Landsat-7 Long-Term Acquisition Plan: Development and Validation

Terry Arvidson, Samuel Goward, John Gasch, and Darrel Williams

Abstract

The long-term acquisition plan (LTAP) was developed to fulfill the Landsat-7 (L7) mission of building and seasonally refreshing an archive of global, essentially cloud-free, sunlit, land scenes. The LTAP is considered one of the primary successes of the mission. By incorporating seasonality and cloud avoidance into the decision making used to schedule image acquisitions, the L7 data in the U.S. Landsat archive is more complete and of higher quality than has ever been previously achieved in the Landsat program.

Development of the LTAP system evolved over more than a decade, starting in 1995. From 2002 to 2004 most attention has been given to validation of LTAP elements. We find that the original expectations and goals for the LTAP were surpassed for Landsat 7. When the L7 scan line corrector mirror failed, we adjusted the LTAP operations, effectively demonstrating the flexibility of the LTAP concept to address unanticipated needs. During validation, we also identified some seasonal and geographic acquisition shortcomings of the implementation: including how the spectral vegetation index measurements were used and regional/seasonal cloud climatology concerns. Some of these issues have already been at least partially addressed in the L7 LTAP, while others will wait further attention in the development of the LTAP for the Landsat Data Continuity Mission (LDCM). The lessons learned from a decade of work on the L7 LTAP provide a solid foundation upon which to build future mission LTAPs including the LDCM.

Introduction

Over the last 33 years and seven Landsat missions, a variety of mission operation approaches have been employed to acquire global Landsat imagery. Inspection of the historical Landsat observation record currently residing in the U.S. Geological Survey's Center for Earth Resources Observation and Science (USGS/EROS) National Satellite Land Remote Sensing Data Archive (NSLRSDA) has shown that during various periods of the Landsat mission history, the observatories were directed toward differing goals (Goward *et al.*, 2006). These efforts varied from manually-specified global observations in early days to much more regional-focused goals when international cooperators became more prevalent, satellite systems failed (such as recorders and antennae) and/or the "commercial" market for Landsat observations declined. It was only with Landsat 7 and the return of the mission to U.S. government operations that the original vision of global monitoring was again pursued under the LTAP concept.

Background

Over the Landsat mission history, most efforts to define image acquisition approaches were *ad hoc*, determined principally by mission operators at:

- The National Aeronautics and Space Administration's Goddard Space Flight Center (NASA GSFC) (1972 through 1978),
- National Oceanic and Atmospheric Administration (NOAA) (1979 through 1985),
- EOSAT (1985 through 2003), and
- USGS/EROS (2003 to present).

There were also times that NASA program managers in charge of large-area missions (such as the NASA Large Area Crop Inventory Experiment (LACIE, 1976 through 1978) and Agriculture and Resource Inventory Surveys through Aerospace and Remote Sensing (AGRISTARS, 1979 through 1982) programs) had a significant impact in defining mission acquisitions. The innovation of an automated long-term acquisition plan originated from discussions in the *ad hoc* Landsat science team, after the mission had been returned to the government in 1992. Implementation of this approach was taken up by the Landsat Project Science Office (LPSO) in 1995 to support Landsat-7 mission operations.

Landsat Scenes and Images

The Landsat scene concept is an interesting artifact of the 34-year Landsat mission heritage. A Worldwide Reference System (WRS) was developed to break the image acquisitions into "bite-size" pieces, suitable for processing with low-capacity computers as well as viewing as imagery by analysts. The Landsat orbit is maintained such that each image collected is centered on a WRS scene location. The WRS system uniquely identifies all locations on the Earth observed by Landsat. Each WRS scene is an area on the Earth of approximately 32,000 km^2, or about 1.5 by 1.5 degrees of latitude and longitude at the equator.

Need for LTAP

There are many users even today (after more that 34 years of operations) who believe that the Landsat system acquires data continuously while in orbit. Given mission and ground

---

Terry Arvidson is with Lockheed Martin, GSFC Code 614.4, Bldg. 33, Rm. G313, Greenbelt, MD 20771 (terry.arvidson@gsfc.nasa.gov).

Samuel Goward is with the University of Maryland, Department of Geography, 2181 LeFrak Hall, College Park, MD 20782.

John Gasch is with Emalico, LLC, 2113 Hallmark Drive, Gambrills, MD 21054.

Darrel Williams is with NASA/GSFC, Code 614, Greenbelt, MD 20771.

Photogrammetric Engineering & Remote Sensing
Vol. 72, No. 10, October 2006, pp. 1137–1146.

0099-1112/06/7210–1137/$3.00/0

TABLE 1. HYPOTHETICAL LANDSAT-TYPE ACQUISITION RATES

Acquisition Strategy	Per Day		Per Year	
	Scenes	Bytes	Scenes	Bytes
Always	3,600	1.1 tera	1.2 mil	400 tera
Day Only	1800	0.6 tera	0.6 mil	200 tera
Day & Land Only	850*	0.3 tera	0.3 mil	100 tera

*Somewhat variable from day-to-day and the land-ocean proportions vary with longitude.

TABLE 2. SENSOR AND SPACECRAFT TECHNICAL CONSTRAINTS THAT MUST ALSO BE CONSIDERED WITHIN THE SCHEDULING ALGORITHMS

Mission Constraint	Constraint Specification
Duty Cycle	• Maximum minutes on/moving time window: 34 mins/100 mins, 52 mins/200 mins, 131 mins/600 mins, 16.7%/23 hrs; at night, 15 mins maximum on time
Communications	• Direct downlink only, no relay capability • Limited number of ground stations at which to downlink data • Only three downlink antennas available
On-board Storage	• 105-scene capacity (at beginning of life) • Maintain sync between on-board storage status and scheduling system
Sensor States	• Must specify gain state for each sensor band for each acquisition (see Markham *et al.*, 2006)
Acquisition Level	• Average of 250 scenes to the U.S. archive daily, averaged over 2 days (was raised to 300 in 2004)
Non-U.S. Stations	• Incorporate imaging requests from International Cooperators (ICs) into available duty cycle after U.S.-archive acquisitions • Manage antenna availability when multiple ICs in view at once

systems constraints, this has never been possible, nor would it be a particularly effective approach to operating such an observatory. For example, a continuously operating Landsat-7 system would generate truly large volumes of data: well over 1.1 terabytes per day or 3,600 images for the Enhanced Thematic Mapper Plus (ETM+) (Table 1). Not only would these data volumes overwhelm existing satellite, telemetry, ground processing, and archival systems, but also it would produce many data of little or no value to the users.

