

CoreWall: A Methodology for Collaborative Visualization of Geological Cores

BY

ARUN GANGADHAR GUDUR RAO
B.S. University of Illinois at Chicago, 2003

THESIS

Submitted as partial fulfillment of the requirements
for the degree of Master of Science in Computer Science
in the Graduate College of the
University of Illinois at Chicago, 2006

Chicago, Illinois

This thesis is dedicated to my grandparents

Maruthi and Shakunthala Repalle
Venkatrama and Savitridevi Gudur

ACKNOWLEDGEMENTS

I would like to thank Jason Leigh, my advisor, and Paul Morin for letting me start this project and receive the opportunities to meet individuals from other fields of study. To Andy Johnson, thank for you valuable input as development went on. Thanks to Frank Rack, for willing to be on my thesis committee and introducing me to many individuals in the geological community, especially on the two-week ride from Panama to Canada. I would like to thank Emi Ito, Anders Noren, Doug Schnurrenberger and Amy Myrbo for guiding me on my first hands-on experiences with cores and continuing to provide me insights into the science when I needed them. For helping me to understand the ocean drilling processes, I want to thank Sean Higgins. To Dave Harwood and Richard Levy, I would like to thank them for adopting this project into their work. For the help in creating posters, brochures, logos, web pages and all other graphic designs, I would like to thank Julieta Aguilera. Thanks to Chris Jenkins and Bill Kamp for their developments of CoreNavigator and CoreCLIP. From CHRONOS, I would like to thank Doug Fils, Cinzia Cervato, Xiaoyun Tang and Josh Reed, for their collaboration on this project.

I received help in making design decisions from my lab-mates Luc Renambot, Venkat Vishwanath, Bob Kooima, Cole Krumbholz, Javid Alimohideen and Dimtry Svistula. A lot of demos were done in order to get community feedback as work progressed, and I appreciate all the help from the faculty and staff of the Electronic Visualization Laboratory. Finally, to Yu-Chung “Julian” Chen and Sangyoon “James” Lee, thanks for helping me finish this phase of work and making strong strides in further developing the project. If it weren’t for your help, I may not have finished in time.

AGGR

TABLE OF CONTENTS

<u>Chapter</u>	<u>Page</u>
1. INTRODUCTION	1
1.1.BACKGROUND.....	1
1.2. EXISTING CORE WORKFLOWS.....	6
1.2.1. INTEGRATED OCEAN DRILLING PROGRAM'S WORKFLOW ...	10
1.2.2. LIMNOLOGICAL RESEARCH CENTER'S WORKFLOW.....	12
1.2.3. ANTARCTIC GEOLOGICAL DRILLING PROGRAM'S WORKFLOW.....	14
1.3. PROBLEMS WITH EXISTING WORKFLOWS.....	17
2. A NEW COMPUTATIONAL APPROACH TO CORE WORKFLOWS.....	20
2.1. CATEGORIES FOR IMPROVING CORE WORKFLOW	20
2.2. COMPARISON TO EXISTING APPLICATIONS	21
3. EMBEDDING THE COREWALL WORK SUITE INTO A GENERALIZED WORKFLOW.....	25
3.1. CORENAVIGATOR AND DATA DISCOVERY.....	30
3.2. CORELYZER, THE CENTRAL ENVIRONMENT	32
3.2.1. CORELYZER ARCHITECTURE.....	33
3.2.2. HOW CORELYZER MEETS DISPLAY SCALABILITY.....	35
3.2.3. HOW CORELYZER MEETS IMAGE DATA SCALABILITY.....	39
3.2.4. ANNOTATIONS.....	41
3.2.5. USER INTERACTION.....	44
3.2.6. CORELYZER AND EXTENSIBILITY.....	45
3.3. CORECLIP, A CORELYZER PLUG-IN TO ANALYZE DOWN-HOLE LOGS	46
3.4. WORKFLOW DATABASE	48
3.4.1. SYNCHRONIZED SESSIONS FOR COLLABORATION	48
4. DETAILED CASE STUDY: ANTARCTIC DRILLING	49
4.1. ANTARCTIC DRILLING SCENARIO	50
4.2. PHYSICAL INTEGRATION INTO THE ANDRILL WORKFLOW	52
4.3. IMAGE AND NUMERICAL LOG MANAGEMENT	53
4.4. CORELYZER AND PSICAT INTEGRATION.....	56
4.5. COMMUNITY RESPONSE TO CORELYZER	57

TABLE OF CONTENTS (continued)

<u>Chapter</u>	<u>Page</u>
5. FUTURE WORK.....	58
6. CONCLUSION.....	61
APPENDIX A: GEOLOGY TERMS USED	62
REFERENCES	63
VITA	68

LIST OF FIGURES

<u>Figure</u>	<u>Page</u>
FIGURE 1 A DIGITAL-LINE SCAN OF A 1.5-METER CORE SECTION RETRIEVED FROM LAKE HVITARVATN, ICELAND. NOTICE THE VISIBLE, ALTERNATING BANDS OF DARK AND LIGHT COLORED MUD. THIS HIGH FREQUENCY BANDING DOES NOT OFTEN OCCUR IN OCEAN CORES. THIS PARTICULAR CORE SECTION WAS RETRIEVED APPROXIMATELY 60 METERS BELOW THE LAKE FLOOR.....	2
FIGURE 2 A MULTI-SENSOR CORE LOGGER (MSCL) AT THE LRC. MULTI-SENSOR LOGS NON-DESTRUCTIVE MEASUREMENTS OF CORES USING VARIOUS SENSORS. THE TERM NON-DESTRUCTIVE IS USED HERE BECAUSE THE CORE SECTIONS ARE NOT PHYSICALLY MODIFIED. SOME OF THE MEASUREMENTS MADE BY SENSORS MOUNTED ON THE MSCL'S, ALSO KNOWN AS MULIT-SENSOR TRACKS (MST), INCLUDE ELECTRICAL RESISTIVITY, POROSITY, DENSITY, MAGNETIC SUSCEPTIBILITY AND ACOUSTIC VELOCITY. 3	3
FIGURE 3 APPROXIMATELY A FOURTH OF THE LIMNOLOGICAL RESEARCH CENTER'S (LRC) REPOSITORY AT THE UNIVERSITY OF MINNESOTA, TWIN CITIES. CURRENTLY THE REPOSITORY IS HOUSING 10 KM OF CORES, WITH EACH KILOMETER PROVIDING APPROXIMATELY 27 GB OF CORE IMAGERY.	3
FIGURE 4 SIMPLIFIED INTEGRATED OCEAN DRILLING PROGRAM'S WORKFLOW SHOWING TIMES WHEN CORES ARE WHOLE (TOP), AND WHEN CORES ARE SPLIT INTO WORKING (LEFT) AND ARCHIVE (RIGHT) HALVES.	7
FIGURE 5 LIMNOLOGICAL RESEARCH CENTER'S WORKFLOW WITH TIMES WHEN CORES ARE WHOLE (TOP), AND SPLIT INTO WORKING (LEFT) AND ARCHIVE (RIGHT) HALVES.....	8
FIGURE 6 ANTARCTIC DRILLING WORKFLOW WITH TIMES WHEN CORES ARE WHOLE (TOP), AND SPLIT INTO WORKING (LEFT) AND ARCHIVE (RIGHT) HALVES	9
FIGURE 7 AN EXAMPLE LIMNOLOGICAL RESEARCH CENTER BARREL SHEET (COURTESY OF THE LRC), WHICH GIVES AN IDEA OF WHAT IS GENERATED BY THE INITIAL CORE DESCRIPTION (ICD) PROCESS. HERE WE SEE MSCL DATA LOGGED NEXT TO A LOWER RESOLUTION CORE IMAGE. IN THIS BARREL SHEET WE SEE LINES INDICATING STRUCTURAL INTERPRETATIONS OF THE EXPRESSION OF FAULTS ON THE CORE, TOGETHER, WITH ANNOTATIONS ON THE SIDE TO FURTHER DESCRIBE CHANGES IN SEDIMENT COMPOSITION AND TEXTURE.	13
FIGURE 8 A COMPLEX DIAGRAM MADE WITH PSICAT (COURTESY OF CHRONOS) PSICAT IS AN APPLICATION USED FOR VISUAL CORE DESCRIPTION (VCD) AND CREATION OF A BARREL SHEET. THE KEY DIFFERENCE BETWEEN A VISUAL CORE DESCRIPTION AND AN INITIAL CORE DESCRIPTION IS THAT THE CENTRAL EFFORT OF CREATING A VCD REVOLVES AROUND PRODUCING AN ABSTRACT VISUAL REPRESENTATION OF A CORE. THE VARIOUS PATTERNS, AND POLYGONS IN THE CENTER OF THE FIGURE ARE REPRESENTATIVE OF CHANGES IN GRAIN SIZE AND COMPOSITION OF THE SEDIMENTARY SEQUENCE RECOVERED IN THE CORE BEING DESCRIBED.	16
FIGURE 9 MUNSELL COLOR CHARTS USED FOR COLOR-CODING FEATURES IN A CORE. COLOR MATCHING IS OFTEN DONE IN LESS THAN FAVORABLE LIGHTING.....	18
FIGURE 10 APPLICATIONS AND WORKFLOW FEATURE COVERAGE. THREE LEVELS OF FEATURE COVERAGE ARE USED HERE TO INDICATE THAT AN APPLICATION PROVIDES LITTLE OR NO COVERAGE OF A FEATURE (SMALL SQUARE), MAKES MODERATE EFFORT TO COVER A FEATURE (MEDIUM-SIZED SQUARE), OR MAKES A POINT TO FULLY COVER A GIVEN FEATURE (LARGE SQUARE).	24
FIGURE 11 GENERALIZED CORE WORKFLOW WITH POST-DRILLING ACTIVITIES (GREEN), SHOWN AS A GROUP OF NODES FOLLOWING CURATION. THE WORKING AND ARCHIVE HALVES CAN BE INTERCHANGED FOR IMAGING AND INITIAL CORE DESCRIPTION OR VISUAL CORE DESCRIPTION WORK.	27

