were found not to be significant, hence the use of the '≈' in (5.4); a boundary term (at 30°N) was, however, included in the calculation.

It has already been seen that $\bar{u}_z$ decreases around solstice, which on its own leads to reduced linear baroclinic growth rates; however, a further suppression would be suggested if condition (5.2) were satisfied. Figure 5.13 shows that large portions of the low atmosphere (outside the dashed region) did satisfy the condition, but at the latitude of strongest temperature gradient it was generally not satisfied near the surface, implying that the most unstable latitudes remained, to some extent, protected from any suppressive influence of the barotropic wind shear. Although the volume of the atmosphere in which the condition was not satisfied decreased around solstice, it remained largely intact at the surface. The seasonal variation of this volume, and of its position with respect to the peak temperature gradient, was very similar in the cloud and no-cloud (not shown) cases, which displayed rather different solsticial pauses. Therefore it would seem that, while the barotropic governor may be limiting baroclinic instability more generally, it is not primarily responsible for the suppression near the surface during winter. The increase in the barotropicity of the atmosphere seen around solstice for τ^*_{MY24} does seem to be important — it was seen consistently in each simulation that produced a solsticial pause, caused either by the presence of clouds or by a dust-enhanced circulation around solstice — but the transition from baroclinic to barotropic conversion in the low atmosphere can be seen in the reduction of vertical wind shear as presented in the previous section, and it does not seem necessary to invoke any additional suppression to explain the observed reduction in eddy activity.

5.4.2 Diabatic heating/cooling

Chang (2001) found from a modelling study that a contribution to the midwinter suppression on Earth could be given by a diabatic dissipation of eddy APE, caused by surface sensible heat flux. Away from winter solstice, he found that stronger generation of eddy APE from condensational heating was enough to overwhelm the dissipation by the surface heat flux, thereby increasing storminess.

Figure 5.14 shows the rate of generation of eddy APE in the northern hemisphere of the two MY24 runs, due to diabatic heating and through baroclinic conversion from zonal-mean APE. The diabatic generation rate is (following Peixoto and Oort (1992)) $\int \Gamma \, \overline{Q'T'} \, dm$, where Q' is the eddy diabatic heating rate and Γ is an inverse static stability parameter defined as $\Gamma = -\left(\frac{\kappa\theta}{pT}\right)\left(\frac{\partial\tilde{\theta}}{\partial p}\right)^{-1}$, where $\tilde{\theta}$ is the domain-mean potential temperature on

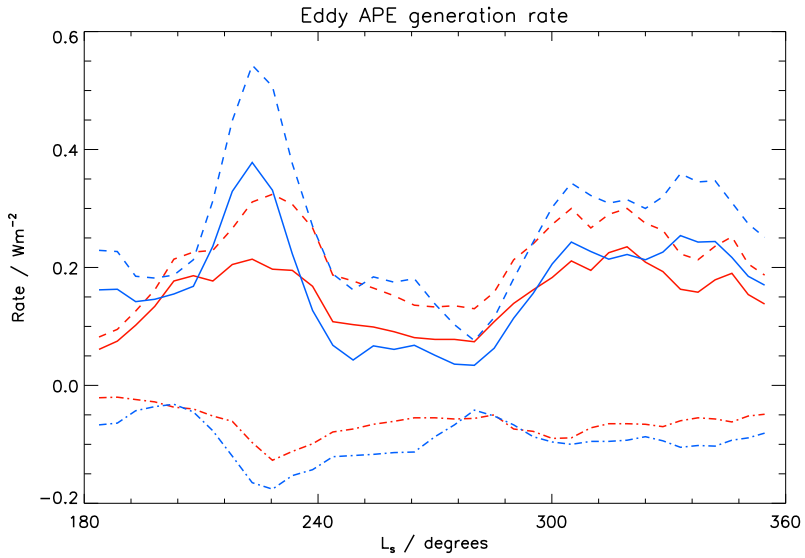

Figure 5.14: Eddy available potential energy generation rates (in $\mathrm{W\,m^{-2}}$), vertically integrated and averaged over 30–90°N, for τ_{MY24} (red) and τ^*_{MY24} (blue). Lines show contributions from baroclinic conversion (dashed) and diabatic heating (dash-dotted), and their sum (solid).

a particular pressure surface. The rate includes contributions from direct heating by dust, CO_2 and water ice clouds, and the surface sensible heat flux. The contribution from latent heat of condensation was several orders of magnitudes smaller, so was neglected. In both cases diabatic effects dissipated eddy APE throughout the autumn/winter period, meaning that perturbations in radiative cooling were positively correlated with positive temperature perturbations, and thus acted to damp baroclinic waves. Since most of the winter polar atmosphere receives little or no solar flux, this reflects the fact that warm regions tend to cool faster than colder regions.

Diabatic destruction of eddy APE acted in opposition to, and primarily varied in response to, its generation baroclinically from zonal APE, which peaked before and after solstice in both cases. The net result (generation − loss) showed seasonal progressions that agreed with the respective near-surface temperature variances, with a lower net APE generation rate over $L_s = 240$–$300°$ for τ^*_{MY24}. Over ~ 240–$270°$, baroclinic conversion was at least as rapid for τ^*_{MY24} as for τ_{MY24}, and the lower net rate was due to greater diabatic dissipation. However, for ~ 270–$300°$, the lower net rate for τ^*_{MY24} was the result of reduced baroclinic conversion. From these domain-average figures, it appears that thermal damping caused by water ice clouds could have been partly responsible for the deeper solsticial pause

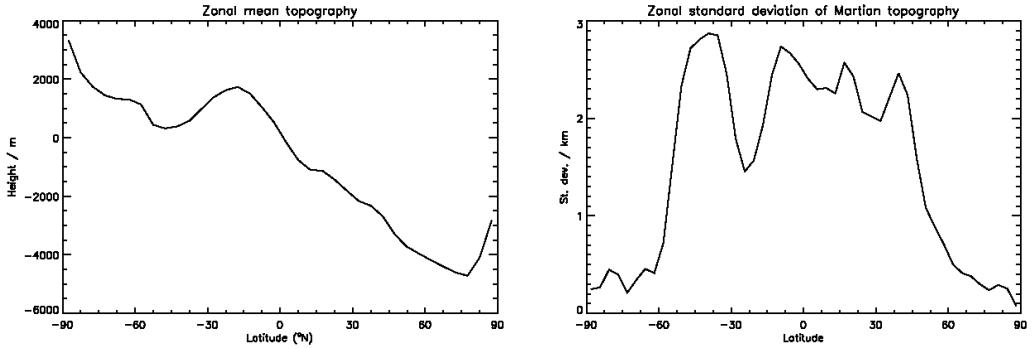

Figure 5.15: (Left) Height of the zonally averaged Martian surface above the areoid (m). (Right) The standard deviation of surface height above the areoid around latitude circles (km).

produced in τ^*_{MY24}.

5.4.3 Topographic effects

Another important consideration which has been ignored thus far is the substantial and complex topography that exists on Mars. Figure 5.15 shows the zonally averaged difference between the Martian surface and the 'areoid' (the Martian equivalent of terrestrial sea level), and its zonal standard deviation. Most of the northern hemisphere slopes downward towards the pole, while the southern polar region slopes upward towards the pole. Blumsack and Gierasch (1972) suggested that the northern hemispheric surface, with the lower boundary sloping in the opposite sense to the isentropes, should reduce baroclinic instability and shift the most unstable mode to shorter wavelengths, with the magnitude of the effect increasing with the steepness of the gradient. The meridional topographic slope is steeper over $\sim$ 35–50°N than it is over 50–75°N, which would suggest that eddy activity in the northern hemisphere should be reduced if and when the baroclinic zone moves south of $\sim$ 50°N. Using the temperature gradients of Figure 5.4(d) as a guide, this progression to or beyond 50°N appears to coincide with the duration of the solsticial pause. A similar argument could be made for the southern hemisphere in the region just north of 60°S, although the meridional variation is more complicated than it is in the northern hemisphere.

The Martian surface is also notably more zonally symmetric poleward of $\pm$60° than in midlatitudes. Yu and Hartmann (1995) found that adding topographic features to a zonally symmetric surface resulted in a transfer of wave energy from transient to stationary waves,

and from high to low frequencies. In their paper, 'high frequency' was defined as those waves with a period less than 7 days — therefore it is possible that most Martian transient waves, having relatively short periods, will be suppressed to some extent by the presence of the large zonal variations in topography that exist equatorward of ±60°.

A critical latitude for suppression of instability can be hypothesised, situated at around 50–60° in each hemisphere. In the northern hemisphere, one or both of the increase in meridional slope and the increase in zonal asymmetry south of 60°N may act to limit transient eddy activity. To investigate which, if either, of these effects exert an influence, two further model runs were performed. One simulation used zonally averaged topography, while the other had the mean meridional slope of the topography removed by subtracting the zonal mean surface height from each gridpoint. Both were forced by the MY24 dust scenario and included radiatively active water ice clouds. The dust scenario was not zonally or meridionally smoothed; however, the scenario is defined relative to a reference pressure level, so dust opacities do not retain a signal of the non-averaged Martian surface, only of heterogeneities that existed in the opacity field at the reference level.

The results, which can be compared to those in Figure 5.4(d), are shown in Figure 5.16. For the zonally averaged topography run, temperature perturbations in the northern hemisphere were increased over those from τ^*_{MY24} throughout autumn/winter, and the solsticial pause was no longer produced; in fact, there was now a solsticial maximum. This was accompanied by a consistently large vertical wind shear (and therefore growth rate) at 60–75°N. The region of peak temperature variance spread further south, as far as 30–40°N, where previously there was a strong $s = 3$ topographic forcing. In this run $s = 3$ eddies were most active over $L_s = 230$–$330°$, during the period of increased dust loading; $s = 1$ and $s = 2$ were also present around solstice and showed peaks post-solstice. However, these variations may have been partly induced by any topographic signal present in the dust scenario used.

A dramatic change was seen in the southern hemisphere, where peak RMS temperatures increased and were now of comparable magnitude to those in the northern hemisphere. Several short maxima and minima were seen but eddies remained fairly strong over $L_s = 30$–$150°$ at 30–50°S, whereas in τ^*_{MY24} they were absent from this region/season. It appears that the zonal topographic asymmetry at these latitudes is a major reason for the weakness

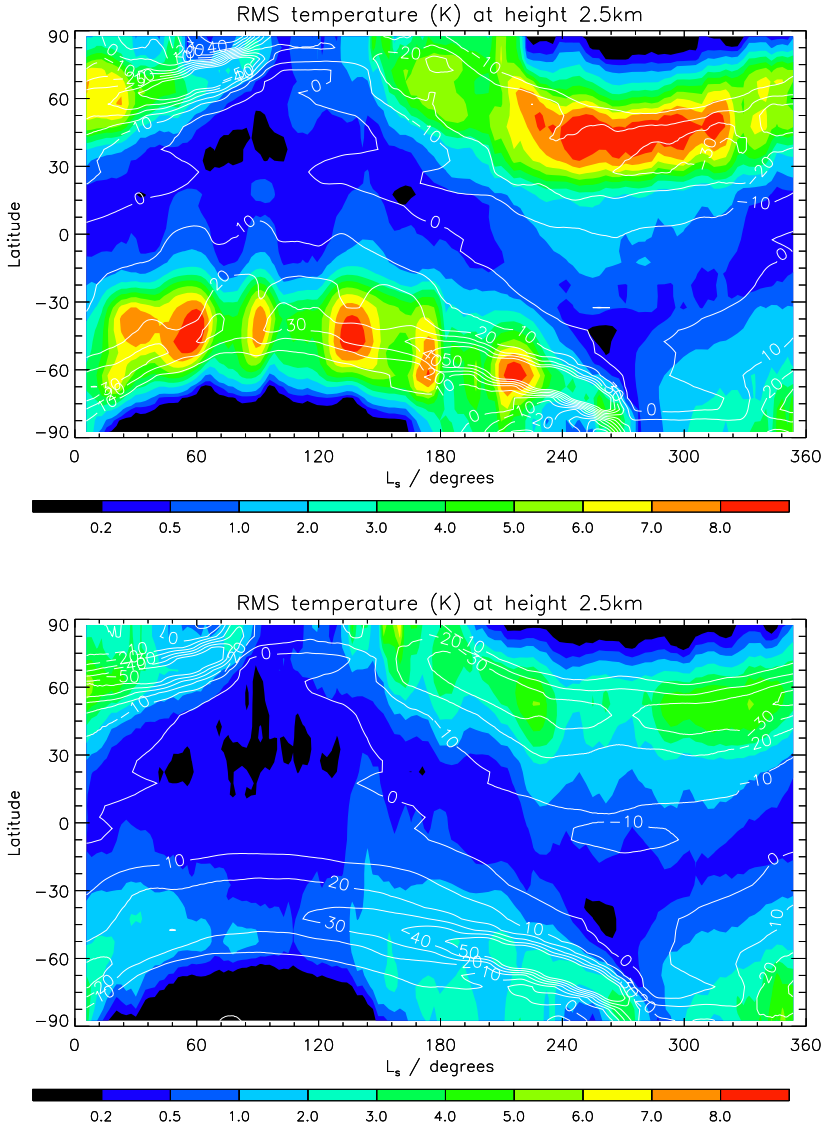

Figure 5.16: Transient RMS temperature (K) at 2.5 km for runs using the MY24 dust scenario with radiatively active water ice clouds, and using (a) zonally averaged topography (b) topography with the mean meridional slope removed. The meridional temperature gradient at 200 m (10^{-6} K m^{-1}) is overlaid in white contours.

of eddy activity in the southern hemisphere, where RMS values $\gtrsim 3\,\mathrm{K}$ are rarely seen north of 60°S. In this simulation with zonally averaged topography, $s = 1$ eddies were weak throughout the autumn/winter period, $s = 2$ peaked at $L_s = 0$–30° and 120–180°, while $s = 3$ were strongest over 30–120°. Peak 200 m Eady growth rates at 60–70°S were much larger than for τ^*_{MY24}, though the peaks in RMS appeared to follow fluctuations in local growth rate at ~ 45°S for much of the autumn/winter period.

The simulation with the meridional topographic slope removed produced a marginally shallower minimum in the northern hemisphere than τ^*_{MY24}, implying that an increased meridional slope at 35–50°N does slightly suppress instability. There was also a general reduction in temperature variance in the NH — with the exception of the pre-solstice peak, Eady growth rates were larger than they were for τ^*_{MY24}, yet weaker waves resulted throughout autumn and winter. This was possibly a result of the decreased surface pressure in the northern hemisphere, as the surface now sat higher relative to the areoid than in τ^*_{MY24}. Baroclinic activity at the most unstable latitude again followed the Eady growth rate where it was strongest, north of 60°N. The southern hemisphere saw very little change from τ^*_{MY24}.

5.5 Summary and relevance to MY24–26

The Martian solsticial pause has been convincingly simulated by the UKMGCM, when run with a realistic dust loading and the inclusion of radiatively active water ice clouds. Both increased dustiness and ice clouds alter the thermal structure of the polar atmosphere, lowering the zonal wind and its vertical shear near the surface at mid- and high latitudes around winter solstice, causing a decrease in baroclinic growth rates. The suppression occurs only very close to the surface and affects eddy activity in the lowest atmospheric scale height, making this midwinter suppression quite different from the equivalent phenomenon on Earth, in which transient eddies are seen to decrease around winter solstice in the north Pacific throughout the troposphere, during certain years. Many of the explanations suggested for the terrestrial suppression, including increased barotropic damping (Deng and Mak, 2006) and an inverse correlation with the strength of the westerly jet (Zhang and Held, 1999), do not apply to the Martian case, but are not needed to explain the observed transient activity. The Martian solsticial pause appears, in fact, to be less mysterious

than the terrestrial suppression, as baroclinicity at a particular location near the surface is actually reduced around solstice. The Pacific solsticial suppression, on the other hand, occurs as a localised effect at odds with an atmospheric baroclinicity that is at a maximum (Nakamura, 1992).

5.5.1 Factors contributing to the solsticial pause

Baroclinic eddy activity near the surface in both hemispheres appears to be accurately predicted by linear instability theory. In all model runs that produced a solsticial pause, and to a lesser extent in the assimilation of MY24, variation in transient eddies was accompanied by changes in Eady growth rate, driven by changes in surface vertical wind shear. Quantitative agreement was found between variations in growth rate and the amplitude of temperature perturbations, though in some situations it was necessary to consider both the growth rate at the location of the eddy maximum and the (often larger) rate elsewhere in the hemisphere. There was evidence that a high wind shear at high latitudes could sustain wave growth at midlatitudes, though the effect was seen not to apply in the southern hemisphere, perhaps due to a larger spatial separation between the two regions.

The barotropic governor effect does not seem to play a significant role in the Martian solsticial pause in either hemisphere. Diabatic dissipation, in the form of radiative damping, appears to play a modest role, in concert with modulations in the local Eady growth rate itself. The Martian topography, however, does suppress eddies particularly strongly around solstice. Specifically, it is the zonal asymmetry of the midlatitude topography which exerts a strong influence on transient activity in both hemispheres. In the case of the southern hemisphere, the removal of this asymmetry increased model RMS temperature during autumn and winter by around a factor of two, bringing magnitudes close to those seen in the northern hemisphere. A difference between the two hemispheres might be anticipated, based on the slightly weaker winter baroclinicity in the southern hemisphere (owing to Mars' orbital eccentricity), but as noted by Barnes et al. (1993), this difference does not seem large enough to explain the significantly weaker transient eddies observed in southern autumn and winter. The effect of removing the zonal asymmetry in the southern hemisphere was to increase Eady growth rates near the surface over $\sim$ 40–80°S — peak growth rates were still located at 60–70°S, but relatively large values reached as far north

as 40°S, producing waves at this latitude around solstice.

In the northern hemisphere, the absence of zonal variations in topography allowed a large surface vertical wind shear to be maintained throughout autumn and winter at 60–70°N, where it would otherwise have tended to reduce around solstice. A secondary peak in wind shear and growth rate at 10–30°N was possibly also important in creating the solsticial maximum seen for this simulation. Further altered-topography runs will be necessary to determine why the zonal asymmetry in topography suppresses waves so much more strongly in the southern hemisphere than in the northern hemisphere, as the variation of this asymmetry with latitude is fairly similar in each of the two hemispheres (Figure 5.15). This is pertinent to the occurrence of cap-edge dust storms (and potentially also larger storms) during the recent history of Mars, since the precession of the solar longitude of perihelion, which occurs with a period of $\sim 50\,\mathrm{kyr}$ (Ward, 1974), is roughly equivalent to the flipping of the topography about the equator every $\sim 25\,\mathrm{kyr}$.

The meridional topographic gradient did not have as strong an influence on model results, as when it was removed, northern hemispheric eddy activity was slightly weakened, when it might have been expected to increase due to the removal of a lower boundary that slopes downwards towards the pole. As a result of the meridional averaging, the northern hemisphere was raised to a higher altitude, which implies a lower CO_2 frost point and greater temperature inversions over the cold polar cap, making the atmosphere more statically stable. It would appear that this had a greater effect on eddies than the change to the slope of the lower boundary. The NH solsticial pause was, however, slightly shallower than for τ^*_{MY24}, presumably because of the removal of the steeper meridional slope at 35–50°N. In the southern hemisphere, the magnitudes of the pre- and post-solstice peaks were essentially unchanged, suggesting that the replacement of a boundary sloping upwards towards the pole (at 50–60°S and 75–90°S) with a flat surface, which should reduce instability, was roughly balanced by the decrease in surface pressure, which may enhance instability due to weakened temperature inversions and static stability over the polar cap.

5.5.2 Interannual variability in surface travelling waves

This insight into the reasons behind the surface suppression in northern midlatitudes made it possible to draw some conclusions regarding the interannual variability seen in northern

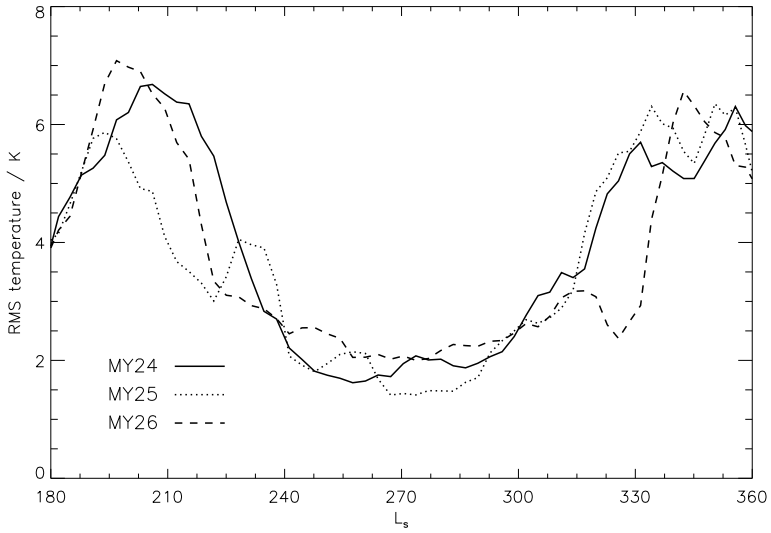

Figure 5.17: Transient RMS temperature at 2.5 km, averaged between 40°N and 80°N, for MY24 (solid), MY25 (dotted) and MY26 (dashed), from the assimilation.

hemispheric travelling waves in the TES data (Figures 5.1 and 5.2), building on ideas discussed in Wang (2007). Figure 5.17 illustrates more clearly the variation in the length and depth of the solsticial pause at the surface in each of the three years.

The earlier reduction in transient eddy activity seen in MY26 compared to MY24 followed the earlier development of the upper maximum and tilt of the westerly jet due to the regional storm at $L_s = 210°$ in MY26. Wang (2007) found that flushing storms, which disappeared upon the onset of the 210° storm, did return briefly over $L_s = 230$–$250°$ which, as can be seen from Figure 5.1, was a period of moderate transient activity along the cap edge. She attributed the eventual cessation of flushing storms at $L_s = 250°$ to the progression into the solsticial window, where the change in atmospheric state makes it impossible to sustain flushing storm activity. The particular aspect of this change that is relevant, according to the results shown in the previous section, is the tilt of the westerly jet, particularly near the ground, which increases naturally around solstice, even without an increase in global dust loading.

The pause seen in τ^*_{low} would suggest that this seasonal change to the jet, when augmented by the radiative impact of water ice clouds, is enough to effectively reduce the vertical wind shear at the surface so as to inhibit low level baroclinic growth, and therefore

that a solsticial pause should occur in each Martian year, regardless of global dust loading. This makes sense particularly with regard to the southern hemisphere where, in contrast to the NH, winter solstice is generally not accompanied by a rise in visible dust opacity much beyond the value of 0.2 used in τ^*_{low} and τ_{low}. The fact that a solsticial pause was seen in the SH for τ^*_{low} but not τ_{low} means that clouds do seem to play a role.

However, some caution is required, as the model currently overestimates polar hood cloud thicknesses. Visible ice opacities in the model north polar hood were typically 0.2–0.5, and reached values of around 1 at their peak at $L_s \sim 240°$, leading to a strong cooling near the ground. Recent results from Mars Climate Sounder (McCleese et al., 2007; Benson et al., 2011) found that typical north polar hood visible opacities are around 0.1–0.3, suggesting significantly thinner clouds than simulated here by the model, particularly during mid-autumn. Related to this, in comparison with the assimilation, it seems that the jet tilt in autumn and winter (shown in Figure 5.6(a)) was overestimated near the ground. Southern hemisphere visible cloud ice opacities were more realistic — 0.05–0.15 along the edge of the polar cap from $L_s = 30°$ to 120° — but a post-solstice thickening in polar hood cloud observed by Benson et al. (2010) was absent, as was a solsticial minimum in ice opacity. If cloud opacities were overestimated, they may have contributed too strongly to the surface cooling and weakening of the near-surface zonal wind in mid- and high latitudes around solstice. This appears to be true in the southern hemisphere, as the simulated pause of τ^*_{MY24} was too severe in comparison to the assimilation, which was arguably reproduced better (in the SH) by τ_{MY24} (compare Figures 5.4(b) & 5.4(d) with Figure 5.1). This would imply that the use of a realistic dust scenario is as at least as important in replicating the observed SH pause as is the action of the clouds in the south polar hood. In the northern hemisphere, the low dust optical depth may have been overcompensated for by the presence of polar hood clouds, so that in reality such low dust loading conditions (should they ever occur on Mars) would not produce such a significant solsticial pause as that seen in τ^*_{low}.

Conversely, a large dust storm may not fully suppress surface transient eddies if it occurs too early in the year. The pre-solstice peak in surface RMS temperature was lower in MY25 than in the other two years, but nevertheless values of $\sim 6\,\mathrm{K}$ were present at $L_s = 200°$, while the global dust storm (GDS) was in progress — it seems that pre-210° is simply too early to force a sufficient tilt in the jet near the ground to properly suppress surface

instability, even with a storm as large as the one in MY25. This fits with the results of simulation τ_{high} (Figure 5.10), in which eddy activity was not reduced until after $L_s \sim 220°$, despite an equally high dust loading during $L_s = 180$–$220°$.

The post-solsticial peak in eddy activity was delayed by $\sim 20°$ in MY26 compared to the previous two years — this can be linked to the dramatic increase in dust loading seen towards the end of that year (the 315° dust storm), which led to a temporary reversal of the seasonal decrease in the magnitude of the polar warming, and associated tilt of the jet, that usually occurs at this time of year. When the late-winter storm occurred, the temperature at 10 Pa, 60°N, which had been declining since $L_s = 270°$, increased again, by ~ 10 K, as the PMOC expanded in response to the increase in dust loading. The tilt that this imparted to the jet allowed a renewed suppression of the near-surface vertical wind shear and therefore of the eddy activity, which had been increasing in line with the other two years until the occurrence of the storm.

Dust lifting in the northern midlatitudes appears to possess a self-limiting quality, whereby the occurrence of a series of flushing storms is likely to inhibit the development of further such storms through the strengthening of the PMOC (Wang, 2007). Periods of flushing storm activity in early autumn and late winter seem certain to occur, however, as outside $L_s \sim 210$–$330°$ it is very difficult to force the tilt in the lower part of the westerly jet that is needed to reduce growth rates. What is less certain is that there must be an annual window where no flushing storms occur. Whether or not these conditions occur in a certain year may depend on there being a sufficiently large dust loading around northern winter solstice, resulting either from the occurrence of a flushing storm or more simply by an increase to the 'background' loading above some baseline level. The minimum simulated by τ^*_{low} suggests that this is actually not necessary (at least not to a visible opacity any larger than 0.2), but as mentioned previously, this run may have overestimated the contribution to eddy suppression of water ice clouds.

Chapter 6

Dust interaction with water ice clouds

With a flexible dust lifting scheme (developed in Chapters 3 and 4) and radiatively active water ice clouds (introduced in the previous chapter), it is possible to use the UKMGCM to study the interaction between dust and water, a primary goal of the project. The key areas in which this interaction is manifested include dust lifting at the cap edges, transport into the winter polar regions and deposition onto the polar caps, cross-equatorial dust transport, modifications to vertical dust mass profiles, and feedbacks on the water cycle. In this chapter, the performance of the model's cloud scheme is analysed using both prescribed and transported dust, and the changes imposed upon dust lifting patterns by the addition of active clouds are shown. Then, the effect of dust scavenging, due to both heterogeneous cloud nucleation and CO_2 snowfall, is introduced. The sensitivity of various climate features to the efficiency of nucleation is investigated, with a view to gaining an insight as to what the true efficiency of the process may be.

6.1 The UKMGCM cloud scheme

The cloud microphysics used in the model were formulated by Montmessin et al. (2004). Briefly, the scheme calculates cloud ice growth rates at each timestep as a function both of the saturation ratio and of the number of dust particles available to act as ice nuclei (IN).

Two tracers, water vapour and water ice, are used, and mass is distributed between the two modes. Ice particles are assumed to be monodisperse, and the mean radius at each gridpoint is calculated by dividing the ice mass present among the number of IN, assuming that each ice particle contains one dust particle at its core. Coagulation of ice particles is neglected, as this is thought to involve only particles smaller than 0.1 μm in radius (Montmessin et al., 2002). Sedimentation velocities are calculated for the ice tracer only, and are adjusted to account for an assumed lognormal size distribution with a prescribed effective variance ν_{sed}. Nucleation is assumed to occur only heterogeneously, i.e. the presence of IN is essential for cloud formation. The cloud ice mass mixing ratio tendency at each gridpoint, dM_c, is calculated, for a timestep of length dt, as

$$dM_c = 4\pi\mathcal{N}r_c^2 dr_c = 4\pi\mathcal{N}r_c\frac{S - S_{eq}}{(R_c + R_d)}dt, \tag{6.1}$$

where $\mathcal{N}$ is the IN number density, r_c is the mean ice particle radius, S is the current water vapour saturation ratio, S_{eq} is its equilibrium value of 1, and R_c and R_d are terms accounting for the heat and molecular diffusion resistances to crystal growth (see Montmessin et al. (2002)). Ice crystal growth is therefore favoured when the saturation S and IN number density $\mathcal{N}$ are large, and becomes progressively easier as r_c increases. Initial crystal formation is allowed only when S exceeds 1.4, which corresponds to a contact angle parameter of 0.95 (Michelangeli et al., 1993). This is higher than the saturation ratio required to begin nucleation calculated by Määttänen et al. (2005) of 1.18, though it was found that results were not particularly sensitive to this choice.

When using prescribed dust opacities, $\mathcal{N}$ must also be prescribed. This is done by assuming for the vertical distribution of dust number density the same Conrath function that is used to distribute the column opacity into each of the vertical layers in the model's radiation scheme (equation (2.1)), and constraining the total column abundance of dust particles using the prescribed column opacity. Dust particle size, also required to calculate r_c, is taken as a constant 1.5 μm in this case. The calculated IN number density is multiplied by a tunable scaling factor, R_s, to account for the unknown fraction of dust particles that become involved in cloud formation (this is later identified as the scavenging efficiency — see §6.3). The setup of the scheme makes it easy to alternatively obtain $\mathcal{N}$ from a transported dust field, if the dust cycle is being simulated. In this case, particle number density and

radius are known at each gridpoint, and it is possible to select as IN only those particles with radii greater than 100 nm, found by Määttänen et al. (2005) to be an approximate critical radius, below which nucleation rates decrease rapidly with decreasing size[1]. The reduced dust number densities are further multiplied by R_s, to give $\mathcal{N}$.

6.1.1 Water cycle results with the MY24 dust scenario

Using radiatively active clouds with the MY24 prescribed dust scenario, and with best-fit settings for ν_{sed} and R_s (Madeleine, 2011) and a value of 0.45 for the visible albedo of surface water ice, produced vapour abundance and cloud opacity fields as shown in Figure 6.1. Cloud opacity at low latitudes compares well with TES data (Figure 1.2), though the aphelion cloud belt (ACB) was centred around 20° too early in solar longitude, and there was too much cloud ice at the equator at $L_s = 180$–$240°$. When polar hood opacities are compared to those measured by MCS, several differences are apparent. South polar hood cloud thickness and location during autumn match well with the data presented by Benson et al. (2010), and a decline in opacity after the peak at $L_s = 40°$ was captured. However, the reformation of the south polar hood over $L_s = 120$–$200°$ was not simulated by the model. In the northern hemisphere, polar cloud opacity may exhibit some interannual variability, dependent on dust storm activity, so the results using the MY24 dust scenario need not necessarily match the MY29 MCS observations; nonetheless, peak model IR ice opacity (~ 1, at $L_s = 230°$) significantly exceeded the maximum value of ~ 0.1 reported by Benson et al. (2011), so it is very likely that the model was overestimating north polar hood cloud thickness at this time. Later in NH winter, cloud optical depths were in agreement with Benson et al. (2011); however, the model north polar hood persisted too long into spring, as MCS observed its disappearance at around $L_s = 20°$, rather than the date of $\sim 60°$ seen here.

Vapour abundances, in comparison with TES (Smith, 2008), were too low by a factor of ~ 2 in this run (even moreso at the south pole in spring), although the qualitative behaviour was largely correct, except for the strange feature seen in northern midlatitudes around $L_s = 210°$, which is not replicated in the observations. The north polar vapour maximum could be increased by using a lower cap albedo (a value of 0.4 would not be

[1]The same phenomenon occurs on Earth: small 'Aitken mode' aerosol particles are known to be less efficient as IN than larger particles (Pruppacher and Klett, 1978).

