Implications of mountain shading on calculating energy for snowmelt using unstructured triangular meshes

A Thesis Submitted to the

College of Arts and Science

in Partial Fulfillment of the Requirements

for the degree of Master of Science

in the Department of Geography and Planning

University of Saskatchewan

Saskatoon

By

Christopher Marsh

Permission to Use

In presenting this thesis in partial fulfilment of the requirements for a Postgraduate degree from the University of Saskatchewan, I agree that the Libraries of this University may make it freely available for inspection. I further agree that permission for copying of this thesis in any manner, in whole or in part, for scholarly purposes may be granted by the professor or professors who supervised my thesis work or, in their absence, by the Head of the Department or the Dean of the College in which my thesis work was done. It is understood that any copying or publication or use of this thesis or parts thereof for financial gain shall not be allowed without my written permission. It is also understood that due recognition shall be given to me and to the University of Saskatchewan in any scholarly use which may be made of any material in my thesis.

Requests for permission to copy or to make other use of material in this thesis in whole or part should be addressed to:

Head of the Department of Geography and Planning
117 Science Place
University of Saskatchewan
Saskatoon, SK Canada
S7N 5C8

Abstract

In many parts of the world, the snowmelt energy balance is dominated by net solar shortwave radiation. This is the case in the Canadian Rocky Mountains, where clear skies dominate the winter and spring. In mountainous regions, solar irradiance at the snow surface is not only affected by solar angles, atmospheric transmittance, and the slope and aspect of immediate topography, but also by shadows from surrounding terrain. Many hydrological models do not consider such horizon-shadows. The accumulation of errors in estimating solar irradiance by neglecting horizon-shadows can lead to significant errors in calculating the timing and rate of snowmelt due to the seasonal storage of internal energy in the snowpack.

A common approach to representing the landscape is through structured meshes. However, such representations introduce errors due to the rigid nature of the mesh, creating artefacts and other constraints. Unstructured triangular meshes are more efficient in their representation of the terrain by allowing for a variable resolution. These meshes do not suffer from the artefact problems of a structured mesh.

This thesis demonstrates the increased accuracy of using a horizon-shading model with an unstructured mesh versus standard self-shading algorithms in Marmot Creek Research Basin (MCRB), Alberta, Canada. A systematic basin-wide over-prediction (basin mean expressed as phase change mass: 14 mm, maximum: 200 mm) in net shortwave is observed when only self-shadows are considered. The horizon-shadow model was run at a point scale at three sites throughout MCRB to investigate the effects of scale on the model results. It was found that small triangles were best suited for this topographic region and that shadow patterns were captured accurately. Large triangles were found to be too easily shaded by the model, created many disjointed

regions. As well, model results were compared to measurements of mountain shadows by timelapse digital cameras. These images were orthorectified and the shadow regions extracted allowing for a quantitative comparison. It was found that the horizon-model produced results within 10 m of the measured shadows, and properly captured shadow transits.

A point-scale energy balance model SNOBAL was run via The Cold Regions Hydrological Model, an HRU based hydrologic model. It was found that in the highly shaded valleys, snowpack ablation could be incorrect by approximately 4 days. Although MCRB was generally not significantly impacted by the over-estimation in irradiance in this study, insight into the horizon-shadowing process was possible as a result of the existing network of radiometers and other meteorological stations at MCRB. Because down-stream processes such as flooding depend on correct headwater snowmelt predictions, quantitative results demonstrating inaccuracies in a modelled component of the surface energy balance can help improve snowmelt modelling.

Acknowledgements

I would like to thank May Guan and Logan Fang for their help with fieldwork. Their company on the long hikes up to Fisera Ridge was invaluable. I wish to thank Kevin Shook for providing advice, to my office mates Chris DeBeer, Matt MacDonald, and Ross Phillips for providing a fantastic work environment as well as providing valuable feedback. Mike Solohub's previous work at Marmot was greatly appreciated. I wish to thank the Numerical Simulation Laboratory at the University of Saskatchewan for access to the computers used for the simulations, as well as to the members Jason Boisvert, Megan Marsh, and Andrew Kroshko of the NSL for their insight into the computational matters of this thesis. My wife, Megan Marsh, for support both in the field and at home. Her ability to listen to hydrology for hours is unmatched. As well I wish to thank my family for their support, for without them I would not have had such an opportunity.

I wish to thank David Gingras, Johannes Korsawe, John D'Errico, Willie Brink, Giuliano Langella, and Oliver Woodford. Although I do not know these individuals, their file submissions to the Matlab File Exchange (FEX) saved me numerous hours and ensured I did not need to reinvent the wheel.

Financial assistance, logistical support or equipment from the IP3 Network funded by the Canadian Foundation for Climate and Atmospheric Sciences, NSERC, MITACS, Environment Canada, Alberta Sustainable Resource Development, Biogeoscience Institute (University of Calgary), and Nakiska Mountain Resort.

Finally, I wish to thank my supervisors Drs. John Pomeroy and Raymond Spiteri for their guidance and their financial support. It was a pleasure and privilege to work with them.

To my wife and family.

Contents

List of Tables

List of Figures

List of Abbreviations

CRHM	Cold Regions Hydrological Model
CSC	Campbell Scientific Canada
DEM	Digital Elevation Model
DSLR	Digital Single Lens Reflex
ESRI	Environmental Systems Research Institute
GIS	Geographic Information System
GRU	Grouped Response Unit
HRU	Hydrological Response Unit
LOF	List of Figures
LOS	List of Symbols
LOT	List of Tables
LTS	Long Term Service
LiDAR	Light Detection And Ranging
MCRB	Marmot Creek Research Basin
MEC	Modélisation Environmentale Communautaire
MESH	Modélisation Environmentale Communautaire-Surface and Hydrology
PRMS	Precipitation-Runoff Modelling System
RAM	Random Access Memory
SAGA	System for Automated Geoscientific Analyses
SCA	Snow Covered Area
SCD	Snow Cover Depletion
SIMD	Single Instruction Multiple Data
SNOBAL	SNOw energy and mass BALance model
SSE	Streaming SIMD Extensions
SWE	Snow Water Equivalent
TIN	Triangulated Irregular Network
UTM	Universal Transverse Mercator
VIP	Very Important Points

List of Symbols

t	Time
Q_m	Energy available for melt ($\mathrm{W{\cdot}m^{-2}}$)
Q_{sn}	Net shortwave radiation ($\mathrm{W{\cdot}m^{-2}}$)
$K_s(t)$	Self-shaded direct-beam irradiance ($\mathrm{W{\cdot}m^{-2}}$)
Q_{ln}	Net longwave irradiance ($\mathrm{W{\cdot}m^{-2}}$)
Q_h	Sensible heat flux ($\mathrm{W{\cdot}m^{-2}}$)
Q_e	Latent heat flux ($\mathrm{W{\cdot}m^{-2}}$)
Q_g	Ground heat flux ($\mathrm{W{\cdot}m^{-2}}$)
Q_p	Precipitation heat flux ($\mathrm{W{\cdot}m^{-2}}$)
$\frac{dU}{dt}$	Rate of change of internal energy ($\mathrm{W{\cdot}m^{-2}}$)
α	Surface albedo
ρ_w	Density of water (1000 $\mathrm{kg{\cdot}m^{-3}}$)
S_0	Solar constant ($\approx$1370 $\mathrm{W{\cdot}m^{-2}}$)
$p(t)$	Zenith mean path transmissivity of the atmosphere
$m(t)$	Optical air mass
$\hat{\mathbf{n}}$	Triangle surface unit-normal vector
$\hat{\mathbf{s}}(t)$	Solar beam unit vector
$\delta(t)$	Solar declination
$Z(t)$	Zenith angle (radians)
A	Slope azimuth
$H(t)$	Hour angle
x', y', z'	Rotated Cartesian coordinate system
x, y, z	Standard Cartesian coordinate system
$\theta_s(t)$	Solar elevation (radians)
$\phi_s(t)$	Solar azimuth, clockwise from North (radians)
$\mathbf{R_x}$	x rotation matrix
$\mathbf{R_z}$	z rotation matrix
Θ_{sn}	Angle between $\hat{\mathbf{s}}(t)$ and $\hat{\mathbf{n}}$ (radians)
K_{meas}	Measured direct-beam irradiance ($\mathrm{W{\cdot}m^{-2}}$)
τ	Measured atmospheric transmittance
ζ	Binary horizon-shaded mask
$K_h(t)$	Horizon-shaded direct-beam irradiance ($\mathrm{W{\cdot}m^{-2}}$)
$\mathcal{O}$	Big-Oh
$\mathrm{Q}_{\Sigma\Delta}$	Cumulative difference in energy ($\mathrm{MJ{\cdot}m^{-2}}$)
h_f	Latent heat of fusion (0.334 $\mathrm{MJ{\cdot}kg^{-1}}$)

Chapter 1

Introduction

High spatial and temporal resolution, physics-based computer-aided simulations are an important tool for improving the understanding of the complex interactions in both the hydrologic system and related aspects of the climate system. A common characteristic of these models is that they require a large number of model parameters such as high-resolution topographic information, vegetation maps, soils profiles, and atmospheric forcing data in order to capture the spatial variability inherent in hydrological systems. Unfortunately, these parameters are often poorly known due to measurement constraints.

When available, these data are often not available for non-research basins, therefore limiting the areas where these models can be applied. Some authors (Sivapalan, 2003; McDonnell et al., 2007) have suggested that even within these research basins, high-resolution modelling does not improve the understanding of processes and that a new paradigm in modelling is required. In contrast however, other authors believe that high-resolution, physically correct models are an important tool for improving the understanding of hydrologic systems (Rigon et al., 2006). The limitations imposed by poor or non-existent data sets are gradually being reduced by ongoing advances in ground, aircraft, and satellite data collection. Satellite remote sensing has increased the ability to obtain certain types of high-resolution spatial data. A decade ago, topographic and vegetation information was not readily available at the required scale for most high-resolution hydrological models; however, with the advent and decreased cost of airborne mapping technologies such as high-resolution Light Detection And Ranging (LiDAR) derived Digital Elevation Models (DEMs), vegetation heights, and leaf area index maps have become more widespread. Al-

though characterizing the spatial variability of a basin's sub-surface properties is more challenging, recent research in applying geophysical techniques such as Ground Penetrating Radar (GPR) and Electrical Resistivity Imaging (ERI) (Berthold, 2004) hold significant promise. These and other advances suggest that in the coming years, a range of appropriate data sets will be available for carrying out large spatial scale (e.g., 2000 km^2), high-resolution hydrologic modelling.

Although an increase in spatial resolution of both models and observations brings with it a potential for increased accuracy and ability to understand the relative importance of various processes, parameters, and physical properties of the landscape, it does raise many research questions. For example, the effect of scale on model performance is sometimes not well known, and various hydrological processes may be important at different scales, both spatially and temporally (Klemeš, 1986). It is clear, however, that such high-spatial and temporal resolution modelling will require an increase in both the efficiency of computer code and computing power in order to improve the governing equation. Such efficiency will allow for faster simulation run times and produce accurate results allowing for increased usage in operational and experimental settings. Given recent advances in obtaining high-resolution spatial data, it is reasonable to concentrate on aspects of the hydrological system that are dependent upon test data and ensure correct prediction.

The surface energy balance, including both turbulent and radiative (shortwave and longwave) fluxes, are spatially variable as numerous studies have demonstrated, e.g., Pomeroy et al. (2003); Pohl et al. (2005). In many environments, such as low- and mid-latitude glaciers and high alpine terrain, shortwave fluxes dominate the surface energy balance, and for snow-covered surfaces it is often the main source of energy for melt (Gray and Male, 1981; Corripio, 2004). Given the importance of shortwave fluxes, correctly estimating the incoming shortwave solar radiation incident to a surface in complex terrain, both spatially and temporally, is necessary to model dependent processes.

There are a number of factors that make it difficult to estimate incident solar radiation across a basin at high resolution. First, as shortwave radiation passes

through a cloud-free atmosphere, the energy partially scatters and is absorbed by atmospheric constituents such as water vapour and other atmospheric gases. The effect of these factors on the quantity of radiation that reaches the earth's surface is referred to as the atmospheric transmittance (Marshak and Davis, 2005). Unfortunately, atmospheric transmittance is difficult to measure and requires knowledge or estimates of all of the atmospheric properties throughout the entire thickness of the atmosphere (Marshak and Davis, 2005). Once additional complexities such as clouds are taken into account, the estimation of the quantity and spatial variability of incoming solar radiation on a horizontal plane becomes more difficult. Another complexity is that shortwave radiation must also pass through any vegetation canopy before reaching the ground. It is also necessary to calculate the shortwave radiation that is incident on slopes of various angles and aspects by taking into account self-shading and shading from surrounding topography. The basic geometric equations for estimating shortwave radiation to a slope are well known and have been applied for many years, e.g., Garnier and Ohmura (1968); Dozier et al. (1981); Dozier and Frew (1990); Frew (1990); Fu and Rich (1999); Chueca and Julián (2004). However for an algorithm to be widely applicable it must be computationally efficient. As a result there still exists a need for improved methods for large scale, high resolution modelling of such a basic process. Therefore, there is a need to develop models that can properly account the spatial variability of shortwave radiation (the appropriate approximation may depend on the purpose of the modelling exercise) and that are able to be coupled with current-generation hydrological models.

The remainder of this thesis is structured into five sections. Chapter 2 is a review of the literature considering the current state of high-resolution modelling in hydrology. It considers some of the limitations that exist in current shadowing models based on structured meshes and outlines some of the reasons unstructured meshes are well suited for hydrological modelling finishing with an outline of the research gap. The hypotheses are formulates and presented. Chapter 3 is an outline of the field methods, the data processing, and data input. The implementation of an unstructured mesh shadowing model is presented, and a surface energy balance

model is described. Chapter 4 describes the results of the shadowing simulations at a point and basin scale and presents the results of the snowmelt energy calculations.

Chapter 2

Background

2.1 Snowmelt processes

2.1.1 Overview

The seasonal melt of snow cover provides water for streamflow that can affect ecosystems far downstream from the source snowpack (Groisman and Davies, 2001). For instance, many western and northern North American rivers derive the majority of their flow from late spring mountain snowpacks that then flow to low-lying regions (Stewart, 2009; Woo et al., 2008). Snowmelt often results in the largest discharge of the year, transporting contaminants that accumulated during the winter and recharging soil moisture reserves (Davies et al., 1987; Gray et al., 2001; Gray and Male, 1981). Snowpack accumulation and melt are of major importance for flood generation in the North American Prairies (Gray and Male, 1981; World Meteorological Organization, 1970). Li and Simonovic (2002) demonstrated that snow accumulation and melt were key variables for describing floods on the Red Deer and Assiniboine Rivers. Capturing the dynamical processes of a hydrological basin is required for proper flood forecasting (Li and Simonovic, 2002). Snowmelt is caused by energy inputs to the snowpack, and generally net shortwave radiation is the largest single input source controlling snowmelt (Male and Granger, 1981). When simulating snowmelt, an incorrect snowmelt prediction causes an error on the simulation day as well as an error later in the melt period (*two* errors total) (Walter et al., 2005).

In areas of complex topography, such as mountainous terrain, snowmelt calculations are complicated due to high spatial variability in energy fluxes (e.g., Marks and

Dozier (1992); Pomeroy et al. (2003)). Terrain slope and aspect have been identified throughout the literature as large contributors to this spatial variability (Carey and Woo, 1998; Dozier et al., 1981; Pomeroy et al., 2003; Ryerson, 1984; Woo and Young, 2004). In a catchment in the Yukon, Canada, Pomeroy et al. (2003) found that on a clear day at midday in the peak snowmelt period, a 20° south-facing slope received 80% more incoming solar shortwave radiation than a similarly steep, north-facing slope. This spatial variability leads to significant differences in snowmelt timing and magnitude throughout the basin and can limit melt on nested glaciers in shaded cirques (Chueca and Julián, 2004), limit snow melt (Carey and Woo, 1998; Pomeroy et al., 2003), cause changes in surface temperature (Pomeroy et al., 2003), and limit or increase photosynthesis and therefore affect subsequent vegetation patterns (Dymond, 2002). Chueca and Julián (2004) demonstrated that the spatial distribution of small cirque glaciers could be primarily attributed to the differences in incoming shortwave radiation, primarily due to shading. The spatial variability of topographically driven shortwave radiation is important in even gently rolling terrain, as shown by Pohl et al. (2006) in a sub-Arctic tundra river basin. Incorporating the effects of topography on shortwave radiation in a modelling framework is required for any of the processes dependent upon it, such as snowmelt, to be correctly modelled. Failure to account for even basic features such as slope and aspect has been shown to be detrimental to snowmelt and stream flow prediction (Dornes et al., 2008a). In modelling a sub-Arctic catchment, Davison et al. (2006) concluded that incorporating surface slope and aspect is required for proper snow-cover ablation and runoff timing. Pohl et al. (2006) identified that the small-scale variability in solar radiation for the melt period was an important contributing factor for the development of patchy snow cover in the tundra and that it also controlled the timing and magnitude of melt water release to the basin. Fundamentally, correct net radiation prediction is a principal contributor of realistic snowmelt models (Munro and Young, 1982).

2.1.2 Surface energy balance

The underlying physics of the snow melt process are governed by the available energy that is input to the snowpack. For a snow pack with an upper snow-air boundary and a lower snow-ground boundary the available energy for melt is given by (Gray and Male, 1981):

$$\mathrm{Q}_m = \mathrm{Q}_{sn} + \mathrm{Q}_{ln} + \mathrm{Q}_h + \mathrm{Q}_e + \mathrm{Q}_g + \mathrm{Q}_p - \frac{dU}{dt}, \tag{2.1}$$

where Q_m ($\mathrm{W{\cdot}m^{-2}}$) is the energy available for snowmelt, $\mathrm{Q}_{sn} = (1-\alpha)\mathrm{K}_s$ is the net shortwave radiation flux absorbed by the snow ($\mathrm{W{\cdot}m^{-2}}$), α is the snow-surface albedo to shortwave radiation (unitless), K_s is the shortwave radiation that has passed through the atmosphere ($\mathrm{W{\cdot}m^{-2}}$), Q_{ln} is the net longwave radiation flux at the snow-air boundary ($\mathrm{W{\cdot}m^{-2}}$), Q_h is the convective or sensible heat flux from the air at the snow-air boundary ($\mathrm{W{\cdot}m^{-2}}$), Q_e is the latent heat flux from evaporation and condensation at the snow-air boundary ($\mathrm{W{\cdot}m^{-2}}$), Q_g is the heat flux from the snow-ground boundary through conduction, Q_p is the heat flux from precipitation ($\mathrm{W{\cdot}m^{-2}}$), and $\frac{dU}{dt}$ is the rate of change of internal energy ($\mathrm{W{\cdot}m^{-2}}$). The total melt energy, Q_m, may then be used to determine a total melt volume via the latent heat of fusion (0.334 $\mathrm{MJ{\cdot}kg^{-2}}$) and the density of water (1000 $\mathrm{kg{\cdot}m^{-3}}$).

The cold content describes the amount of energy required to bring the snowpack up to 0 °C. Once the active layer of the snowpack becomes isothermal at 0 °C, melt can occur with a positive energy balance. A negative energy balance cools the active layer, thus increasing the cold content (Marks et al., 1998). Once melt occurs, melt water infiltrates through the snowpack and can pond at stratigraphic horizons (e.g., created by snow metamorphism) allowing for the formation of flow fingers. As these infiltrate to the the cold lower layers of the snowpack they can re-freeze, resulting in a major source of latent heat to the snowpack (Marsh and Woo, 1984).

Longwave radiation is the component of the electromagnetic spectrum with a wavelength that is approximately 4 μm to 100 μm (Petty, 2006) and is emitted from any object that has a temperature greater than 0 K (Wallace and Hobbs, 2006). In the context of Earth systems, longwave radiation is emitted from the atmosphere

(both towards space and the surface), the surface (including snow on the ground), and vegetation. This process contains a feedback system where the emitted longwave from the earth is absorbed by the atmosphere, thereby warming it, and resulting in a higher longwave output. At high latitudes (e.g., 60° N) , longwave radiation can provide potentially greater amounts of energy to the ground surface than from net shortwave radiation due to the low solar angles (Sicart et al., 2006). As a result, it can affect the timing of the the onset of snowmelt in early spring when the shortwave flux is low (Sicart et al., 2006).

During periods of strong winds and warm humid air, the turbulent fluxes can dominate the snowpack energy budget (Moore, 1983; Marks and Winstral, 2001). The turbulent flux of sensible heat, Q_h, exchanged at the surface of the snowpack is due to the temperature gradient between the snow surface and the overlying air mass. The turbulent flux of latent energy, Q_e, is due to the vapour pressure gradients between the snow surface and the overlying air mass.

2.1.3 Solar shortwave radiation

Solar shortwave radiation (irradiance) is emitted from the sun in the wavelengths of 0.1 μm to 4 μm (Petty, 2006; Wallace and Hobbs, 2006). Incident to the top of the Earth's atmosphere, the sun's incoming rays provide a flux of energy, termed the solar constant, of 1367 W·m^{-2}–1368 W·m^{-2} (Willson, 2003). As this energy is transmitted through the atmosphere, the total flux of energy is reduced due to absorption and reflectance by atmospheric particulates and aerosols. On a flat plane, the radiation incident to the ground is comprised of diffuse- and direct-beam components. Diffuse-beam shortwave radiation is defined as the energy flux produced by any solar photon that has had one or more interactions with any atmospheric constituent, and direct-beam shortwave radiation is any solar photon that has not interacted with any constituent of the atmosphere (Marshak and Davis, 2005). In complex terrain there is a third component: reflected direct- and diffuse-beam from nearby terrain.

Irradiance is an important component of the surface energy balance and net shortwave radiation can affect the spatial variability of ablation of a snowpack in a variety

of ways. Because of the slope and aspect of a surface effects the incident irradiance, the topographic variability present in a basin can greatly affect the spatial variability of surface energy. In addition, ablation (melting, sublimation) and redistribution processes (blowing snow) result in a patchwork of bare ground and snow cover that has a fractal pattern (Shook and Gray, 1996). The geometries of these patches control the contributing areas of snowmelt and runoff (Shook and Gray, 1996). These patchy locations allow for advection from the bare ground, warmed via incident net shortwave radiation due to a low albedo, to the surrounding snowcover thus increasing the rate of ablation (Gray and Male, 1981). DeBeer and Pomeroy (2009) and DeBeer and Pomeroy (2010) identified that snowcovers of various depths complicate estimates of snowcover depletion due to differences in cold content. Therefore locations with limited net shortwave radiation has a greater cold-content, resulting in a prolonged melt. Other small-scale processes can occur. For example when melt water percolates through the snowpack, it may be intercepted by an ice lens that can allow for lateral transport, changing the snow water equivalent (SWE) distribution (Elder et al., 1991). Through these processes, the spatial distribution of irradiance, and subsequently net shortwave radiation, may effectively move SWE between discrete parts of the basin.

Shortwave solar radiation must be modelled for periods of both clear and cloudy skies. During clear sky conditions, only a clear-sky transmittance parameter is required. For cloudy periods, the calculations become much more difficult due to the large spatial and temporal variability in cloud cover, density, and thickness. Frequently, other factors that are related to cloud formation are used to infer and approximate the cloud cover. These may include relative humidity, the presence or absence of precipitation, and daily maximum and minimum air temperatures. Models based on these proxies are generally empirical and involve site specific parameters to account for the large variability. In addition, the solar radiation passing through a vegetation canopy adds additional complexity to a model. There have been many parametrizations and physical models based on theoretical canopies. The effect of canopy transmittance is outside the scope of this review and is not be discussed here.

Incoming solar shortwave radiation (K_s [J· m^{-}2]) to a surface of a given slope and aspect between times t_1, t_2, is given by Garnier and Ohmura (1968):

$$K_s = S_0 \int_{t_1}^{t_2} [p(t)]^{m(t)} \cos(\arccos(\hat{\mathbf{n}} \cdot \hat{\mathbf{s}}(t)))dt, \tag{2.2}$$

where S_0 is the solar constant (W·m^{-2}), $p(t)$ is the zenith mean path transmissivity of the atmosphere, $m(t)$ is the optical air mass, $\hat{\mathbf{n}}$ is a unit vector normal to the slope, and $\hat{\mathbf{s}}(t)$ is a unit vector pointing to the position of the sun. Thus $\arccos(\hat{\mathbf{n}} \cdot \hat{\mathbf{s}}(t))$ is the angle between the sun and the slope. The two vectors are given as:

$$\hat{\mathbf{n}} = \begin{pmatrix} -\cos A \sin Z_n & \sin A \sin Z_n & \cos Z_n \end{pmatrix} \tag{2.3}$$

$$\hat{\mathbf{s}}(t) = \begin{pmatrix} 0 & \cos\delta(t) & \sin\delta(t) \end{pmatrix} \tag{2.4}$$

where A is the azimuth of the slope measured north through east, Z_n is the zenith angle of $\hat{\mathbf{n}}$, and $\delta(t)$ is the sun's declination positive north of the equator and negative south of the equator.

The angle between $\hat{\mathbf{n}}$ and $\hat{\mathbf{s}}(t)$ in Equation 2.2 can then be written in terms of a global coordinate system where

$$\begin{aligned} \cos(\arccos(\hat{\mathbf{n}} \cdot \hat{\mathbf{s}}(t))) = & \\ & \Big\{(\sin\phi \cos H(t))(-\cos A \sin Z_x(t)) - \sin H(t)\,(\sin A \sin Z_x(t)) \\ & + (\cos\phi \cos H(t)) \cos Z_x(t)\Big\} \cos\delta(t) \\ & + \Big\{\cos\phi\,(\cos A \sin Z_x(t)) + \sin\phi \cos Z_x(t)\Big\} \sin\delta(t) \end{aligned} \tag{2.5}$$

where ϕ is the latitude, $H(t)$ is the hour angle measured from solar noon, positive towards the west. An approximation for the optical air mass can be found from $m(t) = \sec Z(t)$ where $Z(t)$ is the zenith angle of the sun (Garnier and Ohmura, 1968). However, this approximation is only valid for $Z(t) \leq 70^\circ$ due to inaccuracies for low sun angles (Haltiner and Martin, 1957). As well, this approximation implies an infinite air mass at the horizon because $\lim_{Z(t)\to 90^\circ} \sec Z(t) = \infty$. Alternatively, precomputed values via a look-up table such as List (1971) may be used, as can empirically fit equations such as that developed by Young (1994) (Equation 2.6).

Young (1994) states that this expression has an error of 0.0037 at the horizon and that the precision in the coefficients is to mitigate roundoff errors; their precision does not indicate the accuracy of the equation or coefficients.

$$m(t) = \frac{1.002432 \cos^2 Z(t) + 0.148386 \cos Z(t) + 0.0096467}{\cos^3 Z(t) + 0.149864 \cos^2 Z(t) + 0.0102963 \cos Z(t) + 0.000303978}. \quad (2.6)$$

The clear-sky atmospheric transmissivity required to solve Equation 2.2 is difficult to obtain via direct measurements, especially over large spatial extents. It should be noted that the terms transmittance and transmissivity are frequently interchanged in the literature, and although similar, are significantly different. Atmospheric *transmittance* is the ratio of the total solar radiant flux that passes through the atmosphere to that incident at the top of the atmosphere. Atmospheric *transmissivity* is the ratio of solar radiant flux that passes through one unit length of atmosphere to the that that would have passed through one unit length of vacuum. The transmissivity and transmittance of the atmosphere are dependent upon the optical properties of the gases, aerosols, and clouds present in the atmosphere (Rasmussen et al., 1998). There is usually a critical lack of vertical resolution of these data (Marshak and Davis, 2005) and although remote sensing platforms are able to view clouds, there is a lack of quantitative error analysis to determine the reliability and accuracy of the retrieved data (Marshak and Davis, 2005). Sophisticated cloud models require large sets of boundary conditions to run and operate at a high vertical resolution and a high temporal resolution (Marshak and Davis, 2005). Although these models are useful, they require significant computational resources; thus there are many empirical relations that attempt to use atmospheric proxies such as relative humidity or temperature to infer cloud cover and atmospheric transmittance, e.g., (Walcek, 1994; Rasmussen et al., 1998; Annandale et al., 2002). For this study, these methods are not used due to the availability of measured data (see Section 3.4); however these methods are included as guidance for when such data are not available.

2.2 Shadows

Shadows occur as a result of surrounding terrain blocking a location from the direct-beam irradiance — that is they are cast from nearby terrain. These shadows are frequently referred to as *horizon-shadows* (e.g., Essery and Marks (2007)). *Self-shadows* form a subset of horizon-shadows that occur when a slope faces away from the sun. In this case, the self-shadow is cast by the terrain only onto itself. These two shadow types lead to terrain cycling in and out of shadow, creating large differences in incident shortwave radiation.