LIST OF FIGURES (continued)

<u>Figure</u>	<u>Page</u>
FIGURE 12 PROTOTYPE CORENAVIGATOR USING GOOGLE EARTH AND KML TO ACCESS DATA ON CORES. EACH DOT ON THE MAP INDICATES A CORE HOLE, AND DATA ABOUT EACH HOLE IS DISPLAYED BELOW THE MAP AND INSIDE THE MAP (FIGURE 13)	31
FIGURE 13 CLOSE UP OF PROTOTYPE DATA BUBBLE FOR A PARTICULAR DRILL HOLE. DISPLAYED IN THE BUBBLE ARE THE NAME, ELEVATION, DRILLING METHOD AND OTHER DATA RELATED TO THE HOLE.	32
FIGURE 14 A SCREENSHOT OF CORELYZER. WE SEE A CORE SECTION AND TWO GRAPHS PLOTTED ALONGSIDE THE IMAGE. HERE A USER HAS SELECTED A GRAPH AND A VERTICAL DASHED LINE APPEARS TO HELP VISUALLY ALIGN THE PLOTS WITH THE RESPECTIVE LOCATION ON THE IMAGE.	34
FIGURE 15 CORELYZER ARCHITECTURAL LAYERS. THE SCENE GRAPH LAYER IS IN C/C++ FOR RENDERING EFFICIENCY. THE USER INTERFACE AND PLUG-IN MANAGEMENT ARE DONE IN JAVA FOR RAPID DEVELOPMENT. EXAMPLE PLUG-INS ARE PLACED ON TOP OF BOTH THE PLUGIN MANAGER AND APPLICATION OBJECTS BECAUSE THEY CAN MAKE USE OF ANY PUBLIC OBJECT IN CORELYZER.	35
FIGURE 16 TWO VERSIONS OF A LINE CROSSING A VERTICAL SEAM. ONE IS DRAWN WITH CONTAINER SPACE DISRUPTION MITIGATION (LEFT), AND THE OTHER WITHOUT (RIGHT). THE VERSION ON THE LEFT REDUCES THE PERCEPTUAL IMPACT THE VERSION ON THE RIGHT CREATES (MACKINLEY 2004).	37
FIGURE 17 CORELYZER DISPLAY CONFIGURATION DIALOG. USED TO DESCRIBE THE SCREEN LAYOUT, SCREEN DIMENSIONS, AND OTHER CHARACTERISTICS TO HELP MITIGATE CONTAINER BASED DISRUPTION AND PROVIDE LIFE-SIZE VIEWS OF CORE SECTIONS.	38
FIGURE 18 DISPLAY SCALABILITY OF CORELYZER SHOWN HERE RUNNING ON TWO 30" DISPLAY MONITORS, EACH AT 2560 X1600 PIXELS, AND CORLYZER RUNNING ON A LAPTOP.	38
FIGURE 19 AVERAGE TIME TO RENDER WITH MIN AND MAX DEVIATION	40
FIGURE 20 DETAILED LOOK AT AVERAGE TIME TO RENDER WITH MIN AND MAX DEVIATION, FOR BLOCK SIZES 256, 512 AND 512 WITH COMPRESSION	41
FIGURE 21 AN ANNOTATION MARKER PLACED ALONG SIDE A CORE SECTION. GIVEN THE VISUAL FEATURE ON THE CORE SECTION IMAGE, THE ANNOTATION IS PROBABLY RELATED TO THE GAP IN THE CORE.	42
FIGURE 22 AN ANNOTATION DESCRIBING A FEATURE VISIBLE IN THE PREVIOUS FIGURE. HERE WE SEE THAT THE USER HAS CATEGORIZED IT UNDER GEOPHYSICS (LOWER LEFT CORNER). AT THE TOP OF THE WINDOW ARE PRIOR ENTRIES MADE, AND AT THE BOTTOM NEW ENTRIES CAN BE CREATED FOR INSERTION INTO THE DISCUSSION.	43
FIGURE 23 SPLICER WITH SOURCE LOGS (LEFT) USED TO CREATE A COMPOSITE LOG (RIGHT)	47
FIGURE 24 A DIAGRAM DETAILING WHICH POINTS OF THE CORE WORKFLOW INVOLVE INTERACTION WITH CORELYZER AND THE PROTOTYPE WORKFLOW DATABASE. INCOMING DATA FROM IMAGES AND SENSORS (DASHED LINES) AND ANNOTATIONS FROM DISCUSSION (DOTTED LINE) ARE INSERTED INTO THE DATABASE, TO BE VIEWED IN CORELYZER WHEN A CLIENT PLUG-IN IS USED TO DOWNLOAD AND BROWSE EXISTING DATA. PSICAT IS USED FOR VISUAL CORE DESCRIPTION (VCD) AND INTEGRATED INTO CORELYZER VIA A PLUG-IN, TO ENHANCE SCIENTIFIC DISCUSSION.	51
FIGURE 25 NUMERICAL LOG ACCESS FLOWCHART. VIEWING DATA IS LESS DIFFICULT THAN ACQUIRING ACTUAL VALUES, IN ORDER TO PRESERVE THE DATA MORATORIUM.	56

LIST OF ABBREVIATIONS

ANDRILL	Antarctic Geological Drilling Program
ANDRILL-SMO	ANDRILL Science Management Office
ASP	Andrill Session Plug-in
CDEX	Center for Deep Earth Exploration
CLIP	Core Log Integration Platform
CWS	CoreWall Work Suite
DOSECC	Drilling, Observing and Sampling of the Earths Continental Crust
DPI	Dots per inch
DSDP	Deep Sea Drilling Project
FDA	Free Draw Area
GLAD	Global Lakes Drilling
ICD	Initial Core Description
IODP	Integrated Ocean Drilling Program
IODP-TAMU	Integrated Ocean Drilling Program, Science Services at Texas A&M University
IODP-USIO	Integrated Ocean Drilling Program - U.S. Implementing Organization
IRB	Internal Review Board
JAMSTEC	Japanese Agency for Marine-Earth Science and Technology
JNI	Java to Native Interface
JVM	Java Virtual Machine
LacCore	National Lacustrine Core Repository
LIMS	Laboratory Information Management System
LRC	Limnological Research Center
MSCL	Multi-Sensor Core Logger
MST	Multi-Sensor Track
ODP	Ocean Drilling Program

LIST OF ABBREVIATIONS (continued)

PPI	Pixels per inch
PSICAT	Paleontological-Stratigraphic Interval Construction and Analysis Tool
PCP	PSICAT Corelyzer Plug-in
SedDB	Integrated Data Management for Sediment Geochemistry
VCD	Visual Core Description
WDB	Workflow Database

Summary

Geological cores are cylindrical bodies containing natural materials, which are retrieved for the sake of understanding the past climates of the Earth. Whether retrieving cores from the world's oceans, lakes, ice sheets or continents, the workflow for studying geological cores follow a common path. However, because of the cultural differences between various scientific drilling projects, creating a system to work with all types of coring operations can be difficult. This is compounded by the fact that certain data fall under restricted-use moratorium, requiring special access. The contribution of this thesis is:

1. Development of a generalized workflow model that merges the core workflow of the three major scientific drilling organizations. A clear definition of the workflow not only aids in the development of effective software but also enables future IT specialists to quickly understand the fundamental requirements of scientific drilling, which is to recover and analyze cores. This workflow could potentially assist researchers who study and develop workflow technologies (Lei et. al. 1997).
2. Mapping of technology, implemented as the Corelyzer application, to this generalized model that is useful to anyone interested in developing technological solutions for the scientific drilling community studying cores.
3. Validation of the efficacy of Corelyzer for a detailed case study- namely the ANDRILL Program.

4. Discovery of new problems that need to be solved, new adjustments that are needed to the generalized workflow model based on findings from the case study.
5. Presentation of a proposed solution comprised of multiple applications, including Corelyzer, to create the CoreWall Work Suite.

The goal of Corelyzer and the CoreWall Work Suite is to provide a common work environment that can be used to improve the core workflow up to and after the Initial Core Description process, including those activities undertaken during post-drilling operations

1. Introduction

1.1. Background

Geological cores are cylindrical bodies of mud, rocks, fossils or ice that record environmental signals that can be interpreted to help understand Earth's past climate and natural variations in other processes. The physical signals recorded in sediment cores are indicators of prior climates, in some cases dating back before the extinction of the dinosaurs (ODP 2003, Nov) (ODP 2003, Apr). These objects are retrieved from the world's continents, lake floors, ocean floors or ice sheets. One of the most pressing questions in paleo-climatology, or the study of the past climates: are the current observed changes in climate part of a natural cycle of climate change, or have human interactions changed the cycle?