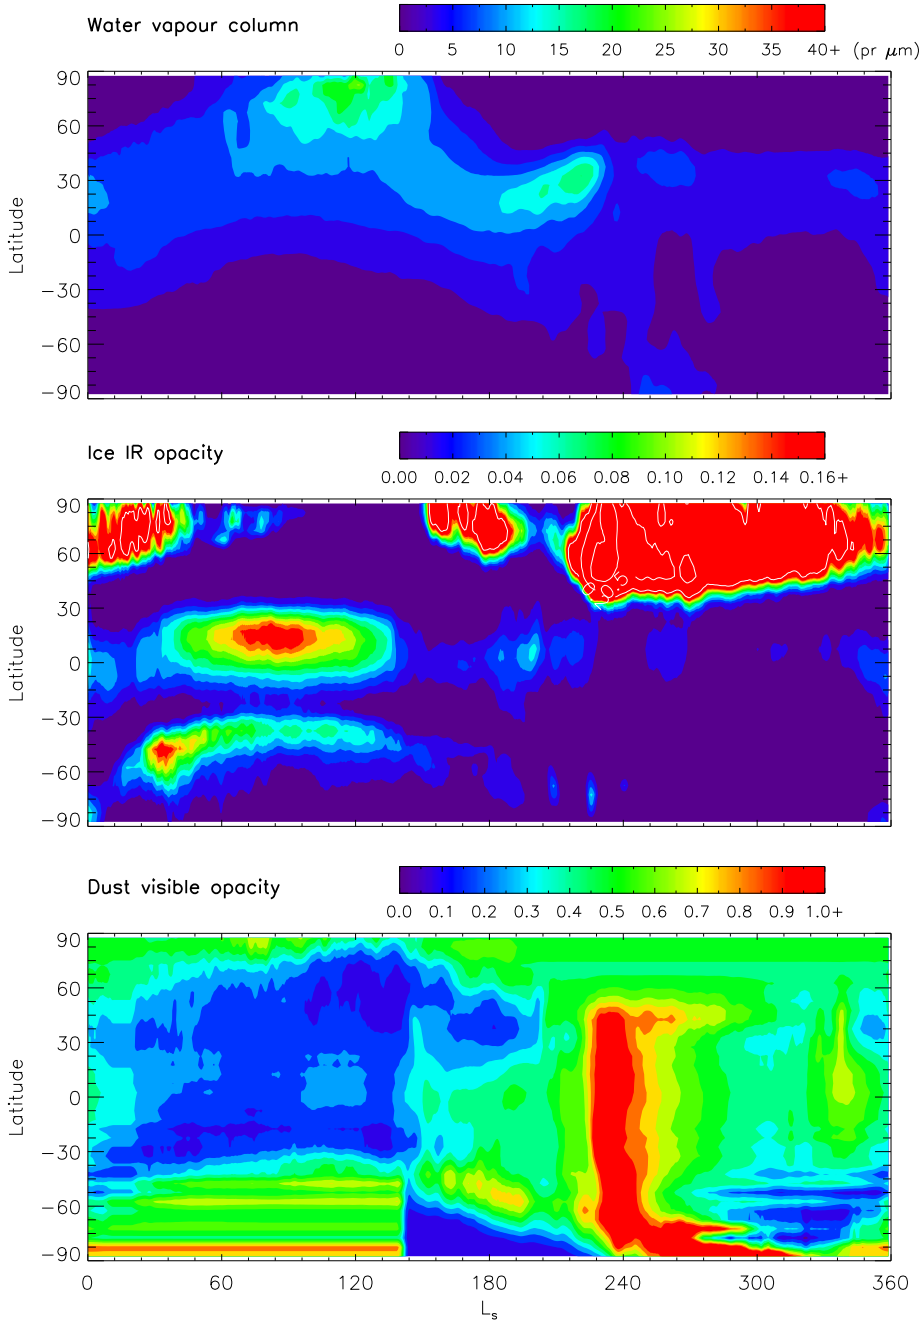

Figure 6.1: Zonal-mean water vapour column abundance (pr μm, top) and water ice cloud 12 micron opacity (middle), for the best-fit water cycle simulation using prescribed (MY24) dust opacity. In the latter, white contours mark opacity levels (0.2, 0.5 and 1) past the limit of the colour scale, which is designed (in both plots) to be visually comparable to the TES observations of Figure 1.2. Visible dust opacity from the MY24 scenario is shown in the bottom plot.

unrealistic); however, the increase in supply of water to the atmosphere would lead to an increase in cloudiness, and therefore further parameter changes may be necessary to retain a realistic ACB.

It must be noted at this point that the variance assumed for the ice particle size distribution when computing sedimentation rates, $\nu_{sed} = 0.45$, was significantly larger than the value of 0.1 used by Clancy et al. (2003) and in the model's radiation code. The larger variance increases the effective ice sedimentation rate and thus reduces the atmospheric ice content: without such a large ν_{sed}, cloud opacities produced were too large. A further reduction to ice abundance comes in the form of R_s, which was set to 0.22; however, this is not in disagreement with any measurements (of nucleation efficiency), since none currently exist. Likely values of R_s are discussed further in later sections.

6.2 Dust lifting with radiatively active clouds

6.2.1 Effect of clouds on surface winds

The addition of radiatively active ice clouds to the model had a significant impact on the thermal structure of the atmosphere, particularly near the surface, enough to affect winter baroclinic wave activity (as seen in the previous chapter), and surface wind strengths at various latitudes. This was explored using two of the simulations performed in Chapter 5 (τ_{MY24} and τ^*_{MY24}, each of which used the same initial conditions and was forced by the MY24 dust scenario); one had no water cycle, and the other did have a water cycle, including radiatively active clouds. The zonal-mean surface drag velocity from the former is shown in Figure 6.2, along with the change due to the addition of clouds. In both cases, winds were strongest in southern spring and summer, in the northern polar latitudes and the southern subtropics. Without clouds, a minimum in windspeed formed in northern high latitudes around $L_s = 270°$, and when clouds were activated this minimum became more pronounced. Cap-edge winds were maintained better in northern spring ($L_s \sim 0$–$60°$) with active clouds. The inclusion of clouds also strengthened winds in the tropics and subtropics near the solstices.

Eddy activity is also of relevance to dust lifting and transport. Since it is the peak wind gusts which often are responsible for bursts of dust lifting, periods of strong wave activity

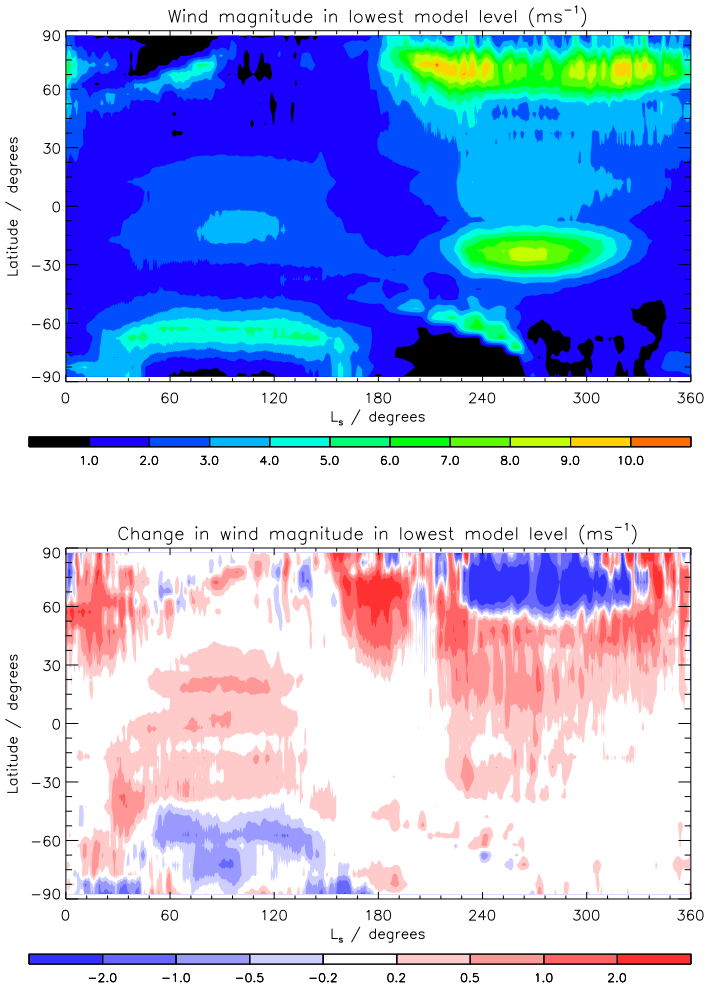

Figure 6.2: Zonally averaged horizontal wind magnitude from the lowest model level ($\mathrm{m\,s^{-1}}$), from the run without clouds (top) and the change due to the addition of radiatively active clouds (bottom).

can result in intense lifting, if a suitable background wind is present; furthermore, it will be shown later that eddy dust fluxes are the dominant means of meridional transport in certain regions. The transient and stationary components of meridional velocity v at a height of 2.5 km above the ground are plotted in Figure 6.3. As in Chapter 5, the transient component has been filtered to include only waves with periods of 1.5–10 sols, and stationary components were calculated using a time window of 20 sols. Differences were seen in transient eddy activity in NH autumn and winter: the solsticial pause was clearly produced when clouds were included, but was relatively weak without clouds. The SH pause in transient RMS v was also much wider and deeper when clouds were included.

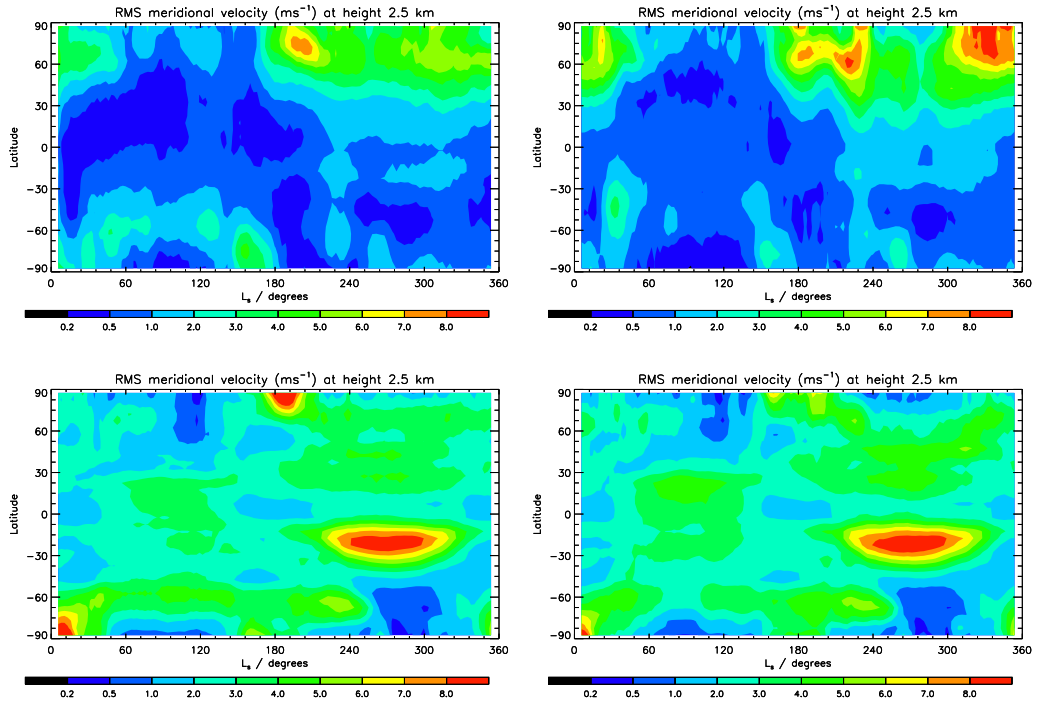

Figure 6.3: The transient (top row) and stationary (bottom row) components of the meridional wind field ($\mathrm{m\,s^{-1}}$) (filtering of each as explained in the text) at a height of 2.5 km above the ground, from runs without (left) and with (right) ice clouds.

Differences in stationary wave amplitudes were more subtle, but a solsticial minimum could be seen at 60°S around $L_s = 90°$ when clouds were used. Such a minimum can also be found in the MY25 reanalysis (not shown), and it was noted that the lower right-hand plot of Figure 6.3 matched very closely with the corresponding plot for MY25 in the SH over $L_s = 0$–$180°$. Southern hemisphere transient waves in the active cloud run over the same period agreed qualitatively with the reanalysis, but were too weak over $L_s = 0$–$60°$ and 120–180°.

Dust lifting was then activated in the model, and the transported dust field used to provide the input IN for the cloud scheme (again with $R_s = 0.22$) instead of the prescribed IN field. For this initial run, the variable threshold scheme was not used. Dust lifting output is very sensitive to the choice of lifting threshold(s), so the results presented in this subsection aim merely to show the general changes to dust lifting patterns and atmospheric transport that radiatively active ice clouds induced, before scavenging effects were considered. The use of a transported dust field to control water ice condensation, while computing the ra-

diative contributions of both dust and ice, made the model considerably more unstable, so some limits were imposed on the cloud and radiation schemes to prevent any unphysical occurrences at isolated gridpoints and to limit the damage arising from numerical errors. Upper and lower limits, of 10^{-10} m and 5×10^{-4} m, were imposed on the dust effective radius, since this (in the two-moment scheme) was calculated using the mass and number mixing ratios, either of which could occasionally fall to zero or close to it without the other doing so, as a consequence of the advection scheme. The lower limit then also acted as a lower limit for the mixed (dust-ice) particle radius, by necessity. Runaway ice crystal growth was a potential problem at specific sites, primarily around Alba Patera and Tharsis, where dust lifting was particularly favoured and where a near-constant supply of IN was provided to the cloud scheme, particularly at low levels. In an active-cloud simulation, the buildup of ice at one gridpoint could lead to very large optical depths, causing the model to crash. A (high) limit on ice mass mixing ratios of 0.002 kg/kg was applied, and cloud ice IR optical depths, as computed by the radiation scheme, were required to be ≤ 1.5 (without further altering mixing ratios).

As before, a pair of year-long simulations were carried out, one with clouds active and the other without. Dust lifting parameters were chosen to give a reasonable 'non-GDS' opacity in both cases, and lifting parameters were identical in the two runs. Differences in wind stress lifting rates were as expected, given the drag velocity and eddy activity behaviour seen previously, and cap-edge dust lifting peaks in both hemispheres could be directly related to variation in transient eddy meridional velocity (not shown). Key changes in wind stress lifting rates when clouds were added included a deeper solsticial pause in the SH, stronger lifting at the NH cap edge — particularly in early autumn — and its persistence in springtime, and a lifting peak at 0–30°N around $L_s = 230°$. The principal link between the presence of clouds and dust lifting was midlatitude eddy activity. The total amounts of dust lifted by wind stress were 2.9×10^{12} kg without clouds and 4.2×10^{12} kg with clouds, so the use of ice clouds in the model strengthened wind stress lifting overall. Dust devil lifting was also slightly increased.

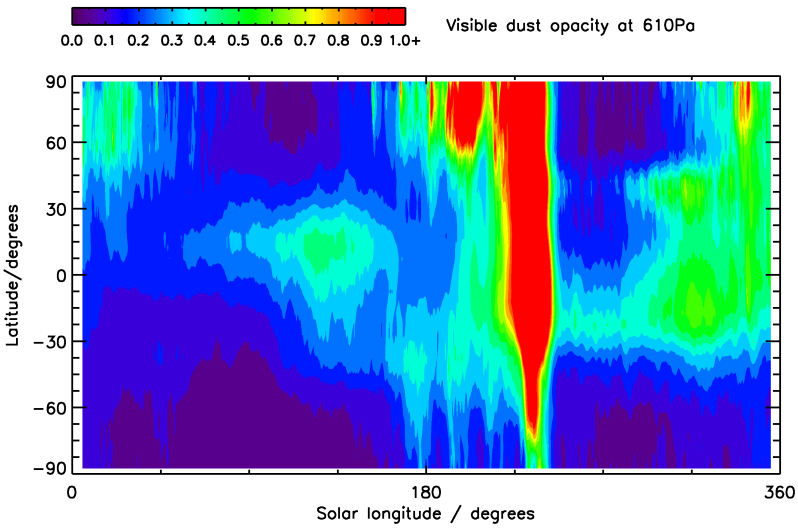

Figure 6.4: Zonally averaged visible dust opacity from the run with active clouds, which assumed infinite dust surface density.

6.2.2 Adjusting dust lifting thresholds

The changes to lifting caused by clouds were, in fact, such that the annual dust opacity cycle became totally dominated by the pre-solstice peak in lifting in the northern hemisphere, as seen in Figure 6.4. Although a peak in global opacity at this time of year is common in observations, the morphology of the 'storm' that occurred in this run was not realistic compared to any observed dust storms, as opacities increased somewhat uniformly across the planet, beginning around 0–30°N. Furthermore, the intrusion of such a large dust mass into the north polar region in autumn is very much at odds with MCS data (see §6.6.1). These results are another example of the inadequacy of an 'infinite surface dust' formulation: lifting becomes dominated by those regions in which surface stresses are largest, even though the regular raising of dust from some of the locations may be demonstrably unsustainable and unrealistic. The development of a quite realistic post-solstice dust loading peak at 300–360° was, however, noted as a positive step resulting from the use of ice clouds (which strengthened winds in the northern midlatitudes during this period).

The variable threshold scheme was therefore used once again. The model was provided with an initial state of uniform baseline effective threshold 0.02 Pa, as before, and was run for several years with ${\zeta^t}_{dec} = 0$, so as to allow thresholds to increase as far as they could before lifting became severely restricted. NH cap-edge lifting actually persisted for

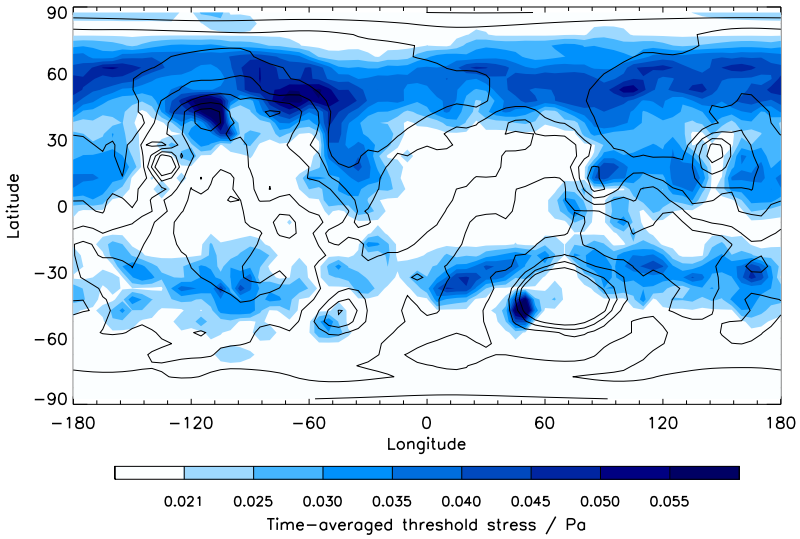

Figure 6.5: The threshold stress field after around ten years of simulation with $\alpha_N = (5–10)\times10^{-5}\,\mathrm{m}^{-1}$, $\zeta^t{}_{inc} = 0.5\,\mathrm{Pa}\,(\mathrm{kg\,m}^{-2})^{-1}$, $\zeta^t{}_{dec} = 0$. The colour scale is the same as in Figure 4.3(a).

several years, pushing thresholds much higher than had been seen previously in the model. The threshold stress field that had formed by end of this procedure is shown in Figure 6.5, which can be directly compared to Figure 4.3(a). The strength of lifting in northern midlatitudes during autumn can be inferred from the huge increases in threshold stress (to more than 0.05 Pa in parts, on the north and east flanks of Alba Patera and in Acidalia) that were caused over several model years. Note that this threshold field, which should be qualitatively equivalent to a map of inverse dust surface density if the possible existence of reservoirs carried over from historic climate regimes is neglected, is still in agreement with the DCI map of Figure 4.3(b) over most of the area for which data are provided; exceptions are in northern Tharsis, where dust is believed to be more abundant than the model would suggest, and at 60°S, where the reverse is true, perhaps due to the lack of model SH cap-edge lifting.

Using this method, pre-solstice dust lifting was reduced to more reasonable levels. Achieving a realistic, variable multi-year opacity cycle like that seen in Chapter 4, however, may require significant changes to the lifting parameters, so was not attempted here. Instead, a series of 'low dust' years, containing some small dust storms, was obtained by continuing the run described here, with $\zeta^t{}_{dec}$ reactivated. The aim in this chapter was not

to examine further the factors affecting dust storm growth and variation, but to investigate on a simpler level the effect of ice clouds on a transported dust field, for which it was required that the spatial and temporal distribution of the dust, pre-scavenging, be as realistic as possible.

6.2.3 Water cycle results with transported dust

Producing a more sensible dust opacity output, the model was now able, for the first time, to convincingly simulate both the dust and water cycles at the same time. The replacement of the prescribed MY24 dust scenario with a transported dust field meant that different forcings were presented to the cloud scheme, in the form of atmospheric temperatures, surface radiative fluxes and IN number densities. There are a number of potential benefits to using the transport method, such as the replacement of the prescribed vertical dust profile with model-generated profiles (including variation in particle size with height) and the replacement of the problematic winter polar prescribed dust opacities, but obtaining realistic cloud opacities is challenging due to the sensitivity of condensation to atmospheric temperature profiles, which are themselves heavily dependent on the transported dust distribution. As expected, there were a number of changes in the typical water vapour and ice outputs — these are shown, along with dust optical depth, in Figure 6.6, which can be compared to Figure 6.1 (obtained using prescribed dust).

Atmospheric temperatures

In order to most accurately reproduce the observed atmospheric state during northern spring and summer, dust devil lifting efficiency α_D was tuned to ensure that dust optical depths — at least at low latitudes — agreed with those of the assimilation dataset (shown earlier in Figure 3.7), as best as could be achieved. Although the model was then able to produce realistic column dust opacities over $L_s = 0$–$180°$, actual atmospheric temperatures were subject to the vertical distribution of the dust mass, and of dust particle sizes (since heating rate $\propto \frac{q}{r}$). Temperatures between latitudes 0° and 30°N, the location of the ACB, are plotted against those from the MY26 reanalysis, over $L_s = 0$–$180°$ (during which period there is little interannual variability between the reanalysis years), in Figure 6.7(a). Temperatures in the lowest 1–2 scale heights agreed well with the reanalysis, except during

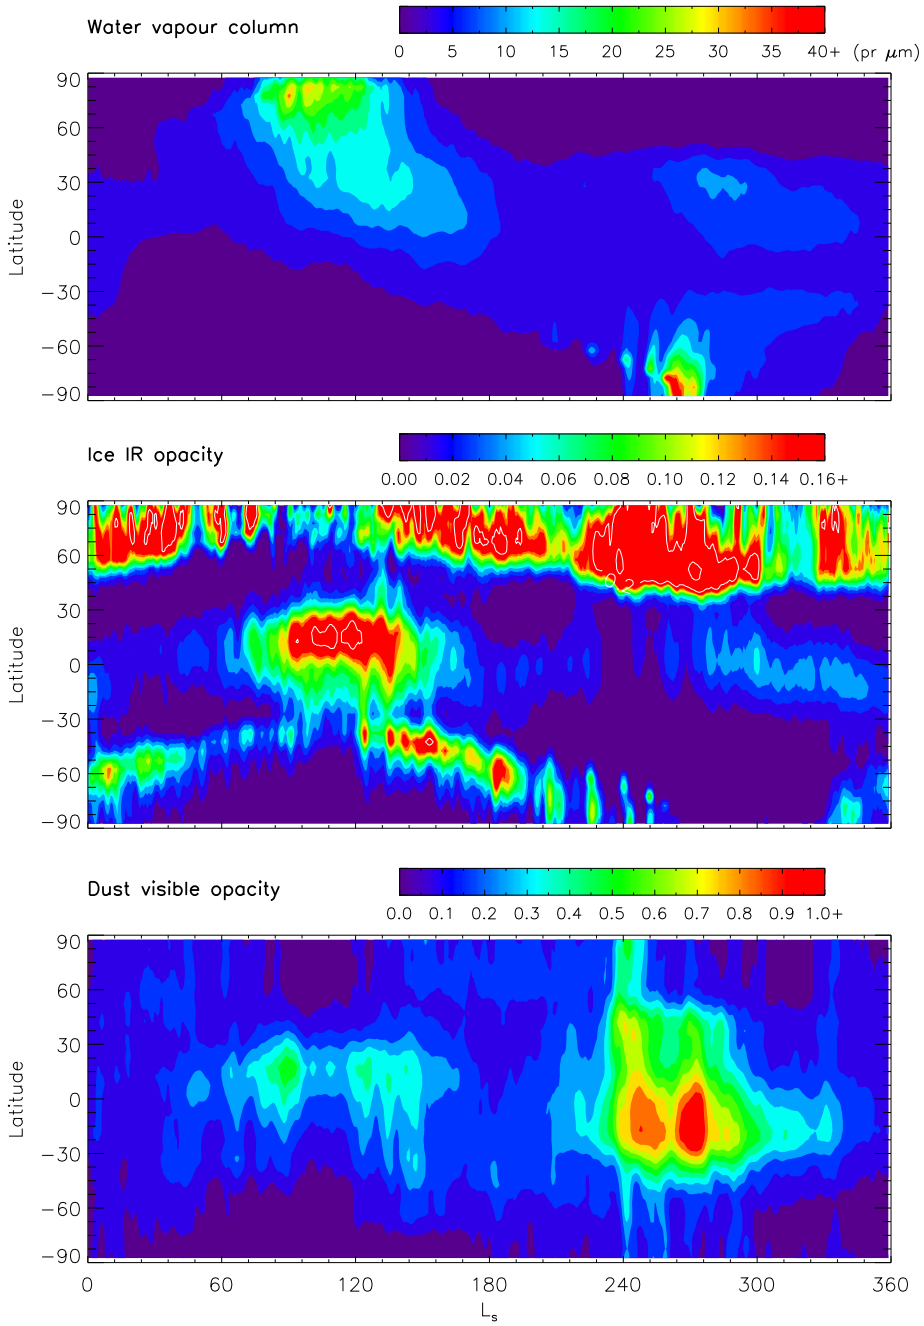

Figure 6.6: As Figure 6.1, but from a well-tuned simulation using active, transported dust and water ice clouds. The opacity fields have been smoothed using a 5-sol moving average.

early spring and in late summer, when increases in small-scale wind stress dust lifting were not adequately captured by the model. At the cloud level (~ 100 Pa), temperatures were on occasion a little low, but higher up (above 10 Pa) a temperature inversion was seen: the atmosphere was much too warm. This was also the case when the dust-transporting results were compared to the earlier run using prescribed dust, in which mass was distributed vertically using a Conrath profile.

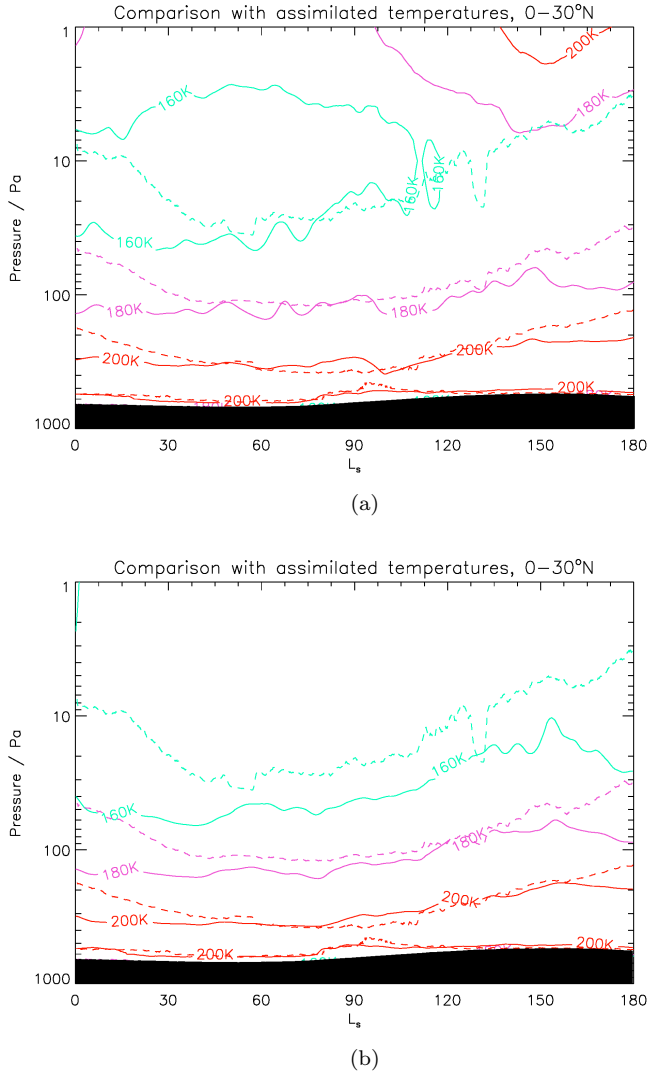

Figure 6.7: (a) Atmospheric temperatures, averaged over 0–30°N, from the model (solid), using size-independent dust scattering assuming a constant effective radius of 1.5 μm, and from the reanalysis of MY26 (dashed). Only three contour levels, at 160, 180 and 200 K, are shown. (b) The same from a model run using size-dependent dust scattering.

The reason for this was that a uniform extinction coefficient, corresponding to a dust

effective radius of 1.5 μm, was being used, regardless of actual particle size. In reality, the extinction coefficient varies slowly over most typical dust sizes, but falls off sharply for particles smaller than about 0.1 μm, so that heating rates will have been overestimated for small particles, such as are found at high altitudes. A more sophisticated scattering scheme, in which coefficients are calculated depending on the effective radius in each gridbox, is available in the newest version of the model physics, and upon its implementation, the high-altitude temperature inversions were removed (Figure 6.7(b)). However, the size-dependent scattering scheme led to further problems with the tuning of the dust lifting model, so was not retained. There was evidence, in both types of run, that dust was not being adequately mixed upwards at low latitudes (seen through the cold temperatures above 100 Pa in Figure 6.7(b), which formed despite column opacities being close to observed values) meaning that with size-dependent scattering, obtaining realistic temperatures at the cloud level required overestimating dust mass and temperature nearer the surface, which had the tendency of prompting further dust lifting. With the uniform scattering coefficient, the decrease in mixing ratio with pressure was partially offset by the overestimation of heating rate (as effective radius decreased), allowing more reasonable temperatures to be retained near the ground. The work throughout this chapter therefore used the simpler size-independent radiation scheme, with the knowledge that the unrealistic temperature inversions were confined above $\sim$ 10 Pa, so that their direct influence on dust lifting and vapour condensation and sublimation should have been fairly small.

The difference between the low latitude dust mass mixing ratio profiles obtained in the transport run and the one used in prescribed dust runs is shown in Figure 6.8, from which it is apparent that dust in the model tended not to be as well mixed up to 30–40 km in altitude as is assumed by the Conrath profile. Also shown are the corresponding model number mixing ratio profiles, including those filtered to include only particles larger than 100 nm, which are used to calculate $\mathcal{N}$, the number density of IN. In the prescribed dust approach, since particle radius is assumed uniform, the number density profile is identical to the Conrath profile. When transporting dust, the decrease in effective radius with increasing height lead to a rapid increase in number density, such that $\mathcal{N}$, around the ACB level ($\sim$ 20 km), was enhanced by as much as a factor of ten over the prescribed case.

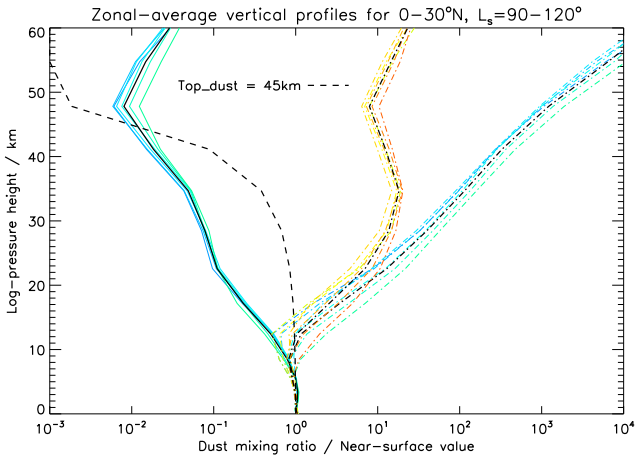

Figure 6.8: Zonal-mean dust vertical profiles from the model, averaged over 0–30°N and $L_s = 90$–120°: mass mixing ratio as solid blue-green lines, number mixing ratio as dash-dotted blue-green lines, and IN number mixing ratio (particles with radii > 100 nm) as dash-dotted orange-red lines. In each case the black solid or dash-dotted line shows the mean profile. The dashed line shows the Conrath profile, using the dust top height assumed by the model to be appropriate for this date and latitude.

Cloud ice and water vapour

Ice opacities in the aphelion cloud belt agreed well with the observed belt optical thickness, using the same ν_{sed} and R_s settings as before. The 100 nm minimum size criterion was applied when calculating $\mathcal{N}$, and it was found that this contributed a reduction to ACB thickness of around 20%. The belt formed later in the year than it did previously: it peaked at around $L_s = 100°$ but covered $L_s \sim 70$–$150°$, which actually agrees fairly well with TES observations, except for the start date. The main reason for the ACB's slow formation appeared to be a lack of water vapour at low latitudes from the beginning of the year onwards (compare the vapour plot to Figure 6.1). The present results were deemed to provide a reasonably accurate reproduction of observed ice abundances, but were obtained through a tuning process that was by no means an exhaustive exploration of the parameter space available: such a search would be lengthy and perhaps ill-advised, given our relatively poor knowledge of the physics governing both Martian cloud nucleation and surface dust emission.

Polar hood ice opacity was significantly improved in the southern hemisphere: the opacity decreased more rapidly at $L_s \sim 40°$, formed a minimum at winter solstice, then built up again post-solstice, from $\sim 120°$. This compares well with MCS data, though

observed IR cloud optical depths were generally lower (0.05–0.1 (Benson et al., 2010)) than those produced by the model. The post-solstice clouds formed due to an increase in vapour supply from lower latitudes, helped — in comparison to the prescribed dust case above — by lower atmospheric temperatures in southern midlatitudes at 100 Pa and below (where the clouds were located). Vapour was carried over the cap boundary by a combination of transient and stationary atmospheric waves; fluxes due to the former increased noticeably after $L_s = 120°$, and were much larger than in the prescribed case. The colder (by up to 10 K near the surface) atmosphere was caused by the lower local dust opacities in the transport simulation: the prescribed opacities of the MY24 scenario are quite possibly too large in southern winter mid- and high latitudes (at least over $L_s \approx 120$–$140°$), due to a lack of TES observational coverage in the region.