The solar radiation incident to the surface is also affected by variations in slope and aspect of the topography and by local terrain shadowing surrounding terrain (Dozier et al., 1981). This shading changes with time of day, time of year, and latitude. This cycling creates large differences in incident shortwave radiation and can play a principal role in, for example, limiting melt on glaciers that are nested in highly shaded cirques (Chueca and Julián, 2004), snow melt (Carey and Woo, 1998; Pomeroy et al., 2003), surface temperature (Pomeroy et al., 2003), and photo-synthesis and subsequent vegetation patterns (Dymond, 2002). In high-relief areas, such as in the Canadian Rocky Mountains, and in low-relief areas with short, steep slopes, the topography can completely block some locations from receiving any direct incoming solar radiation when the sun is low on the horizon. This is common near sunrise and sunset and also for much of the day during the winter months. In fact, much of the spatial variability in snowmelt energies can be attributed to topographic influences (Woo and Young, 2004). Pohl et al. (2005) identified a need to further study the topographic influences on the snowmelt energy balance, particularly the effect of shortwave solar radiation flux on a smaller scale.

When considering a shadowing model that incorporates surrounding topographic features, the most common approach is the hemisphere viewshed/horizon calculation. This algorithm can be described as follows. A structured mesh, rectangular in shape with a fixed spatial cell scale, is taken to define the topography. Then for each point in this mesh, the algorithm computes an upward-looking hemispherical

viewshed by sampling in all cardinal directions from a focal point and then determining the maximum elevation angle of the obscuring terrain (Hetrick et al., 1993). A continuous horizon is created by interpolating between the sampled horizon points. Examples of horizon-based algorithms are those of Fu and Rich (1999) that are incorporated into the ArcGIS® product by Environmental Systems Research Institute (ESRI) under the name SolarAnalyst, as well as those of Dozier et al. (1981), Dozier and Frew (1990), Frew (1990), and Varley et al. (1996). Other algorithms are less common but still used, such as via ray-tracing (Coquillart and Gangnet, 1984) or a sector/hull approach (Stewart, 1998). To determine if terrain is shaded, the horizon in the direction of the sun's azimuth is compared to the solar elevation and the terrain is shaded accordingly. One of the major difficulties with this type of algorithm is choosing a proper interval between sampling points when looking up and computing the horizon. Increasingly, high-resolution data sets are becoming available and this necessitates high-performance algorithms (Tabik et al., 2011) because more terrain elements must be checked. Romero et al. (2008) identified the horizon calculation and subsequent look-up in most shadow-aware models as the most computationally expensive aspect. There have been many attempts at producing increasingly efficient algorithms for the horizon calculation. Essery (2004) developed an alternative approach. By examining slope components for North Wales and the French Alps, it was found that a double-exponential (Laplace) distribution could describe the slopes. Using this, a statistical model was developed that expressed the fractions of the surfaces that are self- and horizon-shaded. This type of approach is suited for sub-grid shadow representation.

An approach that can be used with unstructured meshes is the algorithm of Montero et al. (2009). This is an extension of the parallel point-plane method originally introduced by Clarke (2001). It is an efficient algorithm and one that is easily implemented and modified, making it suitable for inclusion in a hydrology model. The parallel point-plane algorithm determines horizon-shaded areas as follows. A receiver area is an arbitrarily oriented plane defined in 3-space that can potentially be shadowed, and an obstructer area is an arbitrarily oriented plane in 3-space that

can cast a shadow. In this case, the topography is represented by triangles; therefore each triangle is potentially an obstructer and a receiver simultaneously. That is, all parts of the terrain can shade other terrain while being shaded. For this algorithm, the standard x, y, z Cartesian coordinate system is rotated so that the new vertical axis points at the sun. In this rotated configuration, some triangles may be located between the sun and other triangles, and by projecting the rotated coordinate system onto a plane and checking for intersecting triangles, shadow locations can be calculated.

2.3 Terrain representation

Hydrology has had a long history of trying to determining the best method to properly quantify and represent the model domain and to consider the spatial variability in many aspects of the hydrologic system, including solar radiation, as noted above. Correctly representing landscape heterogeneity is key for simulating many spatial processes, including, for example, snow accumulation and snowcover depletion (Grayson and Blöschl, 2000; Dornes et al., 2008b). Previous studies have categorized hydrologic modelling approaches by the way that the spatial representation of the model is conceptualized, with the three primary classes being distributed, semi-distributed, or lumped. Distributed models are appealing because they can potentially capture small scale variability when used at a sufficiently high-resolution, provided the physics are valid for such a scale. This approach involves representing the landscape via a discretizing mesh, where each computational element of the mesh (i.e., a grid cell or triangle) is characterized by elevation, vegetation cover, and soil conditions. Because of this, when the cells are made sufficiently small, high-resolution processes can be captured. One effect of an increase in spatial resolution is that such high-resolution modelling requires an increase in computing power to solve the controlling equations due to the larger number of times the equations must be solved. Although an increase in spatial resolution brings with it a potential for increased accuracy, it adds complexity. For example, it is often not known what an

ideal scale is to correctly model certain processes; different approximating equations and methods may only be applicable at certain scales (spatial or temporal), and various additional parameters (i.e., soil hydraulic conductivity) may also be required at this fine scale. However, such parameters may not be available at the required resolution. Hopkinson et al. (2010) examined modern high-resolution DEMs and their scale effects on glacial ablation, concluding that DEM resolution was important because it affected the underlying terrain representation, causing errors in slope, aspect, and shadow-generating topography. Mckenney (1999) noted that a coarser resolution mesh (100 m versus 20 m) smoothed out the spatial heterogeneity of irradiance.

In contrast, a lumped approach generalizes the model domain into one computational element that characterizes the entire model domain. It considers neither the spatial distribution in input variables nor the spatial variation in parameters (Clarke, 1973). In addition, input variables may be lumped temporally, such as treating rainfall as a daily value (Clarke, 1973). Because of this, lumped models are computationally inexpensive but are unable to capture any high-resolution spatial variability as has been shown on multiple occasions, such as in Pohl et al. (2005); Dornes et al. (2008b). Because the lumped approach does not take into account any spatially variable process, a compromise is often used. This compromise is generally the semi-distributed method, whereby the landscape is discretized by lumping together similar aspects of the basin that are expected to behave in a similar fashion. For example, all north-facing terrain might be considered as a single unit. This allows for some spatial heterogeneity without sacrificing computational speed or requiring a large set of input variables and parameters.

2.3.1 The Hydrological and Grouped Response Units

The Hydrological Response Unit (HRU) conceptualization is a semi-distributed approach that assumes that a basin can be subdivided into uniform elements (response units) that are delineated based on a combination of one or more of elevation, land cover, slope, aspect, soil type, and precipitation (Grayson and Blöschl, 2000). It is assumed that each response unit responds uniformly to a hydrological event; the

internal dynamics of the HRU are small compared to neighbouring HRUs (Flüegel, 1995). HRUs can be at any scale, from small scale such as a hillslope, to large scale, such as a sub-basin (Pomeroy et al., 2007). An example of such a delineation is 'forested, north-facing, and $< 20°$ slope'. Mass may be transferred between response units, thus simulating hydrological events such as runoff, blowing snow, etc. A variety of hydrological models exists that exploit this conceptualization, such as the Precipitation-Runoff Modelling System (PRMS) (Leavesley and Stannard, 1995) and the Cold Regions Hydrological Model (CRHM) (Pomeroy et al., 2007).

The Grouped Response Unit (GRU) conceptualization is a lumped approached utilizing a computational unit that consists of one or more fractional land cover characteristics (Kouwen et al., 1993). Specifically, a GRU is a grouping of all areas within an domain with a similar hydrological response (Soulis et al., 2005) regardless of location (Shaw et al., 2005). GRUs are often used with a grid approach to represent spatial variability within a basin. Each cell thus has associated GRUs and the cell size need not reflect the size of the hydrologic unit being modelled (Cranmer et al., 2001). However, the GRU elements are limited in size to an area that is subject to a uniform hydrological response due to uniform meteorological conditions or where travel times are small compared to the overall basin or the overall meteorological events of interest (Kouwen et al., 1993). The land-cover characteristics of a GRU all contribute to runoff and identical hydrological parameters are used for the same land cover over all elements in the model basin. GRUs are typically based on vegetation based land cover types. A defining characteristic of the GRUs approach is that each group of GRUs must route water directly to a stream. Hydrological models such as WATFLOOD (Kouwen et al., 1993) and the Modélisation Environmentale Communautaire (MEC)-Surface and Hydrology (MESH) (Pietroniro et al., 2007) model both exploit this concept to decrease model runtime and simplify calculations . Topographic information controlling spatial variability of solar radiation is generally omitted from GRU based models. There have been attempts to combine spatially distributed topographic information with the GRU approach, such as Dornes (2009); Pohl et al. (2005). In Dornes (2009), MESH was coupled with the Cold Regions

Hydrologic Model. CHRM was used to produce incoming radiation estimates that approximate the spatial variability of solar radiation. This coupling produced results with an increased accuracy. However, this was a coupling with a HRU model, not a fully distributed model. This suggests that there may be an increase in prediction accuracy when increasingly distributed models are used to predict spatially variable surface energetics.

2.3.2 Structured mesh

A structured mesh, also known as a raster or grid in the GIS nomenclature, is a landscape representation where a structured mesh divides the landscape into uniformly sized cells. A DEM is an example of a structured mesh, where each cell has an elevation value. Raster-based models are common (Tucker, 2001) because their computer representation is trivial to implement using two-dimensional arrays, a feature intrinsic to modern programming languages. Widespread use of rasters, such as satellite imagery, make using rasters a natural choice in models. Thus a raster is intrinsic to many applications.

Despite their widespread use, rasters have a number of significant limitations for usage in hydrological modelling. Tucker (2001) identified many limitations of rasters with one of the most significant being that hydrology does not function on uniform grids. Drainage directions are often constrained to 45° intervals when using structured meshes (Tucker, 2001). By shoehorning hydrological processes into structured grids, geometrical artefacts and other sub-grid variability can be artificially introduced (Tucker, 2001). Landscapes must also be represented by constant-sized cells that are spatially invariant through time. Representing topographic domains with both flat and hilly/mountainous terrain can result in excessive computations and data due to data redundancy (Jones et al., 1994). Because of the rigid, constant nature of a structured mesh, representing a non-rectangular hydrological basin requires buffering around the basin in order to capture its irregular shape with regular cells. This can result in having to either mask areas out, computing values for areas not used in the simulation, or focusing on a 'window' within the basin that can be fully

captured via a raster, neglecting the rest of the basin.

2.3.3 Unstructured triangular mesh

Unstructured triangular meshes, typically referred to as Triangulated Irregular Networks (TINs) in GIS literature, are a way of approximating the topography of the earth's surface and allow for distributed models. Unstructured meshes offer a solution to the problem of a fixed spatial resolution intrinsic to structured meshes. Unstructured meshes are generally comprised of a series of non-overlapping, connected triangles, where each triangle face is of a constant gradient (Chang, 2008). Because the triangles can have variable internal angles, they may have a variable size. Areas of large topographic variability can have a higher density of small triangles in order to capture the spatial variability and areas of relatively homogeneous topography have fewer large triangles.

In contrast to structured meshes, unstructured meshes have a number of advantages in hydrologic modelling. For example, because of the non-uniform spacing of points of an unstructured mesh, meandering features such as streams can be easily represented (Tucker, 2001). When river channel evolution was simulated by Braun and Sambridge (1997), the artificial constraints from using a raster were identified as the cause of creating non-natural channels. This was in contrast to the output from unstructured meshes that produced more natural-looking channels. Tucker (2001) suggested that unstructured meshes devised using unstructured meshes should be preferred over raster DEMs. Unstructured meshes can provide an optimal representation of the terrain using the least number of elements while maintaining conformance to the geometrical and physical properties of the basin (Kumar et al., 2009) to some predefined tolerance. Furthermore, landforms and associated topography need not be represented at the constant spatial resolution of a raster DEM, a resolution that generally corresponds to the highest resolution required to represent one area of the domain of interest (Tucker, 2001). Tachikawa et al. (1994) went as far to say that structured meshes are not appropriate for hydrological models due to their crude movement of water to only the eight adjacent cells (following the commonly

used D8 routing method (O'Callaghan and Marks, 1984)). That is, flow directions in a structured mesh are limited to the horizontal, vertical, and diagonal directions (Nelson et al., 1999). Computation of flow on an unstructured mesh avoids the complication of flow divergence (Zhou et al., 2011). Because of the reduction in spatial information needed to properly quantify the basin due to a more efficient terrain representation (Shewchuk, 1996), Ivanov et al. (2004) found that for an unstructured mesh, 5%–10% of the original number of grid cells could be removed. When representing a non-rectangular hydrological basin, structured meshes necessitate buffering around the basin in order to capture its irregular shape. Structured meshes create artificial scaling and model parametrization problems that lends needless uncertainty to model output. However, terrain resolution is still important in unstructured mesh models. Failing to select a resolution that properly delineates streams, for example, can lead to large discrepancies in hydrologic simulations (Vivoni et al., 2005).

However, unstructured mesh data structures are more complex than those for grids. Because of the differences in terrain representation, using unstructured meshes adds restrictions to the types of algorithms that can be used. For example, in order to incorporate horizon-shading on unstructured meshes, a different approach must be taken from that used in gridded approaches. Although it is sometimes possible to utilize the algorithms designed for structured meshes with unstructured meshes, there are increasing complexities associated with this. For example, the commonly used fixed spatial intervals often used to calculate the horizon can no longer be accurately used with unstructured meshes.

Creating an unstructured mesh from a set of points is referred to as meshing or triangulation. Refining a mesh involves adding triangles, coarsening (or derefining) a mesh involves removing triangles. The resulting triangular mesh is composed of faces, nodes, and edges (Chang, 2008). In the generation of these meshes, a guarantee about the shape of the triangle is necessary for numerical stability if solving partial differential equations (e.g., no thin, sharp triangles). When representing complex topography a valley could be an artefact of a poor meshing algorithm. The triangle shape is affected by the angles at the nodes and it is known that

excessively long, skinny triangles can lead to undesirable behaviour such as numerical instabilities or terrain artefacts (Ruppert, 1995). Therefore, there has been much interest in producing algorithms that produce triangles with certain geometrical characteristics. One example of an algorithm class used for the generation of triangles is the Delaunay refinement. This method satisfies guaranteed bounds on angles, edge lengths, the number of triangles, and the grading of triangles from small to large sizes (Shewchuk, 2002). It is extensively used and implementations can be found in many software packages such as MATLAB® , ArcGIS, and standalone tools like Triangle (Shewchuk, 1996). Examples and further discussion of meshing algorithms can be found in Rivara (1987); Bern and Eppstein (1992); Chew (1993); Ruppert (1995); Shewchuk (1996, 2002).

When converting a structured mesh to an unstructured mesh, it is important to select the points that best describe the topography, and disregard the rest. This allows for reduced data redundancy in the final unstructured mesh. Unfortunately, this process is usually a subjective task and depends on the criteria used as well as the number of points taken to be significant. Outlined by Lee (1991) and referenced elsewhere, such as El-Shimy et al. (2005) and Chang (2008), there are four popular, significant point-selection algorithms: Skeleton, Filter (also known as 'Very Important Points' (VIP)), Hierarchy, and Heuristic. Filter, VIP, Hierarchy and Heuristic share the trait that possible stopping criteria be either a pre-set number of points to be selected or a pre-set tolerance of differences in elevation between the original structured mesh and the final unstructured mesh (Lee, 1991). Slingsby (2002) performed a comparison between these techniques with an eye to hydrological process modelling. He found that the heuristic method was too slow for general use and that the hierarchy/skeleton methods produced the most balanced unstructured meshes which did not over- or under-represent any type of topography. This is in contrast to the VIP method that over-represented complex topography, and under-represented the other areas.

2.4 Research questions

The surface energy balance is frequently dominated by net shortwave radiation fluxes, and incorrectly modelling its behaviour in complex terrain can lead to inaccurate predictions of dependent phenomena. In high-relief environments, high resolution DEMs are required to properly account for horizon-shading. These DEMs generally have large data redundancies that can make them difficult to use in models, given the large number of cells required. This becomes a significant problem when modelling large spatial domains. One method to overcome these problems is to use different terrain representations and algorithms in order to reduce simulation time. Rigorous testing of shadowing algorithms using field collected shadowing data has not been performed. In addition, the effect of unstructured mesh tolerance selection on shadow location prediction has not been qualitatively compared.

2.4.1 Objectives

1. Construct an unstructured-mesh-based shadow model that accounts for topographic shading.

2. Evaluate the shadowing algorithm over a series of spatial domains by comparison with orthorectified imagery of the basin that captures the shadow boundaries. Determine the sensitivity of the spatially distributed radiation output to the mesh resolution.

3. Evaluate the effects of horizon-shadowing on snowmelt by using a CRHM based energy balance snowmelt model (CRHM SNOBAL) using both self- and horizon-shaded irradiance inputs. Measured data will be used as input.

2.4.2 Hypotheses

1. By incorporating an unstructured mesh horizon-shading algorithm, spatially variable radiation can be captured in a more accurate fashion with minimal data overhead.

2. Shadow location prediction is dependent upon topographic data resolution and generally smaller triangles (lower tolerance) produce more accurate results.

3. As long as topographic data sets have sufficiently high resolution so that mountains and valleys are discernible, this unstructured mesh method produces improved hydrological predictions (for variables dependent upon shortwave radiation) than methods that do not account for surrounding topographic shading. By coupling the adaptive mesh model to basin scale models such as CRHM, snowmelt can be predicted more accurately.

Chapter 3

Methodology

3.1 Study site and instrumentation

The research was carried out over the winter of 2010 and the spring of 2011 at the Marmot Creek Research Basin (MCRB) located at 50.96° N and 115.21° W in the Kananaskis River Valley of the Canadian Rocky Mountains (Figure 3.1). The MCRB is currently operated by the Centre for Hydrology, University of Saskatchewan, and was initially established in 1962 by the Canadian Government as an experimental research basin (Golding, 1970). A brief description of the MCRB follows.

The MCRB ranges in elevation from approximately 1450 m.a.s.l to 2886 m.a.s.l. At lower elevations, landcover is generally dense lodgepole pine, Douglas spruce, and sub-alpine fir; tree-line elevations are characterized by larch, spruce, fir, shrubs, and grasses; high alpine zones consists of talus and bare rocks (DeBeer and Pomeroy, 2009; MacDonald et al., 2010). The climatic conditions are dominated by continental air masses where winters are long and cold, averaging $-15\,^{\circ}\mathrm{C}$ between January and March (DeBeer and Pomeroy, 2009). Snow generally covers at least a portion the basin from November to June. Average annual precipitation is about 900 mm, increasing to over 1140 mm at the treeline (Storr, 1967), where about 60%–75% falls as snow, and the melt period is generally April to July (DeBeer and Pomeroy, 2009). DeBeer and Pomeroy (2009) showed that the snowcovered areas (SCA) of the alpine sites such as the Mt. Allan cirque, a north-facing slope, and a south-facing slope were 0.3, 0.4, and 0.1 respectively by mid-June, following a typical winter in 2007.

Measurements conducted over the study period were taken at hydro-meteorological stations in the basin and are shown in Figure 3.1 and summarized in Table 3.1. Field

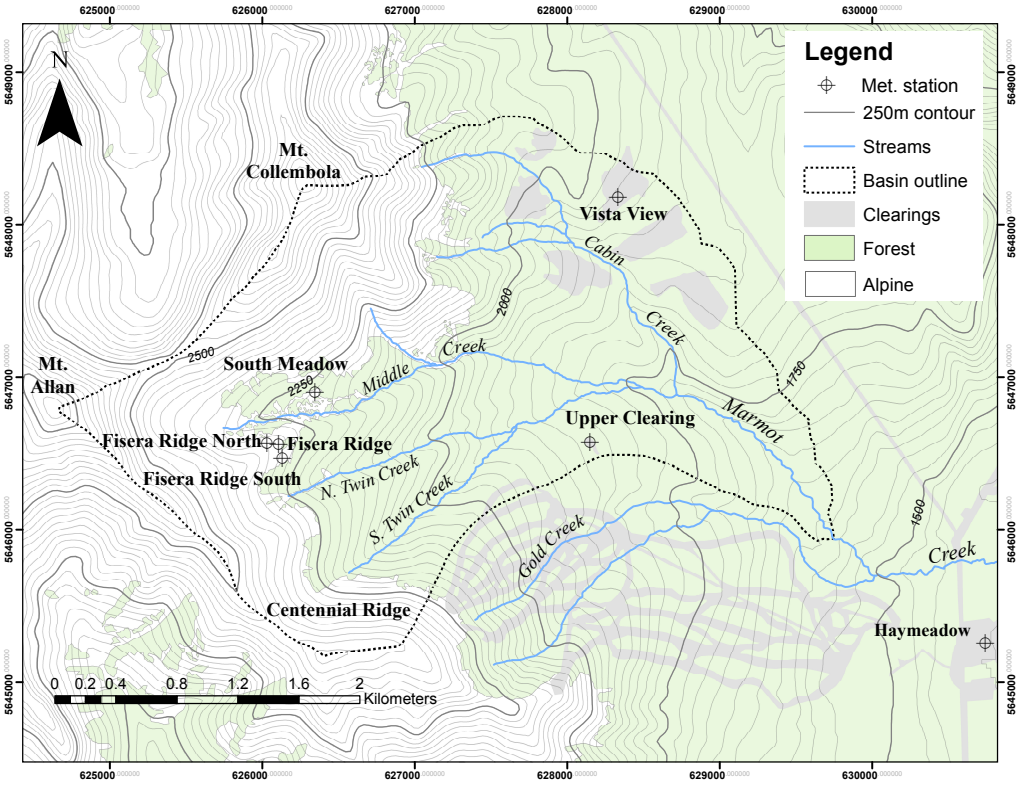

Figure 3.1: Marmot Creek Research Basin, Kananaskis Valley, Alberta at 50.96° N and 115.21° W in the Canadian Rocky Mountains. Locations of meteorological stations used in this study are shown along with major topographical features, stream channels, and man-made clearings.

visits are summarized in Appendix A. Shortwave pyranometers used in this study are the Apogee SP-110s (spectral response: 280 nm–2800 nm, cosine correction of ±1%–±5%, absolute accuracy of ±5%) at Fisera Ridge North and South, South Meadow, and Hay Meadow; and Delta-T SPN1s (spectral response: 400nm–2700 nm, cosine correction of ±2%, resolution 0.6 $W \cdot m^{-2}$) at Fisera Ridge and Hay Meadow. At Vista View, temperature and relative humidity were measured using the Campbell Scientific HMP45C (accuracy: ±0.5 °C and ±2% RH respectively); ground flux was measured using a Campbell Scientific HFT3-L flux plate (accuracy: ±5%); wind velocity was measured using an R. M. Young alpine 05103-45 anemometer (threshold: 1.0 $m \cdot s^{-1}$ accuracy: ±0.3 $m \cdot s^{-1}$). At the Upper Clearing site, precipitation was measured using Geonor T200B accumulation gauge (weight accuracy: 0.1% when full) with alter shield. All instruments were sampled at a frequency of 0.1 Hz, and the

data were averaged and stored every 15 minutes by Campbell Scientific Canada data loggers and Delta-T data loggers. A timelapse camera was installed at Fisera Ridge North, facing Mount Collembola. The camera was a Pentax K110D digital single-lens reflex (DSLR) with a Pentax DA 21 mm F32AL Limited lens that minimizes radial image distortion. The camera was triggered daily by a Campbell Scientific Canada CR200 data logger and images were taken hourly from 08:00–17:00. A Light Detection And Ranging (LiDAR) elevation dataset at 1 m $\times$ 1 m spatial resolution with 0.15 cm vertical accuracy were available from previous research and are summarized in (Hopkinson et al., 2011).

Table 3.1: Listing of the hydrometeorological stations at Marmot Creek Research Basin used in this study, with a list of instruments deployed at each site. The CR data loggers are manufactured by Campbell Scientific Canada, the GP1 data loggers by Delta-T.

Name	Elevation(m.a.s.l)	Instrumentation	Data logger
Fisera Ridge	2318	SP-110, SPN1	CR3000, GP1
Fisera Ridge North	2311	SP-110, camera	CR200X
Fisera Ridge South	2304	SP-110	CR200X
Hay Meadow	1429	SP-110, SPN1	CR3000, GP1
South Meadow	2235	SP-110	CR200X
Upper Clearing	1843	T200B	CR3000
Vista View	1952	HMP45C, HFT3, 05103-45	CR23X

3.2 Orthorectification of timelapse imagery

In order to demonstrate and validate the model against measured shadow locations, a subset of the timelapse images for February 1, 2011 were used. This day was cloud-free, had no meteorological instrument failures, and had an uniform snowcover. The set of hourly images were orthorectified using the algorithm and software of Corripio (2004). Orthorectification transforms georeferenced, oblique photographs into quasi-airphotos by defining a mapping function between the image and the DEM by performing a perspective transform on the DEM (Corripio, 2004). A 1 m $\times$ 1 m LiDAR-derived DEM was used to create a viewshed (the area visible from the point

of an observer [i.e., camera]) looking from the Fisera Ridge North station towards Mt. Collembola was created with ArcGIS 10. Although the optics of the camera and the initial alignment of the camera body are known, the following factors contribute to errors in the orthorectification process. Throughout the season, the camera housing would move and flex (such as during a storm), and these slight changes resulted in the need to use a trial-and-error process to obtain appropriate orthorectification. This process attempts to align the perspective transformed DEM to what is within the view range of the camera (represented by the viewshed). An example of this process is shown in Figure 3.2, where the red points are the perspective transformed DEM demonstrating the alignment process. As is visible, the viewshed does not perfectly reflect reality, leading to an abrupt edge where the viewshed cuts off (right side of image). However, it is clear that the rest of the image generally lines up as expected. Examples of the parameters that can be modified include roll, yaw, and pitch to control camera body misalignments; lateral and vertical alignment to correct for camera target coordinate misalignment.

Once the timelapse images were orthorectified, the shadowed area was extracted from each image via image classification. This classification was done using six manual training groups in ArcGIS 10. Groups were classified as shadow on snow (class 1), snow (class 2), shadow on trees (class 3), trees (class 4), no data (class 5), and sky (class 6). The sky class was needed to mask out the sky above the ridge line as a result of a slight misalignment in the orthorectification process. The classified imagery was then post-processed in Matlab using the Image Processing Toolbox, following Algorithm 1.

The goal of Algorithm 1 is to reclassify all the shadow locations into one class and all the non-shadow locations into another, resulting in a a binary shadow/no-shadow raster (lines 2–8). However, as a result of the original classification, there are holes, regions connected by one pixel, and single unconnected pixels in the classified raster. The algorithm removes these areas (lines 9–10) and fills holes via a three-by-three majority filter, whereby eight of the nine pixels must match (lines 11–14). Noise is further removed via opening, a process that removes small objects such as 2×2

Figure 3.2: Example of the orthorectification perspective transform of the 1 m × 1 m DEM (red points). The perspective transformed DEM is shown on top of the unprocessed timelapse photography take at Fisera Ridge North, Feb. 1st, 2011 at 17:00. The computed viewshed (not shown) causes the sharp edge in the transformed DEM.

pixel groups. The image is subsequently thickened, a process that adds pixels to the exterior of an object unless doing so would cause an unconnected object to be 8-connected. Thickening effectively smooths out the artificially rough edges of the shadow location.

In order to quantify the accuracy of the orthorectified images, a previously obtained air photo of the MCRB (spatial resolution: 1 m × 1 m) was compared to the source LiDAR DEM. This was done by examining the vegetation heights derived from the first-return of the LiDAR dataset. The vegetation extents were compared to the visible vegetation extents in the air photo and an excellent match was found. Fifty points on this air photo were compared to the same fifty points on the orthorectified imagery. Points from one image were selected such that their corresponding points on the other image were clearly matched. For each point, the minimum distance between the orthorectified imagery and the air photo was measured and

Algorithm 1 The shadow extraction algorithm.

Precondition: Set of partially classified images
Postcondition: Set of binary classified shadow areas

```
1: for each Image i do
2:     for each Pixel p in i do
3:         if p = class 1 or 3 then
4:             p ← Shadowed
5:         else
6:             p ← Non-shadowed
7:         end if
8:     end for
9:     i ← 1pxConnected(i)          ▷ Break apart single-pixel connected regions
10:    i ← Clean(i)                 ▷ Remove single-pixel noise
11:    for 1 to 2 do
12:        i ← 3x3Majority(i)       ▷ Eight of nine pixels must match
13:        i ← Clean(i)
14:    end for
15:    i ← OpenImage(i)
16:    i ← ThickenImage(i)
17:    i ← 3x3Majority(i)           ▷ Five or more of nine pixels must match
18:    SaveToDisk(i)
19: end for
```

then this distance was spatially interpolated. The interpolation was done using the ArcGIS software and the kriging spatial interpolation. Kriging is a geo-statistical spatial interpolation that is best known in the earth sciences (Myers, 1994). Kriging interpolation assumes that the distance between known points reflects a spatial correlation that can be used to explain variation within the surface, with this spatially correlated component quantified via a semivariogram. Such an interpolation allows for quantification of the error across the entire orthorectified image. Choosing corresponding points between the imagery and the DEM was difficult at times, with above tree-line points having increased difficulty due to the lack of prominent landmarks. Repeated measurements of the distance between the same point indicated that there was about a $\pm 5\%$ error in this distance measurement.