Simply limiting Landsat acquisitions to only land areas that are sunlit reduces the data volume by approximately 75 percent; however, even with these limits, acquired data volumes would be more than the sensor/data system can afford or tolerate. On average, the Landsat-7 observatory has the potential for observing 850 to 900 sunlit scenes per day. Landsat-7 system developers determined that the observatory could acquire an average of approximately 450 images per day, of which approximately 250 images would be acquired for the U.S. archive. The remaining system capacity would be made available to International Cooperators (ICs).

LTAP Science Drivers

One of the more compelling realizations from the multi-decadal Landsat mission is that the Earth's land areas are better defined by their seasonal variability than the synoptic patterns imaged on any given day (Goward, 1989; Goward and Williams, 1997; Justice *et al.*, 1985). These seasonal variations in land-cover occur primarily as local vegetation responds to variations in climate (Goward and Prince, 1995). Characterizing land-cover and land-cover change from remotely sensed data requires more than one image per year (Townshend *et al.*, 1991).

After the Landsat system was returned to government management in the early 1990s, the *ad hoc* Landsat Science Working Group suggested that the then-future Landsat-7 should acquire "seasonally-refreshed" observations of all the Earth's land areas. This statement was further refined by specifically noting that the long-term goal of the Landsat-7 mission would be to "build and seasonally refresh an archive of global, essentially cloud-free, sunlit, land scenes." It is important to note that this mission statement and the resultant long-term acquisition plan apply solely to the National Satellite Land Remote Sensing Data Archive at USGS/EROS. The acquisitions transmitted to the International Cooperator ground stations for their archives are negotiated between USGS and the ICs. They typically include all imagery for WRS locations within the acquisition circle of the ground station.

The Landsat-7 Long-Term Acquisition Plan

The Landsat-7 LPSO made a concerted effort to systematize and automate the L7 mission goal to build and seasonally refresh a U.S. archive of global, essentially cloud-free, sunlit, land scenes (Arvidson *et al.*, 1999; Arvidson *et al.*, 2001; Arvidson *et al.*, 2002; Arvidson *et al.*, 2000; Goward *et al.*, 1996; Goward *et al.*, 1999). This was done by partnering the Mission Scheduler software with the long-term acquisition plan (LTAP). The LTAP defines the archive-building factors to be considered during scheduling including: location of land scenes, seasonality, cloud avoidance, gain settings, and sun angle. The Scheduler software implements algorithms which balance these factors against the mission constraints (Table 2) and arrive at the "best" 250 candidate images to be scheduled each day[1].

Development Cycle

The LTAP development team initiated work in 1995 and delivered the launch-ready LTAP in late-1998 (Arvidson *et al.*, 2001; Goward *et al.*, 1999). After launch in April 1999, a short-term acquisition plan was used to operate the mission for the first 75 days, to achieve initial spacecraft and sensor assessment, before the spacecraft arrived on the WRS orbit. When on-orbit checkout was completed, the Landsat-7 was placed in operational status as of 29 June 1999. At that time the LTAP took over mission scheduling.

LTAP development continued post-launch, including a three-year effort initiated in 2002 to validate the developed approach. Continual LTAP monitoring during operations, as well as the more detailed LTAP validation reported here, have produced numerous insights into the success of the LTAP-based operations. These insights have led to suggested

---

[1] On 11 May 2004, the daily image acquisition level was raised from 250 to 300. After the failure of the scan-line corrector mechanism on the ETM+ in June 2003, all images acquired had wedge-like data gaps on the left and right edges. As a result, the acquisition strategy was altered to favor acquiring several clear images of a scene within a 2- or 3-cycle time period; these can be merged to form a single image with little or no residual data gaps.

changes in the LTAP which have either already been implemented in LTAP updates or should serve as important considerations for the development of LTAPs for future systems such as LDCM.

LTAP Elements

As stated earlier, the full goal of the LTAP are to build a seasonally-refreshed global archive of sunlit, relatively cloud-free land scenes. Given the mission engineering constraints (Table 2), the LTAP development team hoped to achieve, as a minimum, at least one good image annually for each land WRS location on the entire globe. In addition, we expected to devote most acquisitions to observing locations where substantial land-cover seasonality occurs. The elements of the LTAP consist of land definition, seasonality and cloud avoidance

Land Definition and Niche Locations

Using an atlas (Rand McNally, 1991) and WRS maps (USGS/NOAA, 1982), each WRS scene location was examined to determine if it contained land, no matter how small, shallow coastal waters (to 200 m), and/or regions of permanent sea ice and ice pack. We defined each WRS location as "land" if any or all of these criteria applied.

Additionally, we worked with representatives from the science community to identify and label scenes of interest to various science niches, including: volcanoes, glaciers, reefs, oceanic islands, central Africa scenes for the Global Observation of Forest Cover program, Siberian boreal forest study sites, Earth Observing System (EOS) calibration sites, and sea ice extents (Table 3).

Seasonality

Land seasonality is defined by observed monthly variations in composited normalized difference vegetation index (NDVI) measurements, derived from the NOAA Advanced Very High Resolution Radiometer (AVHRR) (Goward *et al.*, 1994). With the L7 16-day orbital cycle, there are typically two acquisitions acquired during each month of the year. Two alternate monthly acquisition approaches were defined: (a) *acquire-always* at both opportunities, when large changes are observed in the NDVI from year to year, and (b) *acquire-once*, when little change is observed in the seasonality information. Once we had established the baseline global seasonality, we then adjusted this seasonal acquisition scheme for specific additional goals (Table 3).

Cloud Avoidance

The cloud avoidance element is a three-step scheduling process, followed by a post-acquisition analysis that guides future acquisitions. Cloud avoidance consists of:

1. *Historical WRS Average Cloud Cover:* Monthly summary of an ISCCP (International Satellite Cloud Climatology Project) five-year data set (Rossow *et al.*, 1996; Rossow and Schiffer, 1999)
2. *Current Daily Prediction:* Daily NOAA NCEP (National Centers for Environmental Prediction) cloud cover, aggregated to WRS locations.
3. *Priority Evaluation:* Scheduler compares monthly average cloud cover to the daily cloud cover prediction for every potential scene acquisition (Gasch and Campana, 2000). Predictions better than average receive a priority increase and *vice versa.*

Post-acquisition, the scheduler assesses each scene acquired to determine if acquisition was successful, using the automated cloud-cover assessment (ACCA) performed during post-acquisition image processing (see Irish *et al.*, 2006):

- If the ACCA <10 percent, the acquisition is considered successful
- If ACCA >60 percent, the acquisition is considered missed, the next acquisition priority is raised