Changes were also seen in the thickness of the north polar hood. Ice opacities were reduced, generally, from their values in the prescribed dust run; however, there is clearly a strong dependence on global dust loading, so care is required when making the comparison between the two runs. The peak in opacity occurred at $L_s = 230$–$290°$, when the global dust loading was at a maximum and the PMOC was correspondingly strengthened; similarly, the peak in polar hood opacity in the prescribed dust run coincided with the regional dust storm of MY24. The transport run simulated a reasonable near-surface NH solsticial pause in transient RMS temperature — with a min/max ratio, as defined in Chapter 5, of 0.67 — but evidently this did not stop vapour from being transported over the polar cap edge around solstice, to thicken the polar hood.

Another important control exerted by the dust field with regard to the water cycle is on the size of the vapour maxima at the summer poles. This is dependent primarily on the incoming radiation flux reaching the polar surface, and therefore on the dust loading over the pole (surface albedo (unchanged since the prescribed dust results shown earlier) and atmospheric vapour abundance above the polar cap also play a role). Dust cover was lower over the northern summer pole when using transported dust, enabling an increase in total surface flux locally (the increase in shortwave flux outweighed a decrease in longwave flux), and an increase in the vapour maximum to $\sim 30\,\mu$m, closer to the observed value of $\sim 40\,\mu$m (Smith, 2008); however, the year was too dry, in general. In the southern hemisphere, the summer pole was significantly clearer in this particular year, causing an increase in flux.

This increased the vapour maximum, bringing it much closer to observations of the feature, but it still occurred too early in the season ($L_s = 260°$ vs $280°$). The northern midlatitude vapour peak seen previously was no longer produced, presumably due mainly to the lack of a significant rise in pre-solstice dust opacity.

Outstanding issues

The results of the tuned dual-transport simulation were encouraging, and in several respects the vapour and ice behaviour showed improvement over that seen when using prescribed dust. However, several issues with the model's performance remained, the most obvious of which was the continued need for an unrealistically large value of the ice sedimentation variance ($\nu_{sed} = 0.45$, as before) in order to keep ACB optical depths in line with observations. This necessity may have stemmed from a number of model deficiencies, but it suggested most simply that clouds were being formed too readily. Secondly, as is discussed further later in this chapter, the ACB sat too low in the atmosphere for most of its lifetime, and dust mixing ratios also peaked too near to the ground around aphelion. Finally, and most relevant for the dust transport focus of this project, the dust mass loading in both winter polar regions appeared to be too large, again in comparison to recent MCS observations. It was hoped that the introduction of the scavenging interaction could address some of these concerns, since the effect could broadly be expected to make cloud formation more difficult, through a reduction in atmospheric dust mass and number density.

6.3 Introducing dust scavenging

6.3.1 Physical basis and theory

So far, the interaction between ice and dust particles had been in one direction only: dust particles served as ice nuclei for cloud formation but were not directly affected in any way themselves. This was inaccurate, as it meant that IN remained in abundant supply during the onset of ice nucleation, even as the dust particles present became integrated within the ice that was forming. This omission might have had the effect of overestimating ice crystal growth rates, due to overly large IN number densities. Another problem with the method was that the dust column was not reduced during ice nucleation; in reality, it

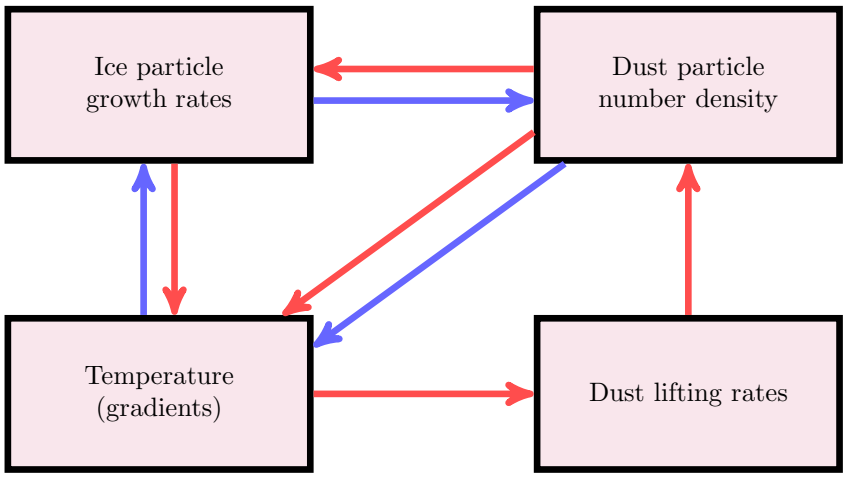

Figure 6.9: A schematic view of the mechanisms involved in the coupling of dust and water ice. A red (blue) arrow from box a to box b indicates that an increase in a will cause an increase (decrease) in b. Arrows of different colour in opposing directions thus indicate a negative feedback, while the continuous loop of red arrows between dust lifting rates, dust particle concentration and temperature signifies the positive feedback on wind stress dust lifting identified in earlier chapters.

can be assumed that a dust particle at the centre of a larger ice crystal will have little or no radiative impact on the surrounding atmosphere, and would thus be 'hidden' from the point of view of the radiation scheme. Finally, a dust particle encased in ice could have a rather different sedimentation velocity to one settling in an ice-free region; therefore, ice nucleation around dust particles may alter vertical dust profiles, with a potential feedback on ice profiles.

The schematic of Figure 6.9 gives an idea of the extent of the coupling between dust and water ice. Ice particle growth rates and dust lifting rates are connected by atmospheric temperature and dust number density. Notice that arrows of both colours are drawn from 'dust particle number density' to 'temperature (gradients)': while temperature is (locally) always increased by the presence of dust, lateral temperature gradients may be increased or decreased depending on the location of the dust heating. The impact on dust lifting rates of reductions in the dust column due to cloud formation is, therefore, unclear at this point.

The removal of aerosol or dust particles from the atmosphere through their interaction with clouds is known as scavenging (e.g. Pruppacher and Klett, 1978), and the subsequent arrival of the particles at the surface is referred to terrestrially as wet deposition, in contrast to dry deposition, in which the particles fall to the surface under their own weight. Dust/aerosol scavenging has received considerable recent attention in terrestrial climate re-

search due to its effect on total aerosol deposition rates (e.g. Sportisse, 2007). It can occur in two different ways. The first, in-cloud scavenging, is the result of heterogeneous cloud nucleation, as described above. On Earth, this can apply to liquid, ice and mixed-phase water clouds. The second mechanism is below-cloud scavenging, which is the removal from the atmosphere of aerosol particles encountered by precipitating liquid drops or ice crystals. On Mars, the situation is simplified somewhat as only ice cloud processes need be considered; however, these are less well understood than the liquid water equivalents, perhaps as they are generally of secondary importance on Earth. Martian dust scavenging can be further divided into removal by H_2O ice and by CO_2 ice, both of which are potentially significant.

In-cloud scavenging by ice clouds can be described simply by a scavenging efficiency R_s, which is the fraction of available dust (or aerosol) particles that act as IN. For an ice mass growth rate of $\frac{dM_c}{dt}$, the rate of change of 'interstitial' (not cloud-borne) dust mass mixing ratio q_{int} — equal in magnitude but opposite in sign to the rate of mass mixing ratio increase within the forming ice crystals — is

$$\frac{dq_{int}}{dt} = -R_s \frac{q_{int}}{M_w} \frac{dM_c}{dt}, \tag{6.2}$$

where M_w is the mass of water vapour in the gridbox. The maximum possible value for R_s is 1, which implies that all nuclei of sufficient size in the ice-forming part of a gridbox will be scavenged. Stier et al. (2005) used such a method in a terrestrial GCM, and used a value for R_s of 0.1 (taken from an experimental study) for stratiform ice clouds. Size-dependent R_s were used for liquid and mixed-phase clouds, but there were evidently insufficient experimental data to use anything other than an R_s that was uniform across the various aerosol size modes used in the model.

Below-cloud scavenging is better understood theoretically, and several expressions have been formulated to describe the efficiency of the process, for specified sizes of precipitate and aerosol particles. It is known to vary strongly with dust particle size, as collection is dominated by Brownian diffusion for small particles and inertial impaction for large particles, leaving a characteristic minimum in scavenging efficiency for particles intermediate in size (Martin et al., 1980). The size dependence of the collection efficiency E can be

modelled by the theoretical form of Dick (1990):

$$E = \frac{mu_t}{6\pi r\eta r_c} + 4Pe^{-1}(1 + 0.4Re^{1/6}Pe^{1/3}) \tag{6.3}$$

where m is the dust particle mass and r is its radius, u_t is the ice or snow particle's (terminal) fall velocity, η is the absolute viscosity of air, Re is the Reynold's number for the ice particle, $Re = \frac{\rho r_c u_t}{\eta}$, and Pe is the Péclét number for the dust particle, $Pe = \frac{2r_c u_t}{D}$, where D is the diffusivity. The first and second terms in (6.3) measure the efficiency of collection via collision and diffusion, respectively. A variation of 2–3 orders of magnitude is predicted for dust sizes of 0.01–10 μm. The collection efficiency is then converted to a scavenging coefficient R_{sbc} (with dimensions of L^2M^{-1}): $R_{sbc} = \frac{\pi {r_c}^2}{M_c}E$, so that the rate of change of interstitial dust mass mixing ratio can be written

$$\frac{dq_{int}}{dt} = -R_{sbc}q_{int}\mathcal{F}_{ice}, \tag{6.4}$$

where $\mathcal{F}_{ice}$ is the vertical ice flux.

The importance of below-cloud scavenging depends very much on the sizes and densities of the particles involved. For water ice crystals on Mars, downward fluxes are so low (typically $\sim 10^{-10}\,\mathrm{kg\,m^{-2}\,s^{-1}}$) that the process can be neglected, even with 100% collection efficiency. CO_2 fluxes in the condensing winter polar atmosphere can, however, be significantly larger, making below-cloud scavenging potentially important. Typical CO_2 snow particle characteristics are not known with much certainty, but reasonable estimates would be a radius of $r_c \sim$ 10–100 μm (Hayne and Paige, 2009) and a density of $\sim$ 100–1000 $\mathrm{kg\,m^{-3}}$ (Smith et al., 2001; Matsuo and Heki, 2009). Martian CO_2 snow particles are therefore smaller and more dense than the terrestrial particles of radius 500 μm and density $\sim 50\,\mathrm{kg\,m^{-3}}$ used by Dick (1990). Both of these differences are such as to give larger collection efficiencies for the Martian snow particles, at least for dust sizes of $\sim 0.1\,\mu$m. In fact, the larger D and u_t values for Mars mean that for a given ice particle size and density, both terms in E are larger, so that the below-cloud scavenging process should be more efficient, by up to an order of magnitude, than it is on Earth. However, if the snow particles are too small, their velocity becomes comparable to the dust particles' fall velocities, reducing the efficiency of impaction scavenging and, in fact, invalidating the theory used: dust particles

of radius 1 micron will only be efficiently scavenged if CO_2 ice particle radii are $\gtrsim 50\,\mu$m.

6.3.2 Implementation in the model

By water ice

The in-cloud scavenging interaction between dust and water ice was modelled by adding another pair of dust tracers, using the two-moment transport scheme, to carry mass and number mixing ratios of scavenged dust. A similar approach was recently used in a terrestrial GCM by Hoose et al. (2008), but in fact the use of extra tracers will be seen to be more important on Mars, because scavenged dust particles spend significant periods of time in the atmosphere, within the cloud ice, rather than being rapidly removed by precipitation, as on Earth. Dust mass was transferred between the two tracer pairs, interstitial and scavenged, during water vapour condensation and ice sublimation. For condensation, dust mass transfer followed equation (6.2), while for sublimation, scavenged dust mass was released in proportion to the fraction of the ice mass that was condensing during the timestep (i.e. $\propto \frac{dM_c}{M_c}$). The latter choice was made assuming that ice particles typically grow over time and scavenge additional dust as they do so, rather than simply nucleating around one ice nucleus initially. Ice particle radii were calculated by dividing the ice mass among the total IN present, both interstitial and scavenged[2]. As before, only dust particles larger than 100 nm were allowed to act as IN, and the resulting number densities were multiplied by the scavenging efficiency R_s to produce the IN density for use by the cloud scheme. Recall that the reference runs above used $R_s = 0.22$, but neglected the feedback on the dust field.

Sedimentation rates were calculated for the mixed dust-ice particles, which contain a dust mass mixing ratio q_{scav}, using their overall density, which can be expressed as

$$\rho_{mixed} = \rho_{ice} + \frac{\rho_{dust} - \rho_{ice}}{\frac{\rho_{dust}}{\rho_{ice}}\frac{q_{ice}}{q_{scav}} + 1} = \rho_{ice} + \left(\frac{r_{scav}}{r_{mixed}}\right)^3 (\rho_{dust} - \rho_{ice}).$$

The same velocities were applied to both the ice tracer and the scavenged dust tracer pair, so that the two fields were transported identically, in both the horizontal and vertical directions. The extent to which the vertical distribution of dust mass is affected by scavenging

[2]It would be more physical to use only the in-cloud IN number, but numerical problems were encountered using this approach, where small numbers of scavenged IN were present. The use of total IN does at least mean that scavenging and non-scavenging runs with the same R_s are consistent, whereas ice particle radii would have differed had they been calculated using scavenged IN only in the former case.

depends on how much the sedimentation velocities for the mixed particles differ from those for interstitial dust particles. For dust particle radii of around 5 μm and above, fall velocities can be approximated as $w_{sed} \propto \rho_p {r_p}^2$, meaning that for scavenged dust to fall more rapidly than interstitial dust requires $\left(\frac{r_{scav}}{r_{mixed}}\right)^2 < \frac{\rho_{mixed}}{\rho_{dust}}$. With $\rho_{ice} = 920\,\mathrm{kg\,m^{-3}}$ and $\rho_{dust} = 2500\,\mathrm{kg\,m^{-3}}$, this is found to be true for all possible values of $\frac{r_{scav}}{r_{mixed}}$, i.e. regardless of the dust:ice mass ratio $\frac{q_{scav}}{q_{ice}}$ in the mixed particle. The enhancement to the fall velocity is greatest when the mass ratio is small: for a ratio of 0.05, the fall velocity is increased by a factor of more than 5. However, dust particles of this size generally have short atmospheric lifetimes anyway, due to their large radii and interstitial fall velocities.

For more common dust sizes of a few microns or less (but larger than 0.1 μm, the minimum size that can be scavenged), fall velocities become better approximated by (see Equation (2.2)) $w_{sed} \propto \rho_p r_p$: this results in a condition for increased dust velocity that is only satisfied if $\frac{r_{scav}}{r_{mixed}} < 0.42$, which corresponds to a dust:ice ratio of less than 0.22. This means that dust will be scavenged downwards only when it forms a relatively small fraction of the dust-ice mixture. Although dust is much denser than water ice, it is the larger radius experienced as part of an ice-rich mixture that gives it a larger fall velocity. When the mass ratio is more than 0.22, the dust particles actually fall more slowly than they would interstitially, as they are accompanied by the lower-density water ice. For a dust:ice ratio of 0.5, fall speeds are decreased by around 10%, while for a ratio of 0.05 they are increased by $\sim 50\%$. Therefore, the effect of scavenging on the dust vertical profile will depend very much on how much dust becomes incorporated into the cloud ice. The fact that large increases in settling velocity are only possible when small amounts of dust are being scavenged would seem to place a limit on the extent of the vertical dust redistribution that may occur through scavenging.

By CO_2 ice

Scavenging by CO_2 required a different approach, as the model does not explicitly simulate CO_2 ice clouds. Rather, condensation rates are calculated at each timestep, and the ice mass that forms is assumed to fall vertically and either resublime at a warmer level or reach the surface as ice, within the timestep. No additional tracers were used for this parameterisation. In-cloud scavenging by CO_2 followed equation (6.2), again using only particles

larger than 100 nm, but the mass lost from interstitial dust was either released again at a lower level, or deposited onto the surface, depending on the behaviour of the falling ice. As with water ice, dust release during sublimation occurred according to the fraction of the ice in the gridbox that sublimed, rather than only upon the complete disappearance of the ice. An additional below-cloud scavenging tendency was calculated as the CO_2 descended through the atmosphere, using equation (6.4). Dust scavenged in this way was assumed to become evenly distributed throughout the snow particles, and was then allowed to resublime at lower levels in the same manner as the dust scavenged by the in-cloud mechanism. All sizes of dust particle were available for below-cloud scavenging.

Choice of efficiencies

The greatest amount of uncertainty in the accuracy of the parameterisations arose from the choices of the scavenging efficiencies, denoted by R_s for in-cloud scavenging and by R_{sbc} for below-cloud scavenging. R_s were varied between 0.01 and 1, in accordance with their definition (a value of zero would have allowed no cloud formation, so was rejected). R_{sbc}, applying only to CO_2 ice, was even less well constrained: as discussed above, the value calculated using equation (6.3) varies over several orders of magnitude depending on the CO_2 snow particle properties used and the size of the dust being scavenged. For simplicity, several possible size-independent R_{sbc} values were used, from $0.1\,\mathrm{m^2\,kg^{-1}}$ to $10\,\mathrm{m^2\,kg^{-1}}$ (corresponding to collection efficiencies of $E \approx 0.001$–0.1), which covered the range of likely values for ice radii of 50–500 microns and dust radii of 0.1–10 microns.

6.4 Testing with passive dust

Radiatively active dust lifting runs are sensitive to changes in the near-surface atmospheric state, so that, as will be seen in subsequent sections, altering the scavenging parameters can dramatically affect dust lifting rates. This then has a knock-on effect on cloud formation rates, due to the changes in IN concentrations in the atmosphere. To remove this complication in the first instance, several year-long simulations were performed using radiatively passive dust, along with active clouds, IN for which were obtained from the prescribed dust opacity scenario used (MY24). Using several illustrative scavenging efficiencies, the spatial distributions of the resulting dust fields were studied, and compared to that of a control

experiment which neglected the scavenging feedback on the dust field. The same lifting settings were used for each run, and were chosen to provide plausible lifting rates, as far as is possible in a passive dust simulation. These simulations therefore provided information primarily on the effect of scavenging on the transport of dust particles, and not on its further influence on the thermal structure of the atmosphere and on dust lifting rates. They also served as a means of estimating the relative importance of each type of scavenging process.

6.4.1 Water ice cloud radiative effects only

Before activating any dust scavenging, a run with passive dust and without active ice clouds was performed, to examine the effect of the clouds themselves on dust transport (following on from some of the results presented in §6.2). As expected, dust mixing ratios in equatorial regions were reduced when clouds were removed, as this weakened dust lifting rates. However, large increases in dustiness were produced in the south polar region during winter ($L_s = 45$–$135°$). This occurred due to greater eddy activity at the cap edge, to the north of the region, which mixed dust down the meridional mixing ratio gradient — set up since no dust lifting occurs over the polar cap — into the polar region. A comparison of the dust mass in each polar region (beyond 60°) for the two simulations is shown in Figure 6.10. In both cases, individual peaks in south polar dust mass over $L_s = 0$–$180°$ could be matched to peaks in transient eddy meridional velocity in the midlatitudes, which led to bursts of dust lifting. The general decrease in south polar dust loading during winter when clouds were used, by factors of $\sim$ 2–10, occurred due to the formation of a deeper solsticial pause in transient eddy activity, as seen earlier in Figure 6.3.

Flux calculations showed that transport across the south polar cap edge occurred mainly via stationary waves, in the lowest 10 km of the atmosphere. Murphy et al. (1995) also found stationary waves to be the dominant transport mechanism for dust in the southern high latitudes. Mixing into the south polar region was particularly inhibited by the addition of water ice clouds at $L_s = 60$–$120°$, around winter solstice, as stationary waves near the surface also became weaker at this time (in addition to the weakened transient waves). This reduction in eddy mixing can be seen in Figure 6.11, which shows (in solid and dashed contours) the residual mean streamfunction $\overline{\chi}^*$ at $L_s = 90°$ for the simulation with active

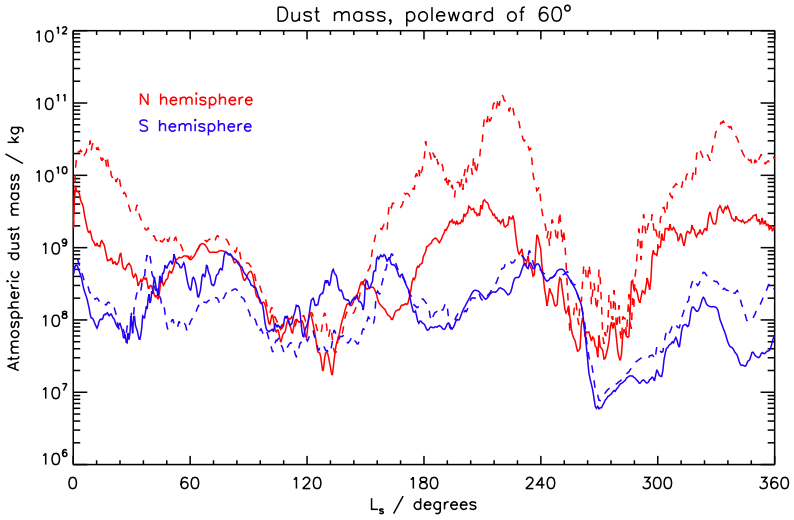

Figure 6.10: Total atmospheric dust mass (in kg, on a log scale) in each polar region, 60–90°N (red) and 60–90°S (blue), for runs without (solid) and with (dotted) active water ice clouds.

clouds. $\overline{\chi}^*$ is defined at a log-pressure height Z' as (Andrews et al., 1987)

$$\overline{\chi}^* = -\frac{2\pi R}{H}\frac{p_0}{g}\cos\phi \int_{Z'}^{{Z'}_T} e^{-Z/H}\left(\overline{v} - {\rho_0}^{-1}\left(\frac{\rho_0\overline{v'\theta'}}{\overline{\theta}_Z}\right)_Z\right) dZ, \tag{6.5}$$

where H is the atmospheric scale height, p_0 is the reference pressure (700 Pa), ρ_0 is the basic density, ${Z'}_T$ is the height at the top of the domain, and a subscript Z denotes a derivative with respect to Z. Tracers should approximately travel along $\overline{\chi}^*$ streamlines, as they are calculated including both mean and eddy transport components. Shown in colours is the difference in $\overline{\chi}^*$ between the runs with and without clouds. It can be seen that the addition of clouds strengthened the PMOC below 30 km, but it weakened the anticlockwise circulation near the ground at 60°S. This reduced the rates of transport of dust to higher latitudes, making the south pole more dynamically isolated at this time. The weakening in the SH included a contribution from the mean circulation (the first term in (6.5)) at ~ 45°S.

Transport into the NH winter polar region occurred through a combination of the zonal-mean circulation (dominant at $L_s \sim 180$–240°) and travelling waves (of greatest importance ~ 270–330°). The two very large peaks in north polar dust loading seen above for the active cloud case formed due to the pre- and post-solstice bursts of NH lifting identified in §6.2.2 as being commonplace in simulations with unlimited surface dust.

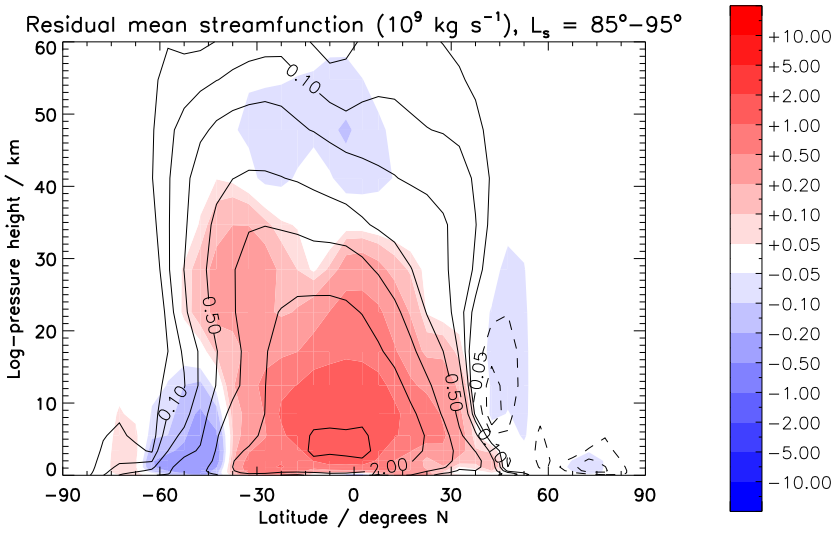

Figure 6.11: Transformed Eulerian mean streamfunction $\overline{\chi}^*$ (in $10^9\,\mathrm{kg\,s^{-1}}$) from L_s = 85–95° of the active cloud case (contours) and the change relative to the no-cloud case (shading). Positive values represent anticlockwise motion.

So, in both hemispheres, it was clear that the radiative influence of water ice clouds can strongly affect the quantities of dust mass in the polar region. In the SH, the presence of clouds reduces the polar dust loading around winter solstice by weakening stationary wave and mean circulation transport near the ground, and by suppressing transient eddies, which reduces the cap-edge dust storm frequency.

6.4.2 Scavenging by water ice

With clouds once again radiatively active, in-cloud water ice scavenging was used for the next simulation, and the results were compared to the control, which neglected scavenging (Figure 6.12). The value used for R_s was 0.22, for consistency with the treatment of IN (calculated from the MY24 dust scenario in both cases) by the cloud scheme, and as the best guess at this stage. Reductions in dust mass were seen in both winter polar regions, which this time was due solely to removal of interstitial dust through scavenging, as lifting patterns were identical. In the SH, the region of dust removal followed the size of the CO_2 cap, and reached ~ 40°S at winter solstice. The actual scavenging took place at the location of the polar hood, along the cap edge, reducing the flux to higher latitudes.

With the efficiency used, scavenging reduced south polar (60–90°S) dust mass by a factor of around 3–10 at $L_s \sim 50$–130°, and by a factor of 2 or less outside of this range.

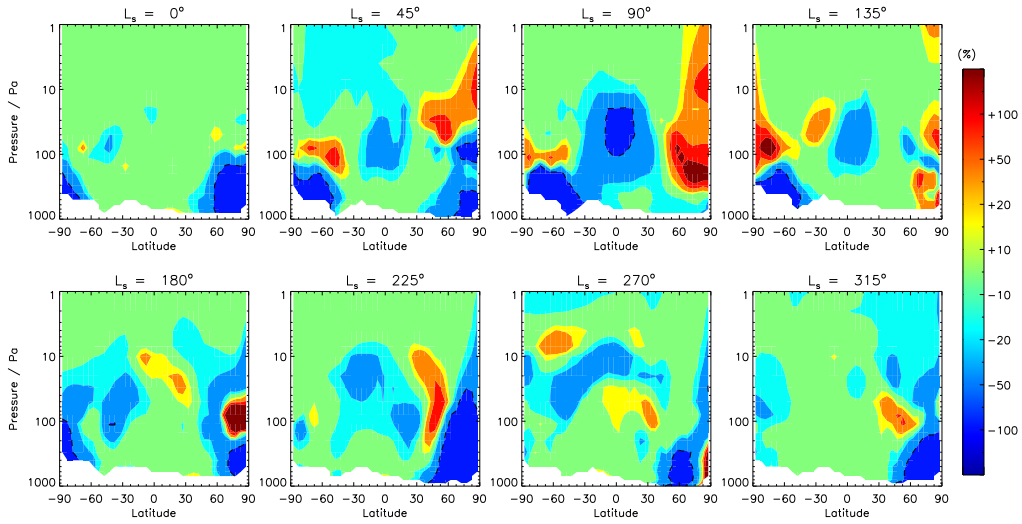

Figure 6.12: Percentage change in zonal-mean dust mass mixing ratio from the run with no scavenging to the run with H_2O ice scavenging with an efficiency of $R_s = 0.22$. Each panel is averaged from 10° before to 10° after the date in L_s shown, except for the first, which covers 0–10°. The lowest contour level shown is actually −50%, and is marked with a dotted line.

So, with an efficiency of 0.22 or larger, scavenging by water ice clouds should be at least as important as the radiative effect of the clouds in keeping the south polar region dust-free around winter solstice. Scavenging in the NH winter polar region was strongest when there was an abundance of dust available to scavenge, which occurred during the two large peaks in lifting, pre- and post-solstice. Scavenging by the aphelion cloud belt can be seen at $L_s = 45$–$135°$. Dust was primarily removed within and above the cloud layer, which sits just below 100 Pa. This, then, is evidence of some 'capping' of the dust column by the ice clouds. Although only a fifth of the dust particles were available for scavenging at each timestep, overall mass reductions in the equatorial region of more than 50% were possible (at $L_s = 90°$). Some increases in dustiness can be seen higher up in the two polar regions over $L_s = 45$–$135°$, which must have been caused by changes in transport patterns induced by the vertical redistribution of dust through scavenging.

6.4.3 Scavenging by CO_2 ice

The final two simulations both neglected water ice scavenging and instead considered scavenging by CO_2 ice clouds. The first of these used only in-cloud scavenging, for which the maximum scavenging efficiency of $R_s = 1$ was chosen. Despite this choice, the effects of

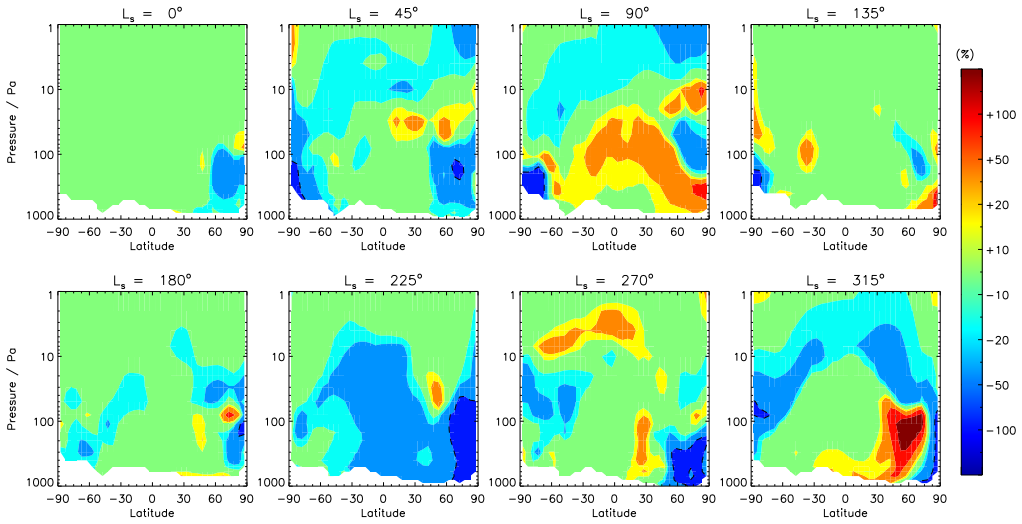

Figure 6.13: As Figure 6.12, but showing the change from the run with no scavenging to the run with below-cloud CO_2 ice scavenging with an efficiency of $R_{sbc} = 1\,\mathrm{kg\,m^{-2}}$.

in-cloud CO_2 scavenging were relatively minor: dust mass reductions occurred only at high latitudes, and did not exceed 10% of the control values.

The second simulation used only below-cloud CO_2 scavenging, using a value of $R_{sbc} = 1\,\mathrm{m^2\,kg^{-1}}$ that is intermediate in the range of possible scavenging coefficients estimated previously. As Figure 6.13 shows, this form of CO_2 scavenging is potentially much more effective at removing atmospheric dust. Large reductions in mass mixing ratio were produced at both winter poles, particularly at the north pole at $L_s = 225°$, where dust was readily available for scavenging. Interestingly, the effects of scavenging extended across the whole planet, even though CO_2 snow is largely limited to high latitudes (except for small amounts of mesospheric equatorial clouds that form above 10 Pa over L_s=0–180°, scavenging and releasing small quantities of dust without greatly affecting the dust mass below). Most of the changes away from the poles come as the result of the removal at the poles of dust that would otherwise have subsequently been advected to lower latitudes. This is seen most clearly at $L_s = 225°$, from which it can be inferred that of the dust raised at the north polar cap edge at this time, some first travels northward, penetrating some way into the condensing polar region, and then moves southward to the equator (and beyond), if it has avoided being scavenged while in the polar region.

6.5 Scavenging simulations with active dust — CO_2

Dust particle scavenging was then activated in a set of fully fledged simulations that included radiatively active clouds and dust. The first number of these were focused on assessing the possible effect of below-cloud scavenging by CO_2 ice. A control simulation, which did not include scavenging, was the same as was presented in §6.2.3. The same initial conditions (including those for threshold surface stress) were then used for the scavenging runs, each of which used a different value for R_{sbc}: respectively, 0.1, 1 and 10 $\mathrm{m^2\,kg^{-1}}$. The impact on dust loading and temperatures away from the polar regions was non-negligible, as suggested by the earlier passive test run. This included a reduction to the summer dust opacity peak at low latitudes of up to 30% (in the strongest case, $R_{sbc} = 10\,\mathrm{m^2\,kg^{-1}}$) and, as a consequence of the less dusty atmosphere, a reduction in north polar hood ice and an increase in the south polar vapour maximum.