The different geological records (e.g. from oceans, lakes, etc.) can leave different kinds of clues and in different ways. Ocean records tend to preserve indicators of basinal to global changes that integrate processes operation on various spatial and temporal scales. These records, when tied together with records from other areas, can provide a global view of changes in Earth's systems. Records preserved in ocean cores can go back to millions of years to the Jurassic period. The Ocean Drilling Program (ODP) recovered cores that extend beyond the Cretaceous-Tertiary boundary, also known as the KT-boundary, (ODP 2003, Apr) and was used to support theories on the extinction of dinosaurs. While cores retrieved from lakes do not go as far back in time, e.g. only 10 million years were recovered from Lake Tanganyika (Cohen 2003), the information retrieved display a record preserving higher frequency of information about environmental changes. Lacustrine cores that are sent to the Limnological Research

Center (LRC) often contain banding created during seasonal changes, as can be seen in Figure 1.



Figure 1 A digital-line scan of a 1.5-meter core section retrieved from Lake Hvitarvatn, Iceland. Notice the visible, alternating bands of dark and light colored mud. This high frequency banding does not often occur in ocean cores. This particular core section was retrieved approximately 60 meters below the lake floor.

The Antarctic Geological Drilling (ANDRILL) program, which is retrieving sedimentary records from below ice shelves and sea ice in McMudro Sound, is expected to reach 40 million years into the past. Their expected goal is to study the environment and climate of Antarctica before and after the development of the East and West Antarctic ice sheets.

One thing in common among the diverse scientific drilling groups examining cores is that large amounts of data are being recovered to analyze and help study the cores retrieved. Efforts are made to gather data from Multi-Sensor Core Loggers (MSCL) (Figure 2), also known as Multi-Sensor Tracks (MST), that process core sections before and after they are split down the length of the core. These MSCL devices are primarily used to acquire non-destructive measurements along the length of the cores using various sensors. The term non-destructive is used here because the core sections are not physically modified. Some of the measurements acquired by the MSTs provide information on electrical resistivity, porosity, density, and magnetic susceptibility.

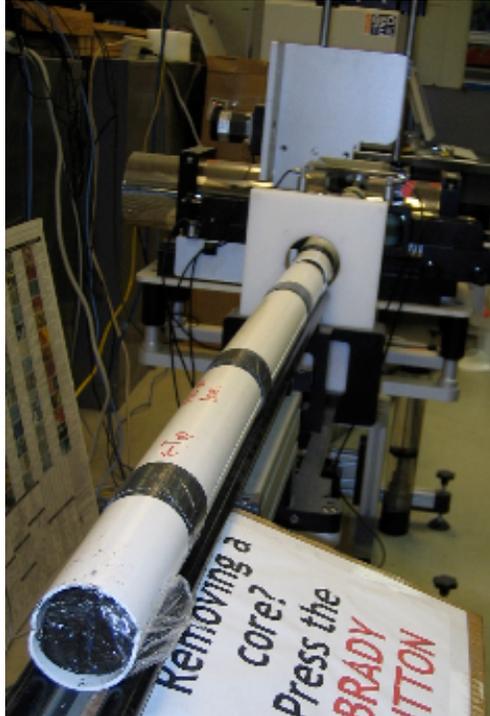


Figure 2 A Multi-Sensor Core Logger (MSCL) at the LRC. Multi-sensor logs non-destructive measurements of cores using various sensors. The term non-destructive is used here because the core sections are not physically modified. Some of the measurements made by sensors mounted on the MSCL's, also known as Multi-sensor Tracks (MST), include electrical resistivity, porosity, density, magnetic susceptibility and acoustic velocity.



Figure 3 Approximately a fourth of the Limnological Research Center's (LRC) repository at the University of Minnesota, Twin Cities. Currently the repository is housing 10 km of cores, with each kilometer providing approximately 27 GB of core imagery.

Another common operation performed over the past several years has been to capture high-resolution, digital line-scan images of the cores. In fact, the amount of imagery created is enough to overstress many commercial and freeware image-related applications, such as Adobe Photoshop. The Limnological Research Center (LRC) houses approximately 10 km of geological cores (Figure 3). The cores are recovered from individual projects that generate tens of meters of core, and the Global Lakes Drilling (GLAD) project that can recover hundreds of meters of core per expedition. The upcoming McMurdo Ice Shelf (MIS) project run by ANDRILL is expecting to recover one kilometer of sediment, most likely to be more lithified than the softer sediment recovered in lake and ocean cores. The Integrated Ocean Drilling Program (IODP), an international successor to the Ocean Drilling Program (ODP) and Deep Sea Drilling Project (DSDP), can recover cores up to the same order as ANDRILL's upcoming project, at each site. Each IODP expedition often includes multiple sites, which can equal thousands of meters of core per expedition. For example, ODP Leg 202 recovered 7 km of core (ODP 2002, Sep). These ocean programs have recovered over 330 kilometers of cores since 1968. All of these cores are stored in repositories that are located in the U.S., Japan and Germany.

Housing the most cores amongst the three organizations, IODP houses hundreds of kilometers of geological cores. IODP has used, and will continue to use and coordinate the operations multiple scientific drilling vessels. The Integrated Ocean Drilling Program U.S. Implementing Organization (IODP-USIO) manages one of the vessels, which is known as the JOIDES Resolution. The Chikyu is the other vessel created by the Japanese component of IODP, the Center for Deep Earth Exploration (CDEX), which is a branch

of the Japanese Agency for Marine-Earth Science and Technology (JAMSTEC). Once operational in 2007, the Chikyu will be able to drill holes in excess of six to seven kilometers. This is enough to drill to the Earth's mantle from the thinnest parts of the ocean crust. Mission specific platforms are used for special expeditions, such as IODP Expedition 302 to the Arctic Ocean, with drilling on Lomonosov Ridge (IODP 2004, Sep).

To understand the magnitude of available data, a calculation can be made as to how much data the cores from a one-kilometer deep hole can create. Numerical data are acquired from MSCL sensors through measurements made of intervals of two centimeters. For 1000 meters of core, each individual physical property being measured could create 50,000 data points, equating to only 1.6 megabytes of floating-point data. The numerical logs are orders of magnitude less in storage size than the image data gathered. The digital line-scan cameras used to create the images of core sections that were split length-wise are often calibrated to create images with a resolution of 100 dots-per-centimeter, or 254 dots-per-inch (DPI). Assuming a total depth of 1000 meters, an average core thickness of 3.5 inches (approximately 9cm) and a resolution of 100 pixels-per-centimeter, we can calculate that a one-kilometer deep hole could create 10,000,000 x 900 pixels of imagery of split-cores, cores that have been split length-wise into two equal halves. If the pixels were assigned values using RGB color space, with unsigned bytes, one kilometer of core can create 27 GB of raw image data. This means that the amount of data from the LRC archives equates to approximately 270 GB of images, for the 10 km of core existing in the LRC repository. The DSDP/ODP/IODP archive equals nearly 9 TB of images.

There does not exist an application that can display so much data in a manner that is capable of helping geologists study cores at this scale, or is able to become an effective part of improving the real-time or asynchronous workflow for studying geological cores. Core workflows include activities during drilling, and after the cores have been curated. In order to understand how to improve the workflow of cores, an understanding of current practices and applications is needed.

1.2. Existing Core Workflows

Once a core has been recovered, by whatever coring means necessary, the core goes through a process, or workflow that results in acquisition of measurements and the curation of the core. Following is a description of the core workflow for the three major science organizations from the recovery of the core to before the post-drilling activity occurs.