TABLE 3. THE BASIC LTAP SEASONALITY WAS MODIFIED BY THESE ADDITIONAL COVERAGE GOALS

Additional Coverage Goal	Seasonality	Cloud Coverage
U.S. (except Alaska)	Acquire always all year long	Ignore predicts
Alaska	Acquire always all year long as long as solar zenith angle (SZA) <85 degrees	Ignore predicts (note – this was changed in 2002 to use cloud avoidance)
Humid tropical rainforest	Acquire always all year long	Use cloud avoidance
USDA campaign sites	Acquire always during specified seasons	Use cloud avoidance; do not acquire if predict is >60% clouds
Volcanoes	Acquire both day and night from 2 to 6 times each year, using science-submitted prioritization	Use cloud avoidance for day scenes
Sea ice	Acquire from 1 to 3 times each year within specified seasons	Use cloud avoidance
Oceanic islands	Acquire twice each year within specified seasons	Use cloud avoidance
Siberia boreal research sites	Acquire always within 9 month period as long as SZA <85 degrees	Use cloud avoidance
EOS Calibration sites	Acquire always	Use cloud avoidance
Reefs	Potential research sites: twice each year at peak/nadir of bleaching; existing research sites: quarterly	Use cloud avoidance
Glaciers and land ice	Acquire glaciers and Greenland once at peak of ablation; acquire Antarctica once during local summer season (changed in 2001 to active areas only)	Use cloud avoidance for glaciers and Greenland; ignore cloud predicts for Antarctica (changed in 2001 to use cloud avoidance)

- For ACCA of 11 percent to 59 percent, a graduated priority increase for the next acquisition is applied.

Thus, the cloud avoidance strategy not only impacts the acquisition priority for the current day but also impacts future priority for this WRS location, until a "successful" acquisition has taken place.

LTAP Results

During the first five years of the L7 mission, the LTAP performed exceptionally well, producing at least annual, nearly cloud-free imagery of most land locations on the Earth (see this issue's cover). For most locations we found that reasonably cloud-free imagery (<30 percent cloud cover) are acquired during the four seasons (e.g., December to February, March to May, June to August, September to November) of the year, although not monthly as we had originally hoped (Plate 1a).

As we inspected the acquired LTAP coverage in more detail, we began to discover some patterns in the acquisitions that we had not anticipated, such as devoting too many acquisitions to repeat coverage in the arid and semi-arid regions of the globe (Arvidson *et al.*, 2001). These insights led us to propose pursuit of a more formal validation of the LTAP, by examining the outcome of the LTAP as represented in the collection of images resident in the U.S. Landsat-7 data archive.

One aspect of the L7 LTAP that we have learned over and over again is that developing and refining the LTAP/Scheduler is entirely non-intuitive. Each time modeling or validation runs are produced and analyzed, we find unexpected results that warrant further investigation. The whole effort has been and continues to be almost entirely recursive. The interactions between demands for science-quality data and the technical constraints of a system such as Landsat-7 create an exceedingly complex mission operations challenge

Validation of the LTAP

From 1995 forward, we periodically conducted qualitative assessments on LTAP performance (Goward and Arvidson, 1998). In 2002, we undertook a definitive, quantitative assessment of LTAP, which ultimately took three years to complete (interrupted by the SLC failure). Every aspect of the LTAP was included in this assessment. We looked at the actual outcome of the LTAP performance, as determined from the archive contents, compared to our desired outcome to determine the effectiveness of the LTAP scheme.

Seasonality

Given the L7 17 percent daily duty cycle, we assumed that all the Earth's land areas would be acquired about once every two 16-day cycles (approximately once per month). To evaluate this outcome, we produced WRS-based coverage maps from the U.S. archive for each 32-day period through the year (Plate 1a and Plate 2a). We also produced seasonal (quarterly) coverage maps, where the seasons are defined as: December to February, March to May, June to August, September to November (Plate 1b and Plate 2b). We also used the USGS/EROS Landsat global visualization (GLOVIS) data exploration tool, available at **http://glovis.usgs.gov**, to qualitatively confirm our coverage analysis. With GLOVIS, we were able to perform visual assessments of specific geographic locations on the quarterly and the 32-day maps by viewing the temporal sequence of imagery that had been acquired under the LTAP for that location.

Overall, as originally implemented, the LTAP takes up to three months to acquire 90 percent of the land scenes, rather than the one month we had anticipated that it would take. However, even over three months (approximately six repeat observation cycles) for many WRS locations, many of the acquired scenes do not meet a criteria of less than 30 percent cloud cover, the minimum acceptable cloud cover for many users. In fact, to achieve this criterion globally requires at least one observation year, particularly for cloud-prone regions of the globe (Plate 1 and Plate 2; also see cover images for this issue).

Coverage Geography

We found that the geographic distribution of acquisitions is not always optimal:

1. *Too many desert scenes.* Qualitative inspection of coverage in GLOVIS as well as statistical assessments shows that the LTAP acquires too many desert scenes and not enough boreal scenes (Table 4). We apparently were not as systematic within the LTAP in restricting desert scene seasonality as we had intended to be. The current plan is to limit acquisition of desert scenes to twice per year: once at high sun angle (>60° elevation) and once at low sun angles (15° to 30° elevation, if this occurs).
2. *Too few boreal scenes.* Statistical analysis and visual browse image inspection also revealed that too few boreal forest images are acquired in the summer and too many during the winter (Table 4). We are not entirely sure why the LTAP acquires too few summer boreal images, but this may be the result of persistent local cloud cover or perhaps less reliable NCEP forecasts during the summer in this region. Our assessment of boreal images acquired (Table 4) indicates that, on average, we acquire two-to-four fewer images per year in the boreal regions than we do in the desert region. This, despite the fact that only about 25 percent of the acquired images are considered "clear" (<10 percent) in the boreal region. To address the summer problem, we now "*acquire-always*" during June to September in the boreal regions. The winter boreal outcome is the reverse, with many clear images acquired but with no change of land surface conditions, as snow cover persists for extended periods of time (see cover images).
3. We also have become uncomfortable with the large number of scenes containing a small portion of land when that same area was also covered by a neighboring scene. Not aware of any group of scientists interested in 100-meter or deeper water, we have already deleted 352 fringe coastal scenes from the LTAP database to conserve instrument resources.