Figure 6.14(a) shows that the annual mean dust wet deposition rate due to CO_2 scavenging increased as R_{sbc} was increased. The relationship between the two was close to linear. In both hemispheres, wet deposition rates decreased in magnitude moving away from the pole; however, the distributions about the poles were asymmetric, especially in the southern hemisphere, where two lobes of high deposition rate were seen, centered at 45°W and 135°E. Dry deposition, by contrast, displayed a wave-1 pattern, with a minimum south of the Hellas basin, at $\sim$ 60°E, 60–90°S, and in the area between the Hellas and Argyre basins at 45–60°S, simply matching the average distribution of dust mass near the ground. The different wet deposition pattern reflects preferred areas for CO_2 condensation, over the lowest 20 km of the atmosphere, which produce enhancements in wet deposition only. It should be noted, though, that the UKMGCM fails to predict the existence of the perennial south polar ice cap; therefore, the simulated climate of the region may be somewhat different to the true climate. In this context, results concerning deposition patterns in the region should be interpreted with caution. In the northern hemisphere, average wet deposition rates were an order of magnitude higher than in the southern hemisphere, due to the greater dust loading over the winter polar region. A wave-1 structure was seen in both wet and dry deposition rates, with a maximum at 0°E due both to colder near-surface temperatures (favouring greater ice condensation) and a similar peak in average autumn/winter dust optical depth. In the northern hemisphere, therefore, the addition of CO_2 scavenging

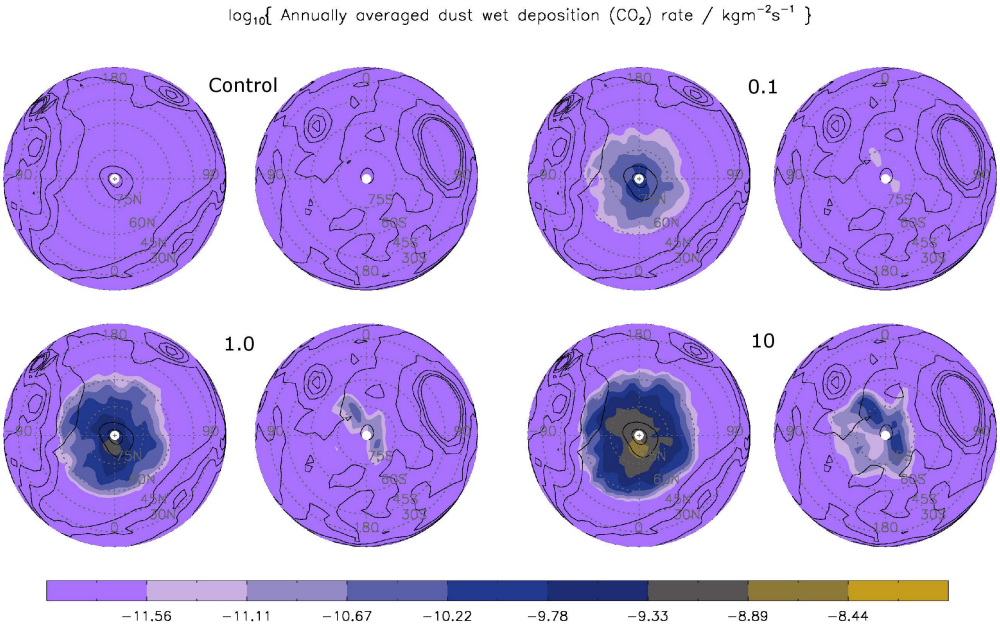

(a)

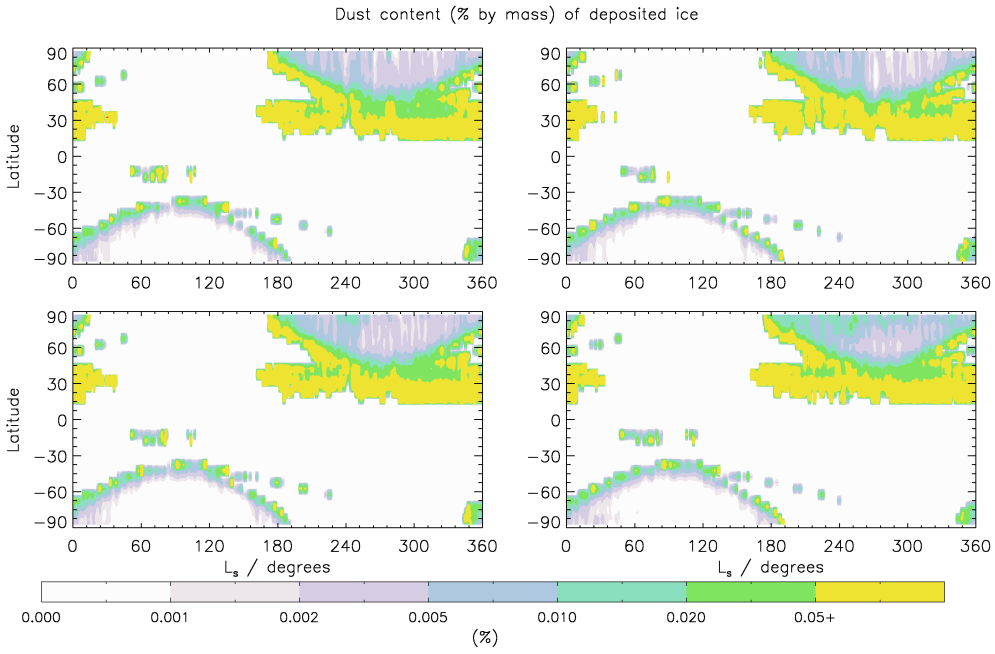

(b)

Figure 6.14: (a) Annually averaged wet deposition rates at the north and south poles from the four runs: the control case (upper left), and R_{sbc} values of $0.1\,\mathrm{m^2\,kg^{-1}}$ (upper right), $1\,\mathrm{m^2\,kg^{-1}}$ (lower left) and $10\,\mathrm{m^2\,kg^{-1}}$ (lower right). (b) Zonally averaged dust content (%) of condensing CO_2 ice, calculated by including both deposition processes, for the same simulations as (a).

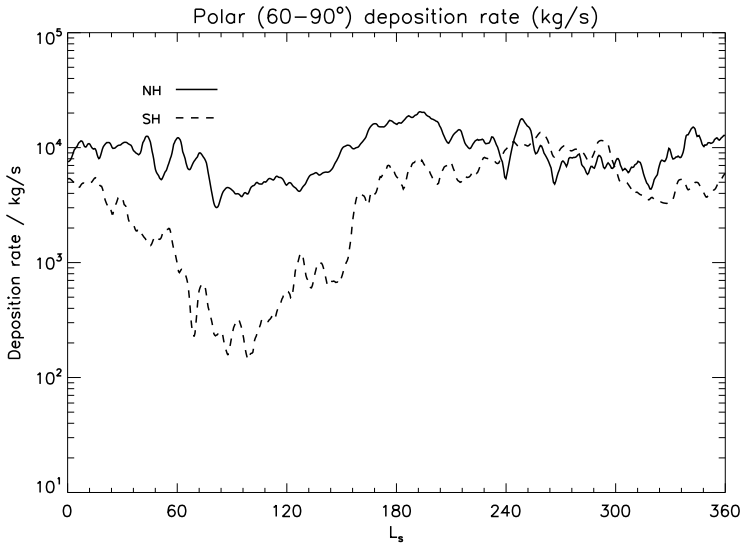

Figure 6.15: Dust mass dry deposition rates ($\mathrm{kg\,s^{-1}}$) over 60–90° in each hemisphere, from the control run.

did not alter the spatial distribution of deposited dust.

Wet deposition occurred significantly more slowly than dry deposition, except in the north polar region when R_{sbc} was large ($10\,\mathrm{m^2\,kg^{-1}}$). This can be seen through the marginal increases in the dust content of sedimenting ice that resulted as R_{sbc} was increased, shown in Figure 6.14(b). The exception was at 60–90°N, L_s = 180–270°, where dust content increased by up to a factor of ten. The plots also show the greater dust fraction produced in north polar cap ice, in comparison to the SH cap, and that dust penetrated more fully to the pole during northern winter (in the SH it was generally limited to the cap-edge region).

Polar dry deposition rates (similar in each of the four runs here, and shown in Figure 6.15) were altered more fundamentally by the inclusion of radiatively active water ice clouds. In comparison to the results of §4.4.2, the magnitude of the south polar deposition rate at its lowest point was unchanged, but rates were increased either side of winter solstice, helped by strong eddy activity during these periods, to create a pronounced solsticial minimum in polar dust deposition that appears to be in better agreement with atmospheric observations. Increases of up to an order of magnitude were seen at both poles throughout most of the rest of the year, signifying dustier polar regions in general. Particularly in the SH, this was a result of increased cap-edge lifting activity, so while the Chapter 4 simulations did feature clearer polar regions, they were also too clear in midlatitudes and therefore lacked

the necessary meridional gradient in dust concentration seen in MCS observations (§6.6.1). The north/south ratio in polar mass deposition was largely unchanged, as rates at both poles, averaged over the year, were increased by factors of $\sim$ 2–5 over those of Chapter 4. The mean dust contents of condensing CO_2 ice were of the order 10^{-5} kg/kg and 10^{-4} kg/kg for the north and south poles, respectively. This is still in rough agreement with the available observations, limited as they are.

6.6 Scavenging simulations with active dust — H_2O

Scavenging by water ice was then, separately, activated in a second set of simulations, again using the existing no-scavenging year as the control. The only difference between each of the simulations initially was the value of R_s: as in the passive case, the initial value used was 0.22, following the control simulation, and further runs were performed using values of 0.1 and 0.5 (as well as an incomplete year with $R_s = 1$). As water ice scavenging takes place over a much larger portion of the planet than CO_2 scavenging, its use affected more significantly the overall simulated climate. Dust lifting was strengthened in the scavenging runs, particularly at $L_s \sim 120$–$180°$, which in turn led to rapid cloud ice growth; intense dust lifting was also possible in the second part of the year, possibly sparked by increases in ice and dust lifting at the edge of the receding south polar cap. The magnitude of the problem increased with increasing R_s, such that the simulation with $R_s = 0.5$ was not particularly useful beyond $L_s \sim 120°$. Therefore, only limited insights could be drawn from the first year of each scavenging simulation through direct comparison to the control run. To complete the analysis, the 0.22 run was continued for three further years, using the variable threshold scheme to allow the surface dust field to adjust to the new climate produced through the inclusion of dust scavenging. Years with $R_s = 0.1$ and 0.5 were then run using the same initial conditions as the final year of the $R_s = 0.22$ series, so that the different values of R_s could be compared with one another.

Near-surface wind stress lifting rates are shown in Figure 6.16, for the control and $R_s = 0.22$ (first year) runs. With scavenging, dust lifting rates were increased at the north polar cap edge at $L_s = 0$–$60°$, at northern low latitudes at $L_s = 80$–$130°$, at the southern cap edge in early spring, and later in the year, when a GDS was produced. Once again, the sensitivity of the dust lifting scheme to changes in the basic model state is clear to see.

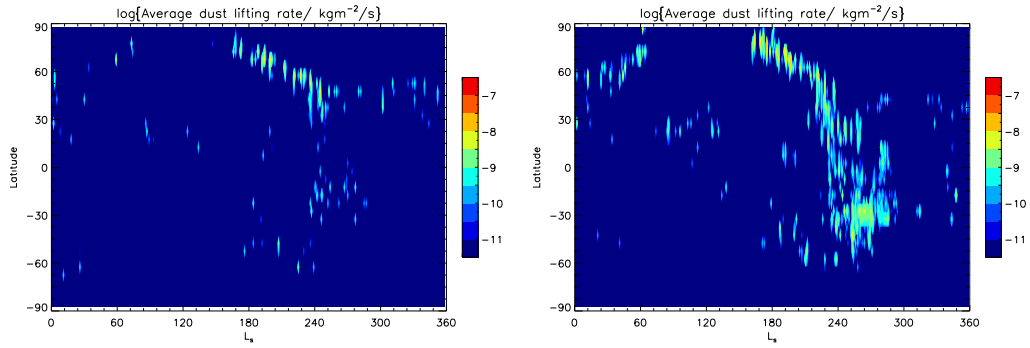

Figure 6.16: Logarithm of zonal-mean dust lifting rate for simulations without (left) and with (right) scavenging, using $R_s = 0.22$ in both cases. Initial conditions, including the threshold stress field, were the same for each simulation.

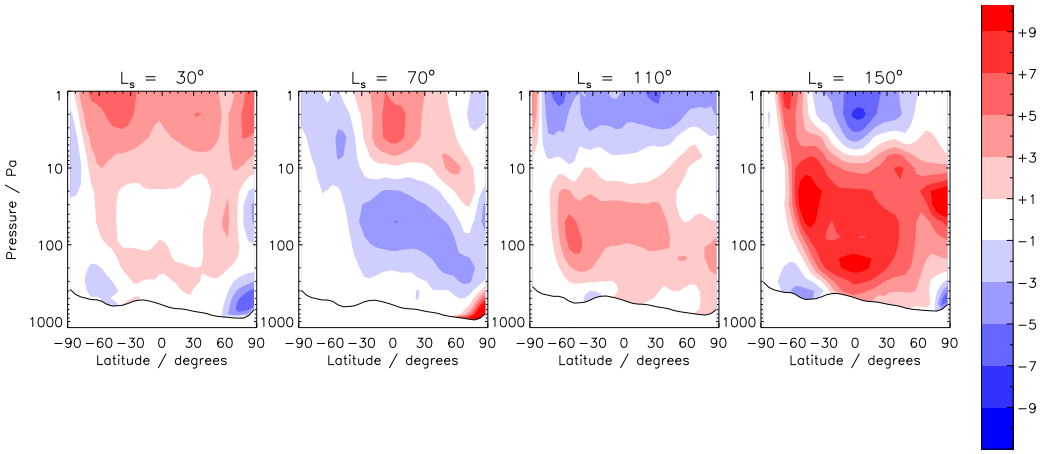

Figure 6.17: Difference in zonal-mean temperature (K) between the control simulation and the same year with scavenging activated, with $R_s = 0.22$.

The increases in lifting at 30°N over L_s = 80–130° were focused at Tharsis and Arabia, apparently in response to the high ice opacities that formed there as part of the aphelion cloud belt. The increased cap-edge lifting on the other hand was attributed to the clearing of dust from the polar regions through scavenging, which reduced temperatures in the polar regions. This is seen in Figure 6.17, particularly at the north pole at $L_s = 30°$, and at 40–70°S at $L_s = 150°$. Meridional temperature gradients were thus increased along the cap edges, which led to stronger travelling waves and enhanced dust lifting. Therefore, in answer to the question raised in Figure 6.9, it appears that the dominant effect of scavenging on dust lifting rates is to increase them, converse to what might have been expected following the removal of atmospheric dust and the associated decrease in temperatures.

These increases in dust lifting rates should be regarded as a real effect of the inclusion of dust scavenging — particularly if the efficiency is relatively high — and although lifting rates were reduced somewhat by the fourth year of the 0.22 run, due to the response of the threshold field, cap-edge lifting was still more intense than in the control simulation. Strong waves along the edges of the autumn/winter polar regions are a natural consequence of the sharp meridional gradients in dust mixing ratio and temperature that evidently exist in reality, and the fact that these gradients were presumably too weak in the southern winter polar region in the work of Chapters 3 and 4 (due to the neglect of both clouds and scavenging) offers a further reason for the weak SH cap-edge lifting produced in those simulations.

After three more years of running with the variable threshold scheme (with resupply turned off), the scavenging simulation with $R_s = 0.22$ acquired a threshold stress field that was able to prevent the large dust storm and, to some extent, the unwanted lifting at $L_s = 120°$, that were seen in the first year. The results of this fourth year, shown in Figure 6.18, thus moved back towards the control case, and were broadly realistic, with the exception of cloud ice at certain points. The dependence of ice opacity at the equator and in northern high latitudes over $L_s = 180$–$360°$ on extra-polar dust opacity is obvious, but polar hood optical depths were nonetheless too high (by perhaps a factor of 5) during the dust storm, and the spurious equatorial cloud ice at $L_s = 270$–$360°$ can be only partly explained by the low dust loading after 300°, which made the atmosphere too cold — temperatures at 270–300° were actually quite realistic. South polar hood ice was also overestimated, by a factor of at least 2 during $L_s = 120$–$180°$, and even moreso with $R_s = 0.5$. These features suggested cloud ice continued to form too readily, i.e. that R_s was too large and/or too much dust was present in cloud-forming regions. Direct comparisons between these dust opacity results and the control simulation of Figure 6.6 should not be made, owing to the different initial conditions and allowing for interannual variability in model dust lifting; one robust difference to note, however, is the lower dust opacity over the autumn/winter polar regions, particularly at the north pole, once scavenging was added.

The north polar vapour maximum seen in Figure 6.18 was larger than that shown earlier for the control run, bringing it closer still in magnitude to the observed vapour peak. The maximum was increased when scavenging was activated due to the radiative shielding (in

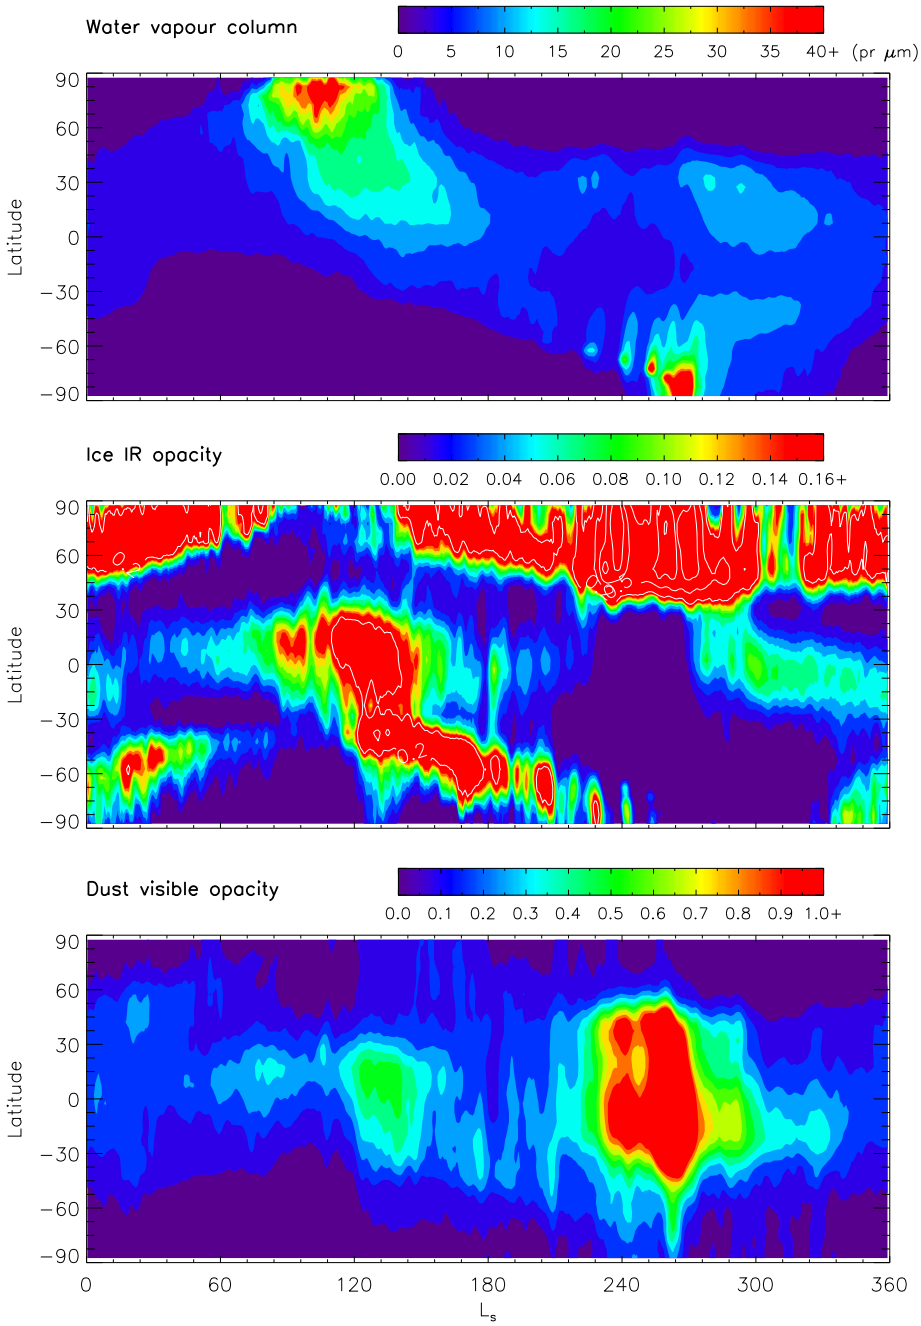

Figure 6.18: As in Figure 6.6, but for the fourth year of the simulation with scavenging activated, with $R_s = 0.22$.

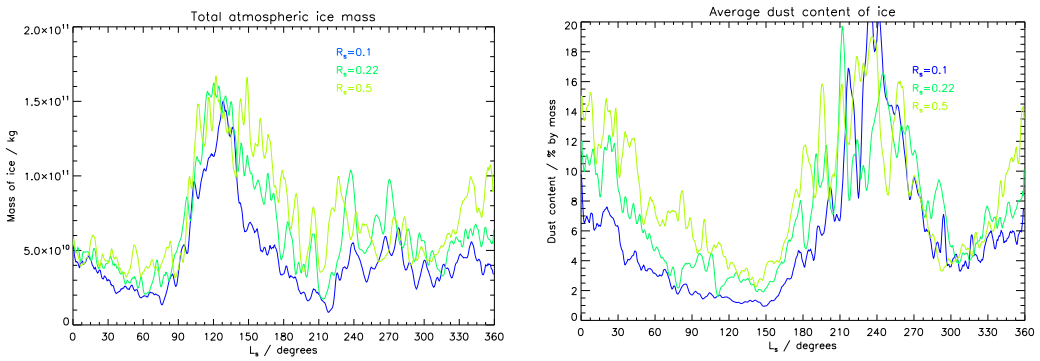

(a) (b)

Figure 6.19: (a) Total atmospheric ice mass (kg) for three different scavenging efficiencies. (b) Globally averaged dust content (as percentage by mass) of mixed particles, weighted by ice mass in each gridbox, for the same simulations.

addition to some accelerated removal) of dust over the pole — though polar clouds were fairly thin at the time — which led to an increase in the shortwave flux received by the cap. As before, the south polar maximum was strongly controlled by the dust storm activity around perihelion, which, as seen already, can be affected by the inclusion of scavenging. In this case, the timing of the main dust event was very similar to the control run, but a small amount of scavenging took place over the south pole around summer solstice, increasing the magnitude of the vapour maximum. The second peak in south polar hood ice from $L_s = 120°$ became increasingly large as R_s was increased, due to stronger eddy mixing of vapour from the tropics.

The sensitivity of the total atmospheric ice mass to the efficiency of dust scavenging used is plotted in Figure 6.19(a). Ice mass growth rates, which in the UKMGCM scheme are proportional to $\mathcal{N}$, were enhanced by the availability of more nuclei as R_s was increased; however, ice mass increased much less than linearly with increasing R_s, as interstitial dust particles were depleted roughly proportional to ${R_s}^2$, and since growth rates were ultimately limited by the availability of supersaturated water vapour. Other potentially important, competing effects occurring as R_s increases include (see Figure 6.9) increasing dust lifting rates, which would increase IN densities but also increase temperatures, probably hindering cloud formation, and decreasing temperatures due to removal of interstitial dust, which would promote further ice growth; also, the radiative effect of increasing cloud thickness (a slight local warming). The net result of all these on ice mass appears to be an increase with

increasing scavenging efficiency, but one which is limited, certainly by vapour availability, and perhaps also by a negative feedback arising from the interactions mentioned.

Globally averaged, mass-weighted dust content of the mixed (dust/ice) particles for each choice of scavenging efficiency is plotted in Figure 6.19(b), where it is seen that dust content increased with R_s: $\mathcal{N}$ increased, allowing more dust to be scavenged, and though ice growth rates also increased, they did so partly at the expense of growth that might have occurred at a later timestep anyway, as S became larger. The increase was less than linear in R_s, by the same token, but dust mass fraction showed greater sensitivity to R_s over $L_s = 0$–$180°$ than did total ice mass. Dust mass content formed a two-peaked pattern in L_s, with one maximum during the dustiest part of the year at 240°, and another at $\sim 0°$ when clouds formed in the presence of dust lifting in the two polar regions.

During $L_s = 0$–$180°$, dust mass content for $R_s = 0.22$ was higher in the southern midlatitudes (up to 15%) than it was in the ACB ($\lesssim 6\%$), due to a greater availability of dust and lower height of cloud formation. This fits with the observation of smaller ice particle radii at 0–45°S than in northern low latitudes by Clancy et al. (2003). The authors also noted that the two types of ice particle displayed distinctly different phase functions — the difference in cleanness of the water ice produced here is a possible explanation for this.

6.6.1 Dust in the polar regions

One of the important developments in recent Martian atmospheric dust observation has been the retrieval of concentrations over both polar regions throughout the whole year (at least to within one or two scale heights of the surface), by Mars Climate Sounder[3] (McCleese et al., 2010). A surprising feature of the data was the remarkably low dust density-scaled opacity (a quantity proportional to mass mixing ratio) in the winter polar regions, particularly at the south pole (60–90°S) around solstice. To date, dust transport models have not been able to adequately simulate the rapid drop-off in dust mixing ratio that occurs from low to high winter latitudes (Heavens et al., 2011). Reasons for this could include overestimated flux towards the pole from the principal meridional overturning cell or from eddy mixing, and the neglect of scavenging. A goal of this part of the project

[3]TES winter polar retrievals have also been made recently from limb profiles (McConnochie and Smith, 2008)

was to assess to what degree the third reason is important, and if it was, how efficient the scavenging process would need to be to produce such clear polar regions.

The MCS nightside data for MY29 are shown in Figure 6.20(a). The data have been collated and binned — at a horizontal resolution that is roughly comparable to the (T31) model resolution, and over 105 pressure levels — as part of an ongoing project within the LMD-UK collaboration to create new dust opacity scenarios for model runs and for use in 3-D data assimilation. The data were also filtered to remove CO_2 ice detections made spuriously by the MCS dust retrieval method, by excluding points at which the temperature lay below the CO_2 frost point at that pressure level. The result of this approach made the clarity of the south pole especially evident — note the reductions in density-scaled opacity of three orders of magnitude or more from low to high southern latitudes around winter solstice ($L_s = 90°$).

The same plot for model data from the control run, though for diurnally averaged data[4], (Figure 6.20(b)) displays significant 'leakage' of dust into both winter polar regions, most obviously near the ground but also noticeably so at the higher altitudes viewed by MCS. Both polar regions were clearest at winter solstice, which, as identified earlier (§6.2, §6.4, and the previous chapter), is at least partly a result of weakened eddy mixing at the cap edge. Dust exclusion from the polar vortices, at 100–50 Pa above the autumn/winter poles, can be observed, signifying the existence of a transport barrier around these regions (there are hints of this phenomenon in the MCS data too, e.g. at the north pole at $L_s = 270°$). When scavenging was used, near-surface dust mixing ratios were reduced near each pole during autumn and winter (Figure 6.20(c)), by an order of magnitude in places (the differences away from the north pole near perihelion were mainly due to interannual variability in dust storm size, since the scavenging year used occurred three years later in the multi-year series than the control run). For $R_s = 0.5$, the winter polar mass loadings approached the levels seen in the MCS data (at least judging from a by-eye comparison).

It is also notable that in both model cases, away from the poles, dust was much more confined to the near-surface than the observations suggest it should be. For example, in northern midlatitudes at $L_s = 45°$, dust opacities were approximately the same between model and observation at 500 Pa, but by 100 Pa model opacities were too low by at least

[4]Using nighttime data only does not greatly alter the results of the comparisons between model and observation here.

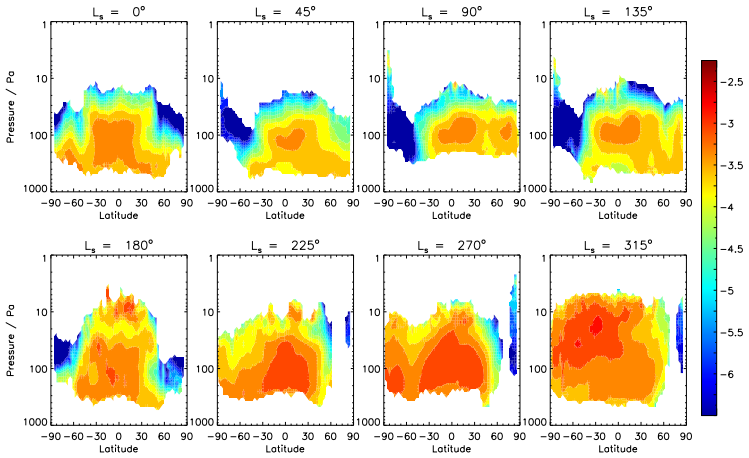

(a)

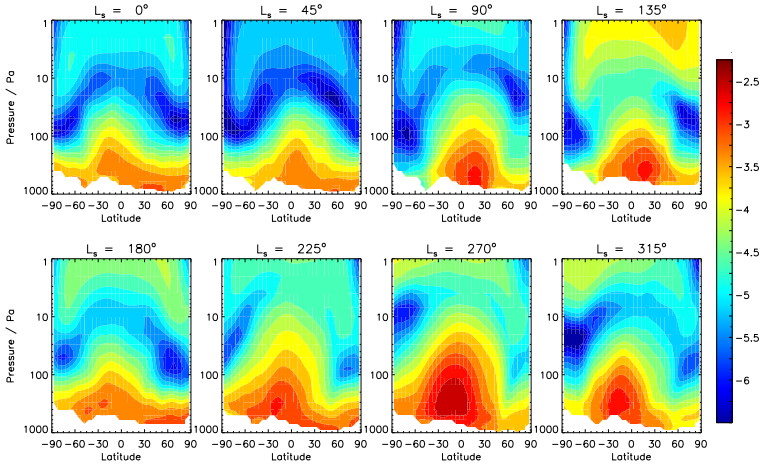

(b)

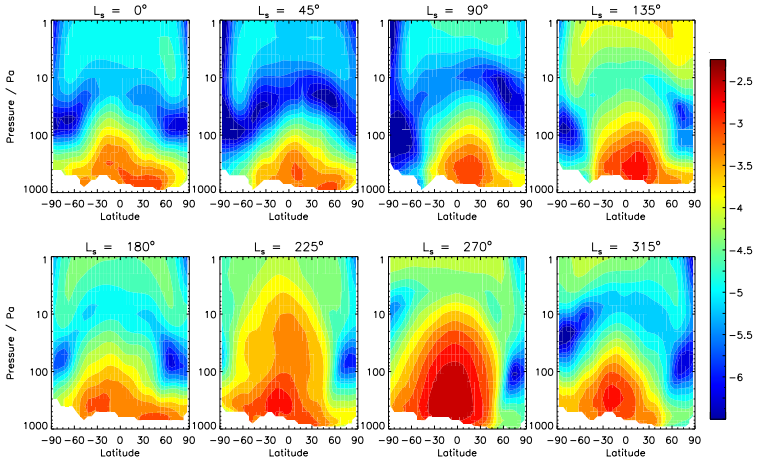

(c)

Figure 6.20: Logarithm of the zonally averaged density-scaled dust opacity ($m^2\,kg^{-1}$) from (a) MCS data from nightside retrievals; (b) the model control run; (c) the scavenging run with $R_s = 0.22$. (b) and (c) show diurnal averages. The lowest contour level shown is -10. No MCS data was available in the white regions of (a).

a factor of ten. No significant elevated dust maxima (Heavens et al., 2011) (seen most clearly in the MCS data at $L_s = 135°$) were produced, at least below 10 Pa; instead dust mixing ratio decreased monotonically with increasing height above the boundary layer, with a scale height of ~ 15 km. The increases in mixing ratio above 10 Pa are, however, surprising, and require comment. They were not present at the beginning of these long simulations, but instead built up over the 10+ model years, apparently as some of the smaller particles carried by the two-moment method became 'trapped' at high levels due to their low sedimentation velocities. While they represent relatively small masses of dust, they are surely unphysical and would have had an unwanted influence on IN densities (as used by the cloud scheme) and atmospheric temperatures (part of the reason for the problem discussed in §6.2.3) at high altitudes. This effect will require attention if any multiannual dust transport simulations are to be performed in future; it may be that it is mitigated by scavenging in the long term, as is hinted at by a comparison of the bottom rows of Figures 6.20(b) and 6.20(c).

To compare the dust content of the model and actual polar regions more quantitatively, since a direct latitude-height difference plot was not suitable due to the different vertical distributions of mass in model and observed data, estimates of the total dust mass in the 60–90° regions were made. The MCS data do not extend all the way to the surface, so it was necessary to use a restricted pressure range; unfortunately, the coverage is such that the largest range that could be used to calculate the 60–90° dust mass over a substantial part of the year was 150–50 Pa. Within this range, the number of datapoints obtained was increased by using either nightside or dayside values wherever only one was available, and taking the average of the two at points at which both were available. The use of this crude day-night mean also makes the results more directly comparable to model data. The outcome of the calculation using this method is shown in Figure 6.21(a) (using only dayside data resulted in broadly similar values, but with sparser coverage, particularly around $L_s = 150$–240°). The purpose here was to illustrate the relative changes in polar mass over the year, rather than estimate absolute totals, since missing data from close to the surface are sure to be vital for the latter. It was observed that in MY29 both polar regions contained a larger dust loading during spring/summer than in autumn/winter, and that the difference between the two solsticial loadings was much greater in the southern

hemisphere, at 3–4 orders of magnitude, than in the north (1–2 orders). The dramatic decline in south polar dust content around winter solstice fits with the solsticial isolation of the polar region mentioned previously. The south pole is affected more strongly than the north pole due to weaker cap-edge dust lifting and eddy activity, a weaker solsticial PMOC, and the larger polar ice cap, which pushes the boundary of the dust lifting region further towards the equator.

It is possible that the results shown, since they cover only a fairly narrow height range, may be missing important dust mass nearer to the surface during the winter minima (this is certainly true, for example, at both poles at $L_s = 0°$). From visual inspection of the available data, this cannot be ruled out; but, as far as can be seen, there is reason to believe that reductions in 150–50 Pa polar mass are representative of reductions in the total polar atmospheric mass. By accepting poorer temporal coverage, the lower limit in the mass calculation could be extended further, to around 350 Pa for the southern hemisphere over 0–180°, and this once again suggested a 3–4 orders of magnitude reduction by winter solstice. In the northern hemisphere, calculating the mass over 250–50 Pa suggested a slightly gentler decrease, of closer to one order of magnitude, from summer to winter than is seen in Figure 6.21(a). The behaviour shown in the plot should therefore be regarded as an approximation; however, this was sufficient for the level of comparison with the model that was undertaken.