3.3 Terrain representation

For use in many types of hydrological research, the spatial domain of interest must be discretized such that a spatial unit, however defined, is connected to other spatial units. A structured mesh, such as that derived from a LiDAR source, may be triangulated such that triangles of variable size represent the topography. Triangle vertices thus have x, y, z components that represent the elevation of a point in some coordinate system, e.g., Universal Transverse Mercator (UTM). A common and powerful triangulation algorithm is Delaunay triangulation; see Shewchuk (2002). Using Delaunay triangulation, the triangulation can be constrained to important characteristics of the model domain such as streams, rivers, and basin delineation, thus guaranteeing their representation. Once complete, a triangulation is essentially a linear interpolant. Thus when creating a triangulation from a structured mesh, a tolerance, i.e., the maximum difference between the triangle's interpolated elevation and the elevation of the structured mesh, must be specified. Therefore, stringent tolerances produce a greater number of triangles, and vice versa. Because computational time increases as a function of number of triangles, choosing a tolerance that is too stringent negates some of the performance increases of an unstructured mesh. As an example (for the MCRB), using the base DEM of 1 m $\times$ 1 m with a tolerance of 1 m results in 206580 triangles whereas a tolerance of 0.5 m results in 544507 triangles. Depending on the modelling requirements, there is a need to find a balance between acceptable error and computational efficiency. The ideal approach would be to use the LiDAR point cloud as the input to the triangulation process.

In order to investigate the effects of the dataset resolution on model performance, sixteen combinations of resolution and tolerance were tested. The base 1 m $\times$ 1 m LiDAR-derived DEM (herein referred to as a 1 m DEM) was resampled to produce a raster with a spatial resolution of 2 m $\times$ 2 m (herein referred to as a 2 m DEM), a 10 m $\times$ 10 m (herein referred to as a 10 m DEM), and 30 m $\times$ 30 m (herein referred to as a 30 m DEM). These structured meshes were used to create the unstructured meshes with tolerances of 1 m, 2 m, 10 m, and 30 m for each DEM resolution. An

area slightly larger than the MCRB area was simulated in order to capture horizon-shadows created from topography outside the basin. The shadow-model model runs were done for February 1, 2011. This day was selected to correspond to the days used for the timelapse imagery, thus allowing for *a priori* knowledge of what the daily shadowing regime.

3.4 Shortwave radiation data

In order to avoid the errors associated with modelled incident shortwave radiation, this study used measured direct- and diffuse-beam irradiance (from the Delta-T SPN1s) as input for the unstructured-mesh-based shadow model that accounts for horizon-shading. Because there were no winter-accessible sites with an unobstructed view of the sky that measured both direct and diffuse irradiance, no one hydro-meteorological station provided a complete record of incoming solar radiation with no shading. For example, the Hay Meadow site is in shadow early in the morning, whereas the Fisera Ridge site is not, with this situation reversing late in the day. In order to provide a complete record of incoming radiation (diffuse- and direct-beam) with no shading, measurements from the SPN1 instruments at the Hay Meadow and Fisera sites were spliced together to create a synthetic dataset spanning October, 2010–June, 2011 utilizing Fisera Ridge values from sunrise to noon and Hay Meadow values until sunset. Each half of the day, from each site, was merged into one data set. Smoothing or interpolation between datasets was not used. Between October, 2010–June, 2011 at 12:00:00 daily, these datasets had a mean difference of 15 $\mathrm{W{\cdot}m^{-2}}$. This synthetic dataset represents the irradiance throughout the day that is as free from topographic shadow contamination as possible.

Diffuse radiation was assumed to be isotropic and corrected for the skyview-factor; the ratio of the diffuse sky irradiance with obstructing terrain to that of an unobstructed sky. This was computed using the the 1 m LiDAR structured-mesh DEM using the algorithm of Dozier and Frew (1990) as implemented in the SAGA-GIS software (SAGA Development Team (saga-gis.org), 2007). The values for each

triangle vertex were extracted from the structured mesh, and the average of each triangle's vertices was used as the skyview for a triangle. Although it is necessary to compute a skyview factor when considering diffuse radiation, this relatively simple hybrid scheme avoids the problems of storing the horizon in every direction for every point and searching through the horizon for every point for every timestep when computing direct-beam shading. Rather, because the skyview is a fixed constant and diffuse shortwave radiation is considered isotropic, there is no need for directional searches. Therefore the skyview factor can be computed once and used for the entire diffuse irradiance calculation. Because diffuse radiation is not affected by shadowing, the shadowing algorithm was not applied to the diffuse component.

Reflected shortwave radiation from the surface can be important in many environments (Marks and Dozier, 1992), especially on cloudy days; however it was not considered in this study and solar irradiance was estimated as the sum of diffuse- and direct-beam components.

3.5 Shadow model development

3.5.1 Model description

The shadowing model presented herein, an extension of the algorithm presented by Montero et al. (2009), determines horizon-shadow locations via Euler rotations; see Fowles and Cassiday (2005) for a treatment of Euler angles. This model is named `Umbra` after the Latin word for shadow. It is an open-source software built upon the Armadillo C++ linear algebra library (Sanderson, 2010) and MATLAB 2011a (see Appendix B for details).

Let $\theta_s(t)$ be the solar elevation angle and $\phi_s(t)$ be the solar azimuth, clockwise from north (radians), at a time t. Let x, y, z be a reference coordinate system with x positive to the east, y positive to the north, and z positive in the vertical direction (corresponding to elevation). This initial configuration is shown in Figure 3.3. Let x', y', z' be a coordinate system obtained by performing an XZX series of Euler

rotations on the coordinate system x, y, z to align the z' axis with the solar vector.

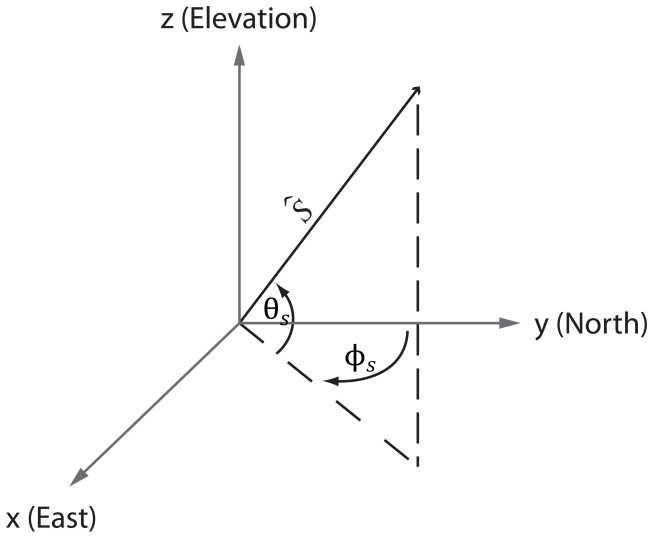

Figure 3.3: The initial coordinate system configuration of x, y, z, where $\theta_s(t)$ is the solar elevation and $\phi_s(t)$ is the solar azimuth (clockwise from north). The x-axis is positive to the east, the y-axis is positive to the north, and the z-axis is positive in the vertical direction (elevation).

For a given angle ψ, let $\mathbf{R}_x(\psi)$ be the x rotation matrix and $\mathbf{R}_z(\psi)$ be the z rotation matrix given as

$$\mathbf{R_x}(\psi) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\psi & -\sin\psi \\ 0 & \sin\psi & \cos\psi \end{pmatrix}, \tag{3.1}$$

and

$$\mathbf{R_z}(\psi) = \begin{pmatrix} \cos\psi & -\sin\psi & 0 \\ \sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{3.2}$$

Then the rotation from x, y, z to x', y', z' is done via the Euler rotations

$$\mathbf{R} = \mathbf{R_x}\left(-\frac{\pi}{2} + \theta_s(t)\right)\mathbf{R_z}\left(-\pi + \phi_s(t)\right)\mathbf{R_x}\left(0\right) \tag{3.3}$$

$$\mathbf{R} = \begin{pmatrix} \cos(-\pi+\phi_s(t)) & \sin(-\pi+\phi_s(t)) & 0 \\ -\cos\left(-\frac{\pi}{2}+\theta_s(t)\right)\sin(-\pi+\phi_s(t)) & \cos\left(-\frac{\pi}{2}+\theta_s(t)\right)\cos(-\pi+\theta_s(t)) & \sin\left(-\frac{\pi}{2}+\theta_s(t)\right) \\ \sin\left(-\frac{\pi}{2}+\theta_s(t)\right)\sin(-\pi+\phi_s(t)) & -\sin\left(-\frac{\pi}{2}+\theta_s(t)\right)\cos(-\pi+\theta_s(t)) & \cos\left(-\frac{\pi}{2}+\theta_s(t)\right) \end{pmatrix}. \tag{3.4}$$

The rotation $\mathbf{R}$ is applied to each triangle's vertex. The result of this rotation is a transformed x', y', z' coordinate system that defines triangles, with respect to the solar vector, further from the sun as having smaller z' values and triangles closer to the sun as having a larger z' values. These rotated triangles are projected onto the $x'y'$ plane and are interpreted as triangles in two dimensions for the purpose of the triangle-triangle intersection test. For the purpose of visualization, the triangles were coloured according to an average of their three vertex z' values. An example of this is shown in Figure 3.4a. The scene is shown with the observer between the sun and the terrain and along the solar vector. Triangles with high values (red) are the closest to the sun, and triangles with low values (blue) are furthest away from the sun. The scene is shown from a different perspective in Figure 3.4b, with the observer looking directly into the basin (approximately west), with the sun on the left of the scene. The ordering of the triangles obtained from the z' values allows for obscuring triangles to be detected.

Each triangle's irradiance was calculated using the standard cosine-correction (Garnier and Ohmura, 1968; Oke, 1987). To calculate the direct-beam irradiance to a triangle (considered uniform over the triangle) the following equations were used. Let $\hat{\mathbf{n}}$ be the unit normal vector to a triangle's surface. Then the angle Θ_{sn} between the surface normal and $\hat{\mathbf{s}}(t)$ is a unit vector pointing to the position of the sun given by

$$\Theta_{sn}(t) = \arccos(\hat{\mathbf{n}} \cdot \hat{\mathbf{s}}(t)) \tag{3.5}$$

where $\hat{\mathbf{s}}(t)$, in a Cartesian coordinate system, is given by:

$$\hat{\mathbf{s}}(t) = \begin{pmatrix} \cos\theta_s(t)\ \sin\phi_s(t) \\ \cos\theta_s(t)\ \cos\phi_s(t) \\ \sin\theta_s(t) \end{pmatrix}. \tag{3.6}$$

The measured direct-beam irradiance data (K_{meas} [W·m^{-2}]) were corrected from flat-plane values to account for the flat-plane solar zenith angle (Garnier and Ohmura, 1968; Oke, 1987) such that

$$S_0\tau = \frac{K_{meas}}{\cos Z(t)}, \tag{3.7}$$

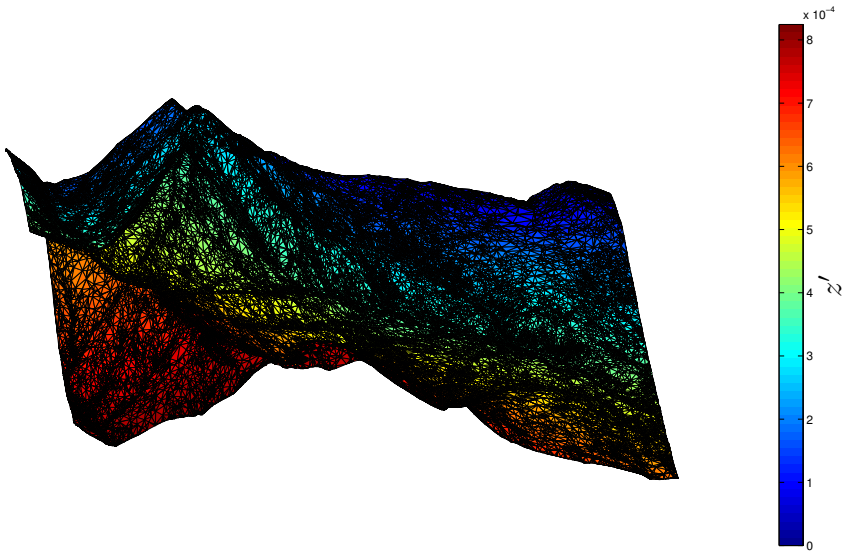

(a) The MCRB viewed along solar vector

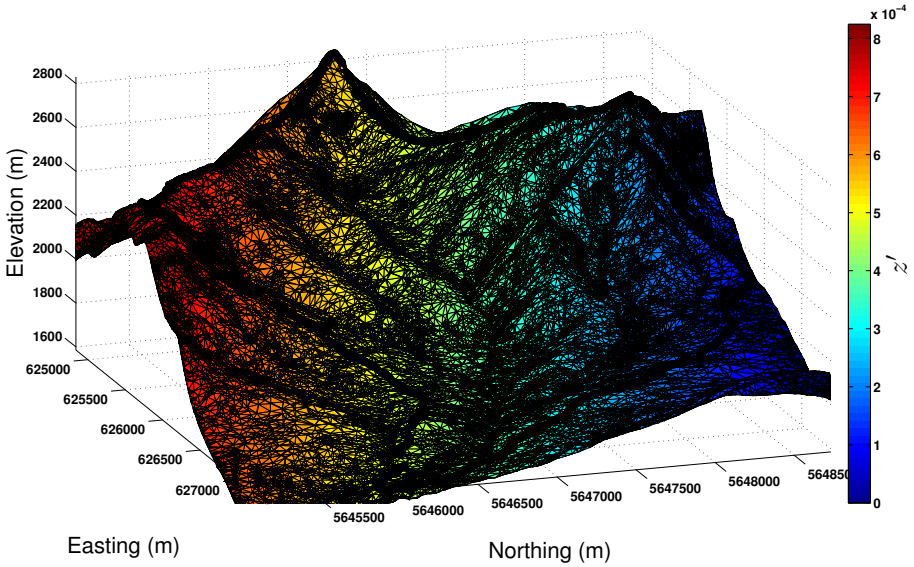

(b) The MCRB viewed in a westerly direction

Figure 3.4: Result of the x, y, z coordinate system rotated to x', y', z' where z' is aligned with the solar vector. The colouring is the average of the z' values for each of the triangles' vertices, where triangles with larger z' values (closer to the sun along the solar vector) are in red and triangles with smaller z' values (farther from the sun) are in blue. The scene is shown from a perspective (a) aligned with the solar vector and (b) looking approximately west towards the MCRB (black outline). Model time: Feb 1, 2011 at 13:30 (local).

where S ($\mathrm{W{\cdot}m^{-2}}$) is the solar constant, $Z(t)$ is the solar zenith angle, and τ is the transmittance of the atmosphere to direct-beam solar radiation. Thus the $S\,\tau$

term captures the measured incident irradiance (including transmittance) that can be used as input to a hydrological model. The self-shaded direct-beam irradiance ($K_s(t)$ [W·m^{-2}]) is then computed as

$$K_s(t) = S\tau \cos \Theta_{sn}(t) \tag{3.8}$$

To account for self-shading, if $\cos \Theta_{sn}(t)$ was less than zero, i.e., the triangle was facing 90° or more away from the sun, the triangle was marked as self-shaded and assigned a direct-beam irradiance of 0 W·m^{-2}.

3.5.2 Horizon-shadow detection

The calculation of the z' values for each triangle gives an ordering to the triangles, allowing for obscuring triangles to be detected. This algorithm is outlined in Algorithm 2 and described in detail below.

For computational efficiency when determining whether two triangles intersect, only triangles with similar x', y' values need to be compared. Because the original algorithm in Montero et al. (2009) used a different mesh type (Rivara 4-T (Rivara, 1987)) than the Delaunay mesh used herein, the shadowing algorithm was modified and adapted accordingly. A three-dimensional bounding box is fit to the terrain. This bounding box is then rotated and flattened to two-dimensions (line 3). This minimum bounding rectangle is segmented into m rows and n columns creating mn regions. Each rotated triangle is assigned to a region (lines 6–10), and if a triangle falls into more than one region, it is considered in each region it intersects. The triangles in each region are sorted in descending order of z' (line 14). The triangles of each region (lines 15–23) are checked in order of largest z' to smallest z' against the remaining triangles in the region. If a triangle lies in front of another triangle (based on their respective z' values) the triangle with the larger z' value is checked for an intersection with the triangle behind it (lines 17–18). To determine intersection between triangles, the vertices and circumcentre of the triangle with greater z' are used as detection points. These points are checked to see if they are contained by the triangle with smaller z'; containment indicates intersection. If intersection is

detected, the triangle with the smaller z' value is marked as shaded with a binary shadow mask (ζ) (line 19), where $\zeta = 0$ when horizon-shaded and $\zeta = 1$ for all other cases. The horizon-shaded direct-beam irradiance ($K_h(t)$ [W·m^{-2}]) was then computed as

$$K_h(t) = \zeta \, K_s(t). \tag{3.9}$$

Algorithm 2 The shadowing algorithm.

Precondition: Tri is a $3 \times N$ matrix where each row are the vertexes of an unstructured triangular mesh of size N.

Precondition: Solar elevation for current time-step > 0.

```
1: function UMBRA(Tri)
2:     N ← Number of triangles in mesh
3:     BR ← CREATEBR(Tri)          ▷ Create 3D bounding rectangle and
                                   ▷ flatten to 2D with m × n regions
4:     for each t in Tri do
5:         Tri(t)' ← EULERROTATION(Tri(t))
6:         for j = 1 to BR'.rows do
7:             for k = 1 to BR'.cols do
8:                 INSERTTRIANGLEINTOBR_{j,k}(Tri'(t))
   ▷ Find region in BR'
9:             end for
10:        end for
11:    end for
12:    for j = 1 to BR.rows do
13:        for k = 1 to BR.cols do
14:            BR_{j,k} ←SORT(BR_{j,k})    ▷ Sort region i,j using z' values
15:            for t_i = 1 to BR_{j,k}.size do
16:                for t_j = t_1 + 1 to BR_{j,k}.size do
17:                    if Tri'(t_i).z' > Tri'(t_j).z' then
18:                        if INTERSECTS(Tri'(t_i),Tri'(t_j)) then
19:                            Tri(t_j) ← Shaded
20:                        end if
21:                    end if
22:                end for
23:            end for
24:        end for
25:    end for
26: end function
```

3.5.3 Computational complexity

Computational complexity, the classification of computational problems according to their inherent difficulty, of the horizon-shading algorithm was determined analytically to be $\mathcal{O}(N \log N)$, where N is the number of triangles in the unstructured mesh. Let N be equal to the number of triangles T in the triangulation. Then lines 4–11 are done in $\mathcal{O}(N)$ and the sorting algorithm (line 14) is $\mathcal{O}(N \log N)$. The lines 15–23 are $\mathcal{O}(N)$.

An empirical verification of this complexity was performed by running the model for one time step (February 1, 2011 at 16:30) for a range of meshes ranging from 851 triangles to 544507 triangles. The results are shown in Figure 3.5, where the black symbols are the measured wall-clock times and the grey dashed line is a curve of best fit of the form $t = a_0 + a_1 N + a_2(N \log N)$. The form of the fit is based on an analytical complexity analysis of the algorithm and has an adjusted R^2 of 0.99994. Although the algorithm is parallelizable, and this fact is exploited when running the algorithm in practice, the complexity analysis and timing information reported are for a serial implementation. The benchmark tests were performed on an Intel Xeon W3520 processor with a clock speed of 2.66 GHz and 16 GB of DDR3 RAM running Ubuntu 10.04 LTS. The Intel C++ compiler version 12.0.4 20110427 was used with level-3 optimization with SSE 4.2 instruction-set compatibility enabled. Timings were taken using the `clock_gettime` function. The values reported are the minimum of fifty runs.

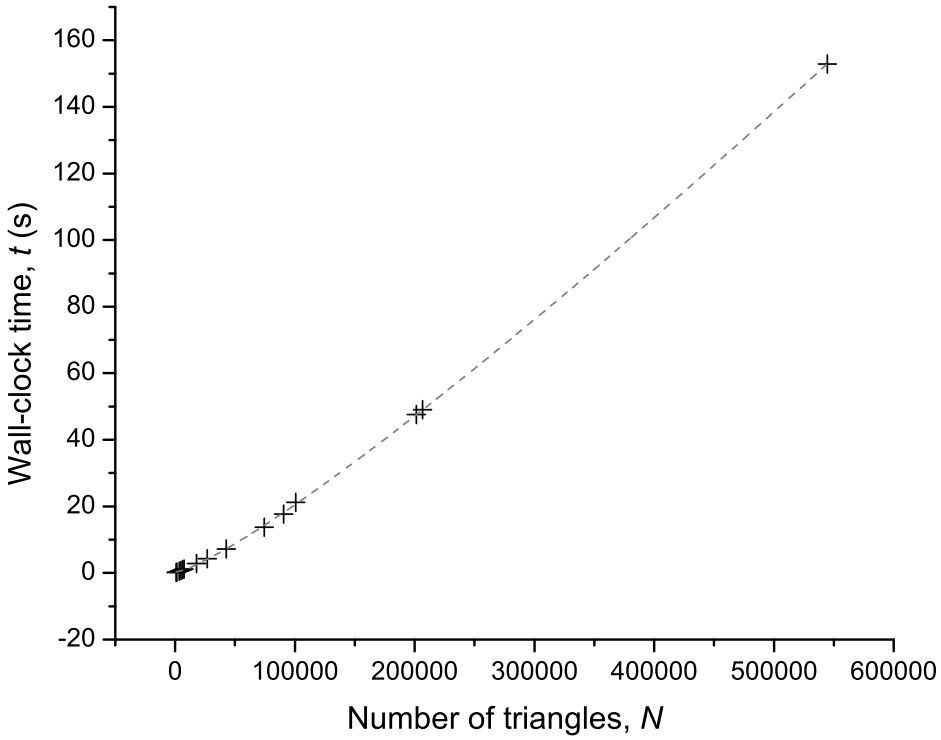

Figure 3.5: Wall-clock times of the horizon-shadow algorithm run for mesh sizes from 851 to 544507 triangles. The grey dashed line is a line of best fit of the form $t = a_0 + a_1 N + a_2(N \log N)$. Wall-clock times are for one model timestep Feb 1, 2011 at 16:30, a time for which there are many triangles in shadow.

3.6 CRHM and SNOBAL

The Cold Regions Hydrological Model platform provides a modular framework to assemble hydrological models (Pomeroy et al., 2007). Within this framework, modules conceptualize a process such as albedo decay, longwave radiation estimation, snowmelt, or infiltration. These modules can be linked together, as appropriate, to create a hydrological model tailored to a basin or hydrological phenomenon in question. Spatial elements are conceptualized by the HRU approach. CRHM is therefore used due to its ability to efficiently assemble a hydrological model and perform sensitivity tests using different inputs.

In this thesis, CRHM is used with the SNOBAL energy balance model for calculating the snow-surface energy balance and melt. SNOBAL is a point scale, two-layer energy balance snowmelt model that simulates each component of the snowcover energy balance (Marks et al., 1998). It is driven using the following measured meteorological inputs: net shortwave radiation ($\mathrm{W{\cdot}m^{-2}}$), incoming longwave radiation ($\mathrm{W{\cdot}m^{-2}}$), air temperature (°C), vapour pressure (Pa), and wind speed ($\mathrm{m{\cdot}s^{-1}}$). Following Marks et al. (1998), Figure 3.6 shows the conceptual structure of SNOBAL, where the arrows denote flux direction.

For the CRHM model runs, two sites within the MCRB were simulated: the South Meadow site and the Deep Valley site. The South Meadow site was chosen because it represents an instrumented south-facing slope that is crossed by a horizon-shadow each day. The Deep Valley site is located outside the MCRB in an adjacent valley. It was chosen as a site that is heavily influenced by horizon-shadows. The two input shortwave irradiances, $K_h(t)$ for horizon-shaded irradiance and $K_s(t)$ for self-shaded irradiance (Equations 3.8 and 3.9), were generated for these sites using the unstructured mesh model with a 1 m DEM and a 1 m tolerance. For input into CRHM, the irradiances are the sum of the direct- and diffuse-beam components. Precipitation measurements were from the forest clearing at Upper Clearing. There were no missing data for this period. The remainder of the meteorological data were from the Vista View site. This site represents a well-drained peat site and is a close

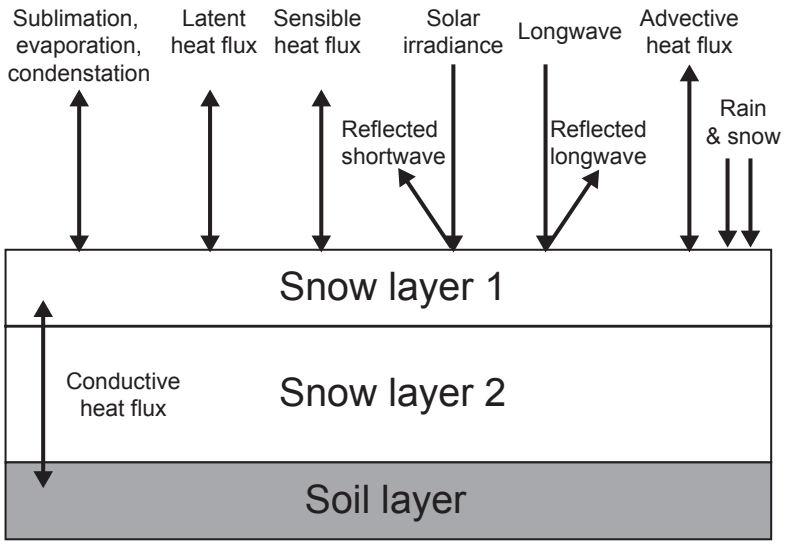

Figure 3.6: Conceptual figure depicting the direction of fluxes in SNOBAL. Arrows denote direction. This figure follows that of Marks et al. (1998).

analogue to the South Meadow site. Because no site visit nor any studies have ever been conducted at the Deep Valley site, knowledge of site specific characteristics are missing. The same meteorology data are used for the South Meadow site as for the Deep Valley site.

In order to simulate the snow accumulation and ablation at a point-scale for South Meadow and Deep Valley, an HRU of size 1 m^2 was created for each, for a total basin size of 1 m^2. CRHM modules used were: `basin`, `global`, `obs`, `intcp`, `longVt`, `pbsmSnobal`, `albedo_Richard`, and `SnobalCRHM#2`.

The `basin` module manages the general parameters of the model by allowing for the definition of the HRU parameters such as area, elevation, and ground slope and aspect. This is a required core module and is the basis for all models. The `global` module calculates the theoretical interval values of flat-plane and slope-corrected direct- and diffuse-beam irradiances for each HRU. Solar geometry is also calculated and used as input for other modules. The shortwave calculations follow those of Garnier and Ohmura (1968). The `obs` module allows for reading the meteorological forcing data. The `intcp` module is a compatibility module for when no canopy interception is required. The `longVt` module calculates terrain-view correct long-

wave radiation based on Sicart et al. (2006). Because this module uses irradiance as an estimate for transmittance, the self-shading irradiance was used as input for the `longVt` module. The `pbsmSnobal` module calculates snow transport and sublimation following Pomeroy et al. (1993). However, it was set to inhibit blowing snow, resulting in a point-scale accumulation and ablation of a snow-pack. The `albedo_Richard` module calculates the surface albedo, including albedo refresh due to snowfall, following Essery et al. (2001). The `SnobalCRHM#2` module is variation#2 of the `SnobalCRHM` module; it uses the observed irradiance and the calculated longwave from the `longVt` module. `SnobalCRHM` is the implementation of SNOBAL following Marks et al. (1998) (as described above). As input, it requires the output all the modules listed above.

For each site, the model was run twice, once using the self-shading $K_s(t)$ irradiance and once using the horizon-shading $K_h(t)$ irradiance. Differences were calculated by subtracting the horizon-shading values from the self-shading values. All other meteorological input values were the same between CRHM runs. Instead of using a measured ground temperature to calculate the ground heat flux via SNOBAL, a measured ground-heat flux was used.

A snowcovered period was simulated from October 17, 2010 to June 14, 2011 at a 15-minute interval, and a snowmelt period was simulated from April 1, 2011 to June 14, 2011 at a 15-minute interval. Finally, a clear-sky day was modelled for February 1, 2011 at 15-minute intervals from sunrise to sunset. This was done to check the performance of the model for a day that did not have any cloud-induced shadows. Surface irradiance ($\mathrm{MJ \cdot m^{-2}}$) was computed assuming a constant irradiance, temporally over the timestep and spatially over each triangle.