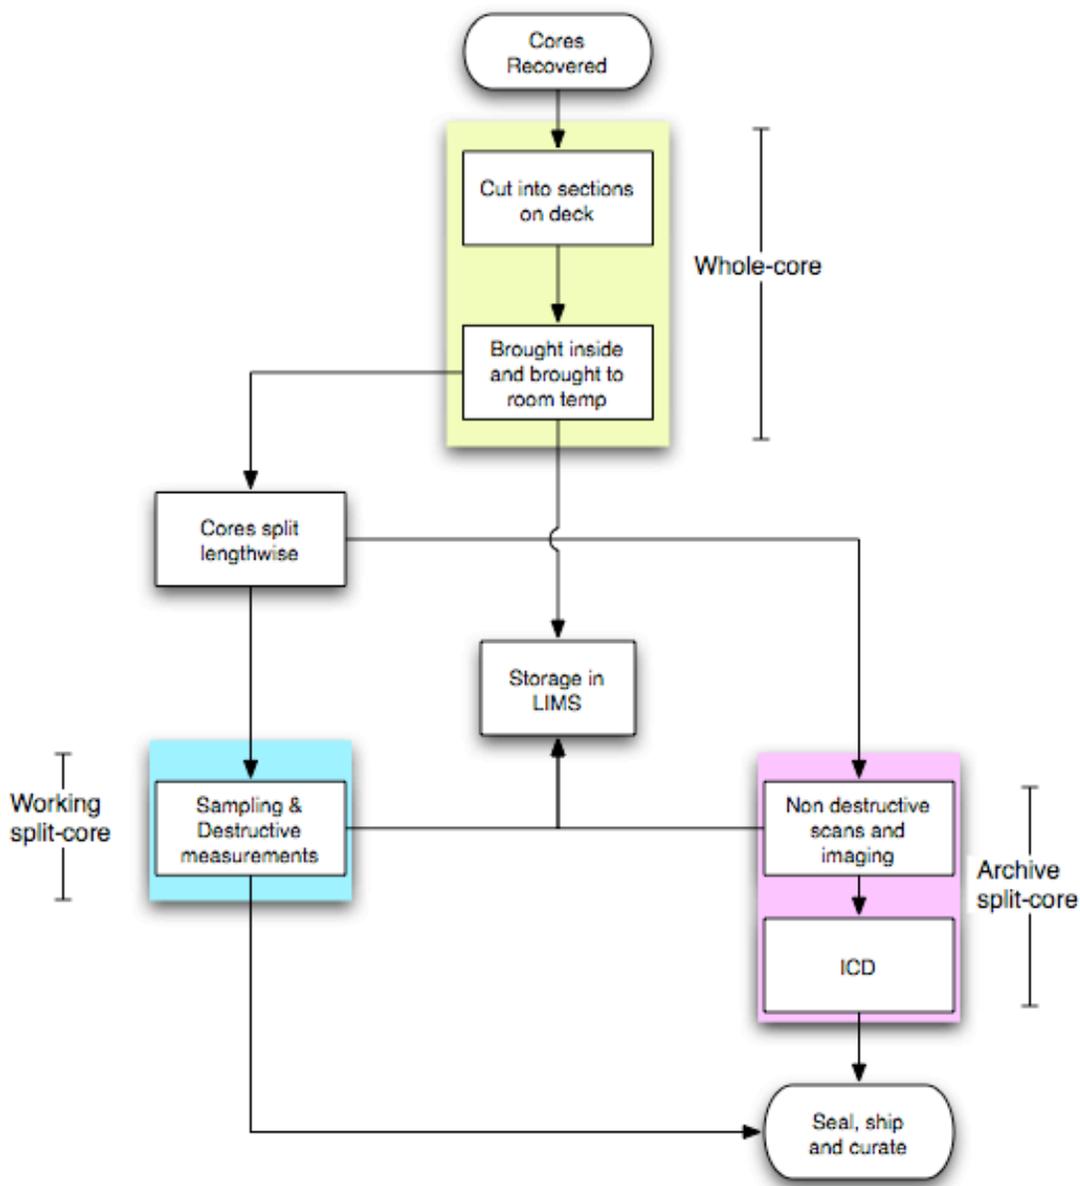


Figure 4 Simplified Integrated Ocean Drilling Program’s Workflow showing times when cores are whole (top), and when cores are split into working (left) and archive (right) halves.

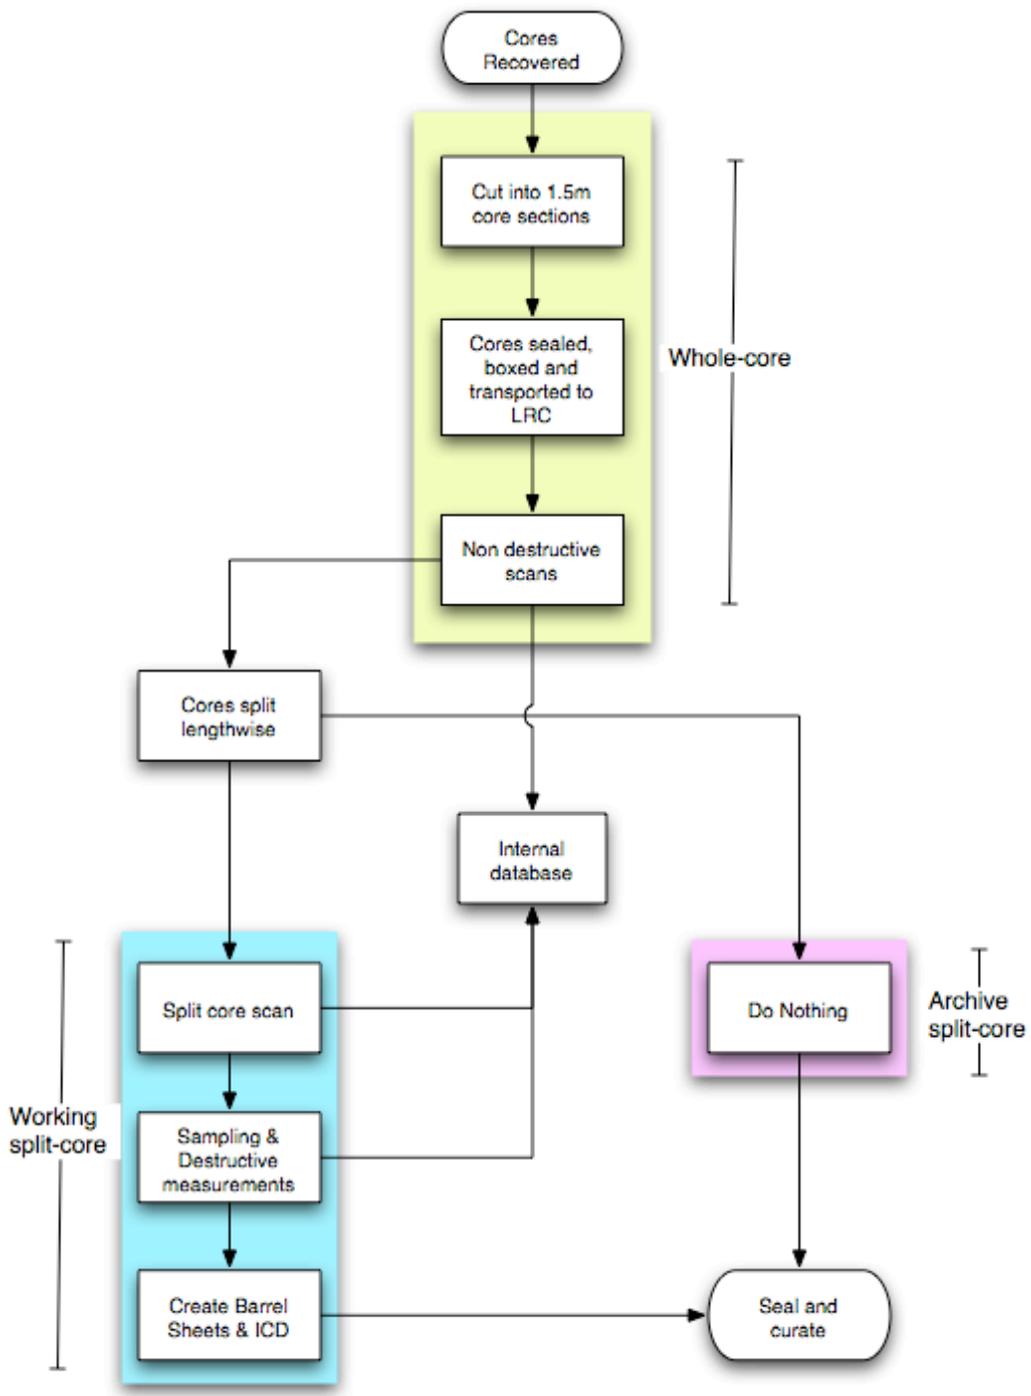


Figure 5 Limnological Research Center’s Workflow with times when cores are whole (top), and split into working (left) and archive (right) halves