TABLE 4. COMPARISON OF IMAGERY ACQUISITION RATES VERSUS CLOUDINESS FOR DESERT AND BOREAL REGIONS (AS DEFINED BY THE DEFRIES GLOBAL LAND COVER ANALYSIS)

Boreal	2001	2002	2003	2004	2005
Avg. acquisitions/ scene	5.6	5.3	3.7	4.8	6.1
Avg. ACCA score (%)	38.1	37	33.3	36.8	37.7
Best ACCA score (%)	5.8	6.5	10.1	11.2	6.1
% clear images (acca<10%)	26.2	28.4	33.6	27.3	27.7
Desert	**2001**	**2002**	**2003**	**2004**	**2005**
Avg. acquisitions/ scene	7.4	7.7	7.9	12.1	12.4
Avg. acca score (%)	14.1	13.9	15.8	12.6	13.9
Best acca score (%)	0.3	0.5	0.6	0.3	0.3
% clear images (acca<10%)	71	70.6	67.6	73.3	70.6

Cloud Avoidance

We found a strong correlation between the assessed and predicted cloud-cover values at the low (0 to 10 percent) and high (80 to 100 percent) ends of the scale (Figure 1). Our analysis shows that predictions less than 10 percent or more than 80 percent are generally accurate. Based on this, we set a threshold of 80 percent predicted cloud cover beyond which scenes were not considered for acquisition. A predicted value of 93 percent for a scene's acquisition is very likely to be proven true by the assessed cloud score after image processing, so we avoid wasting resources by not imaging the scene in the first place.

Correlation for the remaining range (between 10 to 80 percent) varies by latitude under the influence of regional atmospheric instability and related cloud development. This is particularly true during the summer and in the tropics, when and where convective cloud systems dominate and can be difficult to predict. During the winter season, in the mid- to high-latitudes, frontal systems between air masses of differing conditions dominate cloud development and, for now, are easier to forecast. In addition, specific situations (such as fog/cloud banks off western continental coasts and the affinity of clouds for isolated reefs and islands) result in misleading predictions and assessments. These well-behaved problem areas have led to proposals for updates to the ACCA algorithm (Irish *et al.*, 2006) and the Scheduler software to address these situations.

Cloud-Prone Locations

We had speculated that LTAP performance in cloud-prone locations might improve if the Scheduler was instructed to acquire imagery every time for periods greater than one month, and this was done in the tropic regions. To evaluate the effect of this approach and its interactions with cloud avoidance, we produced maps (Plate 2a, 2b, and 2c) showing the best cloud-cover score for each WRS scene, for one month (Plate 2a) and three months (Plate 2b) as well as the impact of cloud cover on acquisition frequency, for three months (Plate 2c).

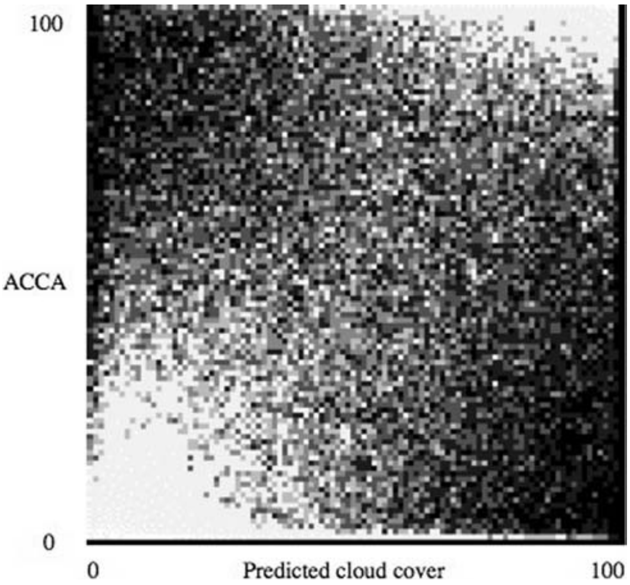

Figure 1. In this scatterplot, the whiter the symbol, the higher the frequency of the paired combination. The strong correlation between assessed and predicted cloud cover at the extremes of the scale have influenced the Scheduler algorithms.

The outcome expected with increased acquisitions is cloud clearing, as noted for the United States, New Zealand, and portions of the Amazon basin. However, what is frequently observed (particularly in persistent cloudy areas such as equatorial Africa, northwest South America, the Aleutian Islands, and far eastern Asia) is that increased acquisition produces less cloudy imagery versus clear imagery, despite persistent acquisitions throughout a full quarter of the year. When such cloud-prone locations are correlated with acquisition repeat frequency (Plate 2b and 2c), the LTAP persistence in acquiring these locations is evident. These are the same locations that international Landsat ground station operators report are difficult to acquire as cloud-free images even though they are acquiring imagery at every opportunity, unconstrained by our cloud-avoidance approach.

These results indicate that the cloud avoidance approach used in the L7 LTAP could benefit from a refined seasonal (monthly) and regional understanding of cloud climatology, such that for some locations and times no imagery would be acquired, freeing up these acquisition resources for other locations where acquisitions would be more successful. This could also assist in specifying urgent areas for radar observations as well. There is more work needed in this area. There is also a compelling need to develop a better understanding of cloud climatology to evaluate whether satellite systems with repeat coverage better than the current 15 to 30 day repeat cycle would produce better global coverage than systems such as Landsat-7 and LDCM.

Value of Cloud Avoidance

The LTAP cloud avoidance in general appears to be successful. In the USGS EROS Landsat-7 archive, accumulated from July 1999 to present, 57 percent of the images have <30 percent cloud cover and 39 percent have <10 percent cloud cover, with an overall cloud-cover percentage of 34 percent since launch (Landsat-7 Monthly Operations Review, May 2006, unpublished).

Assessment of the LTAP cloud-avoidance performance is difficult because we have no cloud information on the scenes we did not acquire. Nevertheless, we do have two alternate assessments of the possible difference cloud avoidance has made: (a) The United States (acquire-always) *versus* the rest of the globe (Table 5), and (b) differences in cloud cover observed when we did not have cloud forecasts to use in the scheduling process for cloud avoidance.

The U.S. *versus* the rest of the world data (Table 5) indicates that we acquired a total of 291,000 scenes during this two-year period. Of these 291,000 images, about 59,000 were from the U.S., including Hawaii and Alaska, leaving approximately 232,000 images acquired over the rest of the globe. Note that the predicted (NOAA) cloud cover in the U.S. exceeds our ACCA estimate by approximately 10 percent. We believe that this bias originates primarily from the ACCA algorithm, which is not effective for cirrus and other thin clouds. For the remainder of the world, the average predicted cloud cover for all possible images was 53 percent, whereas for the images selected for acquisition, the predicted cloud cover was 32 percent and the ACCA observed cloud cover was 34 percent[2]. Assuming the noted

---

[2] This, by the way, seems in conflict with the U.S. bias analysis discussed previously. We assume this has to do with the fact that for the U.S., every possible acquisition is included in the CCp versus CCa comparison in Table 5. For the rest of the globe, only the "best" acquisitions are included, so we have a preponderance (hopefully) of predicts in the 0 to 10 range which are strongly correlated with ACCA, thus lessening the gap between CCp and CCa.

10 percent bias between NOAA predicts and ACCA, the images acquired outside the U.S. have approximately 8 percent lower cloud cover than would have been acquired if cloud avoidance had not been used.