Model polar dust masses within the 150–50 Pa window were generally lower than the MCS equivalents, since dust was more strongly confined to the near-surface. Despite this, the south polar mass at winter solstice was an order of magnitude larger than the MCS data show, and the annual variation in the south polar mass total was less than 3 orders of magnitude. However, the effects of an increased scavenging efficiency were not strongly felt at these heights (except at the north pole, $L_s = 240$–$360°$), as seen in Figure 6.21(b), because the scavenging that occurred did so nearer the ground. A more useful measure of the effect of scavenging in the polar regions was the dust mass between 60° and 90° over the full height of the atmosphere, shown in Figure 6.21(c). Now a decrease in mass with increasing R_s was seen in both winter polar regions, which became significant for $R_s > 0.1$. Increasing R_s from 0.1 to 0.5 reduced SH polar dust mass by a factor of 5 over much of autumn and winter, and had an even greater effect in the NH, reducing polar mass by an

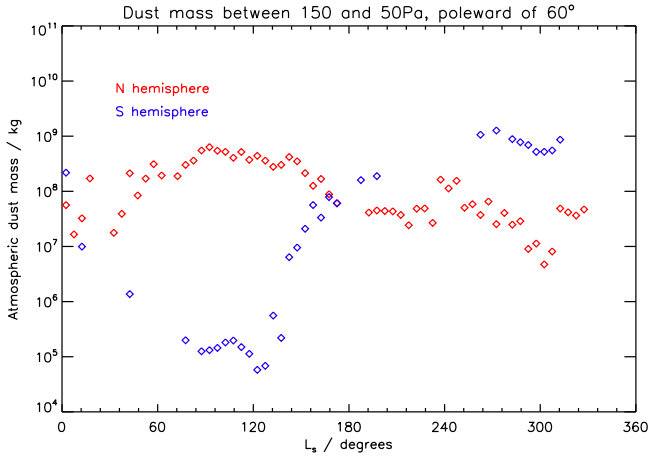

(a)

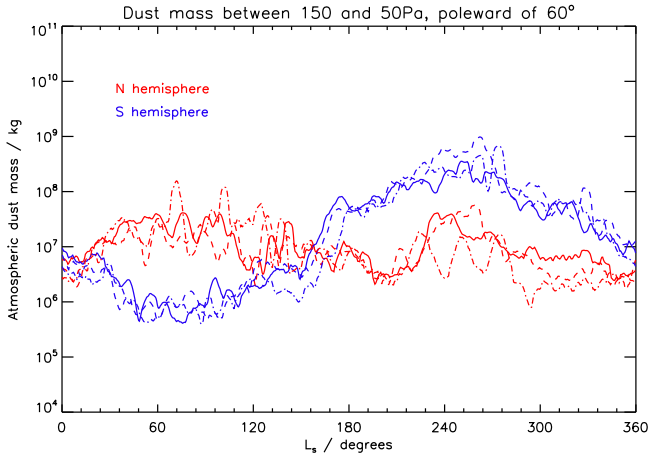

(b)

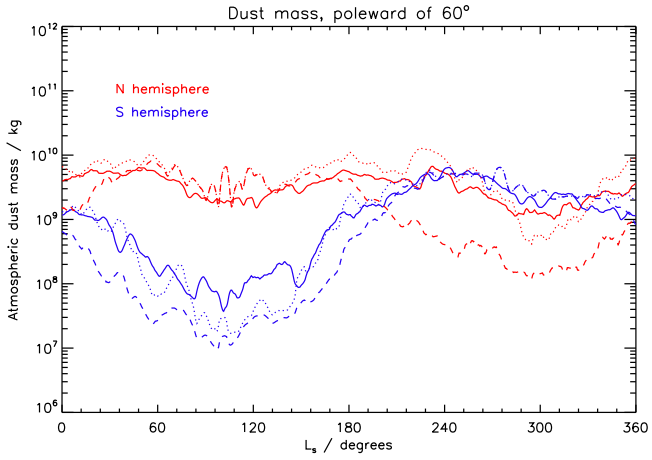

(c)

Figure 6.21: (a) Integrated dust mass (kg) between 150 and 50 Pa at 60–90° in each hemisphere, from the combined dayside and nightside MCS dataset. (b) The same quantity from the model, from runs with $R_s = 0.1$ (solid), 0.22 (dashed) and 0.5 (dash-dotted). (c) Integrated dust mass over the full height of the atmosphere at 60–90° from the model, showing interstitial mass from the runs with $R_s = 0.1$ (solid) and 0.5 (dashed), and total mass for $R_s = 0.5$ (dotted).

order of magnitude during $L_s \sim 240$–$330°$. For efficiencies of 0.5 and 0.22 (not shown, but its behaviour was intermediate between the 0.1 and 0.5 cases), this created the post-180° decline in north polar mass seen in the MCS 150–50 Pa data (though it is possible that a lower efficiency may fit the data better, if the full atmospheric mass declined more gently into autumn/winter, as hinted at above). So, as found in the passive test case, scavenging, if it is efficient, could be of similar importance to the radiative effect of ice clouds in explaining the low dust content of the polar regions around winter solstice.

If scavenging does play a significant role in keeping the winter polar regions dust-free, a further question concerns whether it does so by accelerating the removal of polar atmospheric dust through sedimentation or whether the bulk of the 'missing' dust is merely hidden from view within the cloud ice. The dotted line in Figure 6.21(c) shows that differences in polar mass between the 0.1 and 0.5 cases were largely removed when the in-cloud dust mass was added to the 0.5 interstitial total. This means that most of the extra scavenged dust in the 0.5 run was still present in the atmosphere, which suggests that the largest increases to settling velocities estimated earlier do not apply to a particularly large mass of dust. An exception was in the SH at $L_s \sim 60$–$120°$, when significant amounts of dust were fully removed from the atmosphere by the stronger $R_s = 0.5$ scavenging. This coincided with a minimum in dust content of mixed particles (Figure 6.19(b)), which favoured greater increases to dust settling rates.

Annual mean dust wet deposition rates, for all values of R_s, were largest in those regions where cloud coverage was frequently high; as a result, more dust was deposited as a result of scavenging at the north pole than at the south pole, and wet deposition rates at low latitudes exhibited the zonal wavenumber-3 pattern that exists also in ice opacity (see Figure 6.25), forced by topographic variations. In comparison to dry dust deposition, wet deposition was of relatively low importance in the vicinity of the ACB and at low latitudes in general, but it made an important contribution to the total deposition rates over the winter polar caps: with $R_s = 0.1$, wet deposition accounted for $\sim 40\%$ of the total dust deposition over each cap during periods of CO_2 condensation, and for $R_s = 0.5$ this increased to 80–90%. The total dust mass deposited as a result of scavenging increased with increasing R_s, but did so less than linearly, as summarised in Table 6.1.

R_s	Dry deposited mass / 10^{12} kg	Wet deposited mass / 10^{12} kg	Wet/dry	(Wet/dry)/R_s
0.1	8.84	0.77	0.08	0.80
0.22	8.90	1.33	0.13	0.60
0.5	8.92	2.01	0.18	0.37

Table 6.1: Global dust deposition rates from the scavenging simulations.

6.6.2 Vertical mass distributions

Mass and number profiles

Scavenging has the potential to alter the vertical distributions of both dust and ice, by changing the sedimentation velocities of each and by the removal of dust at the altitude of the cloud layer(s), which itself can reduce local ice growth rates due to a lack of IN. Figures 6.22(a) and 6.22(b) show a series of zonally averaged vertical profiles of dust and ice mass mixing ratio, in the vicinity of the ACB at $L_s = 90°$ for MY29, taken from MCS data, handled as described previously. Features to note in the displayed profiles include the elevated dust maximum at 20–30 km, and the rapid reduction in dust mixing ratio above this, strongly suggesting a capping of dust in the vertical direction by the ice cloud layer that peaks at 30–40 km. A diurnal-mean plot of the same quantities from the control simulation (Figure 6.22(c)) shows that the model, without scavenging, did not produce any such elevated dust layer, and any vertical capping of dust by ice was weak (i.e. there was a slow decline in mass mixing ratio with increasing height within the cloud), if it occurred at all. It is also evident that the ACB formed 10–20 km too low in the atmosphere, and that it peaked much less sharply in the vertical than the MCS cloud layer does. The first of these issues is addressed in detail below. With regard to the second, the increases in temperature above 40 km were responsible for the drop-off in ice mixing ratio that did occur. Otherwise, in the 20–40 km region, cloud growth rates were maintained by a plentiful supply of IN at those altitudes. The prescribed dust simulation shown earlier, which had no upper level temperature inversions, actually produced a much sharper cloud layer, due to lower IN number densities above 20 km (seen from Figure 6.8).

With scavenging activated, using an efficiency of 0.22, dust vertical profiles were only marginally altered, through reductions in mixing ratio at 20 km and above (Figure 6.22(d)).

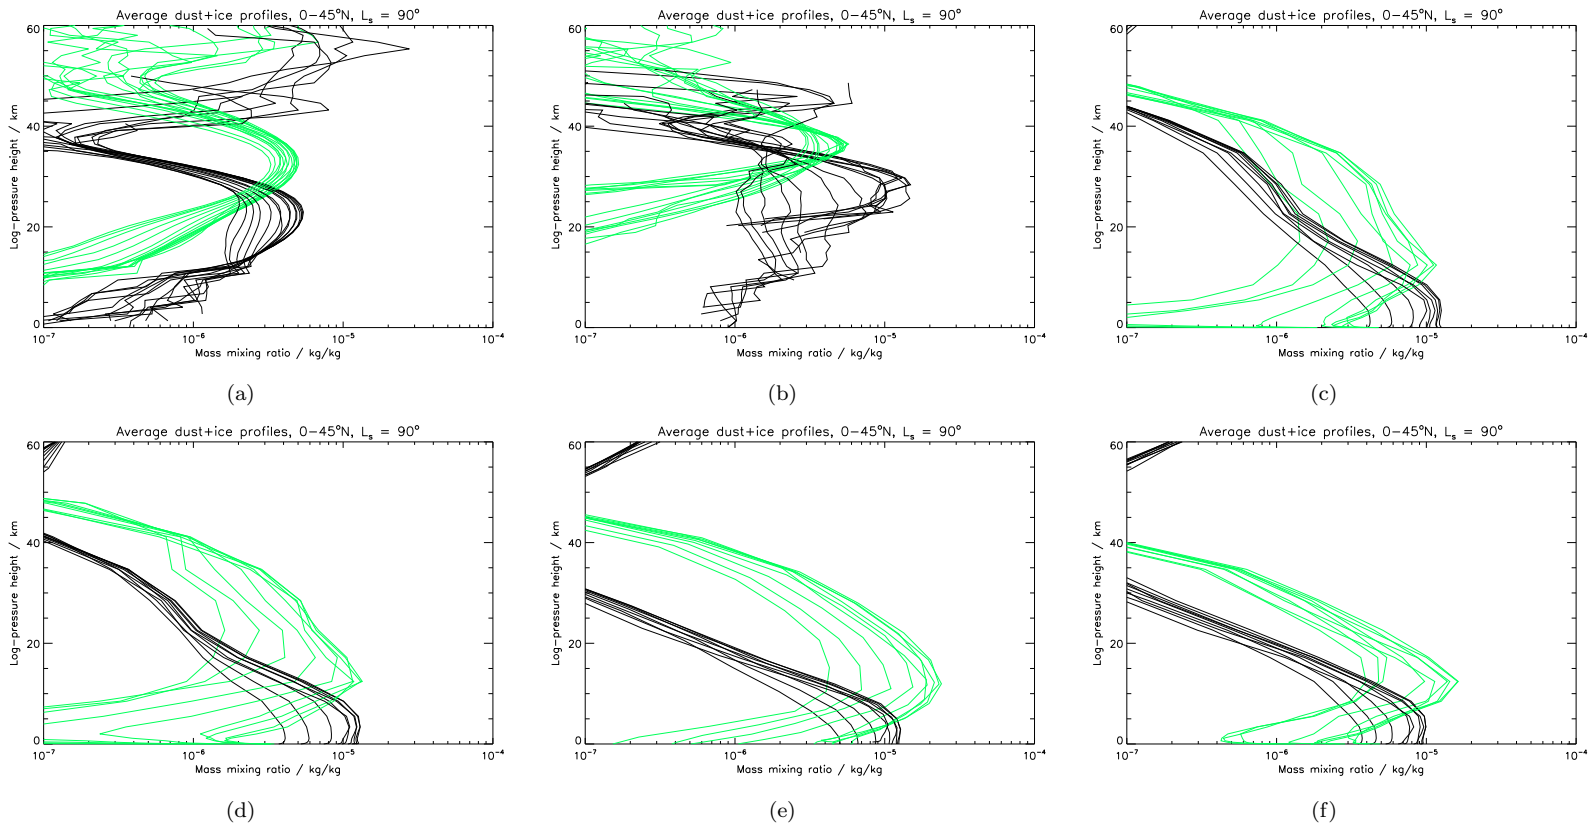

Figure 6.22: Zonal-mean dust (black) and ice (green) mass mixing ratio profiles at northern summer solstice, from (a) MCS MY29 nightside data; (b) MCS MY29 dayside data; (c) the control simulation; (d) the scavenging case with $R_s = 0.22$; (e) the scavenging case with $R_s = 1$; (f) the simplified scavenging simulation, with $R_s = 0.22$, which assumed instantaneous sedimentation of scavenged dust. MCS profiles are plotted against log-pressure height defined relative to a 610 Pa surface.

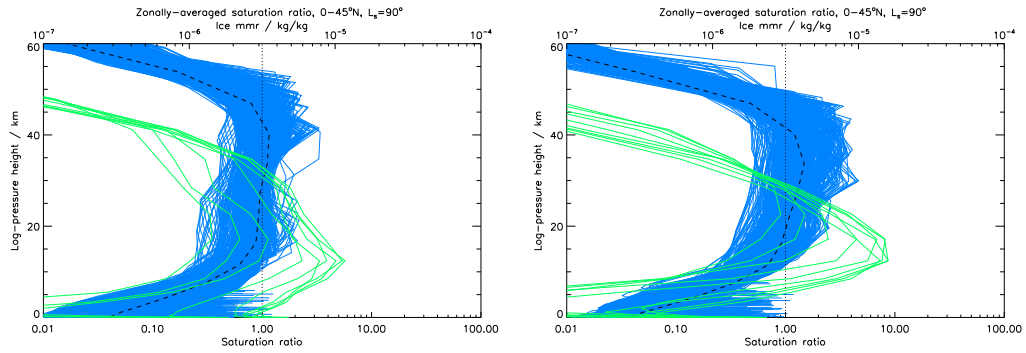

(a) (b)

Figure 6.23: Saturation ratio profiles (blue, bottom axis) from each model point within 0–45°N, each averaged over L_s = 80–100°, with the mean profile shown as a dotted line, from (a) the control run and (b) the simulation with $R_s = 0.5$. Zonally averaged ice mixing ratio profiles from the same region (green, top axis, in kg/kg) show the location of the cloud layer.

There was some sign of a sharpening of the cloud layer in the vertical, though its peak height was unchanged. With more efficient scavenging, these two effects were strengthened (Figure 6.22(e)), but the rapid in-cloud drop-off in dust mass mixing ratio seen in the MCS data was not matched even with the highest possible efficiency of $R_s = 1$.

Figure 6.22(f) shows the profiles produced by a further simulation in which no additional tracers were used to model scavenging; instead, dust scavenged at a given model level was assumed to sediment instantaneously (at least, within a model physics timestep of 30 minutes), and was either released at the first lower level at which ice sublimation was occurring, or deposited on the surface. This approach therefore assumed that dust particles settle very much faster when mixed with ice than they do independently. By comparing Figures 6.22(d) and 6.22(f), each of which used $R_s = 0.22$, it can be seen that vertical redistribution of dust was overestimated in the instant sedimentation case. A scavenging parameterisation without extra tracers may be feasible, but it would need to account for the more modest increases in settling velocity that tend to accompany the incorporation of dust particles into cloud ice.

With a lack of dust nuclei of sufficient size for ice nucleation, ice growth rates decrease, and supersaturated regions can develop in the atmosphere, if the vapour present cannot condense fast enough. Recent analysis of SPICAM data by Maltagliati et al. (2011) has provided evidence that supersaturation does occur on Mars, and in fact is potentially relatively widespread at certain times of year. The authors found that more than half of the

vertical water vapour profiles retrieved in northern midlatitudes in the range $L_s = 50$–$120°$ contained supersaturated portions, at altitudes of 25–55 km, and several had peak saturation ratios of greater than 5, which they noted to be much higher ratios than are observed in the Earth's troposphere. The dust-transporting UKMGCM actually produced supersaturations without scavenging, as Figure 6.23(a) shows: saturation ratios of up to 2–3 were common above the cloud layer, at 35–50 km, in the control run. This must have been a consequence of the slower ice growth rates that occur at low pressures (due to an increase in R_d in the growth rate expression (6.1)). (It should be noted that the supersaturated region would have continued up beyond 50 km, but for the temperature inversion that occurred.) With scavenging activated, larger supersaturations became more common: shown in Figure 6.23(b) is the strong scavenging case of $R_s = 0.5$. Individual (i.e. not time-averaged) profiles in early summer frequently peaked at $S > 5$. The observations of supersaturation are as yet limited, but it is clear that the model was able to produce ratios as large as any that have been seen.

The cause of the increased supersaturations for large R_s was, of course, the removal of some of the IN during cloud formation. This modified the shape of the IN density profiles, as shown in Figure 6.24. Significant reductions were produced in the interstitial population at 20–50 km, in and above the upper portion of the cloud layer, though the particles for the most part remained present at those altitudes, within the cloud ice. Interstitial dust mass was also reduced over the same vertical range, and in this case there was some vertical redistribution seen in the scavenged dust mass, from the upper part of the cloud layer to the lower part.

Cloud layer height

Aphelion belt cloud ice from the control run is shown in both early and late summer in Figure 6.25, along with temperatures and IN number densities. In both cases, a wavenumber-3 pattern in cloud ice formed, as a result of the topographic forcing in northern low latitudes. The cloud layer was initially thickest at 100–200 Pa, and by the end of summer it had moved only slightly upwards, to ~ 100 Pa. The height of the cloud layer is most strongly controlled by the combination of temperature and vapour abundance, and the peaks in ice at both times sat just below the 180 K isotherm, such that the upward shift reflects the

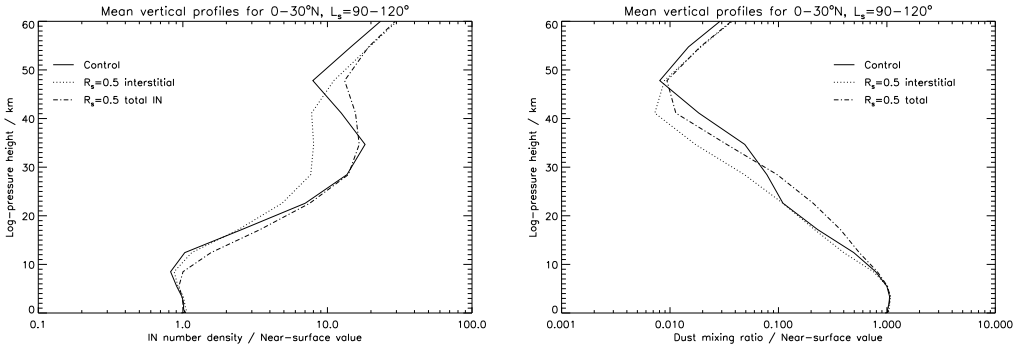

(a) (b)

Figure 6.24: (a) Interstitial IN density profile from the control run, and interstitial and total (interstitial plus scavenged) IN profiles from the $R_s = 0.5$ run. (b) Mass mixing ratio profiles from the same cases.

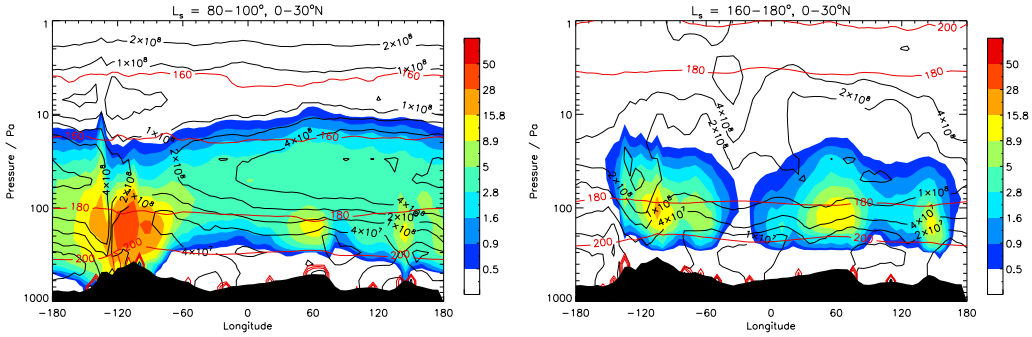

Figure 6.25: Cloud ice mass mixing ratio (ppm, colours), temperature (red contours, at 160, 180 and 200 K) and IN number density (black contours, drawn at $(1,2,4)\times(10^6, 10^7, 10^8, 10^9)\,\mathrm{kg}^{-1}$) from the control simulation at 0–30°N, in early (left) and late (right) summer.

increase in model temperatures that occurred in late summer.

MCS data, on the other hand, show a cloud layer centered — in a crude diurnally averaged sense[5] — at 10–50 Pa in early summer, and rising to $\sim$ 10 Pa or higher by late summer (with some interannual variation between MY28 and MY29) (Heavens et al., 2010). The principal reason for this rise during summer was probably a similar rise in the isotherms around the height of the cloud layer, particularly on the nightside, which pushed the level of saturation up to lower pressures. It was seen earlier that this atmospheric warming was not adequately picked up by the dust-transporting model (see Figure 6.7(a)), which goes some way to explaining the lack of movement of the cloud layer. However, this was not the only problem, as the cloud layer was also too low at the beginning of summer, when

[5]Considering day-night variation in the model does not qualitatively alter the results.

temperatures were reasonably close to reanalysis and MCS values. Clouds in the model tended to form too readily at warmer layers: Heavens et al. (2010) showed an ice layer throughout summer sitting comfortably above the altitude of the 180 K isotherm, while in the model the bulk of the ice formed below this line. The addition of scavenging did not greatly alter the situation, aside from producing a slightly more vertically extended cloud layer in late summer, associated with the warmer atmosphere that existed at this time, for the efficient scavenging cases. Interestingly, the location of the largest cloud thickness in early summer did shift from Tharsis to Arabia (60°E) as R_s was increased, apparently due to changes in the spatial distribution of peak dust lifting rates.

The differences between model and observation in early summer could be due to a number of factors — an overestimation of the amount of water vapour present, either in low levels only or in the atmosphere as a whole, a supersaturation condition for cloud formation that was too lenient, the reduced lapse rate in the middle atmosphere due to the use of size-independent scattering, or an IN profile that was biased towards lower levels. Several separate simulations helped to narrow down the potential causes. The first of these was the prescribed dust run shown earlier, which featured an early summer cloud layer very similar to that of the control (dust-transporting) simulation, with ice mixing ratios peaking below the 180 K isotherm. This run, since it used dust opacities very similar to observations, had a more realistic temperature structure, without an inversion above 10 Pa, yet this did not push clouds any higher in the atmosphere (they did rise to ~ 50 Pa in late summer, following the rise in temperatures that occurred, but the peak of the cloud layer remained below the 180 K isotherm). Another MY24 scenario run was performed using an altered vertical IN profile — instead of the Conrath function used before, a profile was devised to replicate to some extent the shapes produced in transporting runs. It consisted of an increase in number density from the surface to 40 Pa, peaking at ten times the surface concentration, and a decrease from there with decreasing pressure. This did not affect the height of the peaks in cloud ice, though mixing ratios were increased a little above this, at $\sim$ 100–20 Pa, around the IN density maximum. It seems that much more extreme variation in IN number density would have been required to force a genuine change in cloud height. If the suggestion of Montmessin et al. (2002), that there exists a bimodal dust size distribution, is correct, there could be a greater number of particles of radius 100–200 nm

above ~ 30 km than simulated here by the model, which would boost the IN population in the middle and upper atmosphere, but it is unlikely that this could suppress condensation particularly strongly at lower altitudes. Finally, the amount of supersaturation required to initiate cloud formation was, in a transporting simulation that was otherwise identical to the control run, increased from 1.4 to 2.0, but this did not lead to significant changes in ACB height. Recent experimental work has suggested that even higher saturation ratios may sometimes be required for nucleation (Iraci et al., 2010), so it is still possible that underestimated supersaturation in the present case is affecting cloud height.

Vertical transport

None of these alterations came close to producing a cloud layer at as high an altitude as the one observed by Heavens et al. (2010), leaving overestimated water vapour amount as the most likely explanation for the error. Compared to TES data, column vapour abundances in northern low latitudes at this time were, if anything, lower in the model, so the problem must rather have been the distribution of vapour in the vertical. Therefore, once again, weak vertical mixing was the likely issue. To get around this temporarily, in order to determine how the cloud height might be affected by stronger mixing, a short run ($L_s = 90$–$100°$) was performed using altered turbulent diffusion coefficients. Turbulent diffusion is one form of parameterised mixing in the model (the other being convective adjustment) but is usually strong only in the boundary layer (up to a few km above ground) and negligible above this. Diffusion coefficient profiles were artificially altered to increase mixing over ~ 0–30 km by multiplying by a Gaussian function centred around 20 km. It is of course unrealistic to maintain such strong turbulent mixing up to these altitudes, but this adjustment provided a convenient means of altering the tracer profiles; also, it allowed an assessment of whether or not diffusive mixing could possibly explain some of the observed dust and ice profile features.

As expected, it led to dust mass mixing ratios that were well-mixed up to ~ 30 km. Vapour profiles were also altered, but the cloud level was raised only slightly, to around 90 Pa, and this was mainly because temperatures were increased within the mixed region, which increased the holding capacity of the atmosphere at those altitudes (Figure 6.26); in fact, the peak of the cloud layer still sat below the 180 K temperature contour. Temperatures

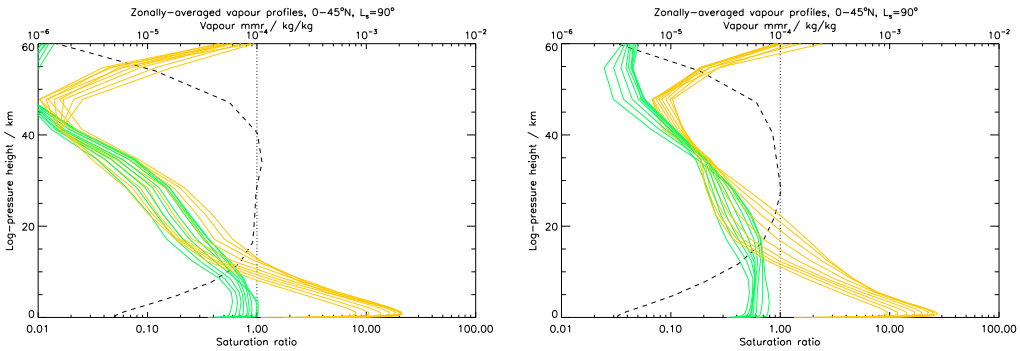

Figure 6.26: Zonal-mean vapour mixing ratio (green, top axis, in kg/kg) and saturation mixing ratio (orange, top axis, in kg/kg) profiles over 0–45° N, along with the saturation ratio averaged over this latitude range (dashed, bottom axis), from the control (left) and enhanced diffusive mixing (right) simulations.

were too high compared to the MY26 reanalysis, due to the exaggerated mixing of dust from the surface at every gridpoint, so the large saturation mixing ratios of Figure 6.26(b) cannot be the true means of avoiding cloud formation below 20 km. To remain below the more accurate saturation curve of Figure 6.26(a), vapour mixing ratios must be reduced below their values in that plot at $\sim$ 10–25 km, and above this the vapour lapse rate must be smaller than the saturation lapse rate. This does not necessarily require the production of an elevated vapour maximum, so a well-mixed vapour profile could be capable of producing a higher cloud layer, but in that situation one would expect dust to be well-mixed as well, which leads to the unrealistic temperature structure produced in the simulation shown in Figure 6.26(b). Therefore, a local vapour enhancement around the cloud level seems more likely. It is clear that diffusive mixing as prescribed in this case could not produce such tracer profiles, and that another vertical transport process is required to explain both the elevated dust layers and the height of the ACB.

A similar enhanced-diffusion simulation with scavenging ($R_s = 0.5$) did not produce a drop-off in dust mass mixing ratio quite as sharp as those of Figures 6.22(a) and 6.22(b), and did not produce any elevated dust maxima below the clouds (Figure 6.27). So, the dust capping of inadequate strength seen earlier appears not to have arisen merely because it was taking place so close to the ground. It may be that the capping seen is predominantly caused by a sharp decline in dust mixing ratio that occurs independently of the presence of an ice layer, as a consequence of the vertical transport processes acting. Furthermore, the scavenging process, even if it is very efficient, does not seem to be capable of generating

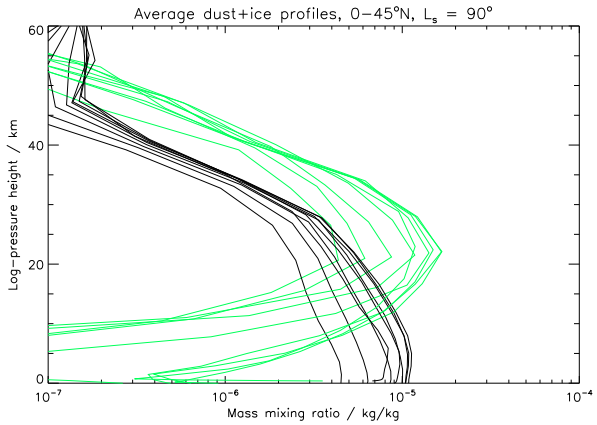

Figure 6.27: Dust and ice mass mixing ratio profiles, as in Figure 6.22, from the simulation with enhanced diffusive mixing and scavenging ($R_s = 0.5$).

elevated dust maxima out of well-mixed dust profiles.

Another possible mechanism for production of these elevated dust and vapour maxima is the thermal topographic circulation identified by Rafkin et al. (2002). Using a mesoscale model centred on Arsia Mons, one of the planet's largest mountains, they simulated strong upslope winds on the flanks of the mountain that transported dust lifted near the base to altitudes of 20–35 km, whereupon winds diverged and the dust was advected away from the mountain. The outflow plumes extended to more than a thousand kilometres in the horizontal direction and created an elevated dust maximum away from the mountain, where mixing ratios below 20 km were much lower than above. This idea is attractive, because the five highest peaks on Mars are all located near the equator, where the elevated dust layers have been observed. None of the topographic peaks lie within 60°W–120°E, which appears to present a problem as the elevated maxima were seen at virtually all longitudes; however, some equivalent circulation could be produced near the large ridge at Arabia Terra, if horizontal advection of dust from one of the mountains were not sufficient. At the very least, topographic circulations offer a means of injecting vapour at lower pressures at the locations of two of the three aphelion cloud ice peaks (over Tharsis and Elysium Mons). Such circulations are not likely to be produced at T31 resolution: in fact, two adjacent volcanoes in the Tharsis range, Arsia and Pavonis Mons, are not even resolved at this resolution. It may be anticipated, however, that total vertical mixing will be strengthened as resolution is increased, due to more effective upslope transport at appropriate locations. A simulation at T42 resolution run over $L_s = 60$–$90°$, using the initial conditions and settings from the

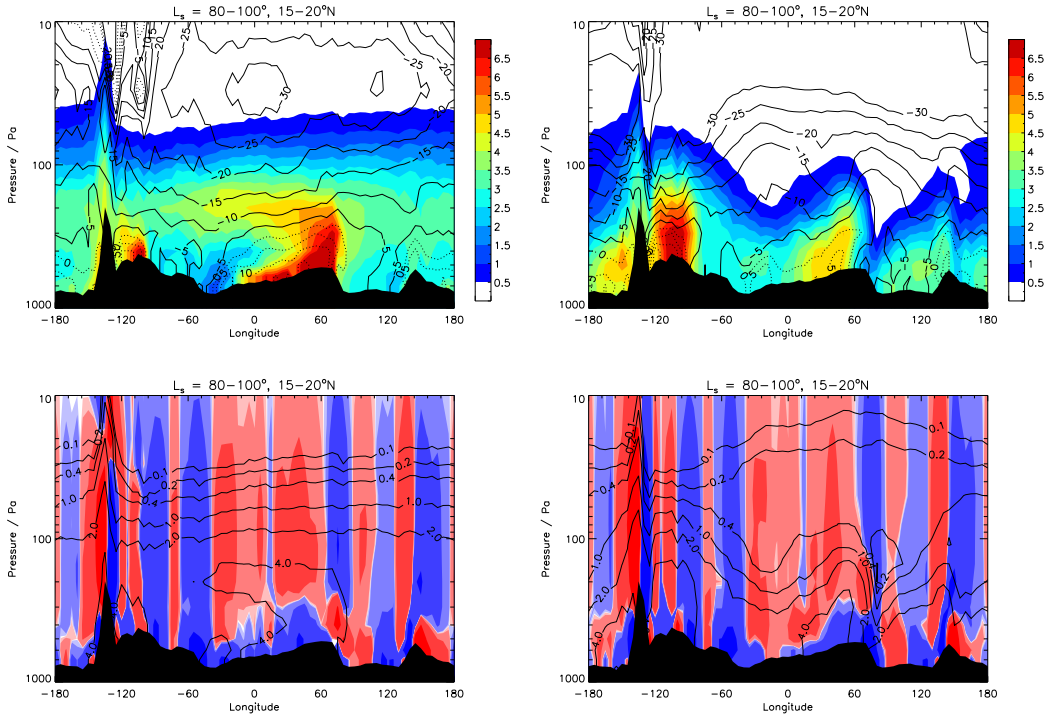

Figure 6.28: Top row: Longitude-pressure views of dust mixing ratio (kg/kg, colours) and zonal wind ($\mathrm{m\,s^{-1}}$ eastward, solid contours; $-30\,\mathrm{m\,s^{-1}}$ is the lowest contour shown) at 15–20°N from simulations using model v3 (left) and v5 (right), showing the difference in the vertical dust distribution. Coordinates shown are in degrees east longitude; Olympus Mons can be seen in the bottom topography at −135°. The run with clouds was dustier at this time than the run without, so mixing ratios in the right-hand plot were multiplied by 0.3 before plotting, to allow a fair comparison between the two plots. Bottom row: the same mixing ratio field as before (solid contours, kg/kg) plotted over the vertical pressure velocity ω (colours, contours are at $\pm(10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}\,\mathrm{Pa\,s^{-1}})$, again from v3 (left) and v5 (right). Blue shading denotes negative values of ω (upward atmospheric motion) and red shading denotes positive ω (downward motion).

control run, confirmed this to be the case. Dust mixing ratios at 0–45°N were increased over the control run above $\sim 10\,\mathrm{km}$, though they still decreased monotonically with increasing height above this level.