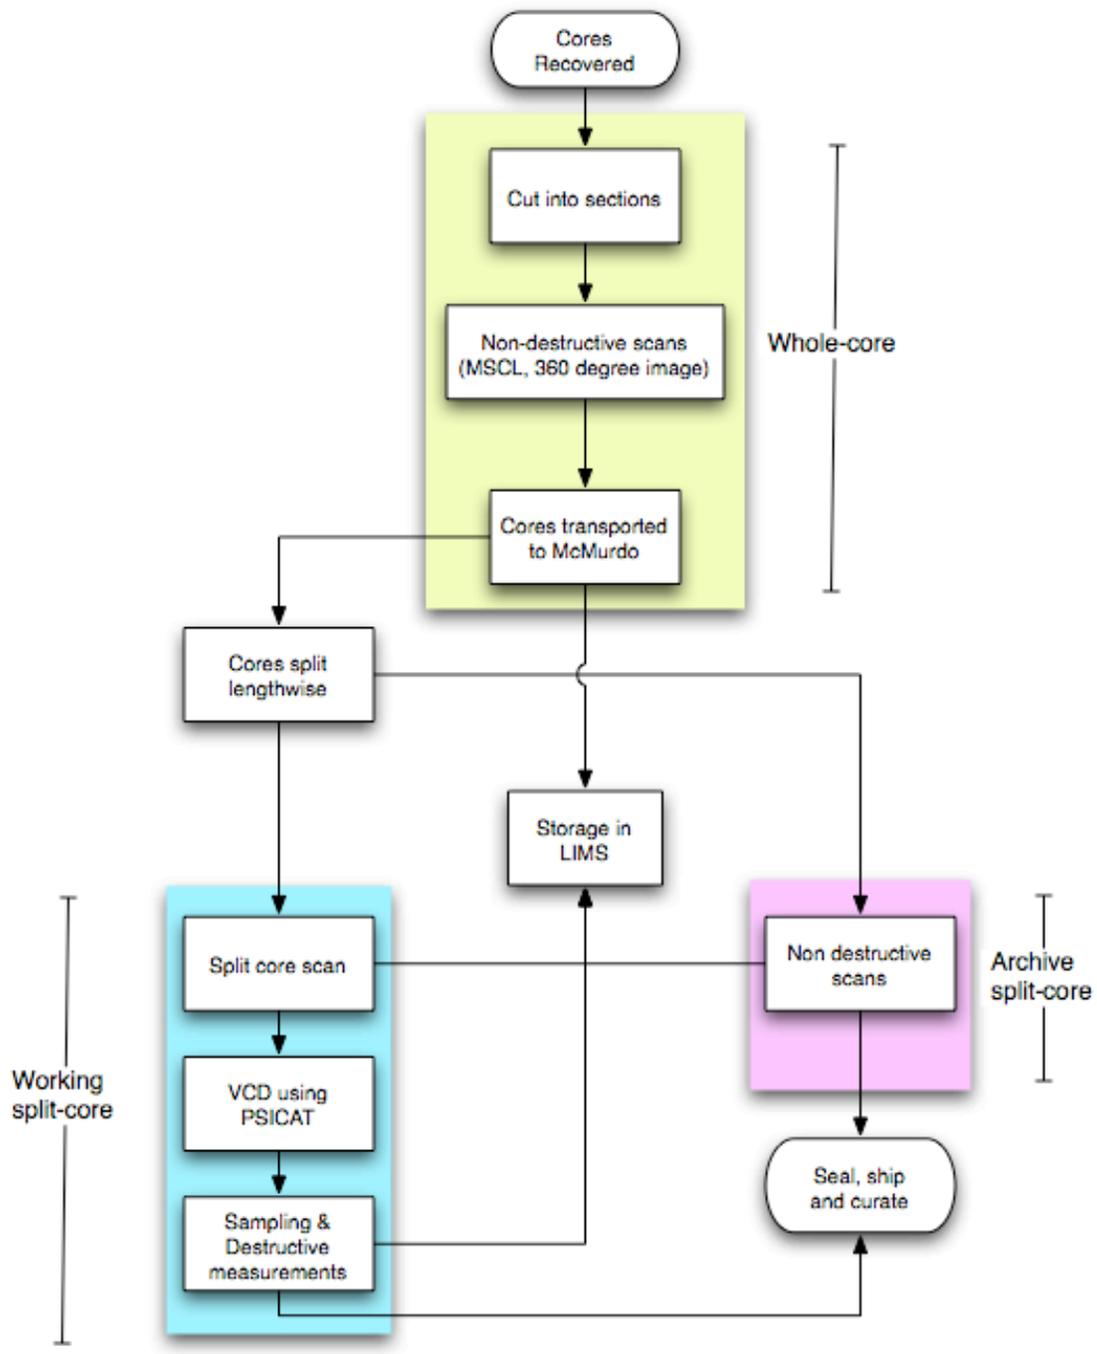


Figure 6 ANTarctic DRILLing Workflow with times when cores are whole (top), and split into working (left) and archive (right) halves

1.2.1. Integrated Ocean Drilling Program's Workflow

Currently, the IODP-USIO is undergoing a transition, with the modification of the U.S. drilling vessel, formerly the JOIDES Resolution. As such, the workflow will potentially undergo alterations based on the final construction of the vessel's science facilities. Individuals in the IODP science team work in 12-hour shifts and may deal with recovery on the order of tens to hundreds of meters of core per shift. As such, many things can go on during 12 hours, which require meetings at shift changes to discuss what happened and compare observations to promote standardization of data acquisition procedures.

When drilling on a typical expedition, a ten-meter core is recovered approximately every thirty minutes. Once recovered, cores are typically cut into seven sections, placed onto racks and brought into the on-board science lab, so that the cores can equilibrate to room temperature. This process can take several hours. By this point the cores are still fully encased in their cylindrical containers (plastic liners), also termed as being whole cores. Once brought to proper temperature, the core sections are sent through the MSCL sensors to acquire non-destructive measurements along the length of the core.

Afterwards, the cores are split length-wise, creating an archive half and a sampling half. Cores that are split in this manner are called split-cores. The archive half of the core is used for additional non-destructive measurements, of magnetic susceptibility (used to indicate concentration of magnetic minerals at a given point in a core), and digital imaging of the split surface for colorimetry and reflectance spectrometry. When an archive half is not being processed by other equipment, the split-

core is usually located at a core description table for Initial Core Description (ICD) by scientists. The ICD process sometimes requires the use of small amounts of material from the core to produce smear slides. Smear slides are microscope slides (glass objects) with small amounts of geologic material smeared onto them, for the purpose of viewing the composition and texture of the material under a microscope. Sometimes smear slides produce images of objects, such as microfossils such as foraminifera or diatoms that are useful in helping determine how to interpret the environmental record preserved by the cores.

The sampling half, also known as the working half, is used for destructive physical properties measurements. The term destructive is used because physical samples are removed from the split-core half and are run through experiments, such as chemical analyses that disaggregate or destroy the sample. These samples are taken at a low resolution, approximately one or two samples per core section, whereas the other measurements made on the archive halves are made at a higher-resolution of approximately one data point every one to two centimeters. While the individual sample processing to gather data occurs, these data are being uploaded into the Laboratory Information Management System (LIMS) for data management and storage during an expedition. When the expedition is completed and the drilling vessel is docked, a permanent on-shore replica of the LIMS database is updated to allow use of the expedition data on a continuous basis. The LIMS is currently used for data management, and at this time is not used to perform Initial Core Descriptions (ICD).

1.2.2. Limnological Research Center's Workflow

Lakes are defined as enclosed bodies of standing water surrounded by land (Cohen 2003). Individuals studying a lake must travel to the lake over land, thus making it unrealistic to build a ship on the scale of the JOIDES Resolution or the Chikyu. This fact alone affects the core workflow of individuals who study lake cores, such as those individuals that are assisted by the Limnological Research Center and National Lacustrine Core Repository (LRC/LacCore).

As seen in Figure 5, the first major difference between studying lake cores and the ocean cores is that the lacustrine cores must be transported to a distant laboratory after acquisition. There can be a long delay, on the order of weeks to months, before the scientists process the cores. Once they have arrived at the repository, the cores go through a workflow that is similar to the one that is performed on board the IODP vessels. However, the cores often arrive at the LRC long before the science crew arrives. Typically, some of the core sections go through the whole process of non-destructive scans, core splitting, more scans and imaging, prior to the arrival of the science crew. This is done to help reduce the amount of time principal investigators have to spend away from their universities, and to help the process of Initial Core Description (ICD) including the creation of barrel sheets. Barrel sheets (Figure 7) are paper or electronic forms that contain core section images at a tenth of the resolution, only a few select numerical logs that follow along the core and space for the scientists to write down notes for core description. The descriptions written on barrel sheets will eventually be used for reports, and as background for analysis and requests for sampling the core in the post-drilling process. In the LRC workflow, the barrel sheets are owned by the science team, and are not currently stored in a database such as the LIMS.

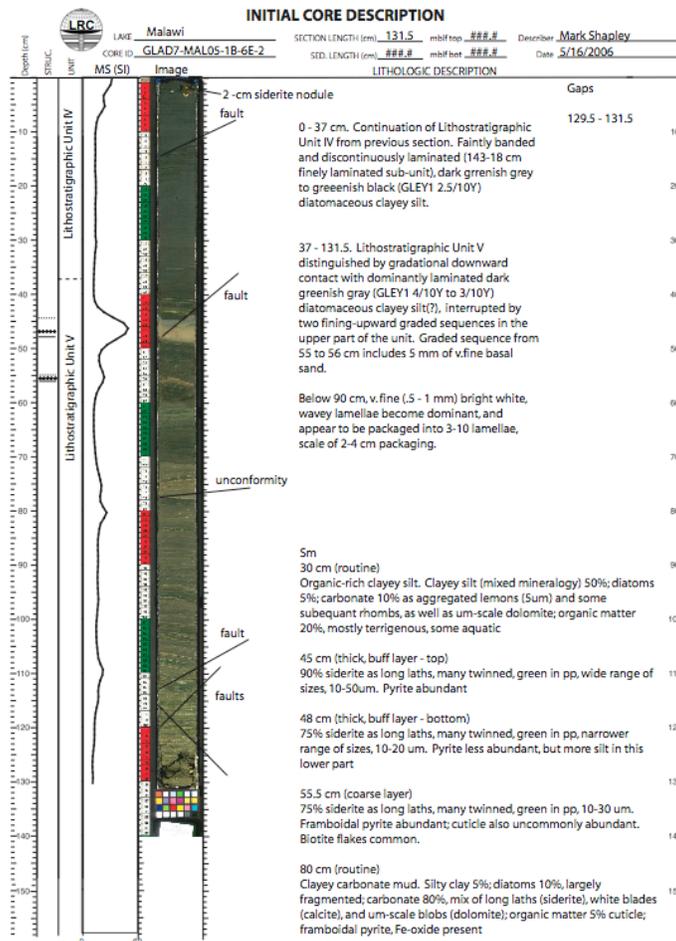


Figure 7 An Example Limnological Research Center barrel sheet (courtesy of the LRC), which gives an idea of what is generated by the Initial Core Description (ICD) process. Here we see MSCL data logged next to a lower resolution core image. In this barrel sheet we see lines indicating structural interpretations of the expression of faults on the core, together, with annotations on the side to further describe changes in sediment composition and texture.