The second evidence that cloud avoidance improves acquisition cloud cover comes from a comparison of ACCA scores for time periods when we were unable to use cloud avoidance because the NOAA forecasts were not available *versus* normal cloud avoidance (Table 5). Of the approximately 291,000 images acquired since July 2002, for approximately 20,000 (7 percent) of the images we did not have cloud forecasts. For the no-forecast scenes, the assessed cloud cover averages 41 percent, whereas for those images where forecasts were available the cloud cover averages 31 percent. This 10 percent difference is comparable to the previously assessed 8 percent noted for the U.S. *versus* the

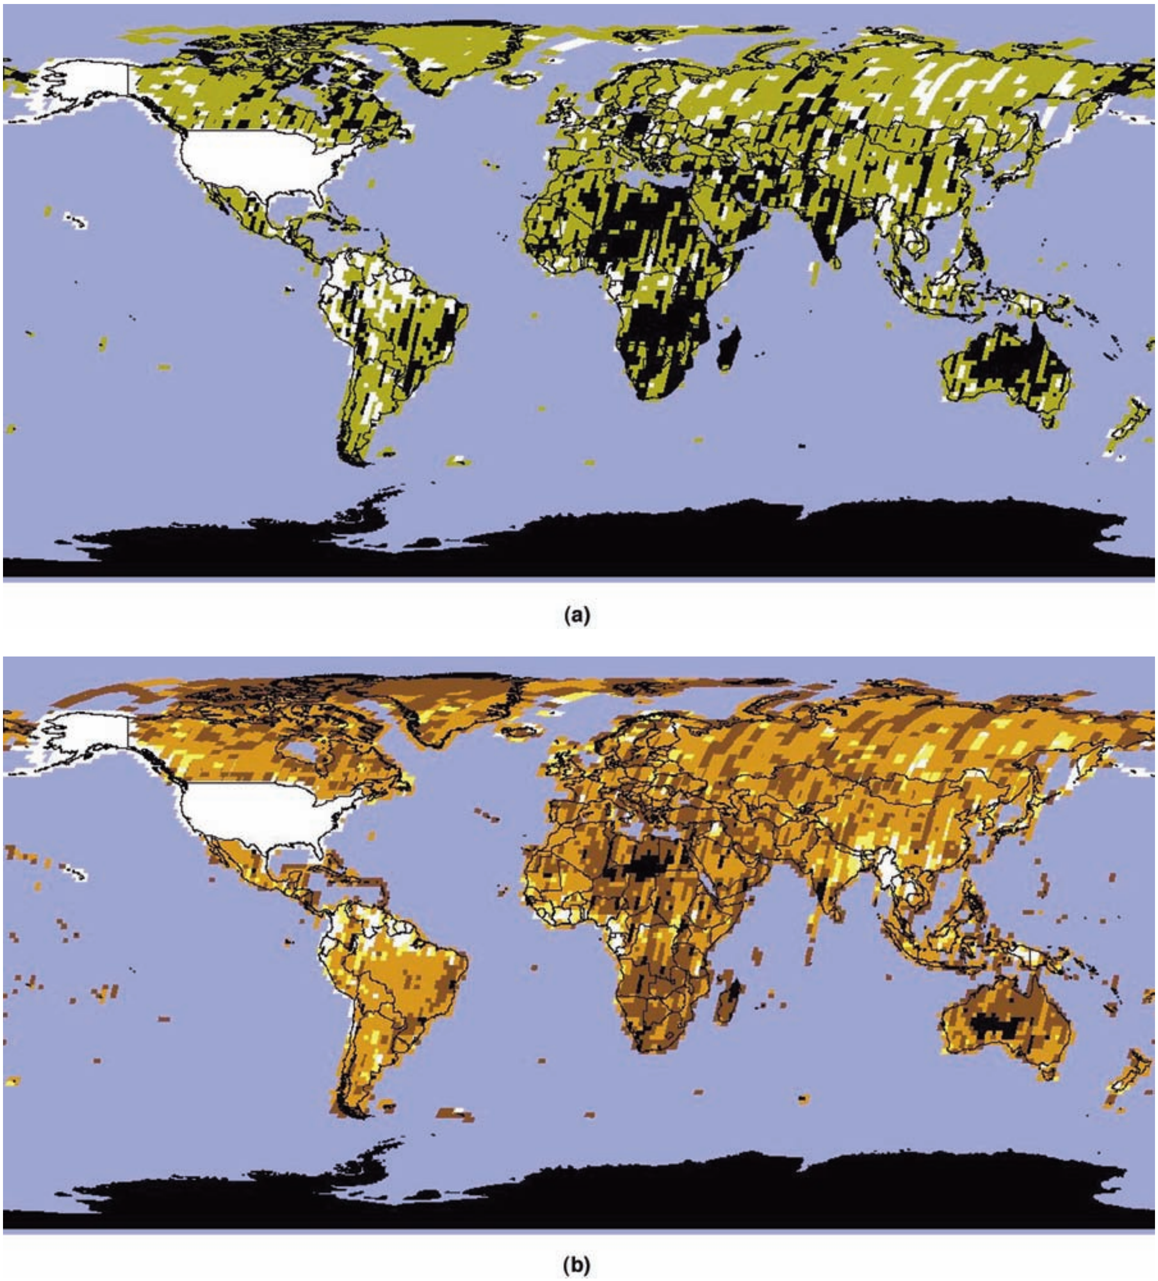

Plate 1. Monthly and Seasonal Acquisition Rates: (a) shows coverage for two orbital cycles in June 2001: green = 1 acquisition, white = two acquisitions, and black = none. Note the lack of acquisitions over southern Africa, Australia, and portions of South America. (b) shows coverage over six orbital cycles, during the Northern Hemisphere summer months of June, July, and August of 2001. Lighter tone indicates more scenes acquired during the time period and black indicates none. We have just about full global coverage during this three-month season. Some gaps remain over the African and Australian deserts and the Pacific islands. Full coverage is achieved at least once each year, and as often as quarterly for many regions.