Upon closer inspection, interesting differences were found between the vertical dust distributions in the current case and in simulations from Chapter 4, which used the older model physics (v3). In the latter case, zonal-mean dust profiles were actually well-mixed up to $\sim 15\,\mathrm{km}$. An example of each, taken from the northern tropics around summer solstice, is shown in the top row of Figure 6.28. In the new model, dust mixing ratios at 100 Pa were very low away from the three main topographic peaks (e.g. at 0–60°W) — it appears that at these longitudes, where wind stress lifting is not occurring, advection from dustier

longitudes is necessary to maintain any significant dust loading above 300 Pa or so. Zonal winds at this height were reduced at these longitudes in the newer model run. Sensitivity tests showed that while the inclusion of clouds was responsible for the shift of dust lifting peak from Arabia to Tharsis, changes in the vertical distribution of dust occurred due to both the switch to the two-moment scheme for transporting a dust size distribution and the use of size-independent scattering in the newer run. In the older model, dust effective radius showed little change from the ground to $\sim$ 100 Pa, at 1.4–1.6 microns — this was because the dust field in this height range was dominated by particles in the 'middle' size bin (1.58 μm), leading to a low effective variance for the distribution, and highlighting a limitation of the use of a small number of dust size bins. When using the two-moment scheme, effective radius was larger near the surface ($> 2\,\mu$m), and though effective radius was similar to the previous case at 100–200 Pa, effective variance was larger (it was fixed at 0.5), implying more rapid sedimentation. Therefore, mass sedimentation rates were larger with the two-moment scheme, from the surface up to 100 Pa and beyond.

Calculation of the vertical (pressure) velocity for each run, plotted in the bottom row of Figure 6.28, shows that at the 100 Pa level (e.g. at 0–40°W) it was this more rapid sedimentation rather than any increase in downward atmospheric motion that caused the clearing of dust in the v5 case. However, differences in vertical velocity near the ground do exist between the two runs, which must primarily be a result of differences in the vertical distribution of dust heating. Further changes in vertical transport were seen when the more accurate size-dependent scattering method was used (lower effective radii near the ground, and larger mass mixing ratios around 100 Pa), indicating that dust heating could play an important role in forcing these mesoscale circulations. (The neglect of size-dependent scattering was therefore not ideal, and introduced errors, but recall (Figure 6.7) that it was necessary to obtain reasonable temperature structure below 10 Pa.)

An important feature seen only in the older run was the development of an elevated dust maximum, seen most prominently west of Arabia Terra. Although the enhancements to dust mixing ratio were slighter than those observed by MCS, and the peak layer formed too low, this nonetheless appears to be evidence that elevated dust maxima could be formed through strong upward transport of dust at specific locations and its subsequent zonal advection to other longitudes. What is particularly interesting is that in this case the feature was

produced at longitudes near the prime meridian, rather than near the much larger peaks at Olympus or Elysium Mons. On the other hand, slope winds were severely under-resolved at the steep mountains, while the gentler slopes of Arabia were better represented at this resolution, and the production of the dust plume was seemingly made possible by the low-level convergence that occurred in the region. Caution is required, as the feature was not seen in any simulations using the two-moment method for dust transport, which should provide more accurate sedimentation rates. The vertical velocity plots suggest that the feature was maintained in part by upward motion in the boundary layer in v3 that was not present in v5 (Figure 6.28, bottom row), indicating the influence of changes in dust heating profiles. Further work is clearly needed to investigate the effects on vertical transport of topography and of the variation of dust mixing ratio and particle size with height (affecting both sedimentation and heating rates).

Cross-equatorial flux

The changes seen in the IN profiles when using large R_s implied that dust transport was affected by the presence of the cloud layer, through a mechanism closely related to the Clancy effect that acts to limit the transport of water vapour out of the northern hemisphere. In both cases, the limiting of the tracer species in the vertical during the aphelion season prevents it from being lifted through the full height of the principal meridional overturning cell, and thus limits its meridional advection, since southward winds at $\sim$ 20–40°N are found only in the upper portion of the PMOC (Clancy et al., 1996). This has implications for the net annual cross-equatorial tracer flux, as a similar effect is not seen in the warmer perihelion season. In the case of dust, since it appears to be insufficiently mixed upwards in the UKMGCM at present, its southward transport around aphelion will have been limited even without any scavenging by clouds. With these simulations, it was possible to analyse the possible effect of scavenging on cross-equatorial transport, while bearing in mind that any capping of the dust column that occurred did so too close to the ground.

The cross-equatorial dust flux, taken as the average of the fluxes in the model latitude bands either side of the equator, was integrated over the height of the model atmosphere and is shown for the control and $R_s = 0.5$ runs in Figure 6.29. In the control run, the overall annual pattern (the red dotted line) was one of transport from summer to winter

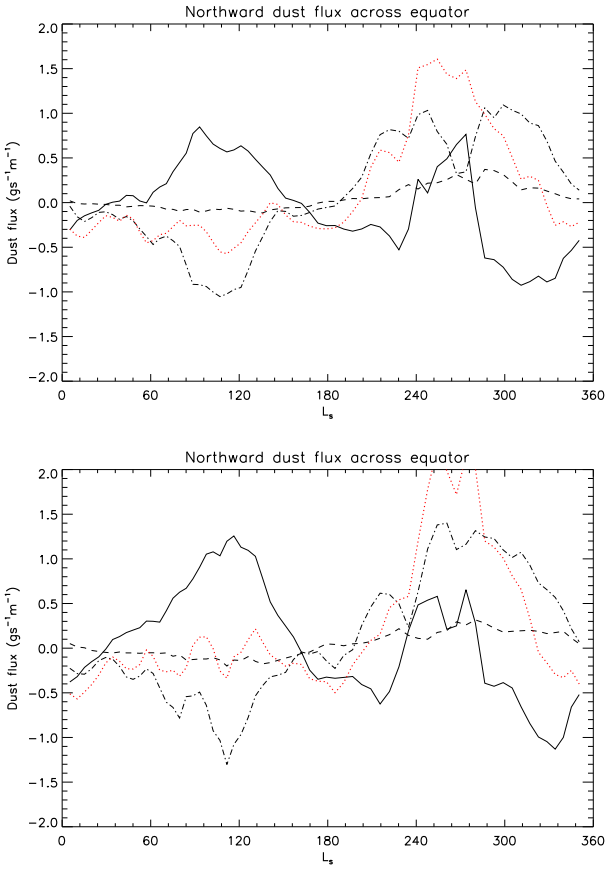

Figure 6.29: Vertically integrated dust mass transport across the equator from the control run (top) and the $R_s = 0.5$ scavenging case, including both interstitial and scavenged dust mass (bottom). Components shown are mean (solid), transient (dashed), stationary (dash-dotted) and total (red dotted).

hemisphere, as expected. However, analysing the different components of this transport (mean, transient and stationary) revealed that the mean circulation, except during the dust storm of 240–280°, acted to move dust in the opposite direction. This transport was outweighed by stationary eddy transport out of the summer hemisphere. The dominant part of the mean transport occurred in the lower branch of the PMOC, in the lowest 5 km of the atmosphere: therefore, this was potentially another consequence of dust being inadequately mixed away from the surface. It is unknown quantitatively what bias this introduced into the net cross-equatorial transport rates analysed in Chapter 4. Although the vertical mixing appears to have been better in those simulations, dust profiles were nonetheless weighted in favour of lower levels with respect to observations, as they did not feature any significant elevated mixing ratio maxima.

The strong variation of average meridional wind with altitude was also the reason for changes in aphelion dust transport that occurred when scavenging was activated. Since scavenging (away from the poles) occurred primarily within the cloud layer, above 300 Pa, the interstitial dust field was depleted solely at the level at which the average flux, due both to stationary eddies and the PMOC, was southward. Meanwhile, the near-surface flux was largely unaltered. This resulted in a weakening of the net southward integrated dust transport over 0–180° as R_s was increased — for interstitial dust only, the transport actually became strongly northward in the $R_s = 0.5$ case, but when the scavenged dust mass was also included, a southward integrated transport was returned, though one that was much (37%) weaker than in the control case. The fact that a weakening was still seen when calculating the flux of the combined dust field shows that a vertical redistribution of dust mass did occur, and acted in the downward direction.

Following this argument, a northward bias in annual dust transport, which becomes stronger as R_s increases, seems inevitable. However, it was not possible to quantify this effect while both dust and water were so vertically limited, and — as stated above — calculations of net inter-hemispheric mass transfer should be treated with caution, as it is unclear how accurately the overall efficiency of the overturning cells has been represented in the model.

6.7 Estimation of H_2O scavenging efficiency

It is clear that many features of the dust and water cycles could be strongly dependent on the efficiency of the heterogeneous nucleation/scavenging process, even though several relevant quantities, such as total ice mass, appear to have a substantially sub-linear relationship with R_s. Efficiency R_s may, in fact, be significantly lower than the range of values studied here: they were chosen, to some extent, to determine the upper limit of modification to the dust field that could be possible through scavenging. Although ice opacities were tuned at the beginning of this chapter with $R_s = 0.22$, it should be recalled that ice mass has been 'artificially' reduced throughout this chapter by using a size distribution variance for sedimentation of $\nu_{sed} = 0.45$.

When this was reduced to the more likely value (for a unimodal ice size distribution) of 0.1, the correct ice opacities were then best reproduced by $R_s \sim 0.02$. However, the actual

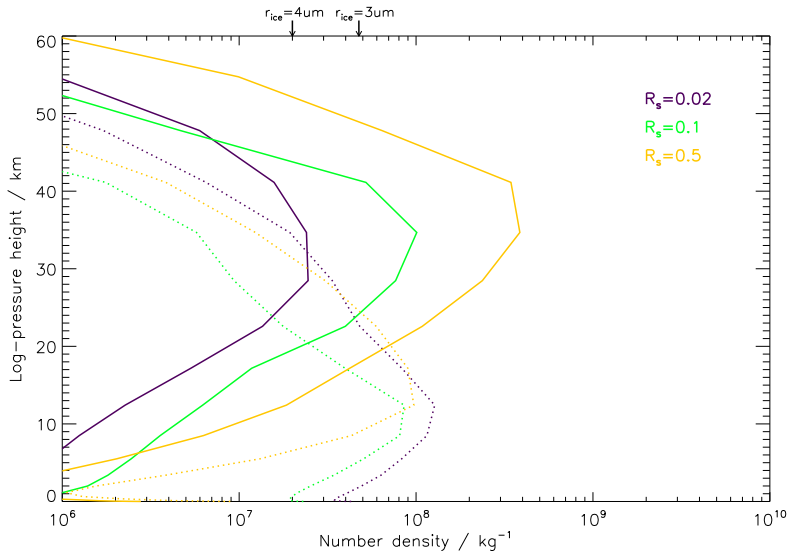

Figure 6.30: Scavenged IN density (kg^{-1}) profiles for different scavenging efficiencies at $L_s = 120°$ averaged over 0–45°N (solid lines). The densities estimated for observed ice particle sizes of 3 and 4 microns are marked. Dotted lines show the scavenged IN density that would be required to produce 4 micron ice particles, given the model ice mixing ratio profiles — sufficient scavenged IN were available where these lines lie to the left of their solid counterparts. Note that the run with $R_s = 0.02$ used $\nu_{sed} = 0.1$, while the other two cases used $\nu_{sed} = 0.45$. This has little impact on the solid lines but does affect the dotted lines: the latter two would be approximately doubled if $\nu_{sed} = 0.1$ had been used, raising the crossing points with the solid lines by ~ 5 km.

ice mass present (increased by roughly a factor of two by the reduction in ν_{sed}, for R_s constant) was probably too large, as opacities were kept artificially low by ice particle radii that were also too large: 10–20 μm within the ACB, which is well in excess of the 3–4 μm sizes reported by Clancy et al. (2003). Ice radii were set by dividing the ice mass among the number of IN present, so to obtain smaller radii required greater IN concentrations at the cloud level. Therefore, while a low scavenging efficiency is needed to keep ice growth rates (and mass) down, this reduces the number of IN that are involved in nucleation, which can lead to the formation of large ice particles, since condensation will inevitably occur as S increases. Comparison among the $R_s = (0.1, 0.22, 0.5)$ cases showed that the smallest particle radii were obtained when R_s was largest, in spite of the increased ice mass produced by this high efficiency (ACB ice radii were reasonable at 2–6 μm in these simulations, but this was helped by the large ν_{sed} used). However, model ice radii are also sensitive to the height of cloud formation (too low at present) due to the variation of IN density with height.

It was possible, using MCS ice mixing ratios together with particle size measurements, to estimate the number density of ice particles in the ACB, which should represent a lower

limit on the in-cloud IN density, assuming that each ice particle contains at least one dust nucleus. For a mean ice radius of 4 microns (Clancy et al., 2003) and a peak mixing ratio in early summer of 5×10^{-6} kg/kg (Heavens et al., 2010), scavenged dust number density should be around $2 \times 10^{7}\,\mathrm{kg}^{-1}$. Plotted in Figure 6.30 (solid lines) are ACB profiles of in-cloud dust number density for the various values of R_s used (it is worth remembering that above ~ 40 km, densities may be affected by the excess high-altitude dust discussed earlier). Scavenged density increased with R_s, due to the increasing numbers of interstitial IN made available in cloud formation. At the model cloud level (15 km), an efficiency of 0.5 was required to attain enough in-cloud nuclei to produce 4 μm ice particles, but at the true (observed) cloud level of ~ 35 km, $R_s = 0.02$ was sufficient. However, it is the number of nuclei in proportion to the ice mass that is more important: the dotted lines in the figure show the nuclei density needed to produce 4 μm ice particles. The high R_s case has the solid line lying furthest to the right relative to the dotted line, indicating that more IN were scavenged for a given mass of ice than for lower values of R_s. $R_s = 0.02$ did not scavenge at a sufficient rate at 15 km, but did so comfortably above 30 km, where the clouds should have formed, and larger R_s produced in-cloud number densities that were much too large above 30 km (albeit assisted by large ν_{sed}). Clearly, the rate of scavenging depended as much on the interstitial IN populations, all of which peaked at 20–40 km, as on R_s. Therefore, any constraints on R_s would be strongly affected by the uncertainty that exists in the model dust profiles. A further complication is that the cloud layer, as seen in the dotted lines in the figure, was too low in each case: if the ice layer was shifted upward by 20 km, there would be insufficient scavenged IN with $R_s = 0.02$, and the $R_s = 0.1$ case would appear to be more realistic (the dotted and solid green curves would roughly overlap at 30–40 km).

The excess in cloud ice produced with $R_s \gtrsim 0.1$ would seem to be strong evidence against an efficient scavenging process. However, cloud formation would have proceeded more slowly had saturation been exceeded at lower pressures, < 100 Pa, where observations show the ACB is actually located. At the very least, the best-fit value of R_s is likely to depend on the particular assumptions used in this cloud scheme; other formulations that produce lower condensation rates may be hypothesised. Therefore, no clear conclusions can yet be drawn regarding the efficiency of the nucleation/scavenging process. In fact, existing terrestrial literature suggests that the number of in-cloud nuclei tends to increase with decreasing

temperature or increasing supersaturation (Meyers et al., 1992). A more complex function for R_s is something that will likely be added to cloud microphysical schemes in future, as more detailed observations of clouds and IN densities become available.

What can be said, based on the work of this chapter, is that for scavenging to exert any significant influence on the dust field would require a moderately high efficiency: $R_s > 0.1$, say. If scavenging is actually weaker than this, then the capping of the dust profile by the ice layer and the clearings produced in both winter polar regions would not occur, and other explanations would be required to justify their lack of presence in the non-scavenging UKMGCM. It is possible that the sharp drop-off seen in the MCS dust profiles is not primarily a result of scavenging: while strong scavenging certainly acts 'in the right sense' and could contribute to a capping of the dust column, inaccuracies in vertical transport of tracers meant that even the use of high R_s values could not properly reproduce the observed relationship between dust and ice vertical profiles. The very clear winter polar regions seem harder to explain without some scavenging taking place, though due to the lack of MCS data at the surface some uncertainty remains over precisely how dust-free these regions are.

Chapter 7

Conclusions

During this project, the dust lifting capability of the UKMGCM has been developed in several areas, and an interaction with atmospheric water ice has been implemented. Significant progress has been made in simulating important aspects of the Martian dust cycle, though the addition of new physics uncovered further complications that will require more attention in future dust lifting and tracer transport schemes.

7.1 Dust lifting, transport and redistribution

The first task involved an update of the original model dust lifting schemes, now a decade old. Actually, in the intervening years, the approach to dust lifting taken by other models has not changed from the method of Newman et al. (2002b), though observations have since clarified the respective roles of dust devil and wind stress lifting in the overall dust cycle, and transporting a particle size distribution rather than a single size has become the norm. Wind stress lifting rates have always been calculated using a cubic function of drag velocity, depending on its value relative to a threshold drag velocity u_*^t, which is heavily tuned to produce appropriate dust emission rates in GCM simulations. The parameterisation described in Chapter 3 also fits into this category; however, the threshold drag velocity is calculated theoretically and then reduced in order to account for subgridscale effects. In this respect it is similar to the method of Kahre et al. (2006). The level of threshold reduction needed is a tunable parameter in the scheme, but the magnitude of this scaling factor must be justified physically, in order to give the parameterisation any

more credibility than a simple prescribed threshold has. The scaling factor value of 0.7 used at T31 resolution appears to be easily justifiable, as it covers the averaging of physical processes occurring on a sub-metre scale over a gridbox that is several hundred kilometres in length (see also §7.2). The scheme also utilised a newer saltation flux formula, following Kok and Renno (2008), and a radiatively active dust size distribution was lifted, using a set of size bins or a two-moment mixing ratio method. Dust devil lifting was handled using the threshold-independent scheme of Newman et al. (2002a), as it provided lifting rates that were more evenly spread over the planet than the threshold-dependent version and was therefore better at maintaining a global dust haze in the absence of any wind stress lifting.

A significant resolution dependence of wind stress dust lifting rates was discovered, and it was found that flushing storms, originating in the northern midlatitudes, were resolved much better at T31 resolution than at the slightly lower T21 resolution. T31 was therefore adopted as the standard resolution for most of the work done, but it is clear that a higher resolution, if it were more computationally feasible, would improve lifting predictions. Topographic flows are very important in generating the high windspeeds needed for dust lifting, and this is one reason why total lifted mass was seen to increase with increasing resolution. This is a significant problem for a dust-lifting GCM, and it may be that higher resolution simulations, at least in particular areas with complex topography, are essential to properly capture some of the flow regimes crucial to the formation of dust storms. It is evident that slope winds in and around the Hellas basin — the initiation site of several major dust storms over the last few decades — were under-resolved by the model at T31 resolution. Cap-edge lifting was also underestimated, particularly in the southern hemisphere, though this was improved by the inclusion of water ice and dust scavenging, through a strengthening of baroclinic waves. The magnitude of the thermal contrast between the ice cap and adjacent ground is perhaps of greatest importance for this type of dust lifting.

Early lifting runs with the above scheme produced some realistic dust storms, typically originating either at Utopia or Chryse Planitia, with a modest amount of interannual variability. Rarely was more than one significant lifting event seen during a model year, which was due in part to an under-resolution of the smaller frontal storms that are likely to be responsible for the majority of wind stress lifting that occurs, especially during years

without a global dust storm. However, a very precise tuning of the lifting parameters (efficiency and threshold) was required to allow the development of dust storms without the problem of runaway growth occurring due to the strong positive feedback that applies to wind stress lifting. The main reason for this was believed to be the use of an unlimited surface dust supply, as this meant that storms were able to continue for extended periods of time, fuelled by sustained lifting from a variety of sources that were activated by the increase in surface windspeeds that can occur as dust opacity increases. While this is not to say that Martian dust storms cannot shut down unless they run out of surface dust, the difficulty in tuning the model to produce large storms that decayed at a suitable rate when unlimited surface dust was used suggested that the cessation of lifting at key sites due to surface dust depletion is a likely mechanism for preventing long-lasting global storms from dominating the perihelion season in each Martian year.

When constraints were placed upon the amount of dust that was available for lifting, runaway untempered storm growth became less of a problem and, more importantly, interannual variability was greatly increased. A wide range of dust storm initiation dates and peak magnitudes were produced, covering a large fraction of the range observed in each, and displaying a greater variety within a single simulation than has been produced in other MGCMs to date (Kahre et al., 2006; Basu et al., 2006). All of these storms decayed fully by the end of the year, in accordance with observations, due to decreasing dust surface density at initial lifting sites, probably helped by some reduction in surface windspeeds due to increasing atmospheric static stability. A low-opacity background dust haze was maintained by dust devil lifting over $L_s = 0$–180°, during which very little wind stress lifting occurred. Several differences remained between these simulations and observations, however, including a lack of late summer ($L_s = 300$–330°) regional storms and early spring (150–200°) SH cap-edge lifting, and the absence of a distinction between regional and global dust storm morphologies. Though the observational record strongly suggests that the largest storms usually begin in the southern hemisphere, almost all the model storms (of any size) began as flushing storms forming at either Chryse or Utopia, some of which grew to global extent when they reached a favourable lifting site in the SH. Dust lifting in the SH was most active at Daedalia Planum, NE Hellas and Noachis, and the temporal variability at these locations was strongly determined by the redistribution of surface dust among them, over

several years.

The form used for the surface dust constraints was an increase in local surface threshold stress (above the baseline value of 0.02 Pa), linearly proportional to the mass of dust removed from the gridpoint. In practice, this was not greatly different from the more straightforward finite dust surface density simulations of Kahre et al. (2005), which failed to maintain lifting beyond several model years of integration due to the depletion of dust at key storm initiation regions, though the constraints were less stringent in the present scheme as lifting was always possible at a given gridpoint, provided winds became strong enough. This scheme also suffered from declining opacities during multiannual runs at first, and a resupply term was needed to avoid this problem. Use of this term, which gradually reduced thresholds back to the baseline value over time even without any dust deposition at the gridpoint, was equivalent to the addition of more dust to the system. While there were some effective dust sinks within the scheme, the threshold changes due to which would require some offsetting, the main reason that this (perhaps) unphysical term was needed was that a net motion of dust from the southern to northern hemisphere was simulated. It is not known on what timescale the cycle of dust transfer between the hemispheres is closed, but the simulations in Chapter 4 implicitly assumed that it was around three years, since each initiation site in the SH could not support more than one GDS before requiring replenishment in subsequent years (and each GDS usually involved at least two of the three main SH sites).

The transport of dust into the NH during GDSs was so strong, relative to the reverse transport during $L_s = 0$–180° and throughout years without a large storm, that southward transport from flushing storms would have needed an enhancement by a factor of 3–4 at T31 resolution, to balance the SH dust loss during GDS years. Such an enhancement was seen at T63 resolution, suggesting that cross-equatorial transport (at least from wind stress lifting) could potentially be balanced over ~ 3 Mars years. At the same time, comparisons of the dust vertical distribution to Mars Climate Sounder (MCS) observations revealed that model dust was being inadequately mixed upwards around aphelion, which has implications for this cross-equatorial transport problem. Transport at aphelion emerges as the net result of southward transport in the upper branch of the PMOC and northward transport in the lower branch, at heights of 0–5 km, so a concentration of dust mass near the surface would

have biased the flux in the northward direction, depriving the southern hemisphere of part of a potential means of resupply. The modelled vertical distribution in the second part of the year appeared to better match the MCS data, which enabled more efficient northward transport and maintained the annual southward bias. Transport of dust lifted by dust devils, which was not affected by the varying surface thresholds, was annually northward, but was strongly dependent on model resolution, and differed significantly from the transport in the NASA-Ames model reported by Haberle et al. (2006), highlighting the uncertainty that exists in both the magnitude and the transport contribution of convective lifting.

7.2 Surface roughness and sub-gridscale variability

The implications of including an accurate, heterogeneous surface roughness length (z_0) field in a dust lifting MGCM were investigated using a recently produced, high-resolution z_0 map. Theory predicts that threshold drag velocity varies significantly with z_0: by up to a factor of four, over the range of roughness values present in the Martian map. All thresholds were raised over the baseline value previously used once z_0 was considered, making dust lifting generally more difficult, and since most of the high-z_0 points were located in the southern hemisphere, lifting from this region was most strongly inhibited, causing a change in the spatial distribution of preferred dust lifting sites. This meant that lifting from the smooth, sloping Tharsis region became especially prominent, causing the formation of storms during NH spring/summer, even when the variable threshold scheme was activated. More importantly, for dust lifting to continue at a realistic level required reducing threshold velocities by more than 70% from their calculated values (up from 30% previously). While a significant amount of this threshold scaling could be justified by considering sub-gridscale wind variability (gustiness) and surface variation, the need for such significant reductions raises the issues of whether or not it is appropriate to use a threshold such as this, which applies to a sub-metre-scale process, across a GCM gridbox. There is reason to believe that this approach is acceptable (Okin, 2005), but the threshold adjustment must be performed with a great deal more care than has been used thus far. Work done in Chapter 4 suggested that the effect of heterogeneities in z_0 on gridbox emission rates depends on both the mean windspeed and the sub-grid z_0 variance, while gustiness may depend on one or more of mean wind magnitude, z_0 and microscale local topography. The heterogeneous roughness

map was not used in subsequent dust lifting simulations, due to the unresolved difficulties described here.

7.3 Dust-ice interaction

Later dust lifting runs also included water vapour and radiatively active water ice clouds, and the dust and water cycles were simultaneously reproduced to a reasonable degree of accuracy. Dust storms were, however, less realistic than before, and the interannual variability seen previously was not retained. It is highly plausible that different tunings applied to the variable threshold scheme could return some of this variability in dust storm characteristics; however, this was not the goal of this part of the project. It was rather to study the influences that the presence of ice clouds exerts on surface wind magnitudes and eddy activity, and therefore on dust lifting, to develop a better understanding of how future dust lifting schemes — which must include, for example, a z_0 dependence — should operate. These influences were most obvious in the autumn-to-spring midlatitudes: cap-edge winds were strengthened in the NH in spring ($L_s = 0$–60°) and weakened at winter solstice, and travelling wave amplitudes, as measured by perturbations in v, were increased either side of northern winter solstice, and decreased at $\sim 60^\circ$S at southern winter solstice. As would be expected, this greatly increased dust lifting intensity at the north polar cap edge in early autumn, and the variable threshold scheme was required to reduce this lifting peak and to push the spatio-temporal distribution in dust lifting towards something more realistic. With respect to the water cycle, the use of transported dust was seen to offer some advantages over the current dust opacity scenario, including the development of a late-winter south polar hood ice peak and reduced north polar hood ice opacities. However, cloud formation (with transported or prescribed dust) was routinely too rapid, and ice opacities needed to be reduced by increasing the sedimentation rates of the assumed ice particle size distribution.

Dust scavenging by cloud ice was included in the model, and it was estimated that only in-cloud H_2O ice and below-cloud CO_2 ice processes are likely to be of much consequence. Scavenging by CO_2 ice was found to increase dust deposition rates at both poles by an amount roughly proportional to the efficiency of the process, R_{sbc}, which is extremely uncertain. Relative to the existing dry deposition rates, these increases were generally

small, except at the north pole in autumn, when the dust content of the CO_2 ice cap was significantly enhanced, by up to a factor of ten, if R_{sbc} was large. Ice cap dust fractions, from both kinds of deposition combined, were an order of magnitude higher in the NH than in the SH, reflecting the greater dust loading over the north pole.

Scavenging of dust by water ice clouds potentially has a more widespread impact on the global dust field, but again, the efficiency of the process R_s is not well constrained. The simulations of Chapter 6 showed increased cap-edge eddy activity and dust lifting rates, which were produced following the clearing of some atmospheric dust from the winter polar regions, where ice clouds had formed. This meant that scavenging, rather than limiting the amount of ice nuclei (IN) available to allow heterogeneous cloud nucleation to proceed, actually increased cloud formation rates in some regions, due to these enhanced dust lifting rates. In this respect, cloud formation and scavenging were sensitive to the surface dust lifting thresholds used, and there was clearly a strong interaction between atmospheric dust and ice abundances. In reality, the strength of this link will depend on the extent to which dust lifting is constrained by the availability of surface dust in key lifting regions, particularly in midlatitudes.

With a high scavenging efficiency ($R_s \geq 0.5$) interstitial dust abundances were reduced within the aphelion cloud belt layer, altering the shape of the vertical dust mixing ratio profiles, but even the $R_s = 1$ scavenging case did not simulate a capping of dust by the cloud layer that was as severe as those observed by MCS. In contrast to scavenging on Earth, dust sedimentation rates are not greatly enhanced by scavenging on Mars, reducing the potential for modification to vertical dust profiles. The situation was complicated by the discovery that the model produced an equatorial cloud layer that peaked at significantly higher pressure than the one observed by MCS, throughout northern summer. The cloud height was not affected by scavenging, and the discrepancy could only partly be explained by temperature differences between model and observations. Since dust mixing ratio also peaked too low in the model — at the surface, rather than just below the cloud layer at 20–30 km — the problem appeared to be inadequate vertical mixing of tracers, affecting both dust and water vapour. While it is possible that boundary layer mixing is currently underestimated, diffusive processes could not produce the elevated dust maxima seen by MCS in the aphelion season. Therefore, some more localised mechanism of entrainment and

transport up to 20–30 km is required, such as the thermal topographic circulation identified by Rafkin et al. (2002). Model vertical mixing was improved when the horizontal resolution was increased, supporting the idea that mesoscale circulations could be responsible for a significant fraction of the net vertical dust transport that occurs, at least during the aphelion season. However, vertical dust distributions were also found to vary with the physical parameterisations used in the model, a result which will require further analysis in future.

Reductions of up to an order of magnitude in winter polar dust loadings at 60–90° were seen for the larger R_s values, caused by scavenging by polar hood clouds (potentially supplemented near the poles by scavenging by CO_2 snowfall). This dust removal appeared to be necessary to match the rapid decline in dust density-scaled opacity with increasing latitude seen in MCS observations, though the lack of data near the surface did not permit a measurement of the total dust content of the polar regions. The clearing of polar dust was the strongest piece of evidence that the scavenging efficiency is actually fairly high (> 0.1); if scavenging is less prevalent in reality, then another explanation is needed as to why the non-scavenging model produced too much dust at the poles. Also, a certain density of scavenged IN is required to provide enough dust cores to keep the effective radius of ACB ice particles down to the sizes of a few microns that have been observed; though uncertainty in the vertical IN profiles made it difficult to gauge how large R_s must be to satisfy this. On the other hand, the amount of cloud ice produced by the model, particularly over the north polar hood, suggested a lower efficiency would be more appropriate. For any reasonable value of R_s, wet deposition due to scavenging by water ice made a non-negligible contribution to the total dust deposition over the polar caps.

7.4 Winter polar meteorology

The autumn and winter midlatitude and polar regions were examined in more detail, in a series of simulations using prescribed dust opacities, to determine the causes and characteristics of the solsticial pause, a midwinter minimum in eddy activity that is the reason behind the cessation of frontal cap-edge dust storms around this time. It was found that both realistic time-varying dust opacities and the presence of radiatively active water ice clouds were necessary for the model to best reproduce the length and depth of the solsticial

pauses observed (in TES data, and through the assimilation of this into the UKMGCM) in each hemisphere. In the model, the occurrence of the pause was linked to a reduction in vertical wind shear near the surface, due both to an increased tilt of the westerly winter jet forced by increased dust loading (in the NH) and by changes to its structure nearer the ground in the presence of the polar hood clouds. This effect was seen to depend strongly on the presence of zonal variations in topography in the midlatitudes of both hemispheres, and in fact topographic suppression was identified as the primary reason for the weakness in baroclinic activity in the southern hemisphere, relative to the NH.