Scientists have typically written down annotations to describe the attributes of the core sections on a sheet of paper. As seen in Figure 7, annotations are written on the right side of the core section image, and lines are drawn to indication geophysical or structural

properties such as interpretations of faults preserved in the cores. As of the summer of 2006, scientists performing the ICD now perform this process electronically by inserting their annotations into a file using an Adobe Illustrator version of the paper barrel sheet. The process of creating the barrel sheets is tedious; requiring the use of Adobe Photoshop to prepare core section images; Microsoft Excel and Sigma Plot to visualize, scale and align the MSCL data; and finally by using Adobe Illustrator to merge the images and logs into a template created to produce the barrel sheet. A workflow difference between the LRC and the IODP workflow is that the split-core imaging and ICDs are done on the working half of the cores rather than on the archive half.

Currently, LRC holds the resulting data in a database that does not provide web access. The primary cause for this structure is due to a lack of I.T. personnel to build and maintain an integrated system. This affects post-drilling operations, and web-based access to data after the data moratorium expires. Future collaborations between IODP and the LRC may provide a means for the LRC to use a LIMS to help store and access data in the future.

1.2.3. Antarctic Geological Drilling Program's Workflow

The workflow for the Antarctic Geological Drilling Program (ANDRILL) is a mixture between the IODP and the LRC workflows. Like the workflow for LRC projects, the core sections need to be transported to McMurdo Station for processing in the Crary Laboratory. However, some measurements and other operations are performed on the whole core sections (i.e. un-split core sections) as soon as possible, at the drill site. These operations include: recovery logging, fracture logging, whole round sampling for microbiology and pore water geochemistry, taking non-destructive Multi-Sensor Core

Logging (MSCL) measurements, and acquiring 360 degree images of specific intervals within a core section.

Once the core sections have been transferred to at the station, within a day or two if weather permits, the typical sets of operations that can be seen in the LRC workflow occur. On ANDRILL's upcoming McMurdo Ice Shelf (MIS) project, the Paleontological-Stratigraphic Interval Construction and Analysis Tool (PSICAT) (CHRONOS 2006) will be in use to perform Visual Core Descriptions (VCD). Visual Core Descriptions are effectively Initial Core Descriptions typically done using an application that allows a user to select annotations visual objects to create an abstract description that includes information about a core's stratigraphy, lithology, geophysical structures, etc. The final output that is similar to a barrel sheet used for ICD work. The key difference between an ICD and a VCD is the inclusion of the visual abstraction of the core. Below (Figure 8) is an example of a VCD created using the PSICAT application. PSICAT is an application created to let users generate VCDs. From this point in the literature, the use of VCD creation will imply the use of a program to aid a user in the ICD process, much like a WYSIWIG editor is used to create web pages.

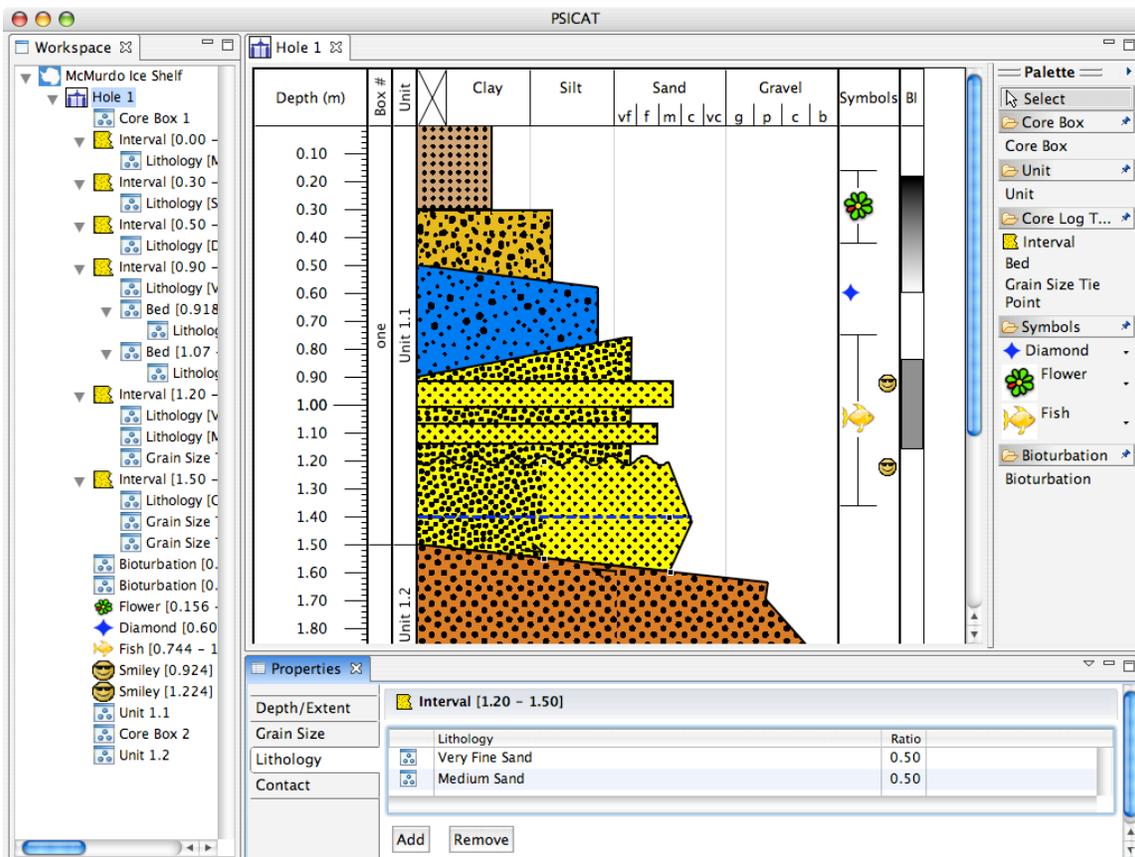


Figure 8 A complex diagram made with PSICAT (Courtesy of CHRONOS) PSICAT is an application used for Visual Core Description (VCD) and creation of a barrel sheet. The key difference between a Visual Core Description and an Initial Core Description is that the central effort of creating a VCD revolves around producing an abstract visual representation of a core. The various patterns, and polygons in the center of the figure are representative of changes in grain size and composition of the sedimentary sequence recovered in the core being described.

Also like the LRC workflow, and unlike the IODP workflow, the VCDs are created when looking at the working half of the core and not the archive half. With the aid of the IODP-TAMU offices, ANDRILL will be making use of a Laboratory Information Management System (LIMS) to store their data on-ice and curing the post drilling phase of the project. Just like the IODP workflow, there are 12-hour shifts and

daily meetings to discuss the past days events and formulate ideas about how the recovered cores should be interpreted.

1.3. Problems with Existing Workflows

There is a recurring problem with all of the existing workflows. In particular, the high-resolution core scans are not being utilized fully by the scientists. The process of color-coding apparent features and intervals on a core section is still done in poor lighting with munsell solid color charts in hand by some groups (Figure 9). Other groups have automated this process using optical sensors and standardized illuminants to provide colorimetry and spectral reflection data in both visible and broader (UV to mid-IR) bandwidths.

Papers forms, called barrel sheets, used for Initial Core Description (ICD) contain core section images at a tenth of the resolution and only a few select numerical logs that follow along the core (Figure 7). All annotations taken down on paper must eventually be transferred manually to digital form to create a final ICD report. Attempts at completely electronic interfaces, such as the use of tablet PCs on the Chikyu vessel coupled with the JCORES data management system, are being made but have deficiencies that prevent them from being practical. A common feature of tablet PCs and Personal Digital Assistants (PDAs) is the ability to recognize hand-written words and sentences. Unfortunately the time taken to perform this task was said to need improvement if regular use of this technique was to become feasible. Other varieties of electronic-based descriptions exist using traditional interfaces, such as LRC's use of barrel sheets created using Adobe Illustrator.



Figure 9 Munsell color charts used for color-coding features in a core. Color matching is often done in less than favorable lighting.

This promotes free form text writing; whereas other stricter data input interfaces cause the ICD process to be more of a classification process. Both options have benefits and undesirable consequences with respect to flexibility and ease of use. Classification based systems need to be continuously updated with new classifications if they hope to prevent becoming outdated, but the benefit of controlling vocabulary and thus preventing a deviation of phrases or names to describe features helps to make descriptions searchable.

Free form entry of descriptions and annotations lies on the opposite end of the spectrum, in that it is incredibly flexible but becomes more difficult to create systems to parse the text. The ability to search through prior preliminary reports and core descriptions is an important part of improving the post-drilling process of analysis of cores, and establishing a systematic core description process.