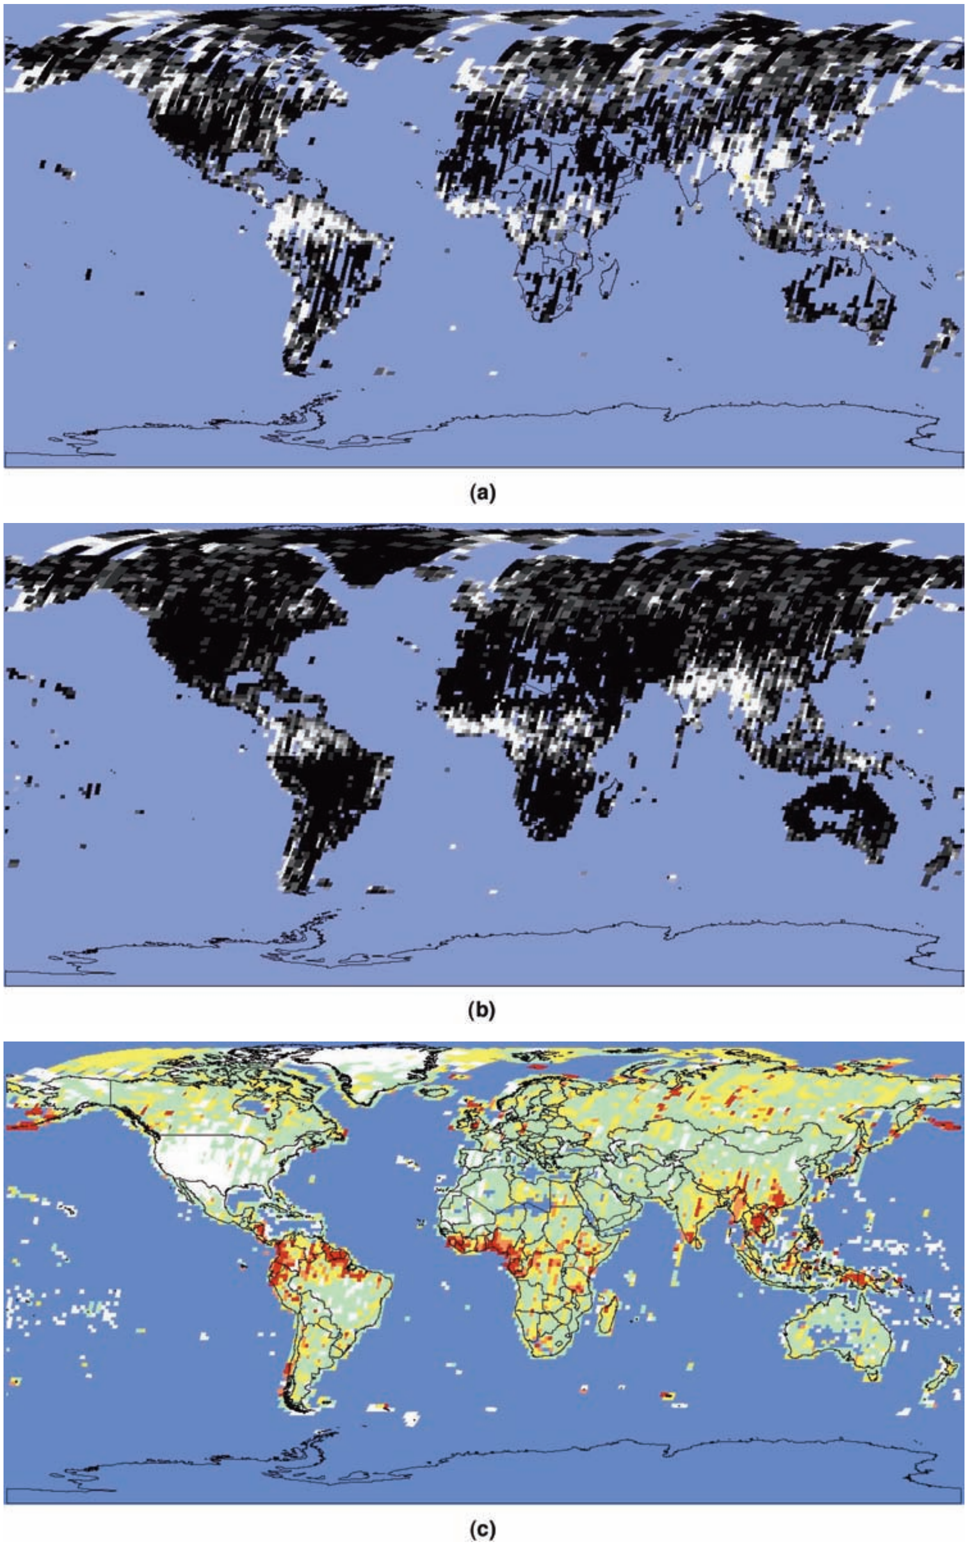

Plate 2. (a) One-Month Acquisition Cloud Cover. Lowest cloud cover acquired over two orbital cycles in June 2001 (orbital cycles 11 and 12). Black is clearest and white is cloudiest. (b) Seasonal Acquisition Cloud Cover. Lowest cloud cover acquired over the six orbital cycles for June, July and August 2001. Black is clearest and white is cloudiest. (c) Repeat Acquisition Rates for a Season. How many times an acquisition failed to meet cloud cover requirements and therefore was targeted for another acquisition. White = 1, light green = 2, yellow = 3, orange = 4, red = 5. Note that many of the highest repeat acquisitions occur for locations where a relatively low cloud cover image is never acquired.

Table 5. Cloud Cover Statistics for Scenes Acquired from 01 July 1999 through 01 June 2003. CCa is the Assessed Cloud Cover Percentage Determined During Image Processing. CCp is the Predicted Cloud Cover Percentage Received from the NOAA National Centers for Environmental Prediction

	All LTAP Scenes	48 U.S. States + HI	Alaska	All Other LTAP Scenes
Acquired images	291080	40892	17689	232455
Avg CCa	36.7	39.9	59.6	34.4
Avg CCp	36.3	48.7	69.2	31.6
% of archive Clear	34.2	33.8	11.0	36.0
All Opportunities	744108	40892	17689	680905
Avg CCp	52.7	48.7	69.2	52.5
Rejected Opportunites	446925			
Avg CCp	63.6			

rest of the world. It is important to note that this reduction in image cloud cover for acquired scenes most likely varies seasonally and regionally, as cloud frequency varies (as noted regionally for Alaska in Table 5).

Niche Communities

In concert with representatives from each identified science niche community, we determined the initial niche seasonality and incorporated it into the LTAP (Table 3). During LTAP validation, we contacted these same representatives and questioned them as to the effectiveness of the LTAP from their science niche point of view (i.e., their satisfaction as data users). In addition, we tabulated and analyzed data sales statistics on a science niche basis. We also analyzed WRS-based maps portraying the acquisition density for each science niche. As a result of these analyses, we made several changes to the niche community acquisition requests, including:

- Volcanoes – Discontinuing night acquisitions except on a request basis
- Sea ice – Discontinuing acquisitions except on a campaign-request basis
- Coral Reefs – Reclassifying high-priority reefs as low-priority reefs, with the associated base-priority value; the reef-mapping community indicated that most of the mapping was complete and they were not in a position to reprioritize acquisitions at this time
- Antarctica – Scaling back acquisitions to concentrate on those areas deemed most active and most cloudy

Lessons Learned

The LTAP has been extremely successful and has demonstrated the value of an automated but easily-modified acquisition strategy. In view of this, the two primary successes of LTAP (seasonality and cloud assessment/avoidance) are included in the operations concept for the next Landsat mission (NASA, 2006).