The need for active ice clouds in order to correctly capture the dynamics of the winter polar regions (at least in this model) illustrates the significant coupling that exists between dust and water in this area of the Martian climate. In the southern hemisphere, the solsticial pause in eddy activity affects not only dust lifting at the cap edge, but also cloud ice formation over the cap, which relies on wave transport of water vapour from lower latitudes, and dust penetration into the polar region, which requires both lifting and eddy mixing over the cap boundary. Therefore, the pause can also be observed in south polar hood ice opacity and in the dust content of the polar region. In the northern hemisphere, the solsticial pause should lead to less near-surface dust over the pole at solstice, but cloud ice formation appears to depend more strongly on the strength of vapour supply from the PMOC. Through their thermal impact, and potentially through scavenging, the polar hood clouds affect cap-edge baroclinic activity and, due to increasing recognition that cap-edge frontal dust storms are responsible for a large fraction of the total lifting that takes place on Mars, this means that understanding and correctly simulating ice clouds is vitally important to dust lifting modelling.

7.5 Future work

7.5.1 Sub-gridscale dust lifting parameterisations

Much work on improving the dust lifting schemes used by Mars GCMs should, and can, be done in the near future. It is my belief that parameterisations should focus on properly accounting for sub-gridscale (microscale) phenomena, which can drastically change the lifting predictions within a model gridbox. As mentioned earlier (in Chapters 3 and

4), these include temporal and spatial variability in windspeed, local variation in surface roughness, and the existence of distinct impact and fluid thresholds on Mars. The first two of these imply that as well as there being some gridpoints at which dust lifting is more common than it is at others, there should also exist areas within each gridbox from which dust is more easily raised than it is from the surrounding region. Because dust lifting is a threshold-dependent process, and because it appears not to be routinely achieved under typical Martian conditions, it is very likely that total emission will be dominated by certain 'hotspots', which is what is observed on Earth, and that the emission from such areas will be significantly greater than would be expected when considering lifting conditions averaged over a larger area. Some of these may be mesoscale regions with distinct circulations forced by topographic variations or otherwise, while others may be extremely localised and rely on favourable combinations of windspeed and surface properties such as roughness length and albedo.

Scaling these small-scale variations in dust emission up to the resolution used by MGCMs is a formidable challenge, but one which could be essential in order to accurately predict the location of key dust sources. Information gained from mesoscale and LES models such as those of Rafkin et al. (2001) or Spiga and Forget (2009) will be invaluable here: mesoscale simulations can tell us more about the kinds of topographic or thermally driven circulations that could be capable of producing particularly large surface winds, and LES models can directly simulate wind gustiness and guide us towards parameterisations that predict its strength as a function of, for example, boundary layer depth or mean windspeed. Results of such simulations should be compared with future in situ measurements of near-surface wind variability, such as will be made by the Curiosity Rover (Gómez-Elvira et al., 2009). Stochastic approaches in GCMs may be useful, in particular for dealing with sub-gridscale variations in z_0 (a stochastic method has already been used here, to simulate gustiness). The hysteresis effect that may result from the large difference between the two lifting thresholds (fluid and impact) should also be studied using LES models. This is something that would make dust lifting more commonplace than previously allowed by theory. It may be possible to observe the effect from a surface vehicle or lander, by recording changes in local wind velocity and correlating them with surface dust movement. It would be hoped that by considering these kinds of phenomena, each of which are likely to vary from place to place,

the spatial distribution of model dust lifting would shift from the results seen when the z_0 map was added, to something more realistic.

7.5.2 Other dust lifting issues

The availability of surface dust is, of course, the other essential factor in predicting the locations from which dust will be most readily emitted — no matter how good the conditions for lifting are at a particular site, it cannot act as a regular dust source if it is not easily resupplied with liftable dust particles. Clearly, this replenishment of surface dust does not occur instantaneously, and so the most active lifting sites in a model that does not consider the dust surface density are likely to be somewhat less active in reality, as they will quickly clear themselves of surface dust. It has been shown in this work that the assumption of an unlimited surface dust supply leads to unrealistic emission results (at least with a lifting formulation at the current level of complexity) whenever ice clouds or the true z_0 distribution are added to the model. A multiannual run using a variable threshold scheme like the one used in Chapter 4 is one way to create a preference for dust lifting at sites with a regular supply of dust, though it does assume that there are no deep reservoirs of dust carried over from historic climate regimes from which dust can be steadily lifted without the need for replenishment. An alternative approach for creating initial surface density conditions would be to use some derivative of the TES Dust Cover Index, though it is not clear how that field could be quantitatively converted into finite surface densities or — acting as a proxy — lifting thresholds. The similarities between the threshold field that was produced after several UKMGCM model years (particularly when clouds were included) and the DCI suggest that the two methods are close to equivalent.

Having set up the surface dust field, it is then important to know how much redistribution occurs on an annual timescale, and how strongly this affects the subsequent dust storm behaviour. The variable threshold simulations of Chapter 4 derived their interannual variability from a continual shifting of dust between several key source regions in the southern hemisphere, and the largest storms were produced only whenever more than one of these areas was well-stocked with dust. One would want to determine to what extent this behaviour actually applies in reality. Some work has been done on this by Szwast et al. (2006), by measuring changes in TES albedo maps, interannually and in comparison to the

Viking era; this approach should be continued with post-TES data. As we are considering the variability of events (GDSs) that occur only once every five Earth years or so, the main issue at present is the amount of data available to analyse.

Our knowledge of dust lifting processes will be improved in future by utilising the technique of data assimilation. In addition to the existing TES reanalysis dataset (used on several occasions during this work), which was produced by assimilating 3-D temperature and 2-D column dust opacity measurements into the UKMGCM, work is underway at Oxford and at The Open University to produce further assimilation products using the 3-D temperature, dust opacity and ice opacity observations made by MCS over the past three Mars years. By running the tracer cycles while assimilating temperature only, it is possible to evaluate the performance of the dust lifting and cloud physics schemes when forced with realistic atmospheric temperatures and wind velocities. More useful is the simultaneous assimilation of the 3-D opacity fields, which corrects the model's prediction of mixing ratios at each gridpoint for which observations are available at that time. By analysing the residuals from the correction step, one can determine where the model is over- or under-predicting dust amounts. Although the lack of MCS data close to the surface will present a problem, it should be possible to obtain quite robust information on where and when surface dust emission is occurring, and at what rate. This can then be used to test the model's lifting parameterisations: I would expect that there will exist many points at which dust emission is required by observations but not predicted by the model, due to relatively low mean surface windspeeds. This technique could, therefore, help to address the sub-gridscale problem described above. The combination of model and observation in this way can also provide much information that may not be available directly from satellite soundings, such as dust particle size distributions, and can 'fill in the gaps' in the observational dataset, including — in the case of MCS — during periods of high opacity.

7.5.3 Vertical transport and dust-cloud interaction

A pressing concern, in the UKMGCM at least, is the discrepancy between simulated and observed tracer vertical distributions during NH spring and summer. Work done in Chapter 6 suggested that vertical mixing was being severely underestimated at T31 resolution, though this improved with an increase in horizontal resolution. The problem affected not

only dust profiles but also vertical atmospheric temperature structure, ACB height and cross-equatorial dust and vapour transport. Getting the vertical distributions correct is of vital importance in simulating the Martian tracer cycles; however, if upward transport occurs primarily through the action of vertical winds forced by topographic slopes, it will inevitably be poorly reproduced by standard $\sim 5^\circ$ GCM simulations, and it is difficult to imagine how it could be parameterised. If particular regions of complex topography can be identified as dominant contributors to the net vertical transport of dust and to the formation of ice clouds at relatively low pressures, it may be possible to conduct GCM simulations with nested areas of increased resolution focused on these regions. However, the same method would also be of value at SH dust storm initiation sites, and if it were extended to include, for example, the Hellas basin as well as areas in the NH, the total nested area may end up covering a large fraction of the planet, which would be computationally demanding.

The coupling of the dust and water cycles depends very much on the efficiency of heterogeneous cloud nucleation, controlled in the UKMGCM cloud scheme by the parameter R_s, whose value is currently unknown. Laboratory nucleation experiments under Martian conditions (pressure, temperature) (e.g. Iraci et al., 2010) should reveal more about the fraction of available dust particles that become activated as ice nuclei, as a function of temperature and/or local supersaturation (i.e. suggesting a non-constant R_s), and confirm whether or not the 100 nm minimum size criterion for IN is in fact valid. Also, any in situ measurement of the dust content of cloud ice particles on Mars would be extremely useful for constraining R_s, as the work of Chapter 6 found that dust content is more sensitive to R_s (particularly in northern spring) than is ice mass/opacity. The likely first step for these measurements would be to look for large dust fractions ($\geq 5\%$, say), allowing an upper limit to be set on R_s through comparison with model predictions. In combination with further theoretical work, this should lead to the development of more accurate cloud microphysical schemes.

7.5.4 Paleoclimate studies

Finally, the ability to model the dust and water cycles in the manner described here will be invaluable for future studies of Martian paleoclimate. Previous attempts at simulating

the climates that existed under different configurations of Mars' orbital parameters, such as obliquity and solar longitude of perihelion, have focused solely on either dust lifting, using unlimited dust surface densities, or water and ice, using prescribed dust opacity conditions. The problem with the former is that lifting schemes tuned to replicate current conditions could produce unrealistic behaviour when applied to different climate regimes. For example, a finding of Newman et al. (2005) was that the stronger global circulation that would be produced in an increased-obliquity situation should lead to more intense dust lifting and more frequent and larger dust storms, perhaps with GDSs occurring every year. However, it is not clear how the finite amount of dust available on the planet's surface would affect this prediction — as shown earlier, the presence of strong surface wind stresses is generally not sufficient to cause continuous, intense dust lifting, and the depletion of source regions must play a role at some point. Using a feedback scheme like the variable threshold method used in this work provides a means of quantifying the effects of finite dust sources, but this would not be straightforward, due to the presence of the resupply term, which could effectively be tuned to allow any GDS frequency. A first guess would be to retain the setting found to work best under present conditions, but as this appeared to be making up for defects in annual cross-equatorial dust fluxes, an application to different climates would be quite uncertain. Nevertheless, it should be possible to roughly estimate the net movement of surface dust under different climate regimes, with the goal of determining the timescale on which the global dust cycle is closed.

For water cycle paleoclimate simulations, which aim to understand the long-term transport of water and the condensation and sublimation of surface ice deposits, the same problem of uncertainty over global dust opacities exists, whether these are prescribed — simply matching current dust loadings or using the results of dust lifting paleoclimate runs — or calculated online by simulating dust lifting and transport. Running the two tracer cycles together (which has not yet been done in this context) is certainly preferable, as it means the direct interaction between the two can be included, albeit with an uncertainty over the scavenging efficiency. If dust transport into and deposition within the polar regions can be accurately modelled for the present-day case (which will require a better knowledge of R_s), it may be possible in future to gather evidence on dust-ice fractions from the polar layered deposits, which have recorded information from a series of past climate regimes,

and use this to determine observationally what kind of dust loadings must have occurred at each stage in Mars' recent history to produce these polar deposition rates. Working backwards, it would be possible to infer past dust emission rates, and thus to understand how strongly the presence of a finite amount of surface dust affected the production of dust storms. The movement of water between deposits at the polar caps and elsewhere may be strongly dependent on the interaction of dust particles with the surface ice, either through changes to ice albedo caused by the presence of dust fractions within the ice or through the formation of dust lag deposits on top of ice caps or glaciers that can inhibit sublimation. The inclusion of dust transport and scavenging in water cycle simulations is therefore important in determining the long-term stability of surface ice deposits and thus in understanding Mars' climate history more generally.

7.6 In summary

To review the work presented in this thesis, recall the important questions identified in the opening chapter as providing the focus for the project:

1. How do dust cover and other surface properties affect major dust storm occurrence, and its interannual variability?

2. Why do cap-edge dust lifting rates vary as they do through autumn and winter, can models reproduce this variation, and is it related to the occurrence of larger dust storms?

3. How important is the interaction between atmospheric dust and water ice clouds, and what effect does this have on dust lifting and transport in the winter hemisphere?

To address the first of these, a variable threshold scheme was added to the UKMGCM to simulate the redistribution of surface dust during storms. This had a significant effect on dust storm production and allowed a much greater degree of interannual variability in long model runs. Temporal change in the surface dust cover at key lifting sites in the southern hemisphere therefore appears to be a plausible explanation for global dust storm interannual variability, if indeed wind stress lifting thresholds are modulated by surface dust availability. The timely return of dust to these sites has not yet been adequately

simulated by the model, but there are signs that this could be possible if vertical dust transport and cap-edge lifting were better represented. Dust lifting was found also to be strongly dependent on the surface roughness length, z_0, and the inclusion of a realistic map of this quantity necessitated severe reductions to calculated lifting threshold values, in order to maintain appropriate global dust lifting rates. It is suggested that attention should be devoted to developing sub-gridscale parameterisations to handle this effect as well as other small-scale effects that may be crucial in controlling dust lifting.

Concerning the second question, seasonal variation in cap-edge and flushing dust storms was best reproduced when radiatively active water ice clouds were included in the simulation. The cessation of dust lifting in northern midlatitudes around winter solstice was linked to a change in the westerly jet structure and reduction in the vertical wind shear close to the surface. This cessation was found to be assisted by the occurrence of a large dust storm at more southerly latitudes, but appears to be forced on a more fundamental level by zonal variations in the surface topography (in both hemispheres).

Lastly, water ice clouds, through their radiative impact, affect dust lifting rates more generally, and their inclusion in the model made the need for a 'finite surface dust' lifting approach more apparent. The direct interaction between dust and water ice, in-cloud scavenging, was found to be capable of reducing polar dust loadings by up to an order of magnitude; however, the efficiency of the process, R_s, is not well constrained. Determining R_s on the basis of model ice opacity — particularly in the aphelion cloud belt — was not possible because of the apparently inadequate vertical mixing of both dust nuclei and water vapour/ice that occurred. The sharp decline in dust mass mixing ratio with decreasing pressure seen at 10–50 Pa by Mars Climate Sounder was not matched by the model, and it appears unlikely that scavenging by water ice is capable of producing these kinds of dust profiles.

Appendix A

Selected list of symbols used

Symbol	Description
L_s	Degrees of solar longitude (0° = spring equinox)
g	Gravitational acceleration ($= 3.72\,\mathrm{m\,s^{-2}}$)
ρ_d	Dust particle density ($= 2500\,\mathrm{kg\,m^{-3}}$)
D_p	Dust particle diameter
w_{sed}	Particle sedimentation velocity
τ	Column opacity
u_*	Drag velocity
u_*^t	Threshold drag velocity
u_{*s}^t	Threshold drag velocity for a smooth surface
ζ	Surface wind stress
ζ^t	Threshold surface wind stress
F	Vertical dust flux
H	Horizontal sand (saltation) flux
α_N	Near-surface wind stress lifting efficiency
α_D	Dust devil lifting efficiency
E	Erodible fraction of surface OR below-cloud scavenging collection efficiency (Chapter 6)
${\zeta^t}_{inc}$	Threshold stress increase per unit dust surface density removed
${\zeta^t}_{dec}$	Threshold stress decrease rate
z_0	Aerodynamic surface roughness length
f_{eff}	Drag partition function ($= \frac{u_{*s}^t}{u_*^t}$)
r_c	Ice crystal radius
M_c	Gridbox ice mass
S	Water vapour saturation ratio
$\mathcal{N}$	Ice nuclei number density
r_{eff}	Dust particle size distribution effective radius
ν_{sed}	Ice crystal size distribution effective variance for sedimentation
R_s	In-cloud scavenging efficiency
R_{sbc}	Below-cloud scavenging efficiency

Appendix B

Calculation of drag partition coefficient f_{eff}

Several different formulations for f_{eff} have been suggested in the literature (Marticorena and Bergametti, 1995; Raupach et al., 1993; Alfaro and Gomes, 1995; Okin, 2008). The forms of Raupach et al. (1993) and Okin (2008) are given in terms of the roughness density, λ, rather than roughness length z_0, so it was necessary to obtain a conversion between these two variables. The roughness density λ, also called the frontal area index, measures the area projected in the direction of the mean wind by the roughness elements, and is defined, for N roughness elements of average height H and width D sitting on a surface area S, as

$$\lambda = \frac{NDH}{S}.$$

Using relationships[1] involving the average height and width of the roughness elements and the rock abundance, ${\zeta_a}^2$, namely:

$$\lambda = 0.3596\frac{\zeta_a}{100}$$

$$H = 0.3596D$$

$$D = \left(1.075\frac{\zeta_a}{100} + 0.201\right) \times 100$$

[1] From Eric Hèbrard, private communication.
[2] Unrelated to surface stress, ζ.

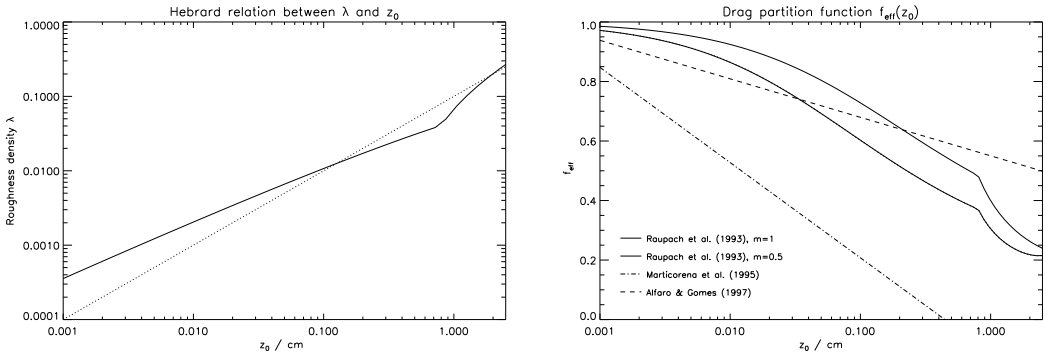

(a) (b)

Figure B.1: (a) The relation between λ and z_0 used here (solid) and the linear relation, following Wooding et al. (1973), $z_0 = \lambda \times 10\,\mathrm{cm}$ (dotted). (b) Several different drag partition functions suggested in the literature. The two Raupach et al. (1993) forms have been converted from functions of λ to z_0 using the relationship shown in (a) (explaining the kinks at 0.8 cm), and the form used for this work was the Raupach, $m = 1$ (thick solid line).

which can be combined to give

$$H = (1.075\lambda + 0.0723) \times 100,$$

together with the empirical relationship between z_0/H and λ[3]

$$z_0/H = \begin{cases} 10^{1.31 \log \lambda + 0.66} & \text{if } \lambda \leq 0.0408 \\ 10^{-1.16} & \text{otherwise} \end{cases}$$

a relation between λ and z_0 was found (shown in Figure B.1(a)):

$$\lambda \approx \begin{cases} 0.487 {z_0}^{0.763} - 0.441 {z_0}^{0.890} & \text{if } z_0 \leq 0.8\,\mathrm{cm} \\ 0.1345 z_0 - 0.0673 & \text{otherwise} \end{cases} \tag{B.1}$$

In practice, this is not greatly different to the use of the form of Wooding et al. (1973), $z_0 = \dfrac{H\lambda}{2}$, and assuming $H \sim 20\,\mathrm{cm}$ everywhere, which is a reasonable estimate of typical rock size on the Martian surface.

The various forms of the drag partition function f_{eff} are presented in Figure B.1(b). The Marticorena and Bergametti (1995) version can be ruled out immediately, as it applies only for $z_0 \leq 0.4\,\mathrm{cm}$, a value that is exceeded over much of the Martian z_0 map. Its

[3]From Eric Hèbrard, private communication.

prediction that no lifting whatsoever can occur for roughnesses $> 0.4\,\mathrm{cm}$ seems untenable for Mars; moreover, non-zero partition function values have been reported experimentally for λ values roughly equivalent to $z_0 > 0.4\,\mathrm{cm}$ ($\lambda > 0.04$; see Gillies et al. (2007) for a summary of field data). The function of Alfaro and Gomes (1995) came from a single experimental study so was not considered as reliable as that of Raupach et al. (1993). The Raupach function was selected for use, as it was a fit to several experimental studies, and was verified by recent measurements by Gillies et al. (2007) over the range $0.01 < \lambda < 0.1$. Both Okin (2008) and Shao and Yang (2005) extended the curve past its original limit of validity of $\lambda = 0.1$; however, below the largest value of z_0 used here ($\approx 2\,\mathrm{cm}$), the correction to Raupach required is small. The Raupach function is

$$f_{eff} = \left[\frac{1}{(1 - m\sigma\lambda)(1 + m\beta\lambda)}\right]^{\frac{1}{2}} \tag{B.2}$$

where $(\beta, \sigma, m) = (170, 2, 1)$ were suggested as providing the best fit to available experimental data. The value of m has been estimated experimentally at various values between 0 and 1, and its value depends on the type of roughness element involved. Vegetation on Earth has a lower m value due to its porosity, whereas solid Martian rocks are likely to require a large m, probably between 0.5 and 1 (Crawley and Nickling, 2003; Gillies et al., 2007). Okin (2008) suggested that m actually varies with λ ($0 \leq m \leq 1$). Though clearly uncertain, the value of $m = 1$ was retained. Use of $m = 0.5$ would result in lower lifting thresholds, but since the curves in Figure B.1(b) are similar in shape, this difference would be absorbed into the scaling factor applied to the threshold (see §4.9).

Bibliography

Alfaro, S. C. and Gomes, L. (1995). Improving the large-scale modeling of the saltation flux of soil particles in presence of nonerodible elements. *J. Geophys. Res.*, 100:16357–16366.

Almeida, M. P., Parteli, E. J. R., Andrade, J. S., and Herrmann, H. J. (2008). From the Cover: Giant saltation on Mars. *Proceedings of the National Academy of Science*, 105:6222–6226.

Andrews, D. G., Holton, J. R., and Leovy, C. B. (1987). *Middle atmosphere dynamics.* Academic Press, New York, NY, USA.

Bagnold, R. A. (1954). *The Physics of Blown Sand and Desert Dunes.* Methuen.

Banfield, D., Conrath, B. J., Gierasch, P. J., Wilson, R. J., and Smith, M. D. (2004). Traveling waves in the martian atmosphere from MGS TES Nadir data. *Icarus*, 170:365–403.

Barker, E. S., Schorn, R. A., Woszczyk, A., Tull, R. G., and Little, S. J. (1970). Mars: Detection of Atmospheric Water Vapor during the Southern Hemisphere Spring and Summer Season. *Science*, 170:1308–1310.

Barnes, J. R. (1980). Time spectral analysis of midlatitude disturbances in the Martian atmosphere. *J. Atmos. Sci.*, 37:2002–2015.

Barnes, J. R., Pollack, J. B., Haberle, R. M., Leovy, C. B., Zurek, R. W., Lee, H., and Schaeffer, J. (1993). Mars atmospheric dynamics as simulated by the NASA AMES General Circulation Model. II - Transient baroclinic eddies. *J. Geophys. Res.*, 98:3125–3148.

Basu, S., Richardson, M. I., and Wilson, R. J. (2004). Simulation of the Martian dust cycle with the GFDL Mars GCM. *J. Geophys. Res. (Planets)*, 109(E18):11006–+.

Basu, S., Wilson, J., Richardson, M., and Ingersoll, A. (2006). Simulation of spontaneous and variable global dust storms with the GFDL Mars GCM. *J. Geophys. Res. (Planets)*, 111(E10):9004–+.

Benson, J. L., Kass, D. M., and Kleinböhl, A. (2011). Mars' north polar hood as observed by the Mars Climate Sounder. *J. Geophys. Res. (Planets)*, 116(E15):E03008.

Benson, J. L., Kass, D. M., Kleinböhl, A., McCleese, D. J., Schofield, J. T., and Taylor, F. W. (2010). Mars' south polar hood as observed by the Mars Climate Sounder. *J. Geophys. Res. (Planets)*, 115(E14):E12015.

Blumsack, S. L. and Gierasch, P. J. (1972). Mars: The effects of topography on baroclinic instability. *J. Atmos. Sci.*, 29(6):1081–1089.

Böttger, H. M., Lewis, S. R., Read, P. L., and Forget, F. (2005). The effects of the martian regolith on GCM water cycle simulations. *Icarus*, 177:174–189.

Cakmur, R. V., Miller, R. L., and Torres, O. (2004). Incorporating the effect of small-scale circulations upon dust emission in an atmospheric general circulation model. *J. Geophys. Res. (Atmospheres)*, 109:7201.

Cantor, B. A. (2007). MOC observations of the 2001 Mars planet-encircling dust storm. *Icarus*, 186:60–96.

Cantor, B. A., James, P. B., Caplinger, M., and Wolff, M. J. (2001). Martian dust storms: 1999 Mars Orbiter Camera observations. *J. Geophys. Res.*, 106:23653–23688.

Cantor, B. A., Kanak, K. M., and Edgett, K. S. (2006). Mars Orbiter Camera observations of Martian dust devils and their tracks (September 1997 to January 2006) and evaluation of theoretical vortex models. *J. Geophys. Res. (Planets)*, 111:E12002.

Capen, C. F. (1974). A Martian yellow cloud - July 1971. *Icarus*, 22:345–362.

Chang, E. K. M. (2001). GCM and Observational Diagnoses of the Seasonal and Interannual Variations of the Pacific Storm Track during the Cool Season. *J. Atmos. Sci.*, 58:1784–1800.

Christensen, P. R. (1988). Global albedo variations on Mars - Implications for active aeolian transport, deposition, and erosion. *J. Geophys. Res.*, 93:7611–7624.

Christensen, P. R., Anderson, D. L., Chase, S. C., Clark, R. N., Kieffer, H. H., Malin, M. C., Pearl, J. C., Carpenter, J., Bandiera, N., and Brown, F. G. (1992). Thermal emission spectrometer experiment - Mars Observer mission. *J. Geophys. Res.*, 97:7719–7734.

Christensen, P. R., Bandfield, J. L., Hamilton, V. E., Ruff, S. W., Kieffer, H. H., Titus, T. N., Malin, M. C., Morris, R. V., Lane, M. D., Clark, R. L., Jakosky, B. M., Mellon, M. T., Pearl, J. C., Conrath, B. J., Smith, M. D., Clancy, R. T., Kuzmin, R. O., Roush, T., Mehall, G. L., Gorelick, N., Bender, K., Murray, K., Dason, S., Greene, E., Silverman, S., and Greenfield, M. (2001). Mars Global Surveyor Thermal Emission Spectrometer experiment: Investigation description and surface science results. *J. Geophys. Res.*, 106:23823–23872.

Christensen, P. R., Jakosky, B. M., Kieffer, H. H., Malin, M. C., McSween, J. H. Y., Nealson, K., Mehall, G. L., Silverman, S. H., Ferry, S., Caplinger, M., and Ravine, M.

(2004). The Thermal Emission Imaging System (THEMIS) for the Mars 2001 Odyssey Mission. *Sp. Sci. Rev.*, 110:85–130.

Clancy, R. T., Grossman, A. W., Wolff, M. J., James, P. B., Rudy, D. J., Billawala, Y. N., Sandor, B. J., Lee, S. W., and Muhleman, D. O. (1996). Water vapor saturation at low altitudes around Mars aphelion: A key to Mars climate? *Icarus*, 122:36–62.

Clancy, R. T., Sandor, B. J., Wolff, M. J., Christensen, P. R., Smith, M. D., Pearl, J. C., Conrath, B. J., and Wilson, R. J. (2000). An intercomparison of ground-based millimeter, MGS TES, and Viking atmospheric temperature measurements: Seasonal and interannual variability of temperatures and dust loading in the global Mars atmosphere. *J. Geophys. Res.*, 105:9553–9572.

Clancy, R. T., Wolff, M. J., and Christensen, P. R. (2003). Mars aerosol studies with the MGS TES emission phase function observations: Optical depths, particle sizes, and ice cloud types versus latitude and solar longitude. *J. Geophys. Res. (Planets)*, 108:5098.

Claudin, P. and Andreotti, B. (2006). A scaling law for aeolian dunes on mars, venus, earth, and for subaqueous ripples. *Earth Planet Sc. Lett.*, 252(1-2):30–44.

Colaprete, A. and Toon, O. B. (2000). The radiative effects of martian water ice clouds on the local atmospheric temperature profile. *Icarus*, 145:524–532.

Colburn, D. S., Pollack, J. B., and Haberle, R. M. (1989). Diurnal variations in optical depth at Mars. *Icarus*, 79:159–189.

Collins, M., Lewis, S. R., Read, P. L., and Hourdin, F. (1996). Baroclinic Wave Transitions in the Martian Atmosphere. *Icarus*, 120:344–357.

Conrath, B. J. (1975). Thermal structure of the Martian atmosphere during the dissipation of the dust storm of 1971. *Icarus*, 24:36–46.

Crawley, D. M. and Nickling, W. G. (2003). Drag Partition for Regularly-Arrayed Rough Surfaces. *Boundary-Layer Meteorology*, 107:445–468. 10.1023/A:1022119909546.

Cutts, J. A. and Lewis, B. H. (1982). Models of climate cycles recorded in Martian polar layered deposits. *Icarus*, 50:216–244.

de Vaucouleurs, G. (1958). World-wide Observations of Mars in 1956. *Leaflet of the Astronomical Society of the Pacific*, 8:1–+.

Deming, D., Mumma, M. J., Espenak, F., Kostiuk, T., and Zipoy, D. (1986). Polar warming in the middle atmosphere of Mars. *Icarus*, 66:366–379.

Deng, Y. and Mak, M. (2006). Nature of the Differences in the Intraseasonal Variability of the Pacific and Atlantic Storm Tracks: A Diagnostic Study. *J. Atmos. Sci.*, 63:2602–2615.

Dick, A. L. (1990). A simple model for air/snow fractionation of aerosol components over the antarctic peninsula. *J. Atmos. Chem.*, 11:179–196. 10.1007/BF00053674.

Drube, L., Leer, K., Goetz, W., Gunnlaugsson, H. P., Haspang, M. P., Lauritsen, N., Madsen, M. B., Sørensen, L. K. D., Ellehoj, M. D., Lemmon, M. T., Morris, R. V., Blaney, D., Reynolds, R. O., and Smith, P. H. (2010). Magnetic and optical properties of airborne dust and settling rates of dust at the Phoenix landing site. *J. Geophys. Res. (Solid Earth)*, 115:0.

Fenton, L. K. and Michaels, T. I. (2010). Characterizing the sensitivity of daytime turbulent activity on Mars with the MRAMS LES: Early results. *International Journal of Mars Science and Exploration*, 5:159–171.

Fisher, J. A., Richardson, M. I., Newman, C. E., Szwast, M. A., Graf, C., Basu, S., Ewald, S. P., Toigo, A. D., and Wilson, R. J. (2005). A survey of Martian dust devil activity using Mars Global Surveyor Mars Orbiter Camera images. *J. Geophys. Res. (Planets)*, 110:E03004.

Fletcher, B. (1976). The erosion of dust by an airflow. *J. Phys. D Appl. Phys.*, 9:913–924.

Forget, F. (1998). Improved optical properties of the Martian atmospheric dust for radiative transfer calculations in the infrared. *Geophys. Res. Lett.*, 25:1105–1108.

Forget, F., Hourdin, F., Fournier, R., Hourdin, C., Talagrand, O., Collins, M., Lewis, S. R., Read, P. L., and Huot, J.-P. (1999). Improved general circulation models of the Martian atmosphere from the surface to above 80 km. *J. Geophys. Res.*, 104:24155–24176.

Forget, F., Hourdin, F., and Talagrand, O. (1998). CO_2 Snowfall on Mars: Simulation with a General Circulation Model. *Icarus*, 131:302–316.

Gillette, D. (1978). A wind tunnel simulation of the erosion of soil: Effect of soil texture, sandblasting, wind speed, and soil consolidation on dust production. *Atmospheric Environment (1967)*, 12(8):1735 – 1743.

Gillette, D. A. (1999). A qualitative geophysical explanation for 'hot spot' dust emitting source regions. *Contr. Atmos. Phys.*, 72(1):67 – 77.

Gillette, D. A., Fryrear, D. W., Gill, T. E., Ley, T., Cahill, T. A., and Gearhart, E. A. (1997). Relation of vertical flux of particles smaller than 10 μm to total aeolian horizontal mass flux at Owens Lake. *J. Geophys. Res.*, 102:26009–26016.

Gillies, J., Nickling, W., and King, J. (2007). Shear stress partitioning in large patches of roughness in the atmospheric inertial sublayer. *Boundary-Layer Meteorology*, 122:367–396. 10.1007/s10546-006-9101-5.

Gomes, L., Rajot, J. L., Alfaro, S. C., and Gaudichet, A. (2003). Validation of a dust production model from measurements performed in semi-arid agricultural areas of spain and niger. *CATENA*, 52(3-4):257–271. Wind Erosion in Europe.

Gómez-Elvira, J., Castañer, L., Lepinette, A., Moreno, J., Polko, J., Sebastian, E., Torres, J., Zorzano, M. P., and Rems Team (2009). REMS, an Instrument for Mars Science Laboratory Rover. In *Lunar and Planetary Institute Science Conference Abstracts*, volume 40 of *Lunar and Planetary Inst. Technical Report*, page 1540.

Greeley, R. and Iversen, J. D. (1985). *Wind as a geological process on Earth, Mars, Venus and Titan.* Cambridge Univ. Press.

Greeley, R., Whelley, P. L., Arvidson, R. E., Cabrol, N. A., Foley, D. J., Franklin, B. J., Geissler, P. G., Golombek, M. P., Kuzmin, R. O., Landis, G. A., Lemmon, M. T., Neakrase, L. D. V., Squyres, S. W., and Thompson, S. D. (2006). Active dust devils in Gusev crater, Mars: Observations from the Mars Exploration Rover Spirit. *J. Geophys. Res. (Planets)*, 111:E12S09.

Greeley, R., Whelley, P. L., Neakrase, L. D. V., Arvidson, R. E., Bridges, N. T., Cabrol, N. A., Christensen, P. R., Di, K., Foley, D. J., Golombek, M. P., Herkenhoff, K., Knudson, A., Kuzmin, R. O., Li, R., Michaels, T., Squyres, S. W., Sullivan, R., and Thompson, S. D. (2008). Columbia Hills, Mars: Aeolian features seen from the ground and orbit. *J. Geophys. Res. (Planets)*, 113:E06S06.