3.6.1 Calculation of phase change mass

The cumulative differences of net solar radiation between the self- and horizon-shading models were converted to a phase change mass per unit area using a constant albedo. Snow in the study area is not affected by wind-borne dust and hence the albedo remains higher than many published values, even during melt periods (Helga-

son and Pomeroy, 2012). Many published low albedo values for snow during melt are contaminated by snow-free area, and the SCA albedo is known to decay at a much slower rate than the areal albedo (O'Neill and Gray, 1973; Male and Granger, 1981; Pomeroy et al., 1998). Pomeroy et al. (1998) show that when SCA corrections are employed, albedo during melt can remain at 0.85. However for this study a constant albedo of 0.8 (Oke, 1987) and a spatially continuous snowcover were assumed. It was also assumed that all net shortwave radiation was applied to snowmelt.

To express the cumulative energy difference ($Q_{\Sigma\Delta}$) between the two shading algorithms to an equivalent phase change mass (mm SWE), the following was used:

$$\frac{Q_{\Sigma\Delta}\,(1-\alpha)}{\rho_w\,h_f} = \text{SWE}, \tag{3.10}$$

where h_f is the latent heat of fusion (0.334 MJ·kg^{-2}), ρ_w is the density of water (1000 kg·m^{-3}), and α is the snow albedo.

Chapter 4

Results

4.1 Shadow location and movement

A series of timelapse images demonstrating the movement of a horizon-shadow across the South Meadow site on February 1, 2011 from 16:00 to 17:00 is shown in Figure 4.1. The South Meadow site is located in the clearing indicated by the red arrow. The image is taken from Fisera Ridge North viewing down the valley, roughly in an easterly direction; see Figure 3.1. The large shadow visible in both images is cast from the Fisera Ridge ridge line. South Meadow is located on a south-facing slope and as a result generally receives relatively high irradiance. A sharp decrease in irradiance at South Meadow is observed once the Fisera Ridge shadow crosses the site. Therefore, in order to determine if the shadowing model can accurately predict the correct timing of shadow movement, it is possible to consider these point-scale irradiances.

Shown in Figure 4.2 is the irradiance measured at the South Meadow station (fine dotted line), the irradiance (from the synthetic spliced dataset) with the horizon-shadow algorithm (dashed line), and the irradiance (from the synthetic spliced dataset) with the self-shading algorithm (solid line). The model results from the horizon-shadows algorithm has a sharp decrease in irradiance at approximately 16:30 when the horizon-shadow from Fisera Ridge crosses the South Meadow site. This decrease in irradiance corresponds well to that shown in the timelapse photos in Figure 4.1. The differences in peak irradiance in Figure 4.2 between measured and modelled are attributed to the fact that the triangle where the South Meadow site is located in the model is not perfectly flat like the measuring pyrometer, resulting

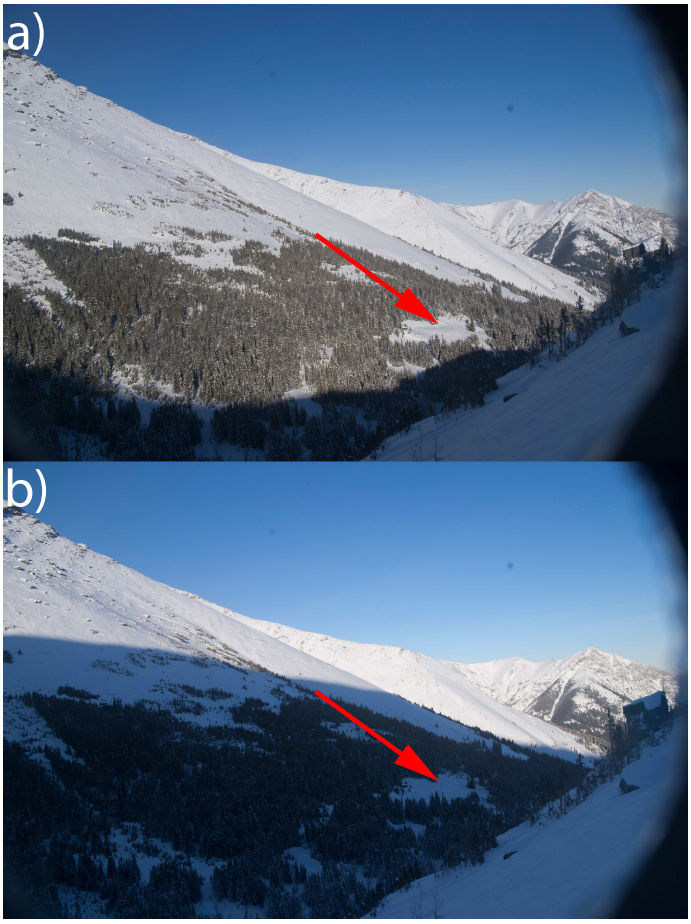

Figure 4.1: Timelapse photographs from Fisera Ridge North showing shadow movement across the South Meadow site (red arrows). Photos taken Feb 1, 2011 at (a) 16:00 and (b) 17:00. Photos taken in an easterly direction.

in slightly higher direct-beam irradiance. As well, because the South Meadow modelled data is a synthetic dataset composed of data from different locations there are different amounts of reflected shortwave radiation. Finally, there are also differences in the actual instruments.

This shadow movement can also be examined at a basin scale. A shadow map, Figure 4.3, was created for the clear-sky day February 1, 2011 at 16:30. Areas in black, grey, and white are the horizon-, self-, and non-shadowed areas respectively, with elevation contours every 100 m. The South Meadow site is indicated. There are misclassifications of some triangles, visible as the lone shaded triangles as a result of imprecisions in the triangle-triangle detection scheme.

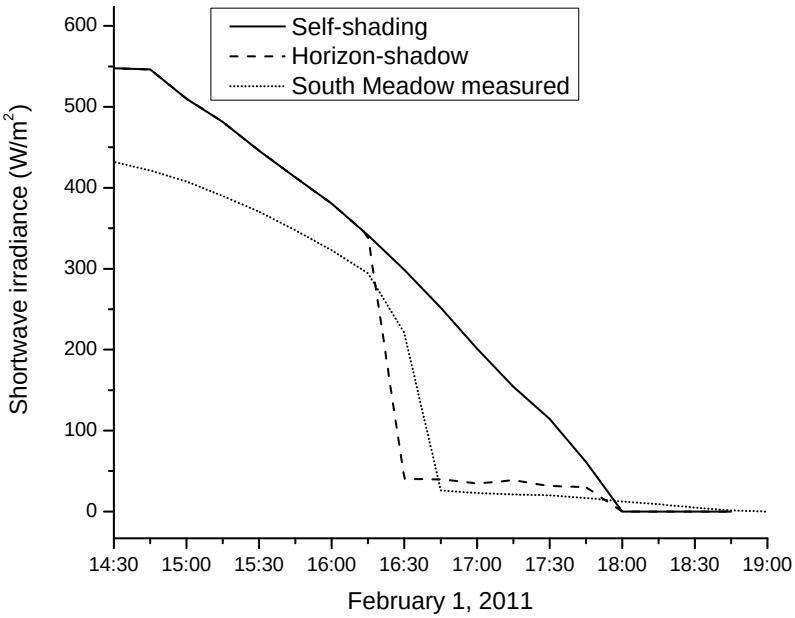

Figure 4.2: Shortwave irradiance for the clear-sky day February 1, 2011 at South Meadow. Measured irradiance (fine dotted line), horizon-shadow corrected irradiance (dashed line), and the self-shading corrected irradiance. Because different instruments were used for South Meadow and the synthetic input data for the model, different amounts of reflected shortwave radiation are present as well as differences in the actual instruments, resulting in the difference between measured and modelled.

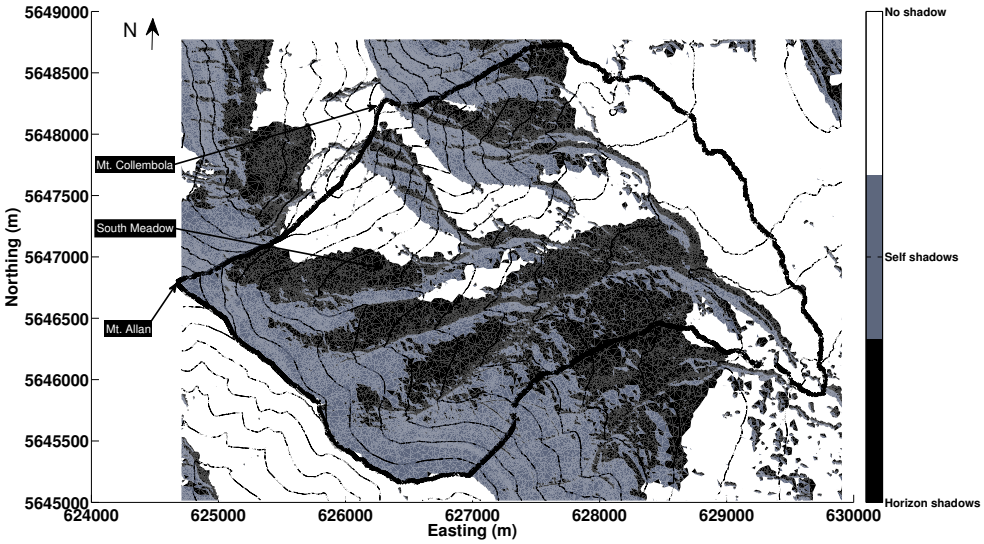

Figure 4.3: Marmot Creek Research Basin with modelled shadow locations. Areas in black, grey, and white are the horizon-, self-, and non-shadowed areas respectively. Model time: February 1, 2011 at 16:30. Contours are every 100 m, and the South Meadow site is indicated. This shows the horizon-shadow passing over the South Meadow site, a phenomenon that is missed with only self-shadows.

4.2 Comparison with orthorectified timelapse photographs

In order to empirically compare the locations of observed shadows with modelled shadow locations, the timelapse photographs from February 1, 2011 were orthorectified and observed shadow locations were classified and extracted for comparison; see Section 3 for details. Figure 4.4 shows the classified orthrectified photographs for February 1, 2011 at (a) 16:00 and (b) 17:00. The red area is the classified shadow area from the orthorectified photographs. The blue area is the modelled horizon-shadow location. Some shadows are predicted in a location where the images do not cover (white areas), so comparison in these locations is not possible. The observed shadow in the bottom left corner of Figure 4.4a demonstrates the close fit of the model. Figure 4.4b demonstrates the model fit one hour later as the South Meadow site is full in shadow. The modelled shadow extents are generally over predicted. This appears to be a result of the triangle-triangle collision algorithm. Because triangle shading is considered binary, a triangle that only has a small intersection with another is completely shaded. If a triangle rests on a uniform slope, such as at South Meadow, the triangle may be large. Thus a larger area can be easily over shaded.

The results of the orthorectification error analysis are shown in Figure 4.5, where the difference (m) between the orthorectified photographs and the aerial photography is shown. Although the aerial photography cannot be considered a perfect truth, these results do show that many of the largest differences between the modelled and observed shadow regions occur in areas that have a large difference between the orthorectified photographs and aerial photography. The maximum and minimum differences are 24.7 m and 4.74 m respectively and an average of 14.5 m (standard deviation 3.46 m). This suggests that utilizing orthorectified photographs for spatial delineation of features has significant uncertainty. However, the uncertainty in matching the points on the aerial photography to those on the orthorectified photographs indicates these differences may be overstated in this analysis. Without

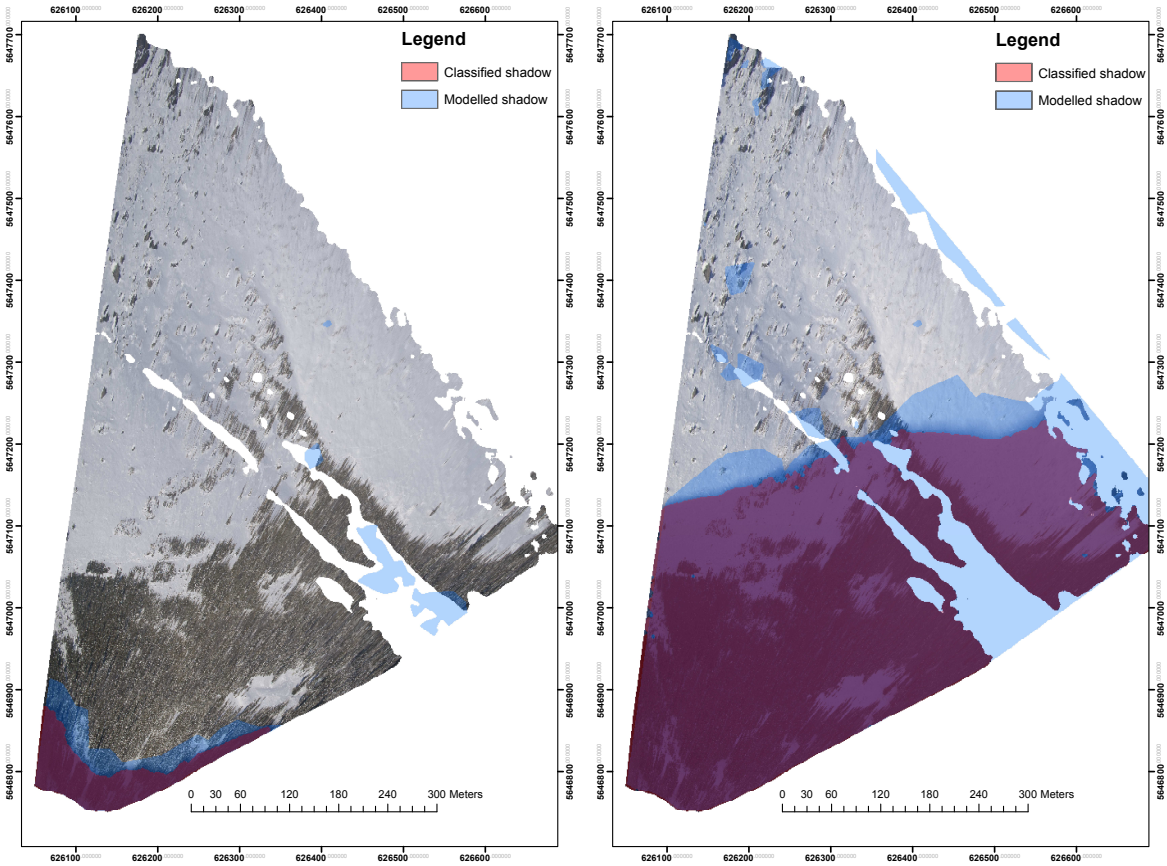

(a) 16:00 **(b)** 17:00

Figure 4.4: Orthorectified timelapse photographs from Fisera Ridge North showing shadow movement across the South Meadow site. Photos were taken February 1, 2011 at (a) 16:00 and (b) 17:00. The red area is classified shadow area from the orthorectified photographs and the blue is from the model.

utilizing precise and accurate differential GPS devices, it is not possible to exactly calculate the differences between orthorectified photographs and reality.

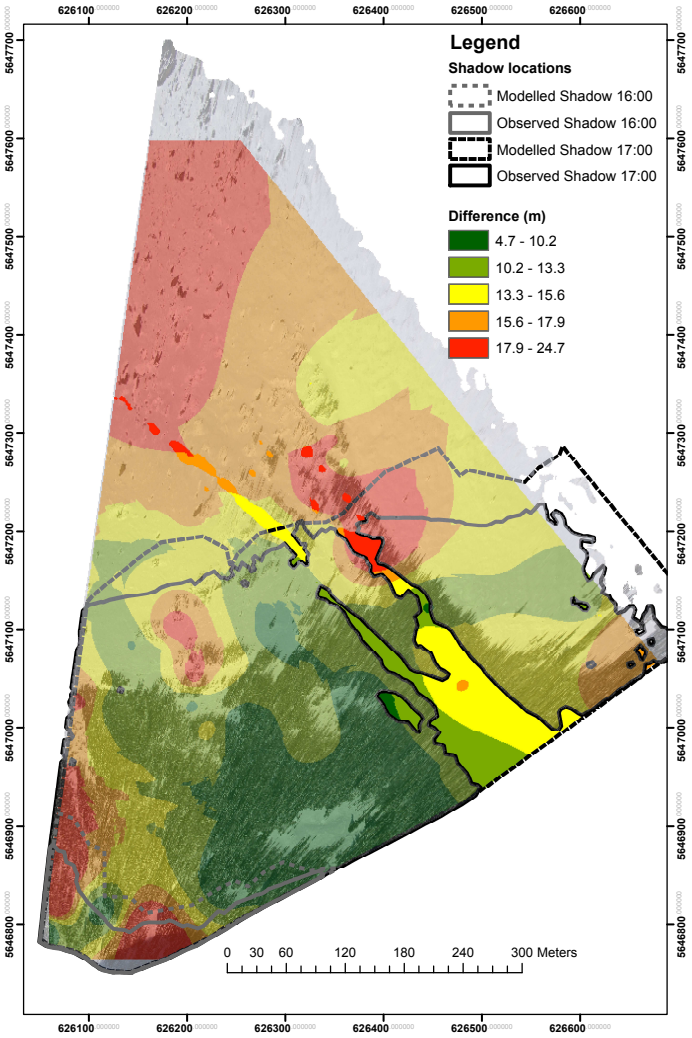

Figure 4.5: Predicted and observed shadow movement across the South Meadow site for February 1, 2011 at 16:00 and 17:00. The horizontal difference (m) between the orthorectified photographs and the aerial photography. The maximum and minimum differences are 24.7 m and 4.74 m respectively and an average of 14.5 m (standard deviation 3.46 m). There is approximately a 5% systematic error in these differences due to uncertainties in selecting matching points.

4.3 Comparison with existing shadowing methods

Although utilizing an unstructured mesh for landscape representation carries many advantages over the commonly used structured mesh, there are certain disadvantages. Because linear interpolants are used, error is introduced into the terrain representation. Thus a comparison with an existing shadowing routine was done us-

ing the ArcGIS 10 Spatial Analyst hillshading algorithm to calculate horizon-shaded shadows. The comparison between the observed shadows from the orthorectified photographs from the Fisera Ridge North site, the unstructured mesh model, and ArcGIS are shown in Figure 4.6 for (a) 16:00 and (b) 17:00.

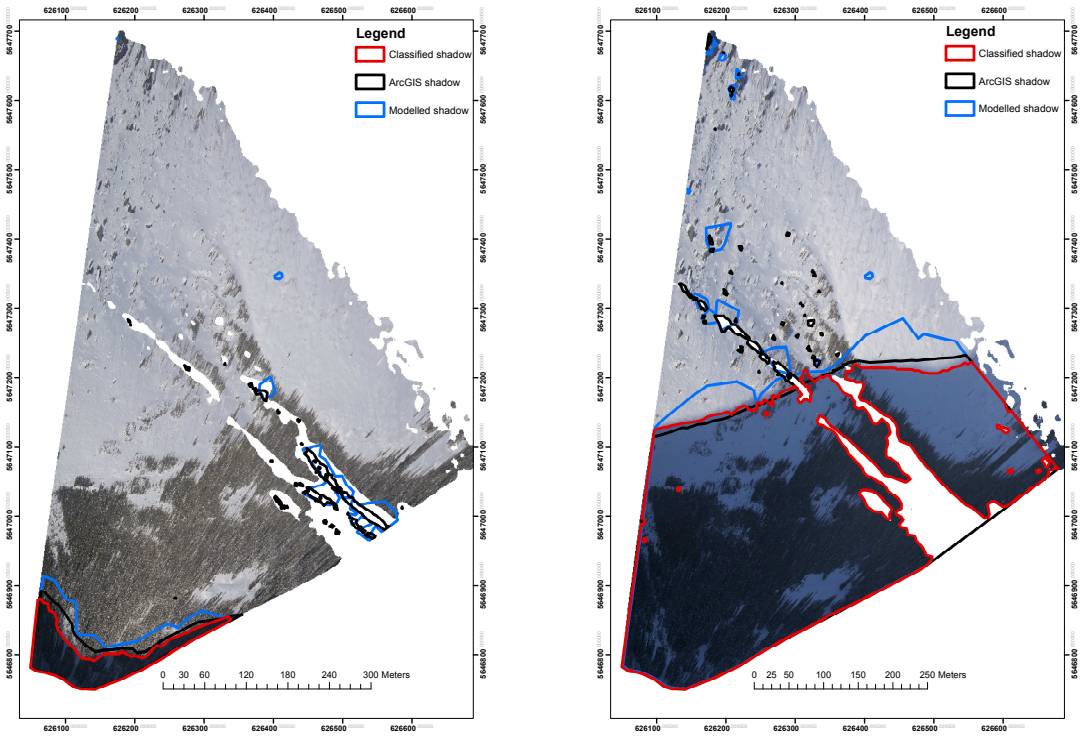

(a) 16:00 **(b)** 17:00

Figure 4.6: Comparison between the classified shadow from the orthorectified photographs (red), the modelled shadow from the unstructured mesh model (blue), and the ArcGIS shadows from the hillshade algorithm (black). Time is February 1, 2011 at (a) 16:00 and (b) 17:00.

In both figures, the shadow from the classified imagery is shown in red, the ArcGIS shadow in black, and the unstructured mesh model in blue. All are overlain on the orthorectified photographs. Both models produce a continuous shadow in an area with no orthorectified photographs (white areas; for example 626500 m E, 5647000 m to 6547100 m N). There is no imagery in these areas because they correspond to valleys, thus obscuring the photographs. Field observations have confirmed that these valleys do have shadows at this time, so it is not unexpected the models are producing shadows in the valley. The main shadow visible in the orthorectified

photographs is in the lower-left corner. ArcGIS does a good job at closely fitting the observed data. The areas where the unstructured mesh has shadow false-positives have a distinct triangular shape. Indeed, it is hypothesized that there are a few big triangles that are overestimated as shadowed due to the limited triangle-triangle collision algorithm. Based on the results presented in Section 4.5, the smaller the triangle tolerance (i.e., the higher the resolution), and thus smaller triangle size, the better the predictive capacities of the model. Thus, due to the inherent higher resolution of a structured mesh (no error due to triangle tolerances), it would be expected that the predicted shadows locations would be more accurate. Indeed, this is what we see. This is further shown in Figure 4.6b. Again it is observed that most of the difference between the models has a triangular nature, leading to further evidence that a better triangle-triangle detection algorithm should be explored in future work, such as that outlined by Devillers and Guigue (2002). This comparison suggests that with a more accurate triangle-triangle collision algorithm, unstructured meshes can perform on par with existing methods. Although this comparison does identify that the ArcGIS method is slightly more accurate at this resolution, the advantages of using an unstructured mesh plus the ability to further refine the detection algorithm suggests that choosing unstructured meshes is desirable.

4.4 Effects of shading on surface energy

4.4.1 Point scale

The point-scale difference of irradiance between the self- and horizon-shading models was averaged on a monthly basis over the course of the seasonal period (October–June). The largest average difference was 4.8 $\mathrm{W{\cdot}m^{-2}}$ during the period of December to February, occurring during periods with the lowest solar angles when the largest shadows are cast by the surrounding topography. Differences in irradiance decreased over the winter and spring through to the summer, with a monthly average difference of irradiance of less than 1 $\mathrm{W{\cdot}m^{-2}}$ in June. Shown in Figure 4.7 is the cumulative

difference of irradiance between the self- and horizon-shading algorithms for the winter season (October–June; dashed line) and the spring (April–June; solid line). The cumulative difference of irradiance accumulated in a linear fashion throughout the winter due to the long horizon-shadows from the low solar angles. With the arrival of spring and larger solar angles, differences became smaller and the cumulative difference in irradiance levelled off. Over the period of October–June, irradiance at the South Meadow site was overestimated by a total of 52 $MJ \cdot m^{-2}$ by the self-shading model. Some of this would be absorbed by the snowpack and contributes to an overestimation of the internal energy of the snowpack (i.e., the snowpack is too warm) over the winter and hence a ripening and melting of snow that are too early in the spring. Even when melt simulations are started in spring from correct initial conditions, there is a small, but potentially important, difference in estimations of snowmelt by ignoring horizon-shading. It is observed that a large portion of the cumulative difference develops in early to mid-April, with diminishing additional differences in irradiance through late May into early June. During this spring period, the cumulative difference in irradiance is 4.5 $MJ \cdot m^{-2}$. Following the conversion to phase change mass as outlined in Section 3.6.1, this translates into a SWE difference of 2.69 $kg \cdot m^{-2}$. These differences introduced when horizon shading is not accounted for results in a cumulative difference in irradiance between models over the season and the melt period. The difference in irradiance due to shadowing is predominantly of similar sign and generally does not average or cancel out over the season.

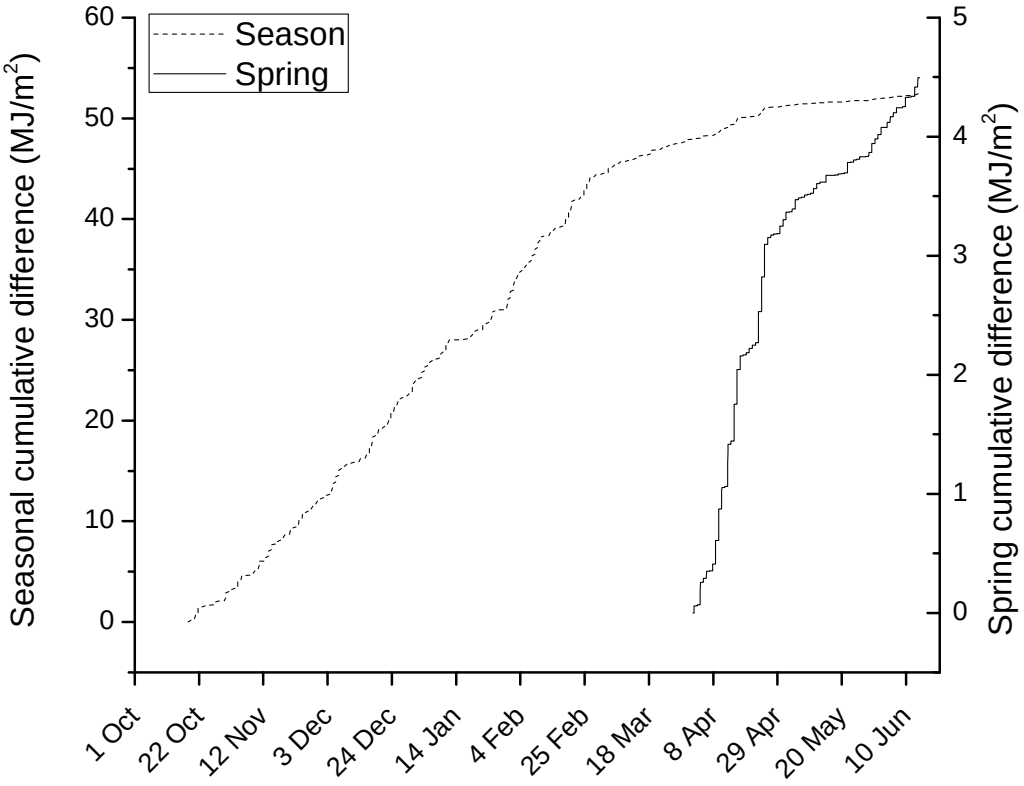

Figure 4.7: Cumulative differences in irradiance between simulations using self-shading and horizon-shading algorithms at the South Meadow site for the season (October 17, 2010–June 14, 2011) in the dashed line and for the spring (April 1, 2010–June 14, 2011) in the solid line.

4.4.2 Basin scale

The cumulative differences of net solar radiation, expressed as phase change mass per unit area, between self- and horizon-shading algorithms for the spring period from April 1, 2011 to June 14, 2011 were considered for the entire simulation domain (206580 triangles). For this period, 99% of the differences are less than 500 kg·m^{-2}, with a peak difference of 1600 kg·m^{-2} predicted over a small fraction of the basin. The mean phase change mass per unit area for the basin was 14.4 kg·m^{-2}.

In Figure 4.8, the seasonal (October–June) cumulative difference of net solar radiation between the two shading algorithms is shown, expressed as phase change mass per unit area. For the majority of this period, low solar angles cast large shadows onto the south-facing slopes, providing the bulk of the difference. The south-facing slope of Mt. Collembola, the slope where South Meadow is located, is a frequent recipient of horizon-shadows from Fisera Ridge. The south-facing slopes of the Mt. Allan cirque, shaded from the surrounding ridge line, are also areas where the irradiance is over-predicted by the self-shading algorithm. The right side of the large 'V' notch on Mt. Collembola is strongly over-predicted using the self-shading algorithm, and timelapse camera observations showed that this large shadow was associated with an abrupt transition between snow and snow-free areas in the notch. In the top left (NW) corner is an area outside of the MCRB. It is a 'U' shaped valley that is almost always in the shadow cast from Mt. Collembola. This topography resulted in a large over-prediction of irradiance throughout the melt period if horizon-shadows are not considered — in fact, most of the largest differences in cumulative difference of net solar radiation between the shadowing models are found in this valley. The south face of the southern ridge line (approximately 6.27 $\times$ 10^6 m E) is shown to have some of the largest differences due to horizon-shadows from the ridge line. The largest differences in net shortwave radiation between shading models for the MCRB are observed to be on east- and south-facing slopes.