We continue to work on the LTAP, making adjustments for operational scenarios (e.g., the need for composite-pair acquisitions instead of single-scene acquisitions after the SLC failure) and rectifying some of the shortcomings observed during the validation effort between 2002 through 2004.

Cloud Avoidance

The NCEP predictive models have changed since 2004 and it would be beneficial for future programs such as LDCM to repeat some of the validations with the new models. The baseline ISCCP data set used to establish the nominal cloud cover was a five-year average. It would be interesting to use a current data set, spanning more than 10 years, to update the nominal values and analyze them further for possible El Niño/La Niña effects.

Landsat Temporal Repeat Coverage

One of the major conclusions of this validation effort, is that the Landsat-7 16-day repeat coverage cannot produce clear (<10 percent cloud) observations for all land areas for each season of the year. The Landsat-7 LTAP mission, prior to the SLC failure, typically could not collect images for many scene locations with cloud cover less than 30 percent once every three months. We found this to be a particular problem in the high northern latitudes, where the growing season is short and dominated by convective, partly-cloudy conditions. The situation is much worse in cloud-prone locations such as the humid tropics, where in some cases it might take years or might not ever be possible to acquire "cloud-free" imagery. Further, with 16-day coverage, it is not realistic to attempt to capture vegetation dynamics during the green-up and senescence cycles. What is not certain is whether more frequent temporal coverage might help to address these cloud-persistence problems.

There is a clear need to conduct better analyses of the regional and seasonal patterns of cloud climatology. At least for some cases, the information needed resides in the U.S. and international cooperator "acquire-always" archives. There are also resources in the AVHRR and Moderate Resolution Imaging Spectroradiometer (MODIS) historical records that can help address these uncertainties. As we move past the LDCM deployment and seek to define future operational mission concepts, a better understanding of these questions will be needed.

NDVI Seasonality

Global Spectral Vegetation Index Observations

Our analysis of vegetation seasonality was from the UMCP NDVI data set which is now over 10 years old. There have been many advances in the development of global, seasonal spectral vegetation index data sets, including from AVHRR, VEGETATION, and now MODIS (James and Kullari, 1994; Justice and Townshend, 2002; Xiangming, 2002). Any new efforts to refine the LTAP seasonality approach should take advantage of these more refined versions of the AVHRR record including starting with the 8-km global acquisition coverage (GAC) files, which adjust for Rayleigh scattering and water vapor absorption and are based on better navigation information. The MODIS sensor is further advanced, producing measurements from quite narrow spectral bands that avoid most water vapor contamination and produce near-daily coverage at 1 km, 500 m and 250 m, in the red and near-infrared bands. In each case, the records of land vegetation seasonality would provide more precise estimates of the regional vegetation dynamics the satellite is most likely to observe at any given time.

Defining Seasonality

Our simple, *binary* analysis of seasonality used in the L7 LTAP could be more sophisticated. The original definition of seasonality throughout the globe prescribed acquisition levels that were too high and not necessarily at the right time of year to detect change. We used NDVI data to determine when change was occurring and tried to capture that change through frequent acquisitions. However, due to the sensor duty cycle and sheer volume of candidate scenes for acquisition, scenes designated as *acquire-always* were in fact acquired once every four or five opportunities, at best. This made it difficult to know where on the upslope (or downslope) of the NDVI curve the image had been captured and the rate of change it represented.

Given these constraints as well as that of the Landsat-7 repeat coverage, we believe that a more useful approach to mission operations would be to focus acquisitions on the growing season and the dormant season, rather than attempting to capture the shorter green-up and senescent periods. This focus would allow us to work harder to produce composite clear views of these landscapes during each dominant season, a significant criteria noted for the development of the USGS National Land Cover Data (NLCD) set with Landsat (Homer *et al.*, 2004).

Our preliminary analysis of the AVHRR NDVI record has led us to conclude that in the future, using the first derivative of the NDVI annual temporal trajectory would provide a clear definition of these seasons and the transitions between these seasons. This suggests the Scheduler system should be supplied with the original or first derivative of the temporal NDVI data sets to evaluate where in the annual seasonal cycles the Landsat will be observing at any given time and location.

Summary and Conclusions

The Landsat-7 long-term acquisition plan (LTAP) was devised as a tool to ensure systematic global acquisitions, building an archive of high-quality cloud-free scientifically-useful imagery. This tool has been quite successful, resulting in a U.S. archive of over 540,000 images providing at least twice-yearly global coverage, with 30 percent or less measured cloud cover for 57 percent of its content (10 percent or less for 39 percent of its content), and with an overall cloud-cover percentage of 34 percent since launch (Landsat 7 Monthly Operations Review, May 2006, unpublished). It should be noted however that the intended goal of seasonally refreshing the archive was not met for some areas of the globe. The LTAP has also proven itself adaptable to changing conditions, including new acquisition requirements, shifting biases in cloud predicts used during scheduling, and the increasing need to preserve sensor resources to ensure continuity and avoid a data gap before the launch of the Landsat Data Continuity Mission.

Importance of Clouds

We have found that the cloud avoidance part of the LTAP strongly dominates all other elements of the system, including seasonality. This might not continue to be the case if we are able to increase temporal repeat frequency of the coverage. However, there is considerable uncertainty now about how far increased temporal repeat coverage will take us toward at least biannual images with 10 percent or less cloud coverage for all land areas of this planet. We have much work yet to do in this direction and it should be done as soon as possible, before a follow-on operational concept to the LDCM is proposed.

Future Landsat-type Systems

Landsat-7's time as our primary land observatory is coming to an end. The U.S. Executive Office of the White House has now stated that Landsat is a national asset and the continuity of Landsat-like data must be pursued (Marburger, 2004 and 2005). The value of an efficient and effective LTAP will increase in future missions as they produce increased data volumes from substantially enhanced observatories.

The LDCM program is conceived as a dedicated satellite with a Landsat-type sensor (Irons and Masek, 2006). As such, the LTAP is included in the operations concept for the LDCM program, with the intent to continue the successful building of a global archive with scientific relevance. We also expect that the LTAP concept will play a pivotal role in the new operational era for Landsat that is currently being defined.