The description and analysis of the cores does not end after the ICD process. Scientists return to their universities and continue to study the numerical logs and images, and perform analyses to create more refined interpretations of the available data. It has often been said that the high-resolution images are best used when a series of core are printed out, tapped together and rolled out into the middle of a long hallway so that an individual scientist can walk along their length and look for patterns.

Searching through previous descriptions and reports, and laying out meters of core images printed out and spread across the floor, requires large amounts of time for a principle investigator (PI). The fact that the data gathered during drilling falls under a limited distribution moratorium does not help the PI's situation. For example, lacustrine cores currently fall under a two-year moratorium starting from the time that the core sections are recovered from the lake floor. This is a policy that is currently being reviewed and reconsidered because of time considerations. For example, from the time of recovery, the cores may take a couple months before they arrive to the LRC for pre-ICD processing and another couple weeks after that for the scientists to arrive; therefore, any extra time spent performing inefficient processes adds to the lag before the PI can begin publishing results. The post-drilling procedure for studying cores is in need for a tool to improve the process.

2. A New Computational Approach to Core Workflows

2.1. Categories for Improving Core Workflow

There is a need for a new software program that can go beyond the capabilities of existing applications. The key attributes for improving the existing core workflow are:

- **Data Scalability:** The ability to view large amounts of high-resolution images and numerical data on the order of gigabytes
- **Display Scalability:** The ability to take advantage of multiple high-resolution displays to display as much of the high-resolution imagery as possible, giving more contextual information.
- **Interactivity:** The ability to interactively browse the data from anywhere with respectable frame rates (approx 10 frames-per-second).
- **Annotations:** A method of keeping track of the analyses and interpretations of cores studied by many researchers.
- **Data Accessibility:** The ease at which first-order data (e.g. numerical logs, images) and second-order data (post-drilling analysis) can be viewed, and managed.
- **Extensibility:** An environment that will balance between speed and the ability to extend applications to include more capabilities that are created with the help of outside developers
- **Multiplatform:** The same program should be usable by a vast majority of users who will have different operating system preferences.

Corelyzer, combined with the whole CoreWall Work Suite (CWS) intends to meet the needs defined by the above categories and more. The entire CWS has not been fully implemented, but Corelyzer has reached a point that is usable by the community and is being implemented in core workflows.

2.2. Comparison to Existing Applications

Several existing applications attempt to help scientists perform core logging and description of cores. AppleCORE is one of the few geological logging programs available for the Macintosh, designed for use in describing subsurface cores. Created in the 1990's before Mac OS X, it has been ported but new features haven't been added. This program provided a logging program on the Mac platform that used standard patterns for displaying lithology and other geological structures. This application was abandoned by broad community of geoscientists because of its lack of flexibility with respect to rock types, sedimentary structures, features, etc. The target audience for the development of the AppleCORE application was the Ocean Drilling Program (ODP), IODP's precursor, and has since become outdated. The application had no features for communicating with outside databases, and had no concept of integrating high-resolution images of cores. Since the practice of gathering these images has become more commonplace, after the program was abandoned.

After struggling with use of Corel Draw for core description during the Cape Roberts Project in Antarctica, the ANDRILL team contacted the CHRONOS group at Iowa State University to develop the Paleontological Stratigraphic Interval Construction and Analysis Tool (PSICAT). PSICAT was developed with the goal to create visuals similar to AppleCORE but to be more flexible and user friendly. The application is stand-

alone, and does not rely on an external database. While PSICAT can upload high-resolution images, because it is written purely in Java it suffers some additional weaknesses, which include the time required to load images and the limit the in number of images that can be loaded because of restrictions due to Java heap sizes. The default maximum heap size is the smaller of one quarter the physical memory or one gigabyte (Sun). This can be made larger with command-line arguments to the Java Virtual Machine (JVM), but PSICAT is still limited to the amount of images that can be fit within the Java heap, preventing it from being able to scale to the amount of data that can be displayed in Corelyzer.

JCORES, another project developed using 100% pure Java, was created for CDEX/JAMSTEC, the Japanese arm of the IODP, and its vessel the Chikyu. JCORES is a very complete package, with many features needed for performing analysis of cores during an expedition. Being a pure Java application, JCORES modules have the same image data scalability and load time issues as PSICAT, and the developers have decided to use scaled down versions of the images for their Visual Core Description (VCD) module. JCORES heavily relies on the use of an ORACLE® database that is designed around the ODP JANUS data model, previously created at IODP-TAMU. This reliance on the database is both its strength and weakness, as the load times for core section images are also exacerbated due to necessary time spent communicating with the ORACLE® database. User feedback from trial runs indicated that using the VCD package seemed to be a tedious task as the interface was complex and had long response times.

The LRC uses several sets of programs (Microsoft® Excel, SigmaPlot, Adobe® Illustrator, etc.) to create a single sheet of paper for individuals to perform the Initial Core Description (ICD). The process alone can consume large amounts of time on some of the large lake drilling projects, such as the Drilling, Observing and Sampling of the Earths Continental Crust (DOSECC) Projects that make use of the Global Lakes Drilling system (GLAD) (DOSECC 2006). This ICD process is arduous and is in need of a stable replacement.

	AppleCORE	PSICAT	JCORES	Corelyzer
Data Scalability				
Display Scalability				
Interactivity				
Annotation				
Data Accessibility				
Extensibility				
Multiplatform				

Figure 10 Applications and workflow feature coverage. Three levels of feature coverage are used here to indicate that an application provides little or no coverage of a feature (small square), makes moderate effort to cover a feature (medium-sized square), or makes a point to fully cover a given feature (large square).

While the above applications are able to excel at core description with the use of standard patterns and basic annotative capabilities, these programs lack the ability to scale to large amounts of images and are mostly fixed in their nature. Corelyzer intends to meet the majority of the emerging community needs, and provide an environment useful for syn-drilling and post-drilling activities. Corelyzer will handle large amounts of numerical data and image data, scale to multiple displays seamlessly, and provide an interactive means to browse and annotate core sections. With its extensibility, Corelyzer can provide a common, effective work environment that can bring in data from heterogeneous databases by creating plug-ins with user-interfaces designed to work with a given database or web-service. In combination with other tools comprising the CoreWall Work Suite, Corelyzer will be able to help access and manage individual studies of geological cores.

3. Embedding the CoreWall Work Suite into a Generalized Workflow

In order to understand the applications of the CoreWall Work Suite, we must first see view a generalized version of the existing workflows. This will illustrate of the general functions, or goals of the applications. In Figure 9, a generalized workflow is diagrammed including a loop that generally describes the post-drilling activities. Most of the workflow is based on a linear process until the post-drilling activities are invoked. The loop of the post-drilling activity would be a task (Lei et. al. 1997) that takes the data created by the linear tasks of the syn-drilling operations to create new data and interpretations of data. Part of the post-drilling task includes upholding rules of data moratorium, indicated by the decision box concerning access privileges. Typically, if data is under moratorium, then the main route to access the data is to contact the core or data

curators, or have the curators send the data electronically if more has been acquired. This system works well in the beginning, for smaller operations and for principle investigators who have data access privileges and who go on to publish preliminary reports, journal articles, etc. Once the data is out of moratorium, the data becomes hosted publicly. At this point the general populace have data access privileges, can perform analysis and collaborate openly. This point of the core workflow can benefit from collaborative tools that are designed to fit the workflow.