For the melt period of April–June 2011, Figure 4.9 shows the cumulative difference of net solar radiation between self- and horizon-shading models, expressed as

phase change mass per unit area. South-facing slopes are again predicted to be areas of large cumulative difference in net solar radiation. The south-facing slope of South Meadow is shown to have a large cumulative difference of net solar radiation, and in general the basin exhibits a similar pattern to that of the seasonal period. However, a notable change is that the north-facing southern ridge (at about 6.27×10^6 m E) has a large cumulative difference of net solar radiation. Generally this area should be self-shaded; however, due to the higher solar angles the triangle faces are not fully facing away from the sun. Although this results in triangles with low values of $\cos\Theta_{sn}$, these values are not so low as to generate self-shadows. Thus during the morning as the sun passes through to the south, the direct-beam irradiances calculated using self-shading are not 0 $\mathrm{W{\cdot}m^{-2}}$, and thus there is a large over-prediction of net solar radiation. Because of these accumulated differences, higher-than-estimated snowmelt would likely be observed in this valley if only self-shading algorithms were used, and the persistence of late-winter snowpacks would not be appropriately modelled. However, these results come with a caveat: the MCRB is dominated by tree-cover, and as such, the sub-canopy irradiance is significantly reduced. Ellis and Pomeroy (2007) demonstrated that mean daily canopy transmissivity values were as low as 0.09 on north-facing slopes and 0.36 on south-facing slopes. Because this study assumes a bare-earth, some of the locations may be impacted less by ignoring horizon-shadows. However, Ellis and Pomeroy (2007) indicated that the forest-cover on south-facing slopes had the greatest variance with respect to incident irradiance. Thus the south-face slopes that are dominated by horizon-shadows may significantly impact the sub-canopy energy balance.

Figure 4.10 depicts the cumulative difference of net solar radiation between the self- and horizon-shading models expressed as phase change mass per unit area for the entire clear-sky day February 1, 2011. This demonstrates the significant effect of the differences in algorithms for even a single day. The deep southern valley at about 6.27×10^5 m East is a source of large net solar radiation differences with average phase change mass values of about 5 $\mathrm{kg{\cdot}m^{-2}}$. Overall, the valley bottoms have over estimated net solar radiation and in many cases greatly so. The South Meadow site

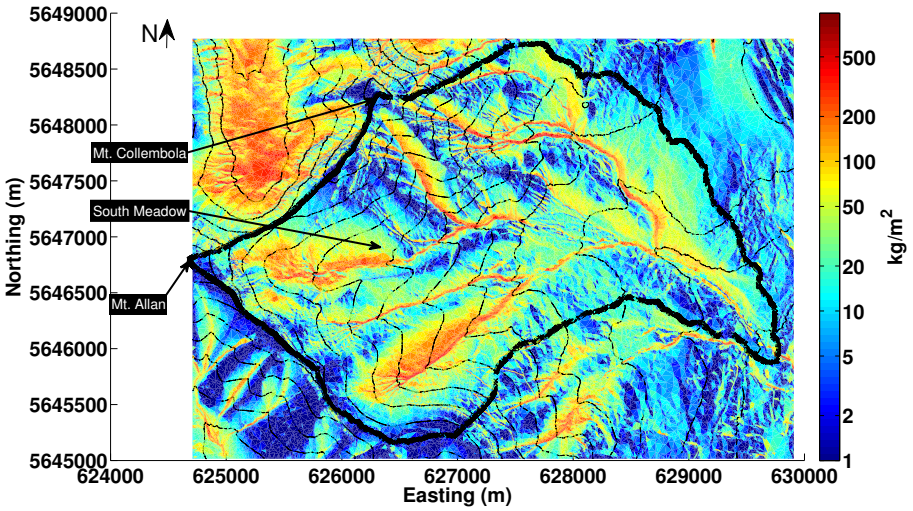

Figure 4.8: Cumulative seasonal (October 17, 2010–June 14, 2011) difference of net solar radiation between self- and horizon-shading algorithms expressed as phase change mass per unit area. An albedo of 0.8 is assumed. Note the logarithmic scale. The Marmot Creek Research Basin is outlined in black and contour lines are every 100 m of elevation.

is located such that although it has a prominent shadow crossing, its accumulated difference is only approximately 2 kg·m^{-2}. The north and south facing sites are areas where limited difference in net solar radiation would be expected, and on this day that remains the case.

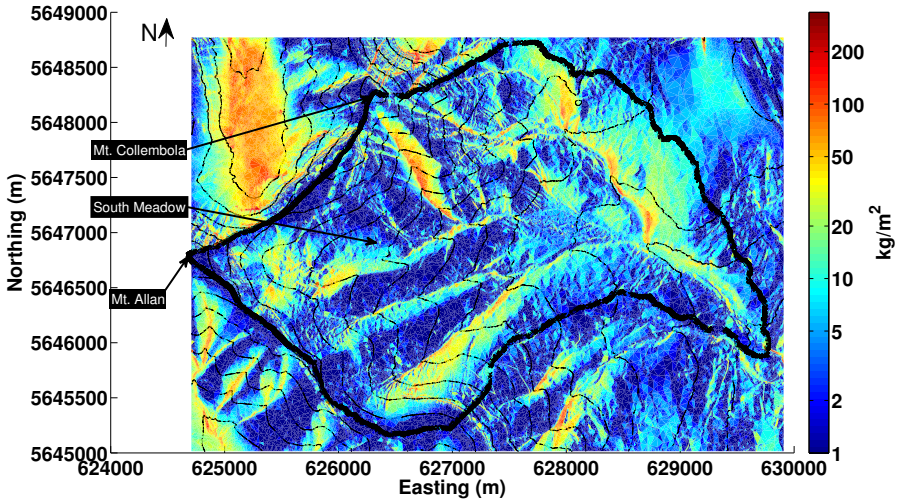

Figure 4.9: Cumulative spring (April 1, 2011–June 14, 2011) difference of net solar radiation between self- and horizon-shading algorithms expressed as phase change mass per unit area. An albedo of 0.8 is assumed. Note the logarithmic scale. The Marmot Creek Research Basin is outlined in black and contour lines are every 100 m of elevation.

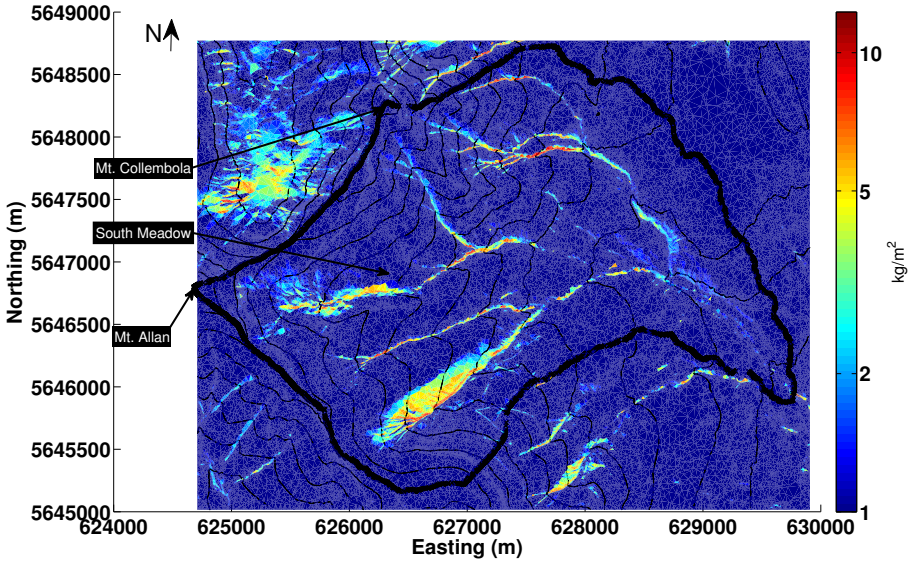

Figure 4.10: Cumulative daily (February 1, 2011) difference of net solar radiation between self- and horizon-shading algorithms expressed as phase change mass per unit area. An albedo of 0.8 is assumed. Note the logarithmic scale. The MCRB is outlined in black and contour lines are every 100 m of elevation.

4.4.3 Comparison with measured SWE

DeBeer and Pomeroy (2009) presented the results from a set of snow surveys conducted in the Mt. Allan cirque over the later winter and spring of 2007. The cirque is on the eastern side of Mt. Allan and is approximately 0.6 km^2 in area. Extensive snow surveys were made on March 29–30, 2007, just before the main snowmelt period. These values are taken to be the pre-melt SWE values and are summarized in Table 4.1. The shadow model for April–June 2011 was applied to the cirque and the cumulative difference of net solar radiation between the self- and horizon-shading algorithms was computed and expressed as phase change mass per unit area. For the slopes of azimuth 0°–40°, the average difference was 19.7 $kg{\cdot}m^{-2}$, and for the slopes of azimuth 140°–180°, the average difference was 16.3 $kg{\cdot}m^{-2}$. For the north- and south-facing slopes, the cumulative difference of net solar radiation converted to phase change mass equates to 6.0% and 11.4% of total pre-melt SWE respectively. The cumulative difference is 66.4% of the total pre-melt SWE for the shallow south-facing slope. This suggests that including horizon-shading is important when using physically based, non-calibrated hydrological models.

DeBeer and Pomeroy (2009) identified that the spatial variability of the surface energetics due to slope and aspect in the cirque had a strong effect on snowmelt rates and snow cover depletion (SCD). Because of the patchiness of SCD, correct prediction of the surface energy balance is required, and DeBeer and Pomeroy (2009) suggest that melt energetics on slope units should be considered further. Although timing issues in SCD were generally attributed to errors in the log-normal fit, it was

Table 4.1: Listing of the pre-melt snow water equivalent for the Mt. Allan cirque following (DeBeer and Pomeroy, 2009)

Survey location	**SWE ($kg{\cdot}m^{-2}$)**	**Average aspect (°)**	**Average slope (°)**
North-facing slope	329.3	20	30
South-facing slope	142.8	160	23
South-facing (shallow SWE)	24.6	-	-
South-facing (deep SWE)	339.8	-	-
All surveys	427.2	-	-

suggested that improvements might also arise from more accurate melt simulations in sheltered areas. Slope units were considered in DeBeer and Pomeroy (2010) where it was shown that consideration of slope-based melt energetics substantially improved the SCD calculations. It is possible that SCD calculations would improve further when coupled with horizon-shading models.

4.5 Effects of terrain resolution on model

As outlined in Section 3.3, the LiDAR-derived 1m × 1m DEM was resampled to produce a raster with a spatial resolution of 2 m × 2 m, a 10 m × 10 m, and 30 m × 30 m. These structured meshes were used to create the unstructured meshes with tolerances of 1 m, 2 m, 10 m, and 30 m for each DEM resolution. An example of these combinations is shown in Figure 4.11 where the base DEM is the 1 m × 1 m with (a) 1 m, (b) 2 m, (c) 10 m, and (d) 30 m triangle tolerances. The MCRB is in dark black, and contours are every 100 m. The triangles are coloured using the mean elevation of the triangle with red corresponding to the highest elevation, and blue corresponding to the lowest elevation. When the tolerance increases the terrain is represented by larger triangles. Consequently, these large triangles produce a disjoint landscape representation, as high elevation triangles are adjacent to lower elevations. There is a lack of smooth transition between elevation bands at these higher tolerances.

4.5.1 Point scale

The effects of changes of spatial resolution and tolerance were tested at a point scale for Fisera Ridge North, Fisera Ridge South, and South Meadow for the clear-sky day February 1, 2011 in order to correspond to the day used for the timelapse photos. The following figures are grouped into four subfigures, one for each base DEM resolution. Within each subfigure, the red line demonstrates the highest-resolution model output — 1 m × 1 m DEM/1m tolerance triangles. All other lines are: black for the 1 m tolerance, blue for the 2 m tolerance, green for the 10 m tolerance, and grey for

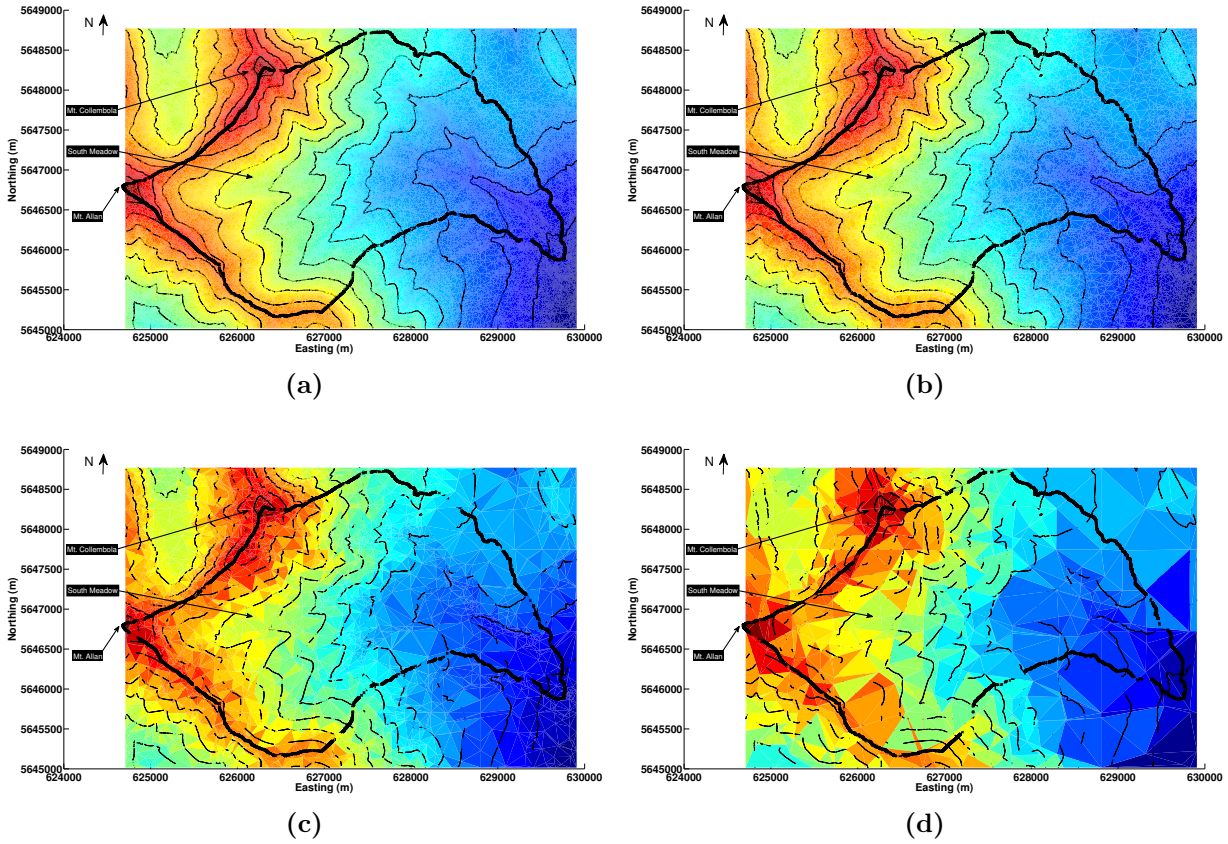

Figure 4.11: Example of four DEM/tolerance combinations. The base DEM is the 1 m × 1 m with (a) 1 m, (b) 2 m, (c) 10 m, and (d) 30 m triangle tolerances. The MCRB is in dark black, and contours are every 100 m. The triangles are coloured using the mean elevation of the triangle. The larger tolerances do not produce fully connected contours and thus there are gaps.

the 30 m tolerance. In order to determine the effectiveness of a particular terrain representation in capturing the shadow movement, the modelled shadowing regime must conform to the observed shadow regime.

South Meadow

At the South Meadow site, as discussed in Section 4.1, a horizon-shadow crosses the site at 16:30. Figure 4.12 shows the sixteen DEM resolution and mesh tolerance combinations showing the differences in timing for the South Meadow site February 1, 2011.

Beginning with the 1 m DEM, the 1 m and 2 m tolerances perform similarly. Although a slight increase in peak irradiance is observed for the 2 m tolerance mesh, the correct timing is present. The 10 m tolerance produces almost 100 $W \cdot m^{-2}$ greater peak irradiance than the 1 m tolerance. In addition, the shadow crossing is roughly 45 minutes too early. The 30 m tolerance mesh fails completely to capture the variability. Because of the large triangles, the site is in shadow almost all day, emerging three times for brief periods before being fully in shadow almost an hour early. The 2 m DEM shows similar traits with the 1 m and 2 m tolerances being indistinguishable. The 10 m tolerance models the shadow crossing about 15 minutes early. Interestingly the 30 m tolerance is not as poor as the 1 m DEM/30 m tolerance mesh combination. It is not currently known why the 1 m DEM/30 m tolerance was the only one to produce this effect. However, the 2 m DEM/30 m tolerance mesh produces a significantly lower peak irradiance than the other models. The 10 m DEM/1 m tolerance mesh performs the best out of this group, correctly producing the shadow timing. The 2 m tolerance, in contrast to the first two DEM groups, does not produce as good of a result as the 1 m tolerance. In fact, the 2 m tolerance performs similarly as the 10 m and the 30 m tolerances. The 30 m DEM variants all perform as well as the 1 m tolerances of the other groups; however the 30 m tolerance exits shadow at 16:48. From all of this, the general trend is that the 1 m tolerance meshes produce the most accurate shadow timings, regardless of the base DEM resolution and the 2 m tolerances produce similar results. Tolerances any larger are likely to have detrimental effects upon the shadow crossing with unpredictable results. This unpredictability demonstrates that there are certain combinations of base DEM and tolerance that appear to fail completely. These failure cases were not anticipated *a priori* and thus demonstrate a need for care if a model is run at larger mesh tolerances. It is generally risky to accept the results of calculations based on a single set of parameters; these results reinforce the better practice of using a number of tolerances and comparing the results.

The following sections show similar results for the other sites in the basin.

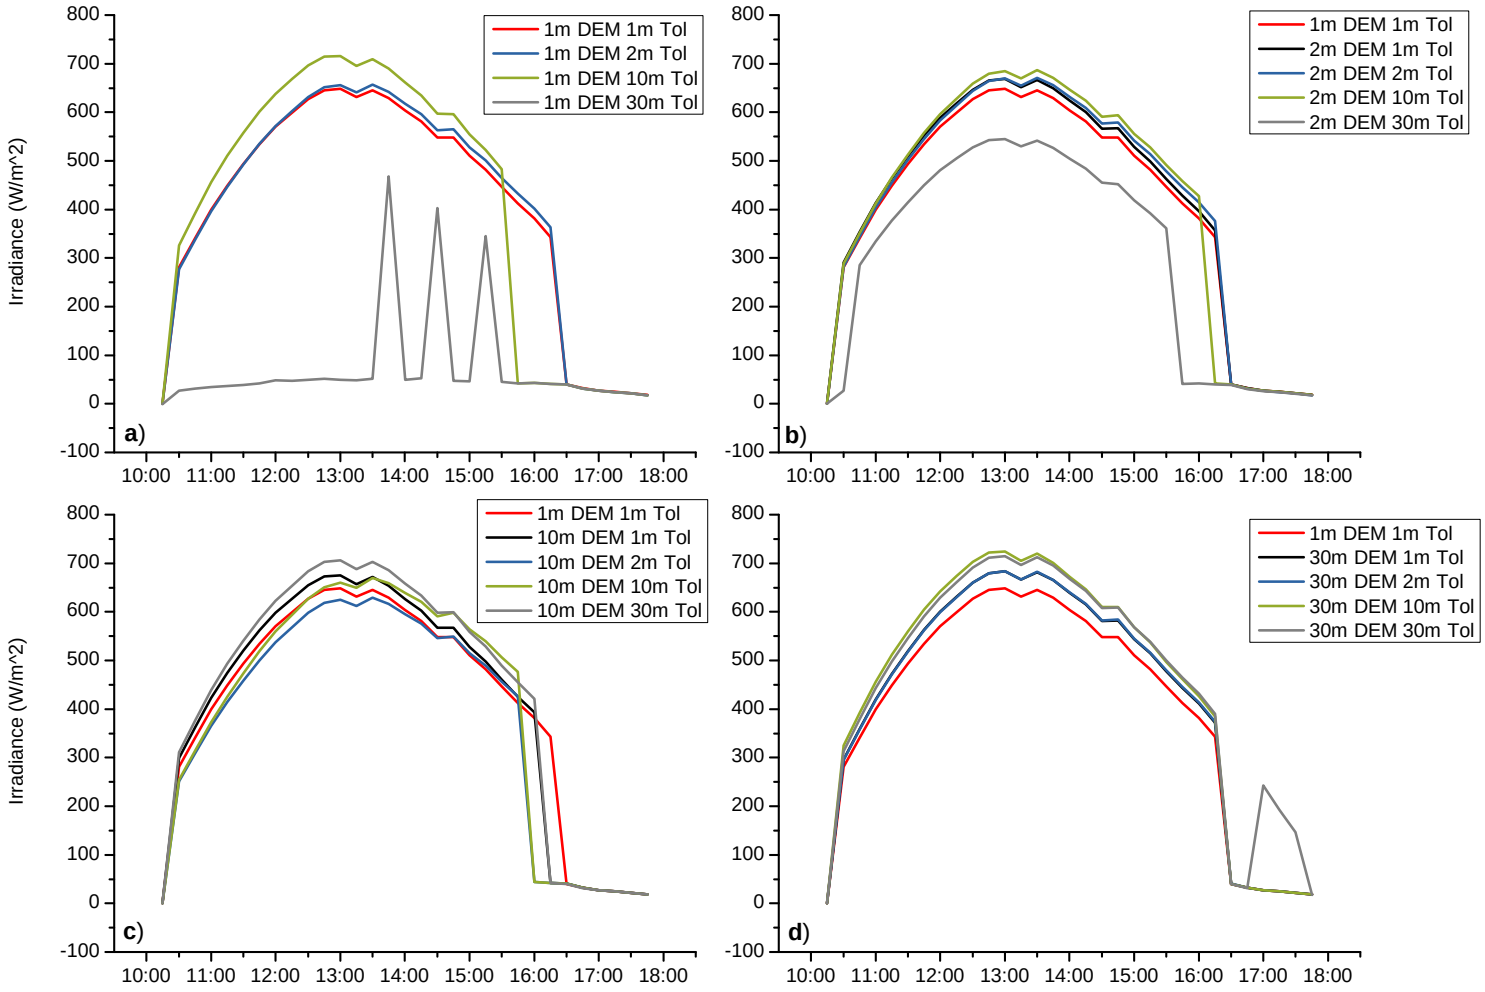

Figure 4.12: Sixteen resolution and tolerance combinations showing the differences in timing for the South Meadow site February 1, 2011. The 1 m × 1 m DEM/1m tolerance is shown in red. Other lines are: black for the 1 m tolerance, blue for the 2 m tolerance, green for the 10 m tolerance, and grey for the 30 m tolerance.

Fisera Ridge North

The Fisera Ridge North site is a north-facing slope that is predominately in shadow, either self- or horizon-shaded, throughout the day. Therefore low irradiance values are expected. Figure 4.13 shows the sixteen DEM resolution and mesh tolerance combinations.

This site is a simpler case, and all the terrain representations perform well expect the 10 m DEM/30 m tolerance mesh and 1 m DEM/30 m tolerance mesh. These two cases fail catastrophically. The 1 m DEM/30 m tolerance mesh fails to be shaded in the morning leading to 140 $\mathrm{W{\cdot}m^{-2}}$ overprediction until 13:30, at which point it returns to being shaded. The 10 m DEM/30 m tolerance mesh begins being properly shadowed, but twice throughout the day returns to a non-shaded state before returning to being shadowed.

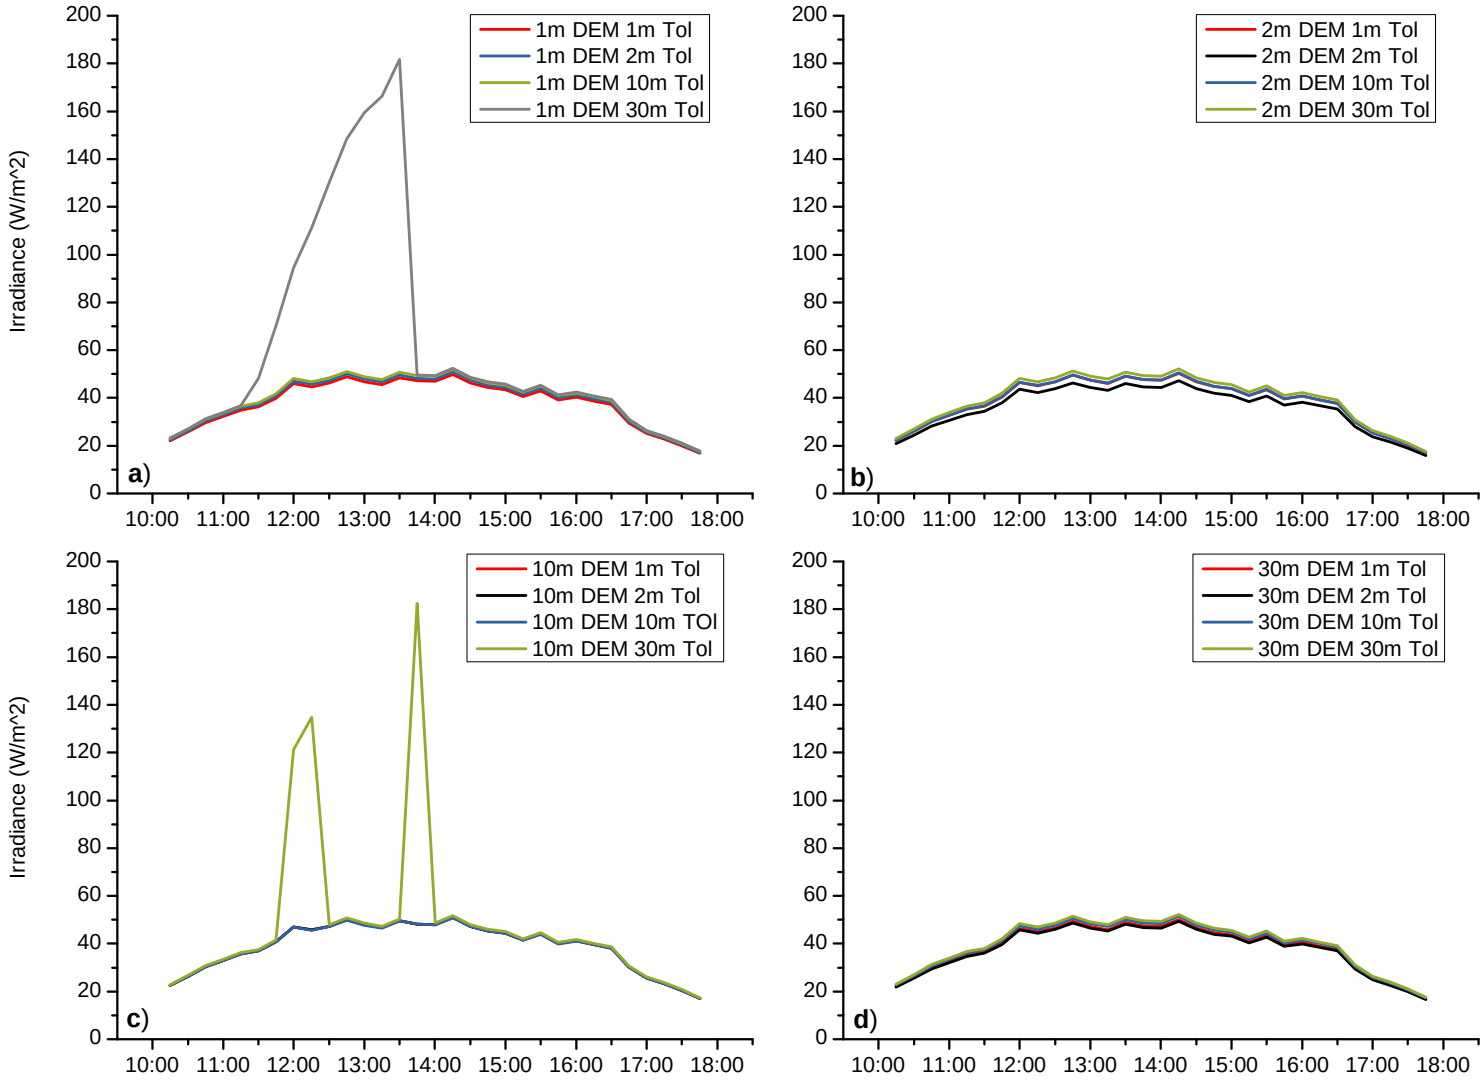

Figure 4.13: Sixteen resolution and tolerance combinations showing the differences in timing for the Fisera Ridge North site February 1, 2011. The 1 m × 1 m DEM/1m tolerance is shown in red. Other lines are: black for the 1 m tolerance, blue for the 2 m tolerance, green for the 10 m tolerance, and grey for the 30 m tolerance.