Haberle, R. M. (1986). Interannual variability of global dust storms on Mars. *Science*, 234:459–461.

Haberle, R. M., Kahre, M. A., Murphy, J. R., Christensen, P. R., and Greeley, R. (2006). Role of dust devils and orbital precession in closing the Martian dust cycle. *Geophys. Res. Lett.*, 33:19.

Haberle, R. M., Leovy, C. B., and Pollack, J. B. (1982). Some effects of global dust storms on the atmospheric circulation of Mars. *Icarus*, 50:322–367.

Hanel, R., Conrath, B., Hovis, W., Kunde, V., Lowman, P., Maguire, W., Pearl, J., Pirraglia, J., Prabhakara, C., Schlachman, B., Levin, G., Straat, P., and Burke, T. (1972). Investigation of the Martian Environment by Infrared Spectroscopy on Mariner 9 (A 5. 2). *Icarus*, 17:423–+.

Hansen, J. E. and Travis, L. D. (1974). Light scattering in planetary atmospheres. *Space Sci. Rev.*, 16:527–610.

Hayne, P. O. and Paige, D. A. (2009). Snow Clouds and the Carbon Dioxide Cycle on Mars. *AGU Fall Meeting Abstracts*, page A11.

Heavens, N. G., Benson, J. L., Kass, D. M., Kleinböhl, A., Abdou, W. A., McCleese, D. J., Richardson, M. I., Schofield, J. T., Shirley, J. H., and Wolkenberg, P. M. (2010). Water ice clouds over the Martian tropics during northern summer. *Geophys. Res. Lett.*, 37:18202.

Heavens, N. G., Richardson, M. I., Kleinböhl, A., Kass, D. M., McCleese, D. J., Abdou, W., Benson, J. L., Schofield, J. T., Shirley, J. H., and Wolkenberg, P. M. (2011). Vertical distribution of dust in the Martian atmosphere during northern spring and summer: High-altitude tropical dust maximum at northern summer solstice. *J. Geophys. Res. (Planets)*, 116:E01007.

Hèbrard, E., Listowski, C., Coll, P., Marticorena, B., Bergametti, G., Määttänen, A., Montmessin, F., and Forget, F. (2012). An aerodynamic roughness length map derived from extended martian rock abundance data. In press.

Heintz, W. D. (1958). On the Martian yellow cloud in 1956. *The Observatory*, 78:203–204.

Hollingsworth, J. L., Haberle, R. M., Barnes, J. R., Bridger, A. F. C., Pollack, J. B., Lee, H., and Schaeffer, J. (1996). Orographic control of storm zones on Mars. *Nature*, 380:413–416.

Hoose, C., Lohmann, U., Stier, P., Verheggen, B., and Weingartner, E. (2008). Aerosol processing in mixed-phase clouds in ECHAM5-HAM: Model description and comparison to observations. *J. Geophys. Res. (Atmospheres)*, 113:7210.

Hoskins, B. J. and Simmons, A. J. (1975). A multi-layer spectral model and the semi-implicit method. *Q. J. Roy. Meteor. Soc.*, 101:637–655.

Hourdin, F. (1992). A new representation of the absorption by the CO2 15-microns band for a Martian general circulation model. *J. Geophys. Res.*, 97:18319–+.

Hourdin, F. and Armengaud, A. (1999). The Use of Finite-Volume Methods for Atmospheric Advection of Trace Species. Part I: Test of Various Formulations in a General Circulation Model. *Mon. Weather Rev.*, 127:822–+.

Hourdin, F., Forget, F., and Talagrand, O. (1995). The sensitivity of the Martian surface pressure and atmospheric mass budget to various parameters: A comparison between numerical simulations and Viking observations. *J. Geophys. Res.*, 100:5501–5523.

Hourdin, F., Le van, P., Forget, F., and Talagrand, O. (1993). Meteorological Variability and the Annual Surface Pressure Cycle on Mars. *J. Atmos. Sci.*, 50:3625–3640.

Iraci, L. T., Phebus, B. D., Stone, B. M., and Colaprete, A. (2010). Water ice cloud formation on Mars is more difficult than presumed: Laboratory studies of ice nucleation on surrogate materials. *Icarus*, 210:985–991.

Iversen, J. D. and Rasmussen, K. R. (1999). The effect of wind speed and bed slope on sand transport. *Sedimentology*, 46(4):723–731.

Jakosky, B. M. and Farmer, C. B. (1982). The seasonal and global behavior of water vapor in the Mars atmosphere - Complete global results of the Viking atmospheric water detector experiment. *J. Geophys. Res.*, 87:2999–3019.

James, I. N. and Gray, L. J. (1986). Concerning the effect of surface drag on the circulation of a baroclinic planetary atmosphere. *Q. J. Roy. Meteor. Soc.*, 112:1231–1250.

James, P. B. and Lee, S. W. (1999). Hubble Space Telescope Observations of Planets and Satellites. *Annu. Rev. Earth Pl. Sc.*, 27:115–148.

Joshi, M. M., Lewis, S. R., Read, P. L., and Catling, D. C. (1995). Western boundary currents in the Martian atmosphere: Numerical simulations and observational evidence. *J. Geophys. Res.*, 100:5485–5500.

Kahre, M. A., Hollingsworth, J. L., Haberle, R. M., and Murphy, J. R. (2008). Investigations of the variability of dust particle sizes in the martian atmosphere using the NASA Ames General Circulation Model. *Icarus*, 195:576–597.

Kahre, M. A., Murphy, J. R., and Haberle, R. M. (2006). Modeling the Martian dust cycle and surface dust reservoirs with the NASA Ames general circulation model. *J. Geophys. Res. (Planets)*, 111:E06008.

Kahre, M. A., Murphy, J. R., Haberle, R. M., Montmessin, F., and Schaeffer, J. (2005). Simulating the Martian dust cycle with a finite surface dust reservoir. *Geophys. Res. Lett.*, 32:20204–+.

Kok, J. F. (2010). Analytical calculation of the minimum wind speed required to sustain wind-blown sand on Earth and Mars. *ArXiv e-prints*.

Kok, J. F. and Renno, N. O. (2008). Electrostatics in Wind-Blown Sand. *Phys. Rev. Lett.*, 100(1):014501–+.

Kuroda, T., Medvedev, A. S., Hartogh, P., and Takahashi, M. (2007). Seasonal changes of the baroclinic wave activity in the northern hemisphere of Mars simulated with a GCM. *Geophys. Rev. Lett.*, 34:L09203.

Langevin, Y., Bibring, J.-P., Montmessin, F., Forget, F., Vincendon, M., Douté, S., Poulet, F., and Gondet, B. (2007). Observations of the south seasonal cap of Mars during recession in 2004-2006 by the OMEGA visible/near-infrared imaging spectrometer on board Mars Express. *J. Geophys. Res.*, 112:8.

Langevin, Y., Douté, S., Vincendon, M., Poulet, F., Bibring, J.-P., Gondet, B., Schmitt, B., and Forget, F. (2006). No signature of clear CO_2 ice from the 'cryptic' regions in Mars' south seasonal polar cap. *Nature*, 442:790–792.

Laurent, B., Marticorena, B., Bergametti, G., Léon, J. F., and Mahowald, N. M. (2008). Modeling mineral dust emissions from the Sahara desert using new surface properties and soil database. *J. Geophys. Res. (Atmospheres)*, 113:D14218.

Lefèvre, F., Lebonnois, S., Montmessin, F., and Forget, F. (2004). Three-dimensional modeling of ozone on Mars. *J. Geophys. Res. (Planets)*, 109(E18):7004–+.

Lemmon, M. T., Wolff, M. J., Smith, M. D., Clancy, R. T., Banfield, D., Landis, G. A., Ghosh, A., Smith, P. H., Spanovich, N., Whitney, B., Whelley, P., Greeley, R., Thompson, S., Bell, J. F., and Squyres, S. W. (2004). Atmospheric Imaging Results from the Mars Exploration Rovers: Spirit and Opportunity. *Science*, 306:1753–1756.

Leovy, C. and Mintz, Y. (1969). Numerical Simulation of the Atmospheric Circulation and Climate of Mars. *J. Atmos. Sci.*, 26:1167–1190.

Lewis, S. R. and Read, P. L. (1995). An operational data assimilation scheme for the Martian atmosphere. *Adv. Sp. Res.*, 16:9–.

Lewis, S. R., Read, P. L., Conrath, B. J., Pearl, J. C., and Smith, M. D. (2007). Assimilation of thermal emission spectrometer atmospheric data during the Mars Global Surveyor aerobraking period. *Icarus*, 192:327–347.

Lindzen, R. S. (1981). Turbulence and stress owing to gravity wave and tidal breakdown. *J. Geophys. Res.*, 86:9707–9714.

Liu, J., Richardson, M. I., and Wilson, R. J. (2003). An assessment of the global, seasonal, and interannual spacecraft record of Martian climate in the thermal infrared. *J. Geophys. Res. (Planets)*, 108:5089–+.

Lopez-Puertas, M. and Lopez-Valverde, M. A. (1995). Radiative energy balance of CO_2 non-LTE infrared emissions in the martian atmosphere. *Icarus*, 114:113–129.

Lorenc, A. C., Bell, R. S., and Macpherson, B. (1991). The meteorological office analysis correction data assimilation scheme. *Q. J. Roy. Meteor. Soc.*, 117(497):59–89.

Lorenz, R. D. (1996). Martian surface wind speeds described by the Weibull distribution. *J. Spacecraft Rockets*, 33:754–756.

Lott, F. and Miller, M. J. (1997). A new subgrid-scale orographic drag parametrization: Its formulation and testing. *Q. J. Roy. Meteor. Soc.*, 123:101–127.

Määttänen, A., Vehkamäki, H., Lauri, A., Merikallio, S., Kauhanen, J., Savijärvi, H., and Kulmala, M. (2005). Nucleation studies in the Martian atmosphere. *J. Geophys. Res.-Planets*, 110:E02002.

Madeleine, J.-B. (2011). *Dust and clouds on Mars: remote sensing, modeling of climate feedbacks and paleoclimate applications*. PhD thesis, Université Pierre et Marie Curie.

Madeleine, J.-B., Forget, F., Millour, E., Montabone, L., and Wolff, M. J. (2011). Revisiting the radiative impact of dust on Mars using the LMD Global Climate Model. *J. Geophys. Res. (Planets)*, 116:11010.

Malin, M. C., Caplinger, M. A., and Davis, S. D. (2001). Observational Evidence for an Active Surface Reservoir of Solid Carbon Dioxide on Mars. *Science*, 294:2146–2148.

Malin, M. C., Carr, M. H., Danielson, G. E., Davies, M. E., Hartmann, W. K., Ingersoll, A. P., James, P. B., Masursky, H., McEwen, A. S., Soderblom, L. A., Thomas, P., Veverka, J., Caplinger, M. A., Ravine, M. A., Soulanille, T. A., and Warr En, J. L. (1998). Early Views of the Martian Surface from the Mars Orbiter Camera of Mars Global Surveyor. *Science*, 279:1681–+.

Malin, M. C. and Edgett, K. S. (2001). Mars Global Surveyor Mars Orbiter Camera: Interplanetary cruise through primary mission. *J. Geophys. Res.*, 106:23429–23570.

Maltagliati, L., Montmessin, F., Fedorova, A., Korablev, O., Forget, F., and Bertaux, J.-L. (2011). Evidence of Water Vapor in Excess of Saturation in the Atmosphere of Mars. *Science*, 333:1868–.

Marticorena, B. and Bergametti, G. (1995). Modeling the atmospheric dust cycle: 1. Design of a soil-derived dust emission scheme. *J. Geophys. Res.*, 100:16415–16430.

Martin, L. J. (1974a). The major Martian dust storms of 1971 and 1973. *Icarus*, 23:108–115.

Martin, L. J. (1974b). The major martian yellow storm of 1971. *Icarus*, 22:175–188.

Martin, L. J. (1984). Clearing the Martian air - The troubled history of dust storms. *Icarus*, 57:317–321.

Martin, L. J. and Zurek, R. W. (1993). An analysis of the history of dust activity on Mars. *J. Geophys. Res.*, 98:3221–3246.

Martin, S. J., Wang, P. K., and Pruppacher, H. R. (1980). A Theoretical Determination of the Efficiency with which Aerosol Particles are Collected by Simple Ice Crystal Plates. *J. Atmos. Sci.*, 37:1628–1638.

Martin, T. Z., Peterfreund, A. R., Miner, E. D., Kieffer, H. H., and Hunt, G. E. (1979). Thermal infrared properties of the Martian atmosphere. I - Global behavior at 7, 9, 11, and 20 microns. *J. Geophys. Res.*, 84:2830–2842.

Martínez-Alvarado, O., Montabone, L., Lewis, S. R., Moroz, I. M., and Read, P. L. (2009). Transient teleconnection event at the onset of a planet-encircling dust storm on Mars. *Annales Geophysicae*, 27:3663–3676.

Matsuo, K. and Heki, K. (2009). Seasonal and inter-annual changes of volume density of martian CO_2 snow from time-variable elevation and gravity. *Icarus*, 202:90–94.

McCleese, D. J., Heavens, N. G., Schofield, J. T., Abdou, W. A., Bandfield, J. L., Calcutt, S. B., Irwin, P. G. J., Kass, D. M., Kleinböhl, A., Lewis, S. R., Paige, D. A., Read, P. L., Richardson, M. I., Shirley, J. H., Taylor, F. W., Teanby, N., and Zurek, R. W. (2010). Structure and dynamics of the Martian lower and middle atmosphere as observed by the Mars Climate Sounder: Seasonal variations in zonal mean temperature, dust, and water ice aerosols. *J. Geophys. Res. (Planets)*, 115(E14):E12016.

McCleese, D. J., Schofield, J. T., Taylor, F. W., Calcutt, S. B., Foote, M. C., Kass, D. M., Leovy, C. B., Paige, D. A., Read, P. L., and Zurek, R. W. (2007). Mars Climate Sounder: An investigation of thermal and water vapor structure, dust and condensate distributions in the atmosphere, and energy balance of the polar regions. *J. Geophys. Res. (Planets)*, 112(E11):5–+.

McConnochie, T. H. and Smith, M. D. (2008). Vertically Resolved Aerosol Climatology from Mars Global Surveyor Thermal Emission Spectrometer (MGS-TES) Limb Sounding. *LPI Contributions*, 1447:9114.

Mellon, M. T., Jakosky, B. M., Kieffer, H. H., and Christensen, P. R. (2000). High-resolution thermal inertia mapping from the mars global surveyor thermal emission spectrometer. *Icarus*, 148(2):437 – 455.

Merrison, J. P., Gunnlaugsson, H. P., Nørnberg, P., Jensen, A. E., and Rasmussen, K. R. (2007). Determination of the wind induced detachment threshold for granular material on Mars using wind tunnel simulations. *Icarus*, 191:568–580.

Meyers, M. P., Demott, P. J., and Cotton, W. R. (1992). New Primary Ice-Nucleation Parameterizations in an Explicit Cloud Model. *J. Appl. Meteorol.*, 31:708–721.

Michelangeli, D. V., Toon, O. B., Haberle, R. M., and Pollack, J. B. (1993). Numerical simulations of the formation and evolution of water ice clouds in the Martian atmosphere. *Icarus*, 102:261–285.

Montabone, L., Lewis, S. R., and Read, P. L. (2005). Interannual variability of Martian dust storms in assimilation of several years of Mars global surveyor observations. *Adv. Sp. Res.*, 36:2146–2155.

Montmessin, F., Forget, F., Rannou, P., Cabane, M., and Haberle, R. M. (2004). Origin and role of water ice clouds in the Martian water cycle as inferred from a general circulation model. *J. Geophys. Res. (Planets)*, 109(E18):10004–+.

Montmessin, F., Haberle, R. M., Forget, F., Langevin, Y., Clancy, R. T., and Bibring, J.-P. (2007). On the origin of perennial water ice at the south pole of Mars: A precession-controlled mechanism? *J. Geophys. Res. (Planets)*, 112(E11):8–+.

Montmessin, F., Rannou, P., and Cabane, M. (2002). New insights into Martian dust distribution and water-ice cloud microphysics. *J. Geophys. Res. (Planets)*, 107:5037–+.

Murphy, J. R., Pollack, J. B., Haberle, R. M., Leovy, C. B., Toon, O. B., and Schaeffer, J. (1995). Three-dimensional numerical simulation of Martian global dust storms. *J. Geophys. Res.*, 100:26357–26376.

Nakamura, H. (1992). Midwinter Suppression of Baroclinic Wave Activity in the Pacific. *J. Atmos. Sci.*, 49:1629–1642.

Neakrase, L. D. V. and Greeley, R. (2010). Dust devils in the laboratory: Effect of surface roughness on vortex dynamics. *J. Geophys. Res. (Planets)*, 115:E05003.

Nelli, S. M., Murphy, J. R., Feldman, W. C., and Schaeffer, J. R. (2009). Characterization of the nighttime low-latitude water ice deposits in the NASA Ames Mars General Circulation Model 2.1 under present-day atmospheric conditions. *J. Geophys. Res.*, 114:E11003.

Newman, C. E. (2001). *Modelling the dust cycle in the Martian atmosphere*. PhD thesis, University of Oxford.

Newman, C. E., Lewis, S. R., and Read, P. L. (2005). The atmospheric circulation and dust activity in different orbital epochs on Mars. *Icarus*, 174:135–160.

Newman, C. E., Lewis, S. R., Read, P. L., and Forget, F. (2002a). Modeling the Martian dust cycle, 1. Representations of dust transport processes. *J. Geophys. Res. (Planets)*, 107:5123–+.

Newman, C. E., Lewis, S. R., Read, P. L., and Forget, F. (2002b). Modeling the Martian dust cycle 2. Multiannual radiatively active dust transport simulations. *J. Geophys. Res. (Planets)*, 107:5124–+.

Nowicki, S. A. and Christensen, P. R. (2007). Rock abundance on Mars from the Thermal Emission Spectrometer. *J. Geophys. Res. (Planets)*, 112:E05007.

Ockert-Bell, M. E., Bell, J. F., Pollack, J. B., McKay, C. P., and Forget, F. (1997). Absorption and scattering properties of the Martian dust in the solar wavelengths. *J. Geophys. Res.*, 102:9039–9050.

Okin, G. S. (2005). Dependence of wind erosion and dust emission on surface heterogeneity: Stochastic modeling. *J. Geophys. Res. (Atmospheres)*, 110(D9):11208–+.

Okin, G. S. (2008). A new model of wind erosion in the presence of vegetation. *J. Geophys. Res.*, 113(F12):F02S10.

Pankine, A. A. and Ingersoll, A. P. (2004). Interannual variability of Mars global dust storms: an example of self-organized criticality? *Icarus*, 170:514–518.

Peixoto, J. P. and Oort, A. H. (1992). *Physics of climate*. New York: American Institute of Physics (AIP).

Phillips, T. E. R. (1924). Mars, Observations of, in 1924. *Mon. Not. R. Astron. Soc.*, 85:179–+.

Pollack, J. B., Colburn, D. S., Flasar, F. M., Kahn, R., Carlston, C. E., and Pidek, D. G. (1979). Properties and effects of dust particles suspended in the Martian atmosphere. *J. Geophys. Res.*, 84:2929–2945.

Pollack, J. B., Haberle, R. M., Schaeffer, J., and Lee, H. (1990). Simulations of the general circulation of the Martian atmosphere. I - Polar processes. *J. Geophys. Res.*, 95:1447–1473.

Pollack, J. B., Leovy, C. B., Greiman, P. W., and Mintz, Y. (1981). A Martian general circulation experiment with large topography. *J. Atmos. Sci.*, 38:3–29.

Priestley, A. (1993). A Quasi-Conservative Version of the Semi-Lagrangian Advection Scheme. *Monthly Weather Rev.*, 121:621–629.

Pruppacher, H. R. and Klett, J. D. (1978). *Microphysics of Clouds and Precipitation.* Dordrecht, Holland ; Boston : D. Reidel Pub. Co.

Putzig, N. E. and Mellon, M. T. (2007). Apparent thermal inertia and the surface heterogeneity of Mars. *Icarus*, 191:68–94.

Rafkin, S. C. R., Haberle, R. M., and Michaels, T. I. (2001). The Mars Regional Atmospheric Modeling System: Model Description and Selected Simulations. *Icarus*, 151:228–256.

Rafkin, S. C. R., Sta. Maria, M. R. V., and Michaels, T. I. (2002). Simulation of the atmospheric thermal circulation of a martian volcano using a mesoscale numerical model. *Nature*, 419:697–699.

Raupach, M. R., Gillette, D. A., and Leys, J. F. (1993). The effect of roughness elements on wind erosion threshold. *J. Geophys. Res.*, 98:3023–3029.

Rennó, N. O., Burkett, M. L., and Larkin, M. P. (1998). A Simple Thermodynamical Theory for Dust Devils. *J. Atmos. Sci.*, 55:3244–3252.

Richardson, M. I. and Wilson, R. J. (2002a). A topographically forced asymmetry in the martian circulation and climate. *Nature*, 416:298–301.

Richardson, M. I. and Wilson, R. J. (2002b). Investigation of the nature and stability of the Martian seasonal water cycle with a general circulation model. *J. Geophys. Res.*, 107:5031.

Rossow, W. B. (1978). Cloud microphysics: Analysis of the clouds of earth, venus, mars and jupiter. *Icarus*, 36(1):1 – 50.

Rover Team (1997). Characterization of the martian surface deposits by the mars pathfinder rover, sojourner. *Science*, 278(5344):1765–1768.

Ruff, S. W. and Christensen, P. R. (2002). Bright and dark regions on Mars: Particle size and mineralogical characteristics based on Thermal Emission Spectrometer data. *J. Geophys. Res. (Planets)*, 107:5127–+.

Shao, Y. and Lu, H. (2000). A simple expression for wind erosion threshold friction velocity. *J. Geophys. Res.*, 105:22437–22444.

Shao, Y. and Yang, Y. (2005). A scheme for drag partition over rough surfaces. *Atmos. Environ.*, 39(38):7351 – 7361.

Smith, D. E., Zuber, M. T., and Neumann, G. A. (2001). Seasonal variations of snow depth on mars. *Science*, 294(5549):2141–2146.

Smith, M. D. (2002). The annual cycle of water vapor on Mars as observed by the Thermal Emission Spectrometer. *J. Geophys. Res. (Planets)*, 107:5115–+.

Smith, M. D. (2008). Spacecraft Observations of the Martian Atmosphere. *Annu. Rev. Earth Pl. Sc.*, 36:191–219.

Smith, M. D. (2009). THEMIS observations of Mars aerosol optical depth from 2002-2008. *Icarus*, 202:444–452.

Smith, M. D., Bandfield, J. L., Christensen, P. R., and Richardson, M. I. (2003). Thermal Emission Imaging System (THEMIS) infrared observations of atmospheric dust and water ice cloud optical depth. *J. Geophys. Res. (Planets)*, 108:5115.

Smith, M. D., Pearl, J. C., Conrath, B. J., and Christensen, P. R. (2000). Mars Global Surveyor Thermal Emission Spectrometer (TES) observations of dust opacity during aerobraking and science phasing. *J. Geophys. Res.*, 105:9539–9552.

Smith, M. D., Wolff, M. J., Clancy, R. T., and Murchie, S. L. (2009). Compact Reconnaissance Imaging Spectrometer observations of water vapor and carbon monoxide. *J. Geophys. Res. (Planets)*, 114:E00D03.

Sorensen, M. (2004). On the rate of aeolian sand transport. *Geomorphology*, 59(1-4):53–62. Aeolian Research: processes, instrumentation, landforms and palaeoenvironments.

Spiga, A. and Forget, F. (2009). A new model to simulate the Martian mesoscale and microscale atmospheric circulation: Validation and first results. *J. Geophys. Res. (Planets)*, 114(E13):2009–+.

Sportisse, B. (2007). A review of parameterizations for modelling dry deposition and scavenging of radionuclides. *Atmospheric Environment*, 41(13):2683–2698.

Stanzel, C., Pätzold, M., Williams, D. A., Whelley, P. L., Greeley, R., Neukum, G., and The HRSC Co-Investigator Team (2008). Dust devil speeds, directions of motion and general characteristics observed by the Mars Express High Resolution Stereo Camera. *Icarus*, 197:39–51.

Stier, P., Feichter, J., Kinne, S., Kloster, S., Vignati, E., Wilson, J., Ganzeveld, L., Tegen, I., Werner, M., Balkanski, Y., Schulz, M., Boucher, O., Minikin, A., and Petzold, A. (2005). The aerosol-climate model ECHAM5-HAM. *Atmospheric Chemistry & Physics*, 5:1125–1156.

Strausberg, M. J., Wang, H., Richardson, M. I., Ewald, S. P., and Toigo, A. D. (2005). Observations of the initiation and evolution of the 2001 Mars global dust storm. *J. Geophys. Res. (Planets)*, 110(E9):2006–+.

Sullivan, R., Arvidson, R., Bell, J. F., Gellert, R., Golombek, M., Greeley, R., Herkenhoff, K., Johnson, J., Thompson, S., Whelley, P., and Wray, J. (2008). Wind-driven particle mobility on Mars: Insights from Mars Exploration Rover observations at "El Dorado" and surroundings at Gusev Crater. *J. Geophys. Res. (Planets)*, 113:6.

Sullivan, R., Greeley, R., Kraft, M., Wilson, G., Golombek, M., Herkenhoff, K., Murphy, J., and Smith, P. (2000). Results of the Imager for Mars Pathfinder windsock experiment. *J. Geophys. Res.*, 105:24547–24562.

Sutton, J. L., Levoy, C. B., and Tillman, J. E. (1978). Diurnal variations of the Martian surface layer meteorological parameters during the first 45 sols at two Viking Lander sites. *J. Atmos. Sci.*, 35:2346–2355.

Szwast, M. A., Richardson, M. I., and Vasavada, A. R. (2006). Surface dust redistribution on Mars as observed by the Mars Global Surveyor and Viking orbiters. *J. Geophys. Res. (Planets)*, 111(E10):11008–+.

Thorpe, T. E. (1979). A history of Mars atmospheric opacity in the southern hemisphere during the Viking extended mission. *J. Geophys. Res.*, 84:6663–6683.

Tillman, J. E. (1988). Mars global atmospheric oscillations - Annually synchronized, transient normal-mode oscillations and the triggering of global dust storms. *J. Geophys. Res.*, 93:9433–9451.

Toigo, A. D., Richardson, M. I., Wilson, R. J., Wang, H., and Ingersoll, A. P. (2002). A first look at dust lifting and dust storms near the south pole of Mars with a mesoscale model. *J. Geophys. Res. (Planets)*, 107:5050.

Tomasko, M. G., Doose, L. R., Lemmon, M., Smith, P. H., and Wegryn, E. (1999). Properties of dust in the Martian atmosphere from the Imager on Mars Pathfinder. *J. Geophys. Res.*, 104:8987–9008.

Toon, O. B., Pollack, J. B., and Sagan, C. (1977). Physical properties of the particles composing the Martian dust storm of 1971-1972. *Icarus*, 30:663–696.

Trenberth, K. E., Branstator, G. W., Karoly, D., Kumar, A., Lau, N.-C., and Ropelewski, C. (1998). Progress during TOGA in understanding and modeling global teleconnections associated with tropical sea surface temperatures. *J. Geophys. Res.*, 103:14291–14324.

Vincendon, M., Langevin, Y., Poulet, F., Bibring, J.-P., Gondet, B., Jouglet, D., and OMEGA Team (2008). Dust aerosols above the south polar cap of Mars as seen by OMEGA. *Icarus*, 196:488–505.

Wang, H. (2007). Dust storms originating in the northern hemisphere during the third mapping year of Mars Global Surveyor. *Icarus*, 189:325–343.

Wang, H., Richardson, M. I., Wilson, R. J., Ingersoll, A. P., Toigo, A. D., and Zurek, R. W. (2003). Cyclones, tides, and the origin of a cross-equatorial dust storm on Mars. *Geophys. Rev. Lett.*, 30(9):090000–1.

Wang, H., Zurek, R. W., and Richardson, M. I. (2005). Relationship between frontal dust storms and transient eddy activity in the northern hemisphere of Mars as observed by Mars Global Surveyor. *J. Geophys. Res. (Planets)*, 110(E9):7005–+.

Ward, W. R. (1974). Climatic variations on Mars. I. Astronomical theory of insolation. *J. Geophys. Res.*, 79:3375–3386.

Warren, S. G. (1984). Optical constants of ice from the ultraviolet to the microwave. *Appl. Opt.*, 23(8):1206–1225.

Whelley, P. L. and Greeley, R. (2006). Latitudinal dependency in dust devil activity on Mars. *J. Geophys. Res. (Planets)*, 111:E10003.

Whelley, P. L. and Greeley, R. (2008). The distribution of dust devil activity on Mars. *J. Geophys. Res. (Planets)*, 113:E07002.

White, B. R. (1979). Soil transport by winds on Mars. *J. Geophys. Res.*, 84:4643–4651.

Whiteway, J. A., Komguem, L., Dickinson, C., Cook, C., Illnicki, M., Seabrook, J., Popovici, V., Duck, T. J., Davy, R., Taylor, P. A., Pathak, J., Fisher, D., Carswell, A. I., Daly, M., Hipkin, V., Zent, A. P., Hecht, M. H., Wood, S. E., Tamppari, L. K., Renno, N., Moores, J. E., Lemmon, M. T., Daerden, F., and Smith, P. H. (2009). Mars Water-Ice Clouds and Precipitation. *Science*, 325:68–.

Wilson, R. J. and Hamilton, K. (1996). Comprehensive model simulation of thermal tides in the Martian atmosphere. *J. Atmos. Sci.*, 53:1290–1326.

Wilson, R. J. and Kahre, M. A. (2010). The role of spatially variable surface dust in GCM simulations of the Martian dust cycle. In *Mars Dust Cycle Workshop*, pages 108–112.

Wilson, R. J., Lewis, S. R., Montabone, L., and Smith, M. D. (2008). Influence of water ice clouds on Martian tropical atmospheric temperatures. *Geophys. Res. Lett.*, 350:L07202.

Wilson, R. J., Neumann, G. A., and Smith, M. D. (2007). Diurnal variation and radiative influence of Martian water ice clouds. *Geophys. Res. Lett.*, 340:L02710.

Wolff, M. J. and Clancy, R. T. (2003). Constraints on the size of Martian aerosols from Thermal Emission Spectrometer observations. *J. Geophys. Res. (Planets)*, 108:5097–+.

Wolff, M. J., Smith, M. D., Clancy, R. T., Spanovich, N., Whitney, B. A., Lemmon, M. T., Bandfield, J. L., Banfield, D., Ghosh, A., Landis, G., Christensen, P. R., Bell, J. F., and

Squyres, S. W. (2006). Constraints on dust aerosols from the Mars Exploration Rovers using MGS overflights and Mini-TES. *J. Geophys. Res. (Planets)*, 111(E10):12–+.

Wooding, R., Bradley, E., and Marshall, J. (1973). Drag due to regular arrays of roughness elements of varying geometry. *Boundary-Layer Meteorology*, 5:285–308. 10.1007/BF00155238.

Yu, J.-Y. and Hartmann, D. L. (1995). Orographic Influences on the Distribution and Generation of Atmospheric Variability in a GCM. *J. Atmos. Sci.*, 52:2428–2443.

Zent, A. P., Haberle, R. M., Houben, H. C., and Jakosky, B. M. (1993). A coupled subsurface-boundary layer model of water on Mars. *J. Geophys. Res.*, 98:3319–3337.

Zhang, Y. and Held, I. M. (1999). A Linear Stochastic Model of a GCM's Midlatitude Storm Tracks. *J. Atmos. Sci.*, 56:3416–3435.

Zimbelman, J. R. (2000). Non-active dunes in the Acheron Fossae region of Mars between the Viking and Mars Global Surveyor eras. *Geophys. Res. Lett.*, 27:1069–1072.

Zuber, M. T., Smith, D. E., Solomon, S. C., Abshire, J. B., Afzal, R. S., Aharonson, O., Fishbaugh, K., Ford, P. G., Frey, H. V., Garvin, J. B., Head, J. W., Ivanov, A. B., Johnson, C. L., Muhleman, D. O., Neumann, G. A., Pettengill, G. H., Phillips, R. J., Sun, X., Zwally, H. J., Banerdt, W. B., and Duxbury, T. C. (1998). Observations of the North Polar Region of Mars from the Mars Orbiter Laser Altimeter. *Science*, 282:2053–+.

Zuber, M. T., Smith, D. E., Solomon, S. C., Muhleman, D. O., Head, J. W., Garvin, J. B., Abshire, J. B., and Bufton, J. L. (1992). The Mars Observer laser altimeter investigation. *J. Geophys. Res.*, 97:7781–7797.

Zuber, M. T., Solomon, S. C., Phillips, R. J., Smith, D. E., Tyler, G. L., Aharonson, O., Balmino, G., Banerdt, W. B., Head, J. W., Johnson, C. L., Lemoine, F. G., McGovern, P. J., Neumann, G. A., Rowlands, D. D., and Zhong, S. (2000). Internal Structure and Early Thermal Evolution of Mars from Mars Global Surveyor Topography and Gravity. *Science*, 287:1788–1793.

Zurek, R. W. (1981). Inference of dust opacities for the 1977 Martian great dust storms from Viking Lander 1 pressure data. *Icarus*, 45:202–215.

Zurek, R. W. (1982). Martian great dust storms - an update. *Icarus*, 50:288–310.

Zurek, R. W. and Martin, L. J. (1993). Interannual variability of planet-encircling dust storms on Mars. *J. Geophys. Res.*, 98:3247–3259.