Adaptation of the LTAP concept to the LDCM mission will not simply be accomplished by transferring some computer code and a couple of technicians to LDCM. We have learned substantially from the decade of effort to develop and validate the Landsat-7 LTAP. Those lessons learned need to be captured and used in the development of an LDCM LTAP that enhances the LTAP performance experienced for Landsat-7.

References

Arvidson, T., J. Gasch, and S.N. Goward, 1999. Pleasing all of the people most of the time: Planning Landsat 7 acquisitions for the U.S. archive, *Proceedings of Pecora 14* American Society for Photogrammetry and Remote Sensing, Bethesda, Maryland, unpaginated CD-ROM.

Arvidson, T.J., J. Gasch, and S.N. Goward, 2000. Building a global, consistent and meaningful Landsat 7 data archive, *Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VI* (S.S. Shen and M.R. Descour, editors), SPIE, Orlando, Florida, Volume 4049, pp. 356–367.

Arvidson, T., J. Gasch, and S.N. Goward, 2001. Landsat 7's long term acquisition plan – An innovative approach to building a global archive, Special Issue on Landsat 7, *Remote Sensing of Environment*, 78:13–26.

Arvidson, T., R. Irish, B. Markham, D. Williams, J. Feuquay, J. Gasch, and S.N. Goward, 2002. Validation of the Landsat 7 long-term acquisition plan, *Proceedings of Pecora 15*, American Society for Photogrammetry and Remote Sensing, Bethesda, Maryland, unpaginated CD-ROM.

Gasch, J., and K.A. Campana, 2000. Cloud cover avoidance in space-based remote sensing acquisition. *Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VI* (S.S. Shen and M.R. Descour, editors), SPIE, Orlando, Florida, Volume 4049.

Goward, S.N., T. Arvidson, D. Williams, J. Faundeen, J. Irons, and S. Franks, 2006. Historical record of Landsat global coverage: Mission operations, NSLRSDA, and international cooperator stations, *Photogrammetric Engineering & Remote Sensing*, 72(10).

Goward, S.N., 1989. Satellite bioclimatology, *Journal of Climate*, 7(2):710–720.

Goward, S.N., and T. Arvidson, 1998. *Science Review of Landsat-7 Long Term Acquisition Plan*, NASA Goddard Space Flight Center, Greenbelt, Maryland, 8 p.

Goward, S.N., R. Dubayah, and J. Haskett, 1996. Landsat-7 long term acquisition plan: Seasonal acquisition strategy, *Proceedings of Pecora 13 Symposium*, American Society for Photogrammetry and Remote Sensing, Bethesda, Maryland, unpaginated CD-ROM.

Goward, S.N., and D.L. Williams, 1997. Landsat and Earth systems science: Development of terrestrial monitoring, *Photogrammetric Engineering & Remote Sensing*, 63:887–900.

Goward, S.N., J. Haskett, D. Williams, T. Arvidson, J. Gasch, R. Lonigro, M. Reeley, J. Irons, R. Dubayah, S. Turner, K. Campana, and R. Bindschadler, 1999. Enhanced Landsat capturing all the Earth's land areas, *EOS*, 80(26):289–293.

Goward, S.N., and S.D. Prince, 1995. Transient effects of climate on vegetation dynamics: Satellite observations, *Journal of Biogeography*, 22:549–563.

Goward, S.N., S. Turner, D.G. Dye, and S. Liang, 1994. The University of Maryland improved Global Vegetation Index data product, *International Journal of Remote Sensing*, 15(17): 3365–3396.

Homer, C., C. Huang, L. Yang, B. Wylie, and M. Coan, 2004. Development of a 2001 National Landcover Database for the United States, *Photogrammetric Engineering & Remote Sensing*, 70(7): 829–840.

Irish, R.R., J.L. Barker, S.N. Goward, and T.J. Arvidson, 2006. Characterization of the Landsat-7 ETM+ automatic cloud cover assessment (ACCA) algorithm, *Photogrammetric Engineering & Remote Sensing*, 72(10).

Irons, J., and J. Masek, 2006. Landsat data continuity mission, *Photogrammetric Engineering & Remote Sensing*, 72(10).

James, M.J., and S. Kullari, 1994. The Pathfinder AVHRR land data set: An improved coarse resolution data set for terrestrial monitoring, *International Journal of Remote Sensing*, 15(17):3347–3364.

Justice, C.O., and J.R.G. Townshend, 2002. Special Issue – The Moderate Resolution Imaging Spectrometer (MODIS): A new generation of Land Surface Monitoring, *Remote Sensing of Environment*, 83(1–2):1–359.

Justice, C.O., J.R.G. Townshend, B.N. Holben, and C.J. Tucker, 1985. Analysis of the phenology of global vegetation using meteorological satellite data, *International Journal of Remote Sensing*, 6(8):1271–1381.

Marburger, J.H., III, 2004. Memorandum from the Office of Science and Technology Policy, *Landsat Data Continuity Strategy*, 13 August.

Marburger, J.H., III, 2005. *Landsat Data Continuity Strategy Adjustment*, Memorandum from the Office of Science and Technology Policy, 23 December.

Markham, B.L., T.J. Arvidson, S.N. Goward, J. Barsi, and P. Scaramuzza, 2006. Landsat-7 long-term acquisition plan radiometry — Evolution over time, *Photogrammetric Engineering & Remote Sensing,* 72(10).

NASA, 2006. Landsat Data Continuity Mission (LDCM) Operations Concept, Version 0.01 June URL: http://ldcm.nasa.gov/procurement/LDCMSys_OpsCon_060601.pdf (last date accessed: 07 July 2006).

Rand McNally and Company, 1991. *The New International Atlas*, The Easton Press, MBI, Inc., Norwalk, Connecticut, 519 p.

Rossow, W.B., and R.A. Schiffer, 1999. Advances in understanding clouds from ISCCP, *Bulletin of the American Meteorological Society*, 80(11):2261–2288.

Rossow, W.B., A.W. Walker, D.E. Beuschel, and M.D. Roiter, 1996. International Satellite Cloud Climatology Project (ISCCP) documentation of new cloud datasets, *WMO/TD-No. 737*, World Meteorological Organization, Geneva, Switzerland, 115 p.

Townshend, J.R.G., C.O. Justice, W. Li, C. Gurney, and J. McManus, 1991. Global land cover classification by remote sensing: Present capabilities and future possibilities, *Remote Sensing of Environment*, 35:243–255.

USGS/NOAA, 1982. *Index to Landsat Worldwide Reference Systems (WRS) Landsats 1, 2, 3, and 4*, Sheets 1–26.

Xiangming, X.S.B., L. Jiyuan, Z. Dafang, and L. Mingliang, 2002. Characterization of forest types in Northeastern China, using multi-temporal SPOT-4 VEGETATION sensor data, *Remote Sensing of Environment*, 82(2–3):335–348.