Fisera Ridge South

The Fisera Ridge South site is a south-facing slope of Fisera Ridge. It receives high irradiance, and like South Meadow, it passes in to shadow late in the afternoon. The shadowing transition is less prolonged than the South Meadow site due to its position high on the ridge line, as opposed to low in the valley. Figure 4.14 shows the sixteen DEM resolution and mesh tolerance combinations.

The 1 m DEM provides similar timings of shadow crossing for the 1 m, 2 m, and 10 m tolerances. Only their peak irradiances differs. The 30 m tolerance, like the other sites, does not properly capture the movement of the shadow and predicts the movement with a lag of about 15 minutes. In addition, the larger tolerance overestimates the flatness of the slope, producing higher irradiances than the smaller tolerances. However, the 10 m tolerance overestimates the steepness of the slope, thus producing significantly lower irradiances. The 2 m DEM has similar performance, although the 10 m tolerance significantly underestimates peak irradiance due to slope overestimation. The 30 m tolerance is once again the weakest and not only overpredicts peak irradiance, it fails to capture the timing of the shadow twice — the site is predicted to move out of shadow at 17:00. The 10 m DEM is once again the first group to have large systematic errors. The 1 m and 2 m tolerances have proper shadow timing; however the 2 m tolerance overpredicts slope resulting in a larger peak irradiance. The 10 m tolerance is modelled as shadowed from 13:00 until 14:30 where it makes a brief exit from shadow before returning to shadow at 14:35. The 30 m tolerance produces poor results; it descends into shadow at 14:30 before rapidly exiting shadow at 14:45. The 30 m base DEM produces good timing results when used with the 1 m and 2 m tolerances. However the 10 m tolerance produces a sharp oscillation throughout the day as the site flips between shadowed and non-shadowed. The 30 m tolerance performs surprisingly well, and this is attributed to the large flat slope being well represented by single, large triangles. However, because the others sites perform so poorly at this tolerance and DEM resolution, this seems to be a rare event.

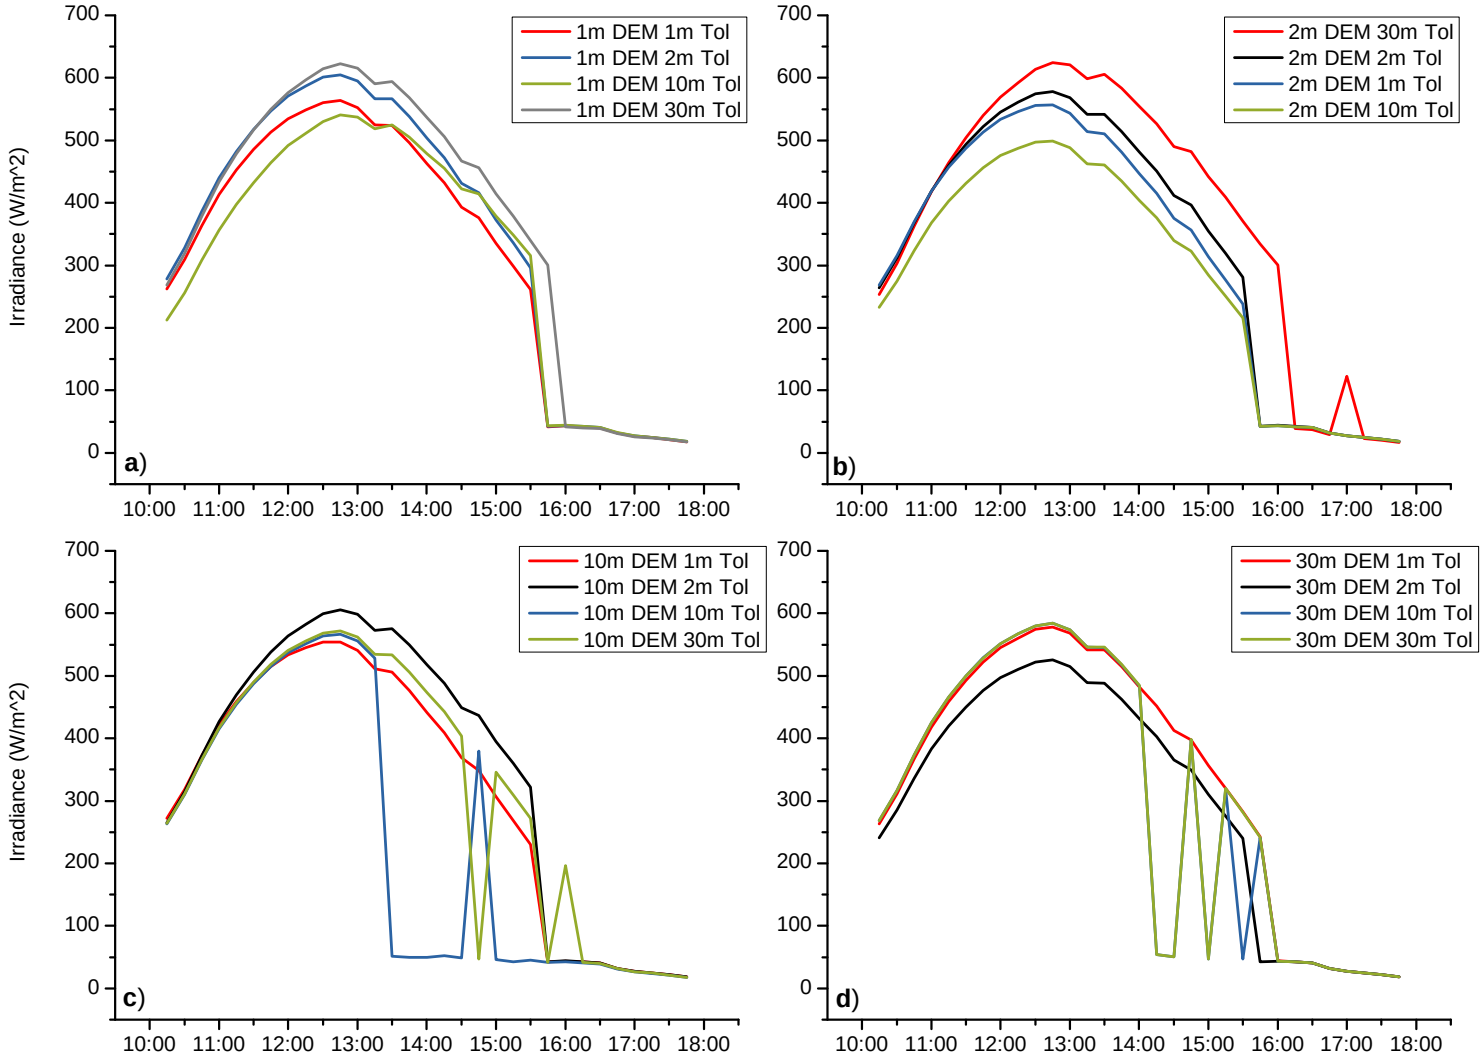

Figure 4.14: Sixteen resolution and tolerance combinations showing the differences in timing for the Fisera Ridge South site February 1, 2011. The 1 m × 1 m DEM/1m tolerance is shown in red. Other lines are: black for the 1 m tolerance, blue for the 2 m tolerance, green for the 10 m tolerance, and grey for the 30 m tolerance.

4.5.2 Point-scale accumulated difference

As demonstrated in Section 4.4, there can be a systematic daily difference between a self-and horizon-shading model that can accumulate. Thus every day there can be a systematic overprediction in irradiance at a point because the self-shading model fails to take into account a horizon-shadow. In some locations this can be a small difference; in others, it can accumulate quickly. As shown in Section 4.5.1, when modelling this difference at different base DEM resolutions and mesh tolerances, the difference between the different DEM resolutions or the triangle tolerance can greatly change the results. Although it was demonstrated that 1 m and 2 m tolerances provide the best point-scale representation, it is instructive to see how this translates to cumulative energy differences over a given temporal period. In this section, the spring temporal period of April 2011 to June 2011 is examined. Given the results from the previous section, it is taken that the 1 m DEM/1 m tolerance mesh is the closest to the 'truth'. All other DEM and tolerance mesh combinations are thus compared to it. In all figures, the line colours are: red for the 1 m tolerance, black for the 2 m tolerance, blue for the 10 m tolerance, and green for the 30 m tolerance.

South Meadow

In Figure 4.15, sixteen resolution and vertical tolerance combinations were compared to the the 1 m DEM/1 m tolerance mesh and the cumulative difference in irradiance was computed. As expected, the 1 m DEM/2 m tolerance mesh performs quite closely to the 1 m DEM/1 m tolerance mesh. The different DEM/tolerance combinations follows the situation of Section 4.5.1, in which the smaller tolerance meshes mostly independent of base DEM resolution, perform increasingly poorly. All of the 10 m tolerance meshes have an 8 $MJ \cdot m^{-2}$ to 15 $MJ \cdot m^{-2}$ cumulative difference in irradiance. However most tolerances are in the 4 $MJ \cdot m^{-2}$ – 6 $MJ \cdot m^{-2}$ range. The 1 m DEM/30 m tolerance mesh and the 2 m DEM/30 m tolerance meshes have a completely different shadow regime. Their cumulative difference in irradiance has a negative sign, indicating that the model is predicting the site to be in shadow much

more than observed. This overprediction of shadowing was observed in Figure 4.12, where the South Meadow site only briefly left shadow in three spiked intervals. This underestimation of irradiance makes these larger tolerances unusable.

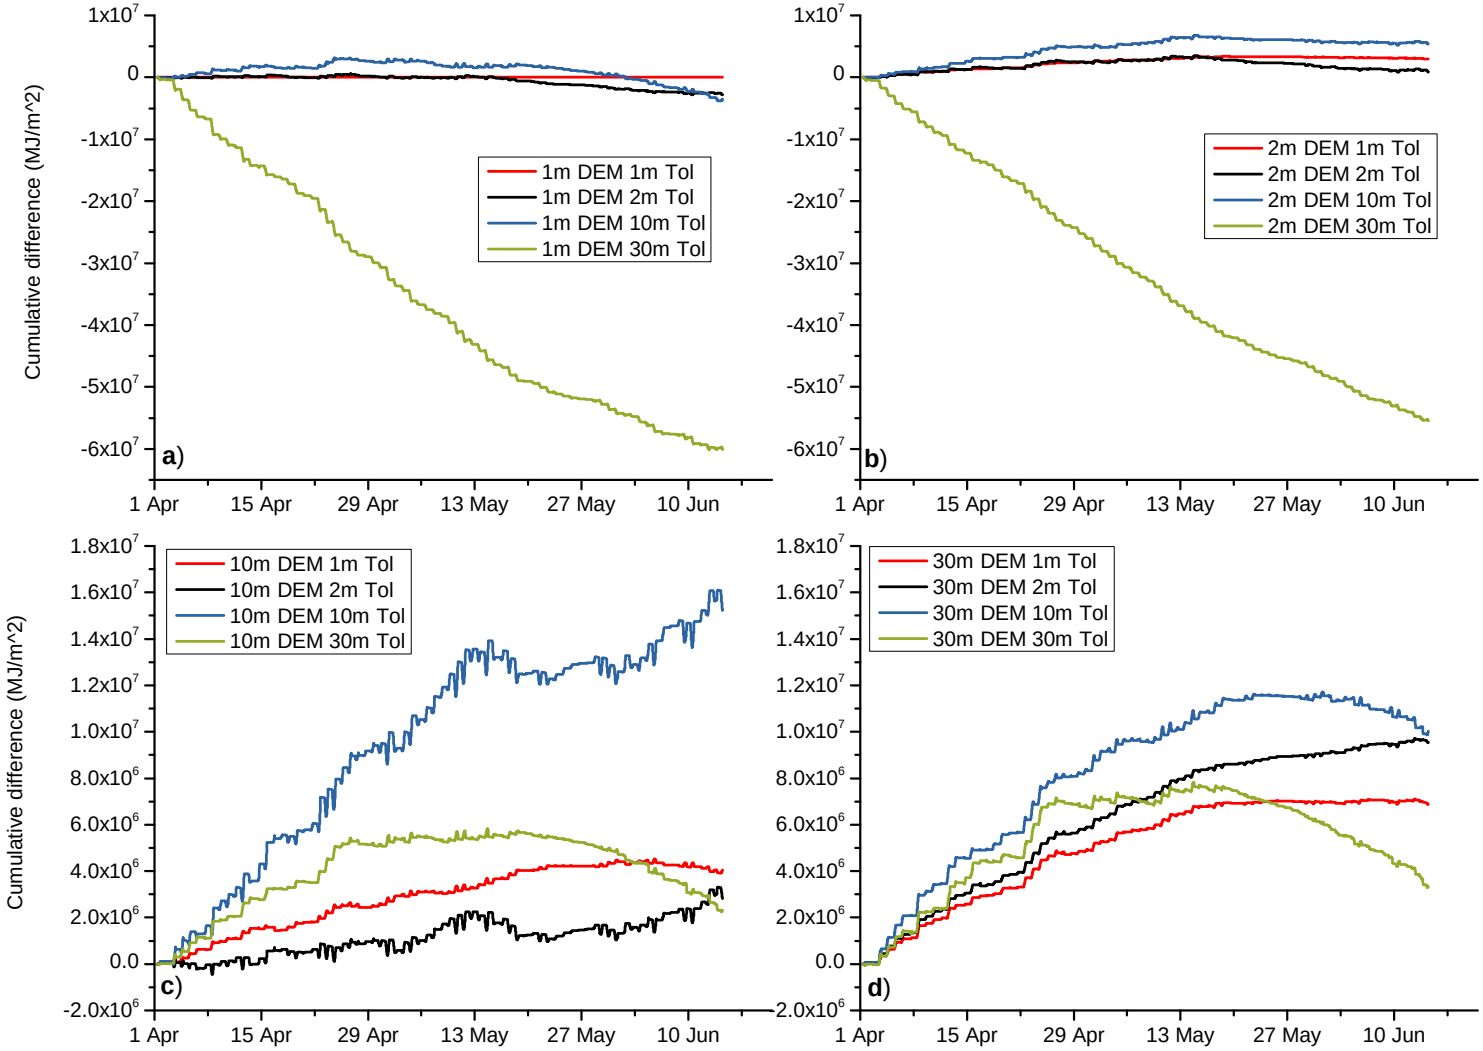

Figure 4.15: Sixteen resolution/tolerance combinations showing the accumulated difference in irradiance for South Meadow for the spring. The 1 m DEM/1m tolerance mesh is taken as true. Line colours: red for the 1 m tolerance, black for the 2 m tolerance, blue for the 10 m tolerance, and green for the 30 m tolerance.

Fisera Ridge North

In contrast to the South Meadow site, no combination closely matches the 1 m DEM/1 m tolerance mesh. This is demonstrated in Figure 4.16. The most similar is the 1 m DEM/10 m tolerance mesh and the 1 m DEM/2 m tolerance mesh performs similarly to 2 m DEM/1 m tolerance mesh. Again, a similar pattern to the other sites is observed where the larger tolerance meshes have the largest cumulative differences. The largest over-predictions are caused by the four 30 m tolerances and the 30 m DEM/30 m tolerance mesh has the largest difference of 175 $MJ \cdot m^{-2}$. The two largest under-predictors are the 30 m DEM/2 m tolerance mesh and (surprisingly) the 2 m DEM/2 m tolerance mesh.

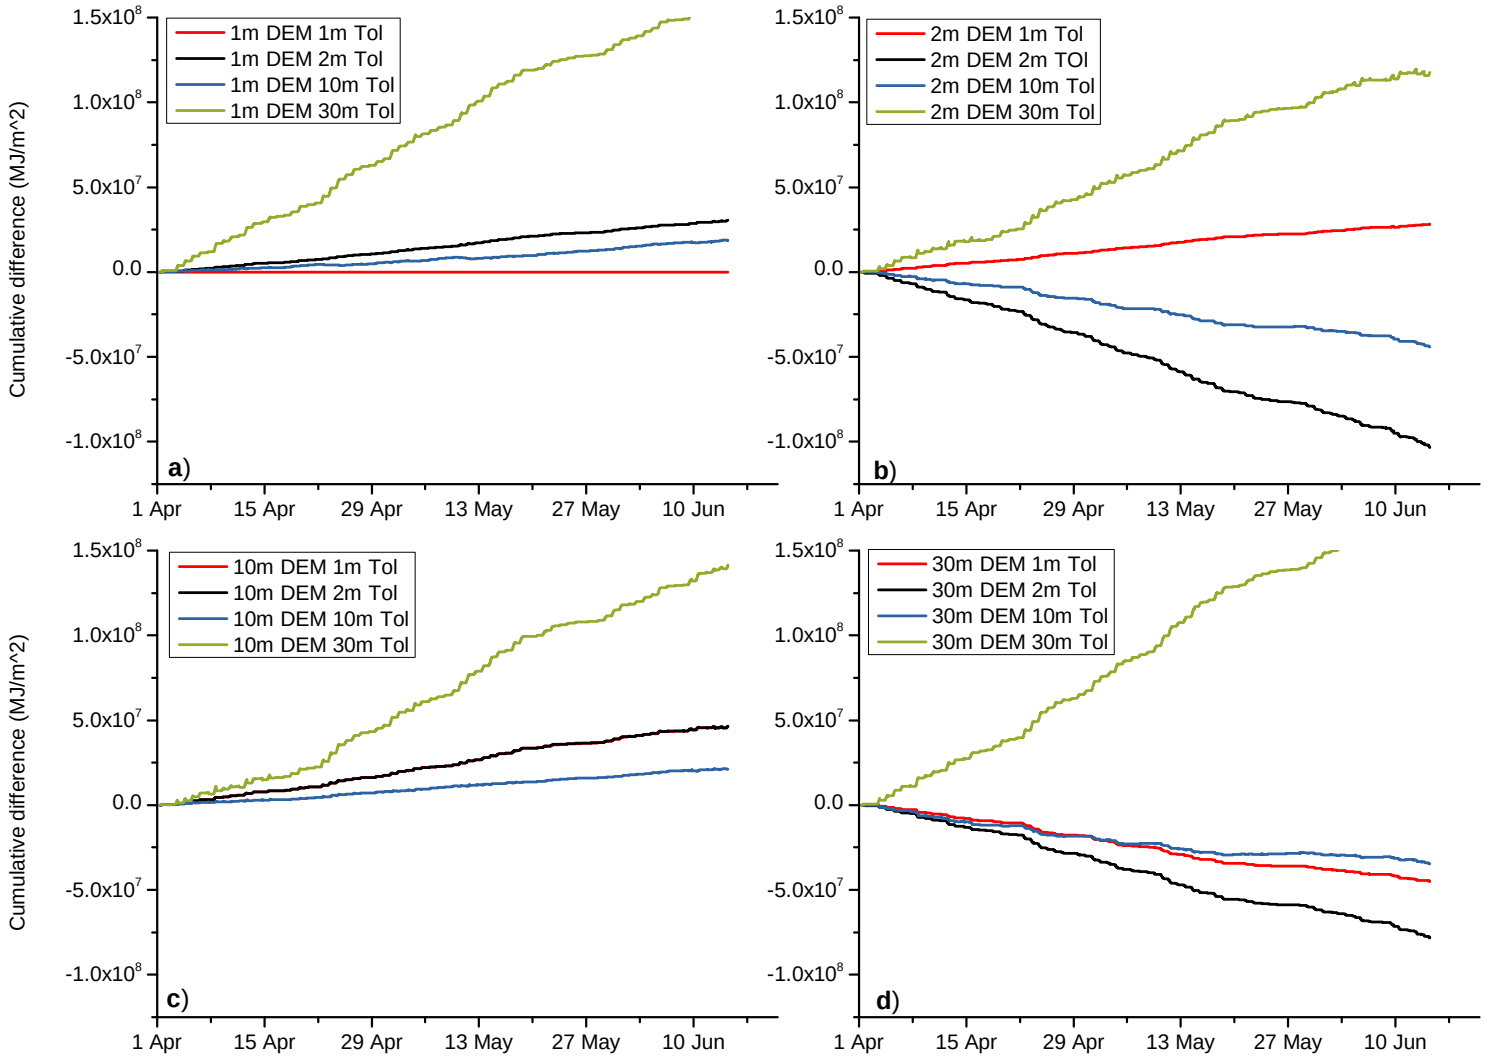

Figure 4.16: Sixteen resolution/tolerance combinations showing the accumulated difference in irradiance for Fisera Ridge North for the spring. The 1 m DEM/1m tolerance mesh is taken as true. Line colours: red for the 1 m tolerance, black for the 2 m tolerance, blue for the 10 m tolerance, and green for the 30 m tolerance.

Fisera Ridge South

Shown in Figure 4.17, Fisera Ridge South shows the most variance in cumulative differences in irradiance. Surprisingly it is the 10 m DEM/30 m tolerance mesh that performs closest to the reference solution. The general pattern observed at the other sites where the 1 m and 2 m tolerances perform well does not hold here. The 10 m DEM/2 m tolerance mesh and the 1 m DEM/1 m tolerance mesh perform almost equally well and accrue the largest cumulative difference in irradiance of 40 MJ. The 30 m DEM/10 m tolerance mesh and 10 m DEM/10 m tolerance mesh have a similar linear increase in accumulated difference in irradiance; however one is an overestimation and the other an underestimation. The 1 m DEM/30 m tolerance and 2 m DEM/30 m tolerance meshes both have a shift in mid-May where the accumulated difference in irradiance decreases after an initial accumulation. The interpretation of these results is as follows. For the Fisera Ridge South site, the horizon-shadows that cross it are cast from the reasonably distant Centennial Ridge (see Figure 3.1). Small changes in the mesh can lead to large differences in shadowing when calculating shadows from far away terrain. Thus for a site like Fisera Ridge South, these small shadowing differences can quickly cause a site to have a reasonably difference set of shadow regimes for different mesh combinations. The 30 m tolerance meshes have been so far shown to produce poor results and this tolerance produces poor results for this site. It is hypothesized that as the sun moves through the sky, the large triangles eventually make a sudden change in where shadows are predicted.

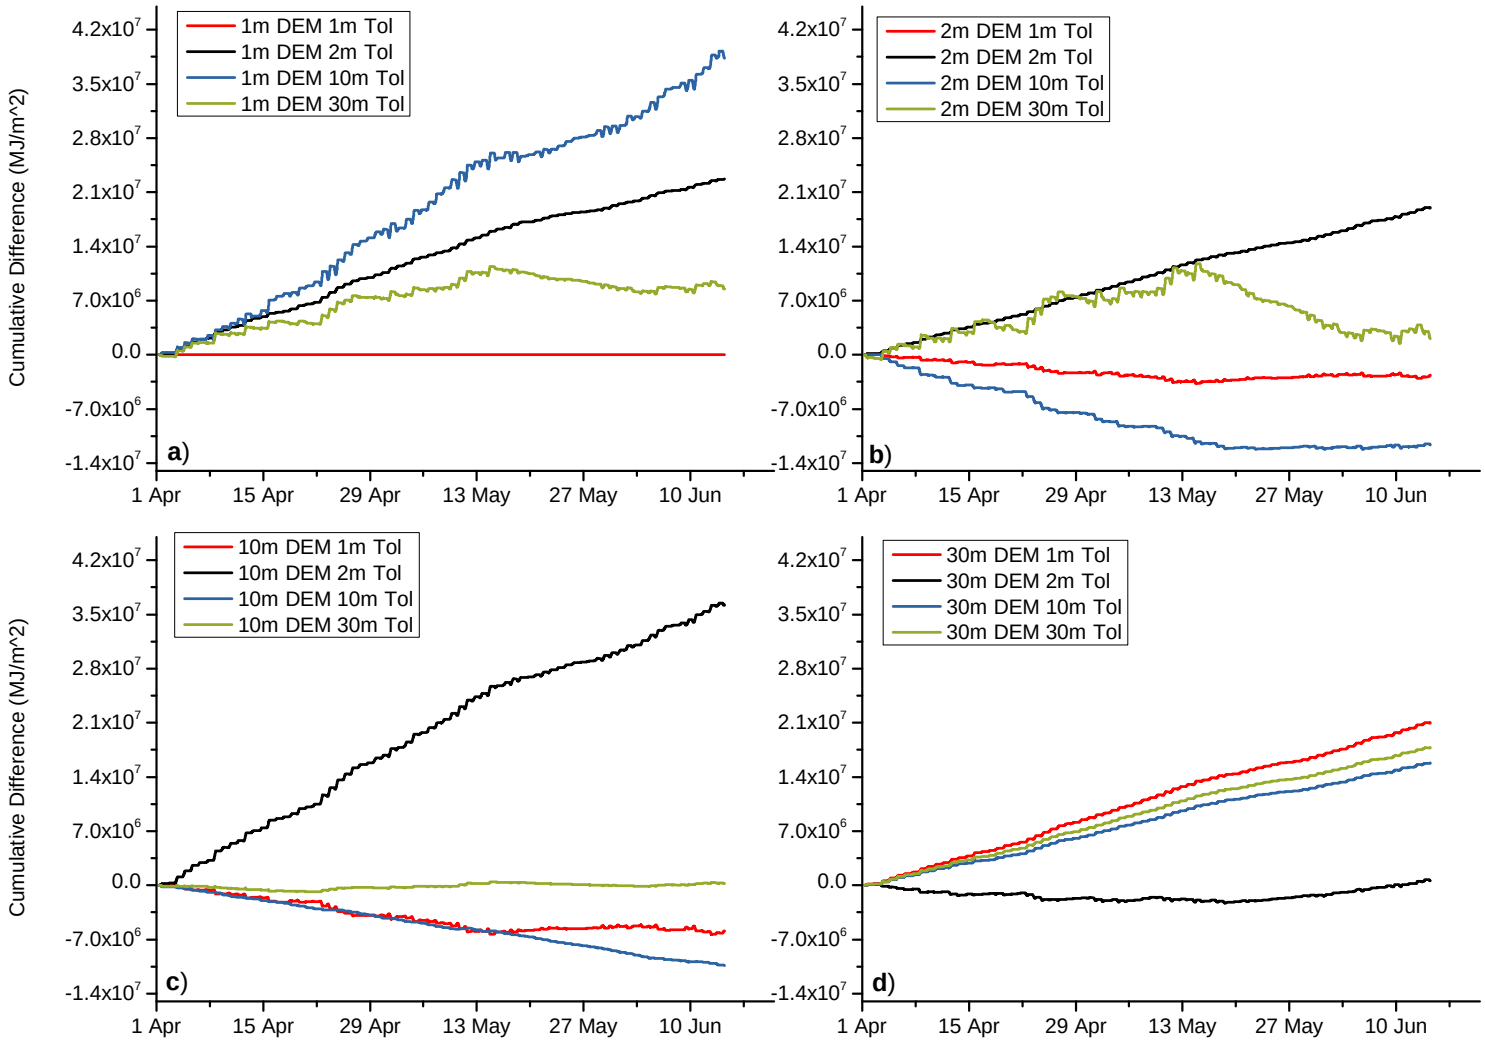

Figure 4.17: Sixteen resolution/tolerance combinations showing the accumulated difference in irradiance for Fisera Ridge South for the spring. The 1 m DEM/1m tolerance mesh is taken as true. Line colours: red for the 1 m tolerance, black for the 2 m tolerance, blue for the 10 m tolerance, and green for the 30 m tolerance.

4.5.3 Basin-scale accumulated difference

The point-scale results show that the mesh tolerances have the largest effect on the timing and peak irradiance. This is not unexpected because larger tolerances result in larger, flatter triangles. Flatter triangles have less cosine-correction for self-shading and thus result in higher peak irradiances. This is easiest to visualize by comparing the cumulative difference of net solar radiation expressed as phase change mass for different resolution and tolerance combinations. Figure 4.18 shows the 1 m DEM and 1 m, 2 m, 10 m, and 30 m tolerance meshes expressed as phase change mass per unit area. From Figure 4.18d it is clear why such poor performance is observed when using such a large tolerance. The triangles are so large as to be immediately identifiable, and because of this the basin is a patchwork of net shortwave radiation differences. Areas of large net solar radiation difference are adjacent to triangles of lower net shortwave radiation difference. This is in comparison to Figure 4.18a where individual triangles are difficult to see, and there are graduations between areas of differing net shortwave radiation differences.

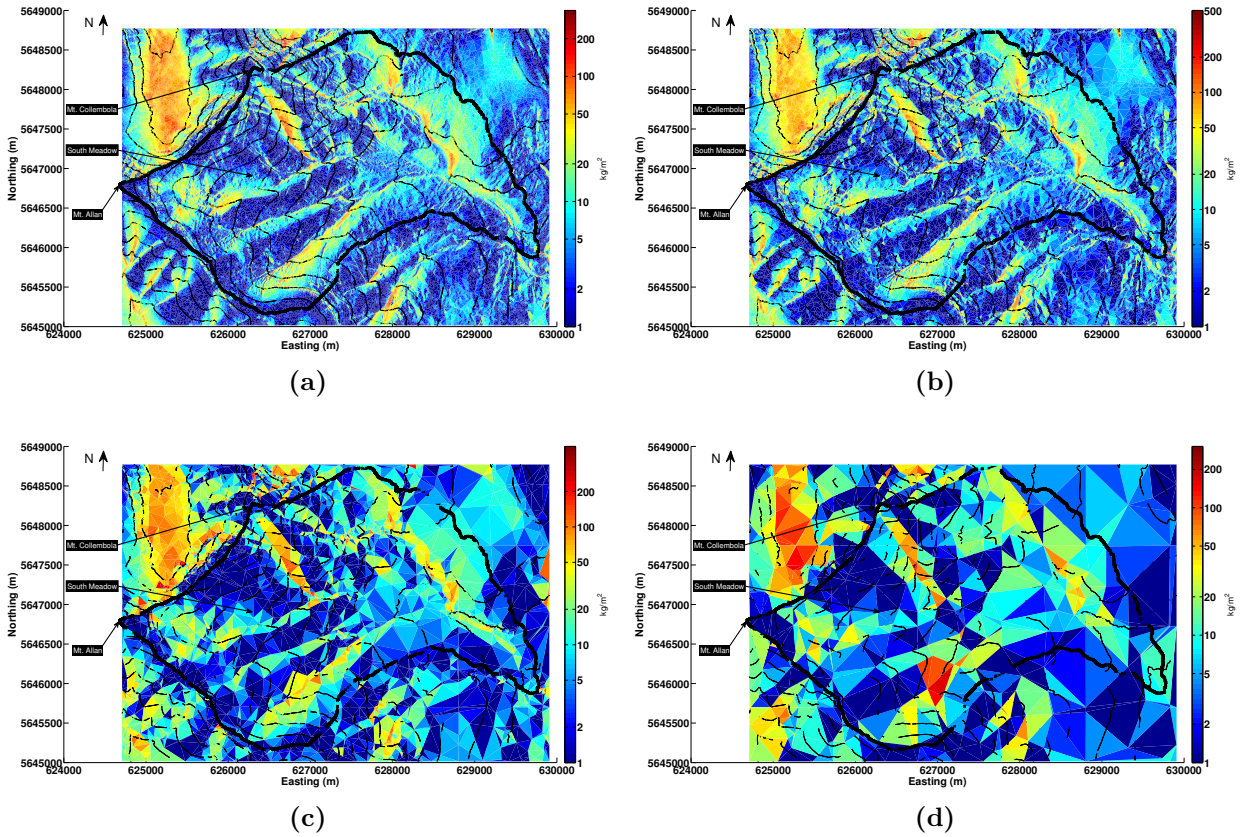

Figure 4.18: Cumulative differences of net shortwave radiation expressed as phase change mass for the MCRB for April 1, 2010–June 14, 2011 for the 1 m DEM with the four mesh tolerances; (a) 1 m, (b) 2 m, (c) 10 m, and (d) 30 m. The MCRB is in dark black, and contours are every 100 m. The larger tolerances do not produce fully connected contours and thus there are gaps.

4.6 CRHM SNOBAL modelling

4.6.1 South Meadow

For the South Meadow site, the model predicts snow accumulation beginning November 7, 2010 and full ablation by June 9, 2011 for both Q_{si} model inputs – one that is corrected for self- and horizon-shadows, and one that is only corrected for self-shadows. The predicted SWE using the two Q_{si} values demonstrates a small difference in SWE, with only a two-hour timing difference in final snowpack ablation. The modelled SWE is illustrated in Figure 4.19 where the top plot shows seasonal

accumulation and ablation, and the bottom plot shows the difference between the two modelled SWE curves. The red line (horizon-shading) is not visible in the top plot because of the limited difference between the two curves.

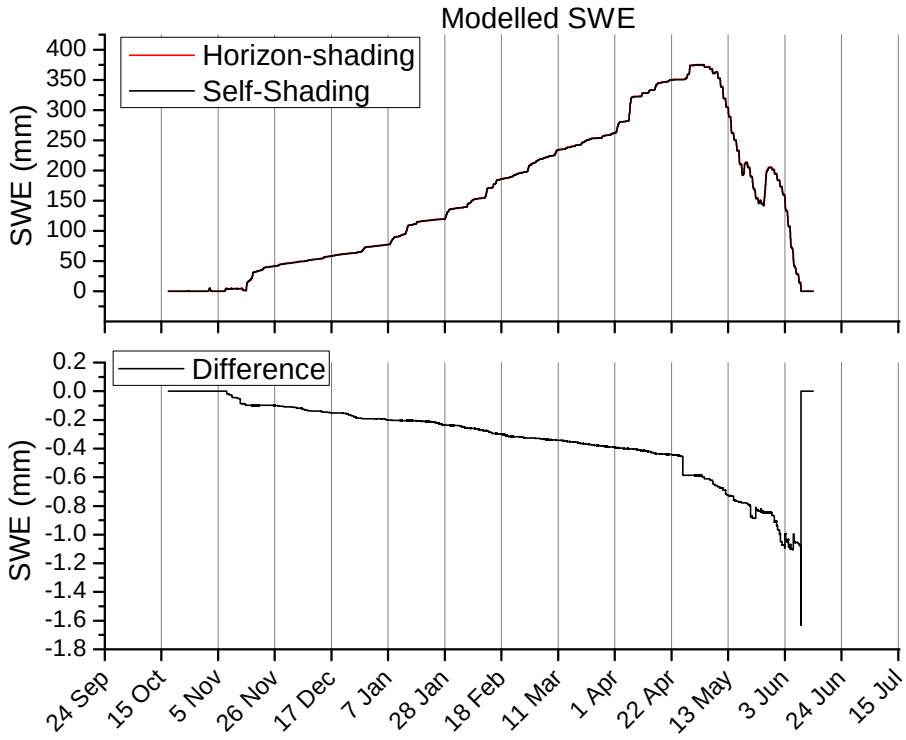

Figure 4.19: Top: Modelled SWE (mm) for October 2010–June 2011 for South Meadow using the two irradiance inputs from the self- and horizon-shadow models. Bottom: Modelled SWE (mm) difference between the two irradiance inputs from the self- and horizon-shadow models. The red line is not visible because of the limited difference between curves.

Between November 15, 2010 and April 1, 2011, there is a gradual divergence of the modelled SWE with a maximum difference of 0.3 mm. To explain this, the model's energy balance was de-constructed for this time, and the cumulative difference ($MJ \cdot m^{-2}$) for each component (net shortwave radiation [Q_{sn}], net longwave [Q_{ln}], sensible heat flux [Q_h], latent heat flux [Q_e], ground heat flux [Q_g], precipitation heat flux [Q_p], and change in internal energy storage [dU/dt]) was calculated. The winter period cumulative energy balance is shown in Figure 4.20.

The winter period energy balance, as shown in (Figure 4.21), demonstrates that there were no differences between the modelled precipitation or ground fluxes. The

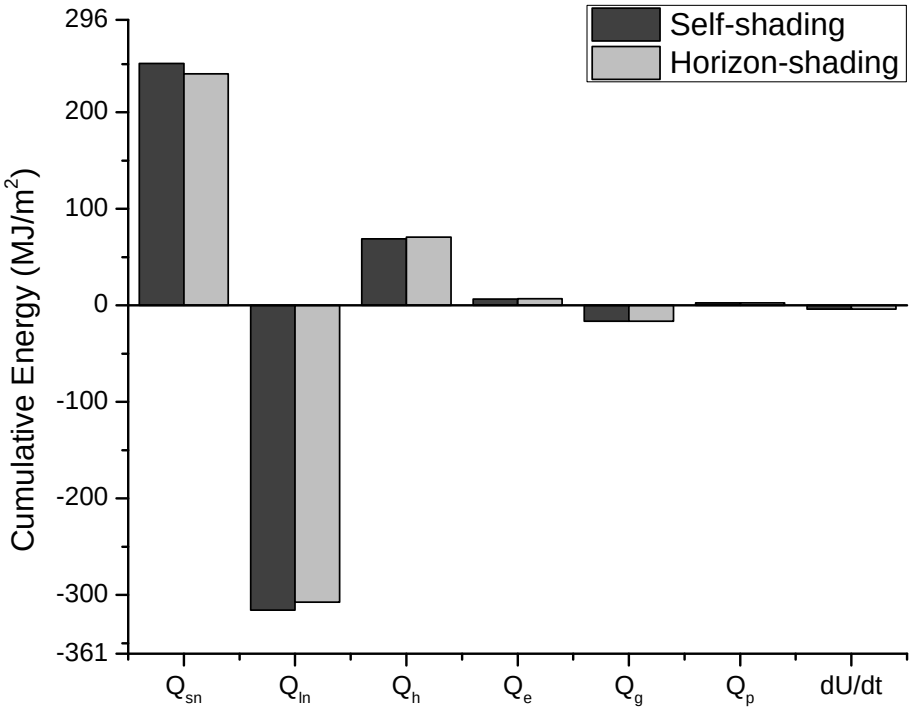

Figure 4.20: Modelled cumulative energy balance for the winter period (November 15, 2010–April 1, 2011) for South Meadow using the two irradiance inputs (self- and horizon-shading).

difference in net shortwave radiation is 10.6 MJ·m^{-2}; because the change in internal storage is small (0.06 MJ·m^{-2}) and there is no melt during this period the difference in Q_{sn} is partially balanced by increased longwave loss. The self-shaded snowpack, due to its higher surface temperature, emitted 8.05 MJ·m^{-2} more energy. Because of this, the horizon-shaded snow pack had a slightly increased sensible heat gain than the self-shaded snowpack. Finally, the difference in the latent heat flux explains the small divergence in SWE throughout the winter. The self-shaded snowpack experienced an increase in latent heat loss due to requiring energy for sublimation. The energy difference in the latent heat flux (0.7 MJ·m^{-2}) indicates that 0.3 mm SWE sublimated between November 15, 2010 and April 1, 2011, exactly matching the SWE divergence. This difference in latent heat is too small to be visible in Figure 4.21 due to the scale used.

Between April 1, 2011 and June 14, 2011, the largest difference in both modelled SWE and surface balance between the two models occurs. The difference is 1.6

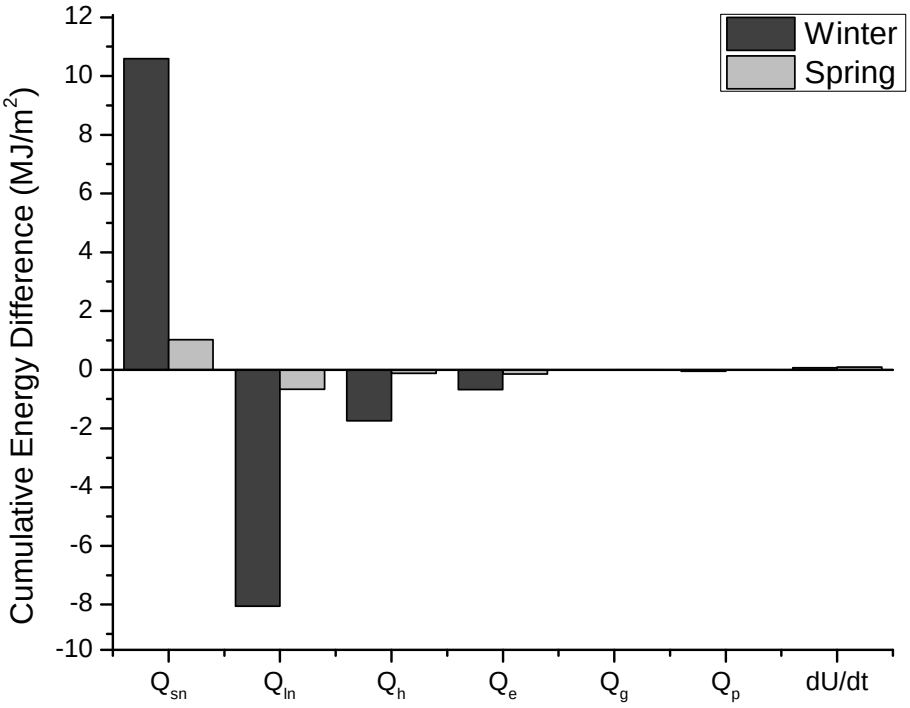

Figure 4.21: Modelled cumulative energy balance difference at South Meadow between the self- and horizon-shading models for the winter period (November 15, 2010–April 1, 2011)[dark grey] and for the spring period (April 1, 2011 – June 14, 2011) [light grey].

mm SWE occurring on June 9, 2011. That is to say, the horizon-shading model has 1.6 mm more SWE than the self-shading model, a negligible difference. The cumulative difference, for each component, between the two shading algorithms for the spring and winter periods is shown in Figure 4.21 with the winter period in dark grey and spring period in light grey. Throughout the spring period, the small difference in cumulative energy results in small SWE differences. The phase change mass calculations suggest that 2.69 kg·m^{-2} could be over-melted given the extra input energy (via Q_{sn}, thus overestimating total SWE by 1.2 mm.

4.6.2 Deep Valley

The Deep Valley site is highly shaded, and as such the first snowfall events of the year are different than that of South Meadow. This early snow-fall period is shown in Figure 4.22. Although the main accumulation period does not start until November

8, 2010, the first accumulation event on November 2, 2010 is of interest. Using the self-shading model, the snowpack is predicted to have fully ablated by November 7, 2010, whereas when using the horizon-shaded model, the snowpack persists without any melt.

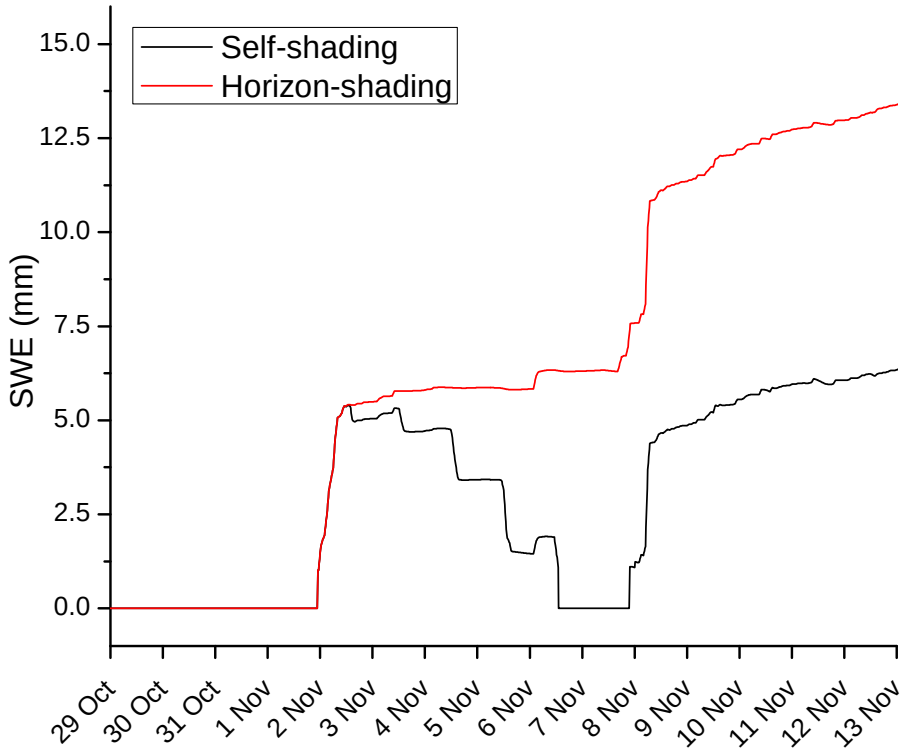

Figure 4.22: Two early snowfall events at the Deep Valley site showing the self-shading model ablating the snowpack three to five days early.

Between November 15, 2010 and April 1, 2011 there is again a gradual divergence in SWE, similar to the South Meadow site, with a maximum divergence of 10.3 mm SWE. However, 7.4 mm of this is due to the initial snowfall not ablating; thus the systematic decline of SWE is 2.8 mm. To explain this, the energy balance was deconstructed for this time period. The cumulative energy balance for each component is shown in Figure 4.23, with the cumulative difference of each component shown in Figure 4.24.

Because it is winter, there is assumed to be no melt. The cumulative net shortwave radiation difference is 59.8 MJ·m^{-2} and the cumulative difference in internal snowpack energy is 0.16 MJ·m^{-2}. The self-shaded snowpack loses much of the extra

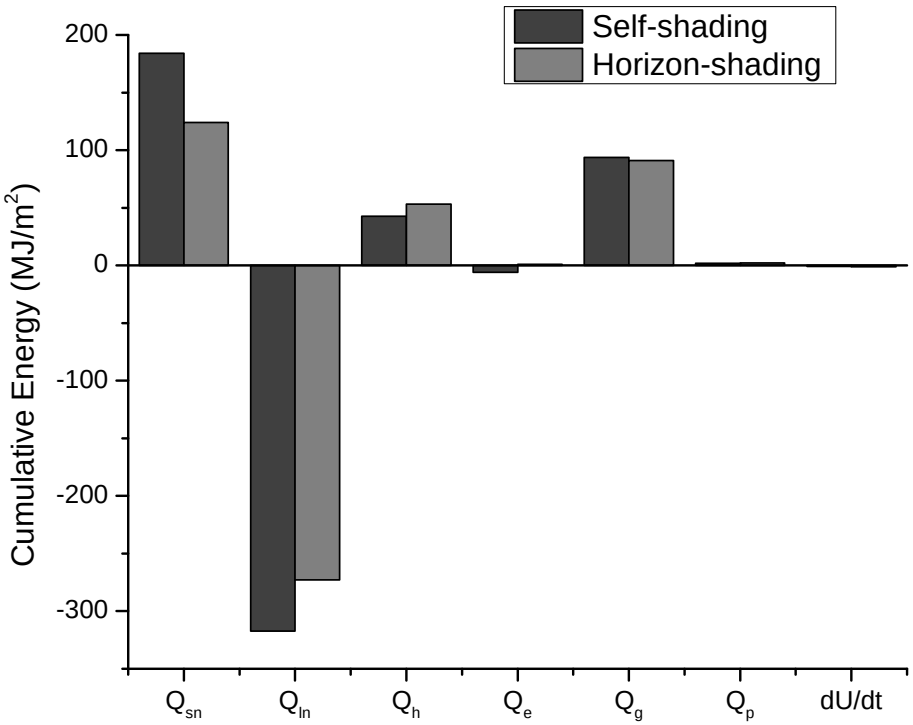

Figure 4.23: Modelled cumulative energy balance for the winter period (November 15, 2010–April 1, 2011) for Deep Valley using the two irradiance inputs (self- and horizon-shading).

energy through longwave radiative loss (44.5 MJ·m^{-2}) and through sensible heat loss (10.4 MJ·m^{-2}). The self-shaded snowpack, like the self-shaded snowpack at the South Meadow site, loses more energy to latent heat due to sublimation than the horizon-shaded snowpack. The cumulative difference for latent heat is 6.93 MJ·m^{-2}, equivalent to 3.07 mm SWE. This agrees well with the observed SWE decline throughout this period.

Unlike the South Meadow site, the Deep Valley site has different accumulation and ablation patterns, shown in Figure 4.25. The horizon-shading model has a predicted peak SWE of 380.4 mm versus the self-shading model value of 368.4 mm. Ablation begins at the same time for both models with the horizon-shading model having a consistently larger SWE.

Here, the final ablation for the self-shading model is June 8, but for the self-shading model, it is June 10. The shading in the valley site has caused a two-day ablation difference. As well, there are significant timing differences throughout

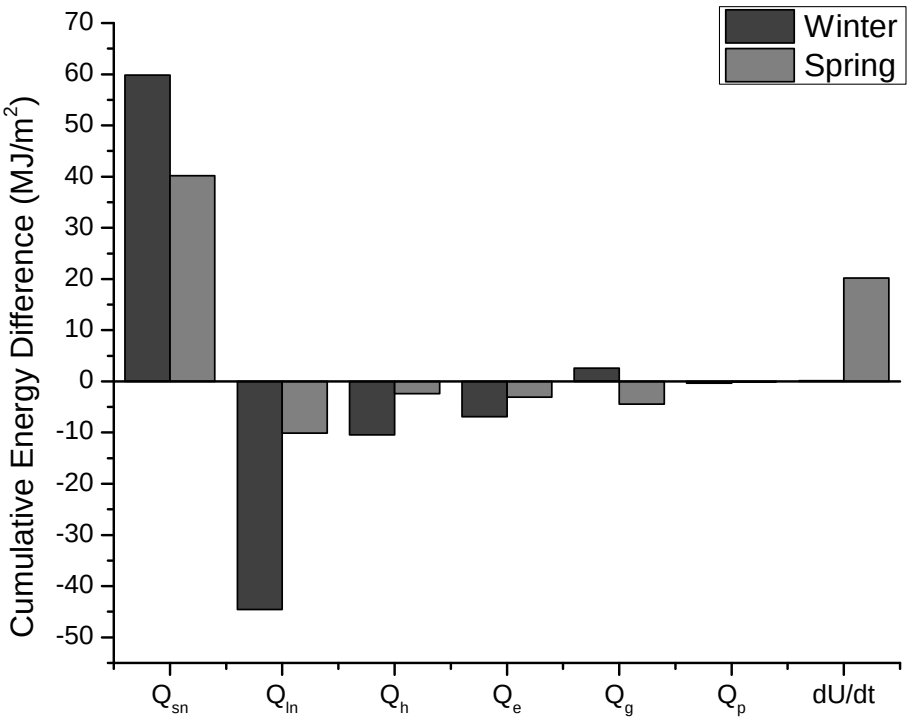

Figure 4.24: Modelled cumulative energy balance difference at Deep Valley between the self- and horizon-shading models for the winter period (November 15, 2010–April 1, 2011)[dark grey] and for the spring period (April 1, 2011 – June 14, 2011) [light grey].

the spring, leading to a peak SWE difference of 93 mm SWE. Phase change mass calculations (using modelled albedo) suggest that 114 mm SWE could melt as a result of the extra energy incident with the self-shading model for the spring. This underestimation suggests that associated feedback mechanisms associated with melt, such as the turbulent fluxes and ground flux, play an important role at this site leading to greater ablation throughout the melt period leading to the self-shading site having an increased SWE difference. It is worth noting that the reduced longwave emission and smaller shading may be important for the atmospheric energy balance, but it is outside the scope of this thesis to examine this issue in detail.

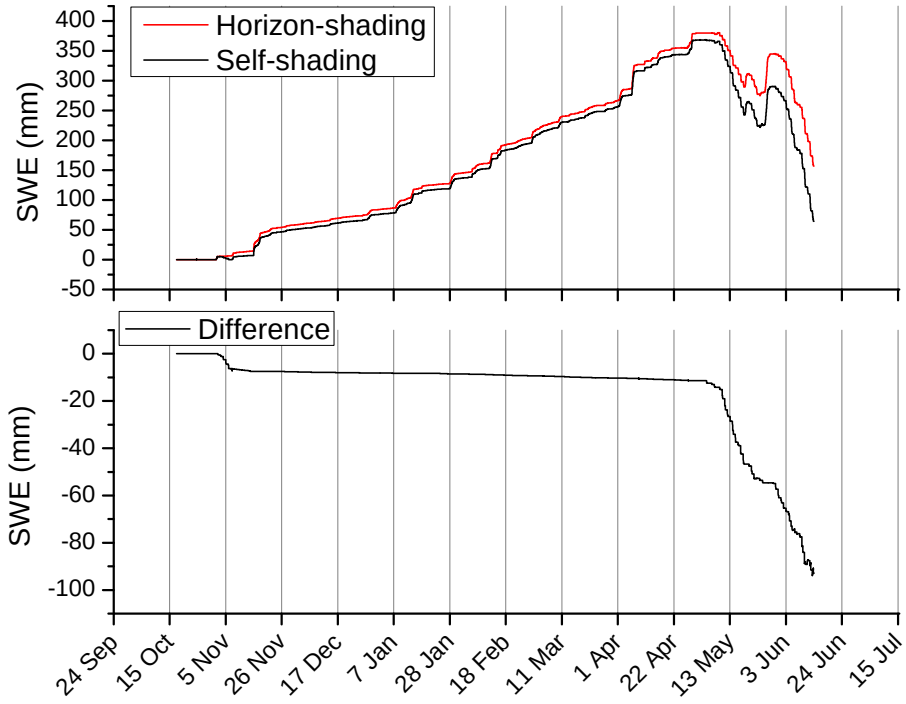

Figure 4.25: Top: Modelled SWE (mm) for October 2010–June 2011 for Deep Valley using the two irradiance inputs from the self- and horizon-shadow models. Bottom: Modelled SWE (mm) difference between the two irradiance inputs from the self- and horizon-shadow models.

Chapter 5

Conclusion

In hydrological models, a common approach to representing the landscape is through structured meshes. However, such representations introduce errors due to the rigid nature of the mesh, creating artefacts and other constraints. On the other hand, unstructured meshes may offer a number of advantages, which are explored in this thesis. An unstructured mesh was used to discretize a Canadian Rockies complex terrain site, the Marmot Creek Research Basin, using a 1 m $\times$ 1 m LiDAR-derived DEM as input. In order to capture the horizon-shadows of surrounding topographic features, the parallel point-plane method used by Clarke (2001) and expanded upon by Montero et al. (2009) was implemented to model shadows in complex terrain. A new coordinate system, aligned with the solar position, was created by rotating the standard coordinate system using Euler angles. Self- and horizon-shading shadow predictions were compared at a point scale to measured irradiance observations at the South Meadow site in Marmot Creek Research Basin on the clear-sky day February 1, 2011. It was found that the shadow passage over the South Meadow site at 16:30 was correctly modelled by the horizon-shadow algorithm. The shading effects of surrounding mountain topography on irradiance in Marmot Creek were substantial. Overall, there was a systematic overprediction of irradiance when not considering horizon-shadows.

The model output from the shadow model was compared to observed data using orthorectified timelapse photography from February 1, 2011. These photos allowed for a quantitative comparison with observed shadow locations. Observed shadow locations were extracted from the timelapse photos using supervised classification within ArcGIS and post-processed using Matlab. It was found that the shadow

model performed well, closely matching the observed shadow locations. However, some locations were over represented in the model due to the binary shading algorithm used. Although some of the large triangles were only partially shaded, due to the lack of sub-triangle representation within the model, the entire triangle had to be marked as fully shaded. The output from the shadow model was compared to that generated using a structured mesh (1 m $\times$ 1 m) generated using existing algorithms via ArcGIS. The output from ArcGIS was a closer match to the observed shadow locations; however it was using an higher resolution mesh, and thus it was expected it would perform well. The error in the orthorectified imagery was quantified by comparing distances between points on an airphoto and the orthorectified image and spatially interpolating these differences for the entire orthorectified image extent. The maximum and minimum differences were found to be 24.7 m and 4.74 m respectively with an average of 14.5 m (standard deviation 3.46 m). Repeated tests suggests that there is an approximate error of 5% in selecting corresponding points on each image.

Modelled over the season from October to June, the difference in irradiance between self- and horizon-shading resulted in a total cumulative difference of 52 MJ$\cdot$m^{-2} and, in the snowmelt period of April to June, 4.5 MJ$\cdot$m^{-2}. This corresponded to a difference equivalent to a phase change mass of 2.69 kg$\cdot$m^{-2} if all of the net shortwave radiation difference were applied to phase change as snowmelt, assuming a continuous snowcover and an albedo of 0.8. At the basin scale for the snowmelt period of April–June, the mean triangle potential snowmelt mass difference was 14.4 kg$\cdot$m^{-2}, under the same assumptions. However, these results come with a caveat: the MCRB is dominated by tree-cover, and as such, the sub-canopy irradiance is already significantly reduced. Ellis and Pomeroy (2007) demonstrated that mean daily canopy transmissivity values were as low as 0.09 on north-facing slopes and 0.36 on south-facing slopes. Because this study assumes a bare-earth, some of the locations may be impacted less by ignoring the horizon-shadows. However, Ellis and Pomeroy (2007) indicates the forest-cover on south-facing slopes had the greatest variance. Thus the south-face slopes dominated by horizon-shadows may still have significant

impacts on the sub-canopy energy balance. When the phase change mass is compared to pre-melt SWE values from 2007 following DeBeer and Pomeroy (2009) for the Mt. Allan cirque, this mean cumulative difference is 6.0% and 11.4% for north- and south- facing slopes, respectively, of total pre-melt SWE. The cumulative difference is 66.4% of the total pre-melt SWE for the shallow south-facing slope. At the basin scale for the spring of April to June, 2011, the peak cumulative difference was 1600 kg·m^{-2} despite the fact that 99% of the basin had a cumulative difference of 500 kg·m^{-2} or less. The deep valleys are chronically over predicted by self-shading methods, resulting in snowpacks that ablate too early in the spring. However, these differences are dependent upon the geometry of the river basin, and further studies should examine multiple basins to determine the importance in different terrains.

The effects of the spatial resolution of the DEM and the tolerance of the unstructured mesh on irradiance and shadow prediction were examined. The model was run at three point locations (Fisera Ridge, Fisera Ridge North, and South Meadow) for February 1, 2011. The 1 m $\times$ 1 m spatial resolution LiDAR-derived DEM was downsampled to produce a 2 m $\times$ 2 m, a 10 m $\times$ 10 m, and a 30 m $\times$ 30 m DEM. Unstructured meshes with vertical tolerances of 1 m, 2 m, 10 m, and 30 m were created from each of these DEMs. This generated a total of sixteen combinations of resolution and tolerance. It was found that smaller tolerances, regardless of the DEM resolution, had better predictive capacity than a higher-resolution DEM with large tolerances (large triangles). This was partially attributed to the triangle-triangle intersection algorithm. Because the large tolerances produced large triangles, it was possible for triangles to be shaded too aggressively, as only a small portion of the triangle was in shadows. Further research should focus on improving the triangle-triangle detection as well as sub-triangle variability to enhance the shadow prediction. An ability to capture some sub-triangle variability would help to mitigate this problem.

The CRHM SNOBAL model was run from within the CRHM modelling environment for October, 2010 to June, 2011 for the South Meadow and the Deep Valley sites. It was found that at South Meadow there was a 0.3 mm SWE loss to sublimation throughout the winter and a maximum difference in SWE of 1.6 mm at the

end of the melt period. At the Deep Valley site, the first snowfall of the season (7.4 mm) fully ablates with the self-shading model but remains with the horizon-shaded model. Throughout the winter, there is a 2.8 mm SWE loss due to sublimation in the self-shaded snowpack. Although final ablation is not modelled due to insufficient model input, extrapolation suggests a final ablation timing difference of 2–4 days. Although snowmelt was found to be influenced by these topographic shadows, the MCRB as a whole is a poor basin for fully exploring the impact of neglecting horizon-shadows. This is due to the low, less rugged topography present in this basin. Because most of the MCRB is forested, the generally small overprediction of irradiance due to neglecting horizon-shadows may be overstated, due to the low transmissivity to irradiance of the forest canopies. Further work should examine basins that are characterized by steep sides, much like the Deep Valley site. However, these results do demonstrate that a highly shaded basin could be significantly impacted. Incorrect snowmelt timing throughout a variety of headwater basins could have a potentially large cumulative effect when predicting downstream processes such as floods.

As the triangle tolerance is increased, single large triangles begin to represent large portions of the basin and these triangles are treated as hydrologically heterogeneous. There are clear parallels between this and the HRU conceptualization. Of interest is where these two landscape representations, unstructured meshes and HRUs, begin to converge. Although it was found that low-tolerance unstructured meshes performed best at a point scale, the overall spatial patterns of the high tolerance meshes suggests that perhaps they many be sufficient for large-scale modelling. If a form of sub-triangle representation was included, high tolerance unstructured meshes may be able to complement or replace the HRU approach by providing a landscape representation that captures spatial variability while still requiring less parametrisation than high spatial resolution (e.g., 1 m $\times$ 1 m) models. By exploiting fully nested unstructured meshes, where a triangle may be composed of sub-triangles, multi-scale processes may be more fully captured, allowing for a merging of an HRU approach with that of a distributed model.

References

Annandale, J., N. Jovanovic, N. Benadé, and R. Allen (2002), Software for missing data error analysis of Penman-Monteith reference evapotranspiration, *Irrigation Science*, *21*(2), 57–67, doi:10.1007/s002710100047.

Bern, M., and D. Eppstein (1992), Mesh generation and optimal triangulation, *Computing in Euclidean geometry*, *1*, 23–90.

Berthold, S. (2004), Integrated hydrogeological and geophysical study of depression-focused groundwater recharge in the Canadian prairies, *Water Resources Research*, *40*(6), doi:10.1029/2003WR002982.

Braun, J., and M. Sambridge (1997), Modelling landscape evolution on geological time scales: a new method based on irregular spatial discretization, *Basin Research*, *9*(1), 27–52.

Carey, S., and M.-k. Woo (1998), Snowmelt hydrology of two subarctic slopes, southern Yukon, Canada, *Nordic Hydrology*, *29*(4), 331.

Chang, K.-t. (2008), *Introduction to Geographic Information Systems*, 450 pp., McGraw-Hill, New York, New York, USA.

Chew, L. (1993), Guaranteed-quality mesh generation for curved surfaces, in *Proceedings of the ninth annual symposium on Computational geometry*, edited by C. Yap, pp. 274–280, ACM, San Diego, CA, USA.

Chueca, J., and A. Julián (2004), Relationship between solar radiation and the development and morphology of small cirque glaciers (Maladeta Mountain massif, Central Pyrenees, Spain), *Geografiska Annaler: Series A, Physical Geography*, *86*(1), 81–89.

Clarke, J. A. (2001), *Energy Simulation in Building Design*, 2nd ed., 362 pp., Butterworth-Heinemann, Oxford, UK.

Clarke, R. (1973), A review of some mathematical models used in hydrology, with observations on their calibration and use, *Journal of hydrology*, *19*, 1–20, doi: 10.1016/0022-1694(73)90089-9.

Coquillart, S., and M. Gangnet (1984), Shaded display of digital maps, *IEEE Computer Graphics and Applications*, *4*, 35–42.

Corripio, J. G. (2004), Snow surface albedo estimation using terrestrial photography, *International Journal of Remote Sensing*, *25*(24), 5705–5729, doi:10.1080/01431160410001709002.

Cranmer, A., N. Kouwen, and S. Mousavi (2001), Proving WATFLOOD: modelling the nonlinearities of hydrologic response to storm intensities, *Canadian Journal of Civil Engineering*, *28*, 837–855, doi:10.1139/cjce-28-5-837.

Davies, T. D., P. Brimblecombe, M. Tranter, S. Tsiouris, C. E. Vincent, P. Abrahams, and I. L. Blackwood (1987), The removal of soluble ions from melting snowpacks, in *Seasonal Snowcovers: Physics, Chemistry, Hydrology*, edited by H. G. Jones and W. J. Orville-Thomas, pp. 337–392, D. Reidel Publishing Company, Dordrecht, Holland.

Davison, B., S. Pohl, P. Dornes, P. Marsh, A. Pietroniro, and M. MacKay (2006), Characterizing snowmelt variability in a land-surface-hydrologic model, *Atmosphere-Ocean*, *44*(3), 271–287.

DeBeer, C. M., and J. Pomeroy (2010), Simulation of the snowmelt runoff contributing area in a small alpine basin, *Hydrology and Earth System Sciences*, *14*(7), 1205–1219.

DeBeer, C. M., and J. W. Pomeroy (2009), Modelling snow melt and snowcover depletion in a small alpine cirque , Canadian Rocky Mountains, *Hydrological Processes*, *23*(18), 2584– 2599.

Devillers, O., and P. Guigue (2002), Faster triangle-triangle intersection tests, *Tech. rep.*, Institut National de Recherche en Informatique et en Automatique, Sophia-Antipolis, France.

Dornes, P. (2009), An approach for modelling snowcover ablation and snowmelt runoff in cold region environments, Doctor of philosophy, University of Saskatchewan.

Dornes, P. F., J. W. Pomeroy, A. Pietroniro, S. K. Carey, and W. L. Quinton (2008a), Influence of landscape aggregation in modelling snow-cover ablation and snowmelt runoff in a sub-arctic mountainous environment, *Hydrological Sciences Journal*, *53*(4), 725–740.

Dornes, P. F., J. W. Pomeroy, A. Pietroniro, and D. L. Verseghy (2008b), Effects of spatial aggregation of initial conditions and forcing data on modeling snowmelt using a land surface scheme, *Journal of Hydrometeorology*, *9*(4), 789–803, doi: 10.1175/2007JHM958.1.

Dozier, J., and J. Frew (1990), Rapid calculation of terrain parameters for radiation modeling from digital elevation data, *IEEE Transactions on Geoscience and Remote*, *28*(5), 963–969.

Dozier, J., J. Bruno, and P. Downey (1981), Faster solution to the horizon problem, *Computers and Geosciences*, *7*(2), 145–151.

Dymond, C. (2002), Mapping vegetation spatial patterns from modeled water, temperature and solar radiation gradients, *ISPRS Journal of Photogrammetry and Remote Sensing*, *57*(1-2), 69–85, doi:10.1016/S0924-2716(02)00110-7.

El-Shimy, N., C. Valeo, and A. Habib (2005), *Digital terrain modeling: Aquisition, manipulation, and applications*, 257 pp., Artech House, Inc., Norwood, MA.

Elder, K., J. Dozier, and J. Michaelsen (1991), Snow accumulation and distribution in an alpine watershed, *Water Resources Research*, *27*(7), 1541–1552.

Ellis, C., and J. W. Pomeroy (2007), Estimating sub-canopy shortwave irradiance to melting snow on forested slopes, *Hydrological Processes*, *21*(19), 2581– 2593.

Essery, R. (2004), Statistical representation of mountain shading, *Hydrology and Earth System Sciences*, *8*(6), 1045–1050, doi:10.5194/hess-8-1047-2004.

Essery, R., and D. Marks (2007), Scaling and parametrization of clear-sky solar radiation over complex topography, *Journal of Geophysical Research*, *112*, 1–12, doi:10.1029/2006JD007650.

Essery, R., M. Best, and P. Cox (2001), MOSES 2.2 Technical Documentation, *Tech. Rep. August*, Hadley Centre, Berks, UK.

Flüegel, W. (1995), Hydrological response units (HRUs) to preserve basin heterogeneity in hydrological modelling using PRMS/MMS – case study in the Bröl basin, Germany, in *Modelling and Management of Sustainable Basin-scale Water Resource Systems (Proceedings of a Boulder Symposium, Jul 1995y)*, vol. 231, pp. 79–88, IAHS Publ. no. 231.

Fowles, G. R., and G. L. Cassiday (2005), *Analytical Mechanics*, 7th ed., 544 pp., Thomson Brooks/Cole, Belmont, USA.

Frew, J. E. (1990), The image processing workbench, Doctor of philosophy, Dept. Geography, University of California Santa Barbara.

Fu, P., and P. M. Rich (1999), Design and implementation of the Solar Analyst: an ArcView extension for modeling solar radiation at landscape scales, in *Proceedings of the 19th Annual ESRI User Conference*, pp. 1–33, San Diego, USA.

Garnier, B. J., and A. Ohmura (1968), A method of calculating the direct shortwave radiation income of slopes, *Journal of Applied Meteorology*, *7*(5), 796–800, doi: 10.1175/1520-0450(1968)007<0796:AMOCTD>2.0.CO;2.

Golding, D. L. (1970), Research results from Marmot Creek experimental watershed, Alberta, Canada, in *Symposium on the results of research on representative and experimental basins*, pp. 397–404, IASH-UNESCO, Wellington, New Zealand.

Gray, D. M., and D. H. Male (1981), *Handbook of Snow: Principles, Processes, Management, and Use*, 776 pp., Pergamon Press, Toronto, Canada.

Gray, D. M., B. Toth, L. Zhao, J. W. Pomeroy, and R. J. Granger (2001), Estimating areal snowmelt infiltration into frozen soils, *Hydrological Processes*, *15*(16), 3095–3111.

Grayson, R., and G. Blöschl (2000), *Spatial patterns in catchment hydrology: observations and modelling*, 408 pp., Cambridge University Press.

Groisman, P., and T. Davies (2001), Snow cover and the climate system, in *Snow Ecology: an Interdisciplinary Examination of Snow-covered Ecosystems*, edited by H. Jones, J. Pomeroy, D. Walker, and R. Hoham, chap. 1, pp. 1–44, Cambridge University Press, Cambridge, UK.

Haltiner, G. J., and F. L. Martin (1957), *Dynamical and physical meteorology*, 470 pp., McGraw-Hill, New York.

Helgason, W., and J. W. Pomeroy (2012), Characteristics of the near surface boundary layer within a mountain valley during winter, *Journal of Applied Meteorology and Climatology*, *51*(3), 583–597, doi:10.1175/JAMC-D-11-058.1.

Hetrick, W., P. Rich, and S. Weiss. (1993), Modeling insolation on complex surfaces, in *Thirteenth Annual ESRI User Conference*, vol. 2, pp. 447–458.

Hopkinson, C., L. Chasmer, S. Munro, and M. Demuth (2010), The influence of DEM resolution on simulated solar radiation-induced glacier melt, *Hydrological Processes*, *24*(6), 775–788, doi:10.1002/hyp.

Hopkinson, C., J. Pomeroy, C. DeBeer, C. Ellis, and A. Anderson (2011), Relationships between snowpack depth and primary LiDAR point cloud derivatives in a mountainous environment, in *Remote Sensing and Hydrology 2010*, IAHS Publ. 3XX, Jackson Hole, Wyoming, USA.

Ivanov, V., E. Vivoni, R. Bras, and D. Entekhabi (2004), Preserving high-resolution surface and rainfall data in operational-scale basin hydrology: a fully-distributed physically-based approach, *Journal of Hydrology*, *298*(1-4), 80–111, doi:10.1016/j.jhydrol.2004.03.041.

Jones, C. B., D. B. Kidner, and J. M. Ware (1994), The implicit triangulated irregular network and multiscale spatial databases, *The Computer Journal*, *37*(1), 43–57.

Klemeš, V. (1986), Operational testing of hydrological simulation models, *Hydrological Sciences Journal*, *31*(1), 13–24, doi:10.1080/02626668609491024.

Kouwen, N., E. Soulis, A. Pietroniro, J. Donald, and R. Harrington (1993), Grouped response units for distributed hydrologic modeling, *Journal of Water Resources Planning and Management*, *119*(3), 289–305.

Kumar, M., G. Bhatt, and C. Duffy (2009), An efficient domain decomposition framework for accurate representation of geodata in distributed hydrologic models, *International Journal of Geographical Information Science*, *23*(12), 1569–1596.

Leavesley, G., and L. Stannard (1995), The precipitation-runoff modelling system–PRMS, in *Computer Models of Watershed Hydrology*, edited by V. Singh, Water Resources Publications, Highlands Rand, CO.

Lee, J. (1991), Comparison of existing methods for building triangular irregular network, models of terrain from grid digital elevation models, *International Journal of Geographical Information Science*, *5*(3), 267–285.

Li, L., and S. P. Simonovic (2002), System dynamics model for predicting floods from snowmelt in North American prairie watersheds, *Hydrological Processes*, *16*, 2645–2666, doi:10.1002/hyp.1064.

List, R. J. (1971), *Smithsonian Meteorological Tables*, 350 pp., Smithsonian Institution Press, Washington, DC.

MacDonald, M. K., J. W. Pomeroy, and A. Pietroniro (2010), On the importance of sublimation to an alpine snow mass balance in the Canadian Rocky Mountains, *Hydrology and Earth System Sciences*, *14*(7), 1401–1415.

Male, D. H., and R. J. Granger (1981), Snow surface-energy exchange, *Water Resources Research*, *17*(3), 609–627, doi:10.1029/WR007i001p00184.

Marks, D., and J. Dozier (1992), Climate and energy exchange at the snow surface in the alpine region of the Sierra Nevada, 2, Snow cover energy balance, *Water Resources Research*, *28*(11), 3043–3054.

Marks, D., and A. Winstral (2001), Comparison of snow deposition, the snow cover energy balance, and snowmelt at two sites in a semiarid mountain basin, *Journal of Hydrometeorology*, *2*(3), 213–227, doi:10.1175/1525-7541(2001)002<0213:COSDTS>2.0.CO;2.

Marks, D., J. Kimball, D. Tingey, and T. Link (1998), The sensitivity of snowmelt processes to climate conditions and forest cover during rain-on-snow: a case study of the 1996 Pacific Northwest flood, *Hydrological Processes*, *12*(March), 1587–1587.

Marsh, P., and M.-K. Woo (1984), Wetting front advance and freezing of meltwater within a snow cover: 1. Observations in the Canadian Arctic, *Water Resources Research*, *20*(12), 1853–1864.

Marshak, A., and A. B. Davis (2005), *3D Radiative Transfer in Cloudy Atmospheres*, 686 pp., Springer-Verlag, Berlin, Germany.

McDonnell, J. J., et al. (2007), Moving beyond heterogeneity and process complexity: a new vision for watershed hydrology, *Water Resources Research*, *43*(7), 1–6, doi:10.1029/2006WR005467.

Mckenney, D. W. (1999), Calibration and sensitivity analysis of a spatially-distributed solar radiation model, *International Journal of Geographical Information Science*, *13*(1), 49–65, doi:10.1080/136588199241454.

Montero, G., J. Escobar, E. Rodriguez, R. Montenegro, and E. Rodríguez (2009), Solar radiation and shadow modelling with adaptive triangular meshes, *Solar Energy*, *83*(7), 998–1012, doi:10.1016/j.solener.2009.01.004.

Moore, R. D. (1983), A comparison of the snowmelt energy budgets in two alpine basins, *Archives for Meteorology, Geophysics, and Bioclimatology Series B*, *33*(1-2), 1–10, doi:10.1007/BF02273986.

Munro, D. S., and G. J. Young (1982), An operational net shortwave radiation model for glacier basins, *Water Resources Research*, *18*(2), 220–230, doi:10.1029/WR018i002p00220.

Myers, D. (1994), Spatial interpolation: an overview, *Geoderma*, *62*(1-3), 17–28, doi:10.1016/0016-7061(94)90025-6.

Nelson, E., N. Jones, ASCE, and R. Berrett (1999), Adaptive tessellation method for creating TINS from GIS data, *Journal of Hydrologic Engineering*, *4*(1), 2–9.

O'Callaghan, J. F., and D. Marks (1984), The extraction of drainage networks from digital elevation data, *Computer Vision, Graphics, and Image Processing*, *28*(3), 323–344.

Oke, T. (1987), *Boundary layer climates*, second ed., 435 pp., University Press, Cambridge, UK.

O'Neill, A., and D. Gray (1973), Solar radiation penetration through snow, in *UNESCO-WMO-IAHS Symposia on the Role of Snow and Ice in Hydrology*, pp. 227–249, IAHS Press, Wallingford, UK.

Petty, G. W. (2006), *A First Course in Atmospheric Radiation*, 2nd ed., 472 pp., Sundog Publishing, Madison, Wisconsin, USA.

Pietroniro, A., et al. (2007), Development of the MESH modelling system for hydrological ensemble forecasting of the Laurentian Great Lakes at the regional scale, *Hydrology and Earth System Sciences Discussions*, *11*(4), 1279–1294.

Pohl, S., B. Davison, P. Marsh, and A. Pietroniro (2005), Modelling spatially distributed snowmelt and meltwater runoff in a small Arctic catchment with a hydrology land-surface scheme (WATCLASS), *Atmosphere-Ocean*, *43*(3), 193–211, doi:10.3137/ao.430301.

Pohl, S., P. Marsh, and A. Pietroniro (2006), Spatial-temporal variability in solar radiation during spring snowmelt, *Nordic Hydrology*, *37*(1), 1–19, doi:10.2166/nh.2005.026.

Pomeroy, J. W., D. Gray, and P. Landine (1993), The prairie blowing snow model: characteristics, validation, operation, *Journal of Hydrology*, *144*(1-4), 165–192.

Pomeroy, J. W., D. M. Gray, K. R. Shook, B. Toth, R. L. H. Essery, A. Pietroniro, and N. Hedstrom (1998), An evaluation of snow accumulation and ablation processes for land surface modelling, *Hydrological Processes*, *12*(15), 2339–2367, doi:10.1002/(SICI)1099-1085(199812)12:15<2339::AID-HYP800>3.0.CO;2-L.

Pomeroy, J. W., B. Toth, R. J. Granger, N. R. Hedstrom, and R. L. H. Essery (2003), Variation in surface energetics during snowmelt in a subarctic mountain catchment, *Journal of Hydrometeorology*, *4*(4), 702–719, doi:10.1175/1525-7541(2003)004<0702:VISEDS>2.0.CO;2.

Pomeroy, J. W., D. M. Gray, T. Brown, N. R. Hedstrom, W. L. Quinton, R. J. Granger, and S. K. Carey (2007), The cold regions hydrological model: a platform for basing process representation and model structure on physical evidence, *Hydrological Processes*, *21*(19), 2650–2667, doi:10.1002/hyp.

Rasmussen, L., H. Conway, and S. Ferguson (1998), Estimation of atmospheric transmittance from upper-air humidity, *Solar energy*, *62*(5), 359–368.

Rigon, R., G. Bertoldi, and T. M. Over (2006), GEOtop: a Distributed Hydrological Model with coupled water and energy budgets, *Journal of Hydrometeorology*, *7*(3), 371–388, doi:10.1175/JHM497.1.

Rivara, M.-C. (1987), A grid generator based on 4-triangles conforming mesh-refinement algorithms, *International Journal for Numerical Methods in Engineering*, *24*(7), 1343–1354, doi:10.1002/nme.1620240710.

Romero, L., S. Tabik, J. Vias, and E. Zapata (2008), Fast clear-sky solar irradiation computation for very large digital elevation models, *Computer Physics Communications*, *178*(11), 800–808, doi:10.1016/j.cpc.2008.01.048.

Ruppert, J. (1995), A Delaunay refinement algorithm for quality 2-dimensional mesh generation, *Journal of Algorithms*, *18*(3), 548–585, doi:10.1006/jagm.1995.1021.

Ryerson, C. (1984), Mapping solar access in high relief areas for regional energy planning, *The Professional Geographer*, *36*(3), 345–352.

SAGA Development Team (saga-gis.org) (2007), System for automated geoscientific analyses (SAGA GIS).

Sanderson, C. (2010), Armadillo : an open source C ++ linear algebra library for fast prototyping and computationally intensive experiments technical report, *Tech. rep.*, NICTA, St Lucia, Australia.

Shaw, D. A., L. W. Martz, and A. Pietroniro (2005), A methodology for preserving channel flow networks and connectivity patterns in large-scale distributed hydrological models, *Hydrological Processes*, *19*(1), 149–168, doi:10.1002/hyp.5765.

Shewchuk, J. (1996), Triangle: engineering a 2D quality mesh generator and Delaunay triangulator, in *Applied Computational Geometry Towards Geometric Engineering*, edited by G. Goos, J. Hartmanis, and J. van Leeuwen, pp. 203–222, Springer Berlin / Heidelberg, doi:10.1007/BFb0014474.

Shewchuk, J. (2002), Delaunay refinement algorithms for triangular mesh generation, *Computational Geometry*, *22*(1-3), 21–74, doi:10.1016/S0925-7721(01)00047-5.

Shook, K., and D. M. Gray (1996), Small-scale spatial structure of shallow snowcovers, *Hydrological Processes*, *10*, 1283–1292.

Sicart, J. E., J. W. Pomeroy, R. L. H. Essery, D. Bewley, and D. Dewley (2006), Incoming longwave radiation to melting snow : observations, sensitivity and estimation in northern environments, *Hydrological Processes*, *20*(17), 3697– 3708, doi:10.1002/hyp.

Sivapalan, M. (2003), Process complexity at hillslope scale, process simplicity at the watershed scale: Is there a connection?, in *EGS-AGU-EUG Joint Assembly*, vol. 17, p. 7973, Nice, France.

Slingsby, A. (2002), An object-orientated approach to hydrological modeling using triangular irregular networks, Masters of science, University of Edinburgh.

Soulis, E., N. Kouwen, A. Pietroniro, F. Seglenieks, K. Snelgrove, P. Pellerin, D. Shaw, and L. Martz (2005), A framework for hydrological modelling in MAGS, in *Prediction in Ungauged Basins: Approaches for Canada's Cold Regions*, edited by C. Spence, A. Pietroniro, and J. Pomeroy, chap. 7, pp. 119–138, Canadian Water Resources Association.

Stewart, I. (2009), Changes in snowpack and snowmelt runoff for key mountain regions, *Hydrological Processes*, *23*(1), 78–94, doi:10.1002/hyp.

Stewart, J. (1998), Fast horizon computation at all points of a terrain with visibility and shading applications, *IEEE Transactions on Visualization and Computer Graphics*, *4*(1), 82–93, doi:10.1109/2945.675656.

Storr, D. (1967), Precipitation variations in a small forested watershed, in *Proc 35th Annual Western Snow Conference 18-20 April*, pp. 11–17, Boise, Idaho.

Tabik, S., L. F. Romero, and E. L. Zapata (2011), High-performance three-horizon composition algorithm for large-scale terrains, *International Journal of Geographical Information Science*, *25*(4), 541–555, doi:10.1080/13658810903149995.

Tachikawa, Y., M. Shiiba, and T. Takasao (1994), Development of a basin geomorphic information system using a TIN-DEM data structure, *Journal of the*

American Water Resources Association, *30*(1), 9–17, doi:10.1111/j.1752-1688.1994.tb03268.x.

Tucker, G. (2001), An object-oriented framework for distributed hydrologic and geomorphic modeling using triangulated irregular networks, *Computers & Geosciences*, *27*(8), 959–973, doi:10.1016/S0098-3004(00)00134-5.

Varley, M. J., K. J. Beven, and H. R. Oliver (1996), Modelling solar radiation in steeply sloping terrain, *International Journal of Climatology*, *16*(1), 93–104, doi:10.1002/(SICI)1097-0088(199601)16:1<93::AID-JOC992>3.3.CO;2-K.

Vivoni, E. R., V. Y. Ivanov, R. L. Bras, and D. Entekhabi (2005), On the effects of triangulated terrain resolution on distributed hydrologic model response, *Hydrological Processes*, *19*, 2101–2122, doi:10.1002/hyp.5671.

Walcek, C. (1994), Cloud cover and its relationship to relative humidity during a springtime midlatitude cyclone, *Monthly Weather Review*, *122*(6), 1021–1035.

Wallace, J. M., and P. V. Hobbs (2006), *Atmospheric Science: An Introductory Survey*, 2nd ed., Elsevier Inc, Burlington, MA.

Walter, T. M., E. Brooks, D. McCool, L. King, M. Molnau, and J. Boll (2005), Process-based snowmelt modeling: does it require more input data than temperature-index modeling?, *Journal of Hydrology*, *300*(1-4), 65–75, doi:10.1016/j.jhydrol.2004.05.002.

Willson, R. C. (2003), Secular total solar irradiance trend during solar cycles 21–23, *Geophysical Research Letters*, *30*(5), doi:10.1029/2002GL016038.

Woo, M., W. R. Rouse, R. E. Stewart, and J. M. R. Stone (2008), The Mackenzie GEWEX study: a contribution to cold region atmospheric and hydrologic sciences, in *Cold Region Atmospheric and Hydrologic Studies Volume 1: Atmospheric Dynamics*, edited by M.-k. Woo, chap. 1, pp. 1–22, Springer-Verlag, Berlin, Germany.

Woo, M.-k., and K. Young (2004), Modeling arctic snow distribution and melt at the 1 km grid scale., *Nordic Hydrology*, *35*(August 2003), 295–307.

World Meteorological Organization (1970), *Guide to hydrometeorological practices*, wmo-no. 16 ed., 168 pp., Geneva, Switzerland.

Young, A. (1994), Air mass and refraction, *Applied Optics*, *33*(6), 1108–1110.

Zhou, Q., P. Pilesjö, and Y. Chen (2011), Estimating surface flow paths on a digital elevation model using a triangular facet network, *Water Resources Research*, *W07522*(7), 1–12, doi:10.1029/2010WR009961.

Appendix A

Field work

Field work was completed by the candidate at the MCRB is shown in Table A.1.

Table A.1: List of fieldwork performed by the candidate.

Date	Duties
August 19–20, 2010	Site survey
October 11–20, 2010	Installation of instruments, erect South Meadow
January 27–February 9, 2011	Site visit, data download
April 14–18, 2011	Site visit, data download
May 19–26, 2011	Site visit, data download
June, 2011	Site visit, data download, removal of SPN1
August, 2011	Site visit, data download, removal of South Meadow

Appendix B

Code

The code used for this study is a mix of C++ and Matlab and the code is split into two sections:

`Umbra` is the shadowing model, written in C++ and Matlab. It provides the ability to read from the input meteorological files, to create the triangulation, generate the shadow locations, and write the output.

`libmaw` (Matlab Application Wrapper) is the Matlab/C++ interface that provides a link between Matlab and `Umbra` written in C++ and Matlab. It is built upon the open source linear algebra library Armadillo[1]. The `libmaw` library allows for simple interaction with the Matlab engine from within C++ and handles all the memory allocations and deallocations via the use of the Boost[2] `smart_ptr` library.

The code listing is too large to include here. However, it is publicly available from GitHub at github.com/Chrismarsh[3]. All code is licensed under the MIT[4] or GPL[5] open source licenses. Numerous Matlab File Exchange (FEX) submissions were used to avoid re-implementing methods. These are available in the GitHub repository.

[1] http://arma.sourceforge.net/
[2] http://www.boost.org
[3] http://www.github.com/Chrismarsh
[4] http://www.opensource.org/licenses/MIT
[5] http://www.gnu.org/licenses/