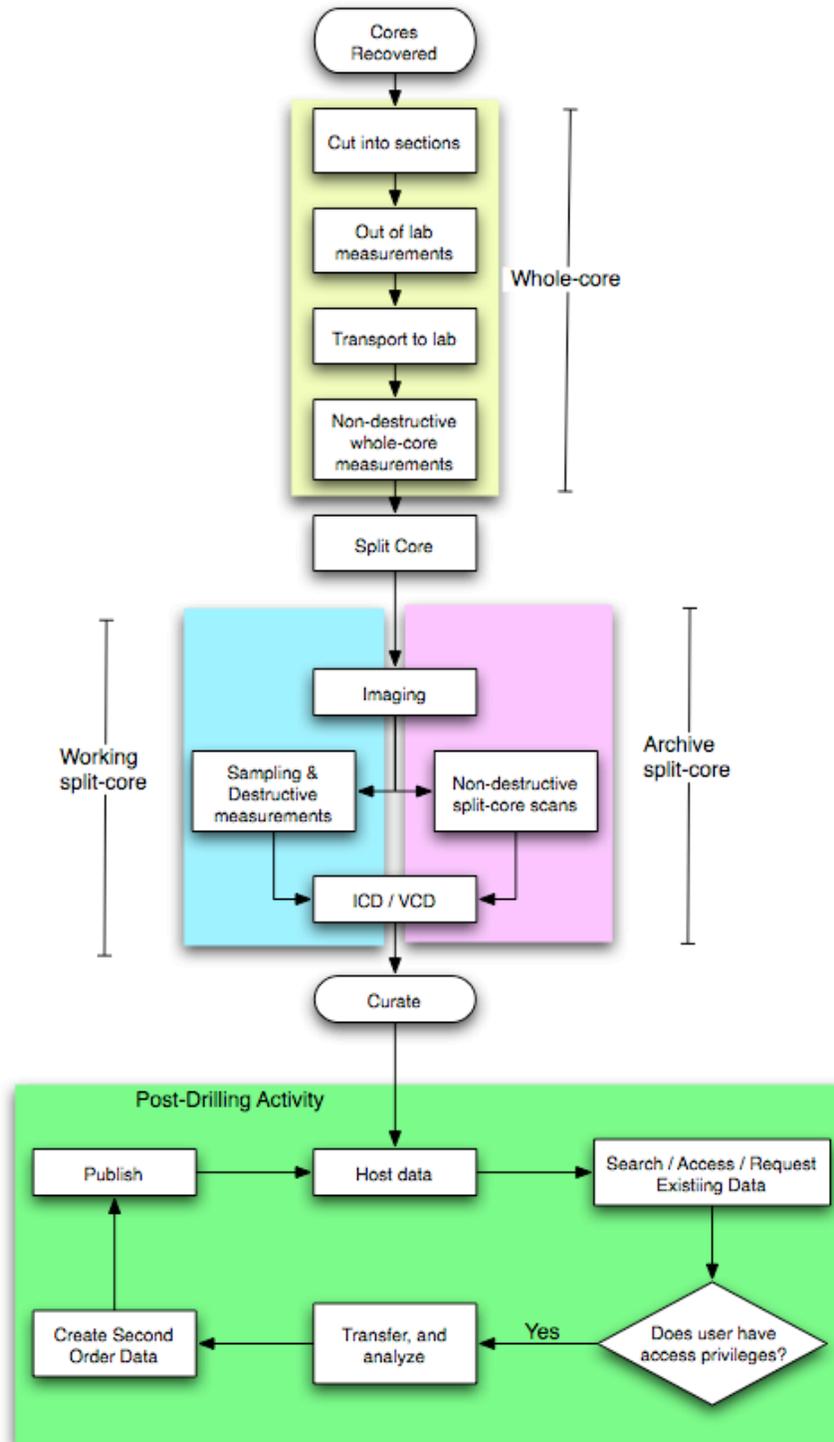


Figure 11 Generalized Core Workflow with Post-Drilling Activities (green), shown as a group of nodes following curation. The working and archive halves can be interchanged for imaging and Initial Core Description or Visual Core Description work.

Within the CoreWall Work Suite (CWS), applications are in development to discover, analyze, and manage data concerning geological cores. There are 4 parts of the CWS and each has a different visual/data integration or exploration function. The following sections will describe the purpose of each application, according to the initial proposal for developing and implementing the CWS.

CoreNavigator- A data discovery tool that has the ability to browse data using a 2- or 3D “GIS-like” (geographic information system) context. This tool provides a visual way of comprehending cored stratigraphic datasets, integrated with seismic and overlaying oceanographic data where these exists. This tool will be particularly useful for individuals to “discover” data from past drilling activities. Within the workflow, this tool is intended to primarily be used in searching, and requesting of data in the post-drilling or pre-drilling proposal development tasks. This tool is intended to be a means to bring data into Corelyzer.

Corelyzer- Designed to be used during ICD/VCD and second order data creation, this is the primary visual integration workspace that allows users to view depth-registered data and images from any compatible databases. Corelyzer is easily expandable through the development of custom plug-ins. Because individual user communities will always want to integrate new functionality and features into the CoreWall Suite, a plug-in structure has been developed that allows programmers to extend the functionality of the base software on their own. Corelyzer can also pull data from the Web,

which was demonstrated at the CoreWall Workshop (Washington, DC; May 2006) with a CHRONOS Web Data Portal plug-in. In the workflow, annotative capabilities in Corelyzer could be used for ICD/VCD work, syn-drilling operations. For pre-or-post drilling tasks, Corelyzer can retrieve data from known locations, such as those discovered using Core Navigator, and make use of plug-ins to perform analysis, create second order data.

CoreCLIP (Core-Log-Integration-Platform)- The revised SPLICER (stratigraphic composite builder) and SAGAN (core-log mapping) software that was previously two separate UNIX applications used mainly by ODP. CoreCLIP will provide a standalone application and a Corelyzer-ready plug-in with many new features, such as the ability to use images to assist users in building composite logs, an integrated text parser to improve data entry, and add new depth correlation algorithms to improve composites. This tool would be used syn-drilling operations to help make drilling decisions after preliminary lab measurements (e.g. MSCL) are taken. Another point of use would be in post-drilling tasks to create new interpretations of composite cores, which would become second order data. The second order data could be used to correlate among geographically separated sites or be analyzed as time series of environmental data with addition on an appropriate age model.

Workflow Database- A local working database(s) that will be designed to allow for local interaction with data, interaction with web services and other databases, data synchronization between multiple CoreWall setups/databases, and for remote collaborations where data/images are being shared. To help with on-site and off-site collaboration using Corelyzer, a workflow database would be used to help manage communication. This tool would fit entirely within both syn-drilling and post-drilling tasks.

3.1. CoreNavigator and Data Discovery

For individuals that are new to the techniques of studying geologic cores, data discovery and searching is one of the first actions taken. Individuals need information about any previous cores that have been recovered, where they come from, how many cores at a given location, why and when they were gathered, etc. Searching for this information can be difficult for individuals that don't know where to go to access these data. CoreNavigator is an application in the CoreWall Work Suite that is designed to meet this need. By taking advantage of existing technology, such as Google Earth Enterprise (Google 2006) and the Keyhole Markup Language (Google 2006), an integrated web-services oriented viewer can be developed to provide an interactive environment, rich with context about the cores available.

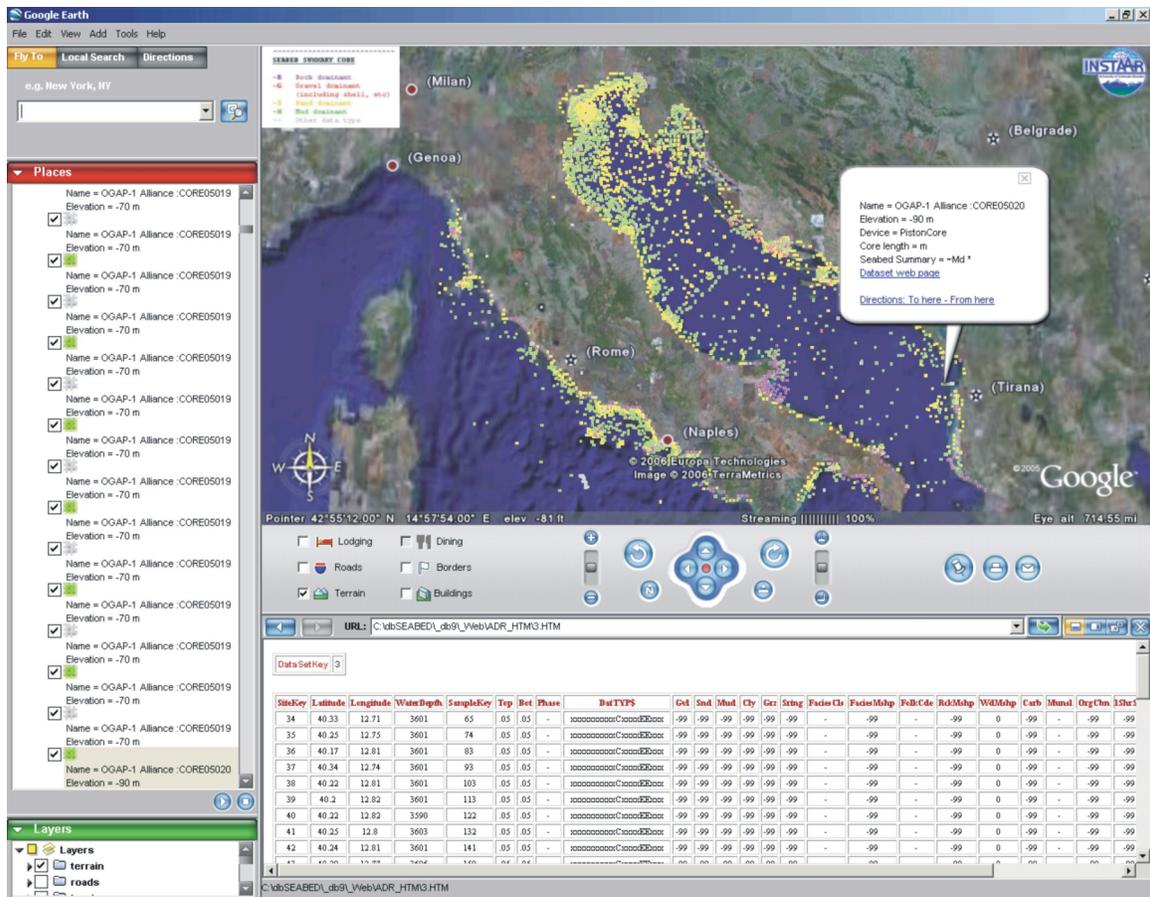


Figure 12 Prototype CoreNavigator using Google Earth and KML to access data on cores. Each dot on the map indicates a core hole, and data about each hole is displayed below the map and inside the map (Figure 13)

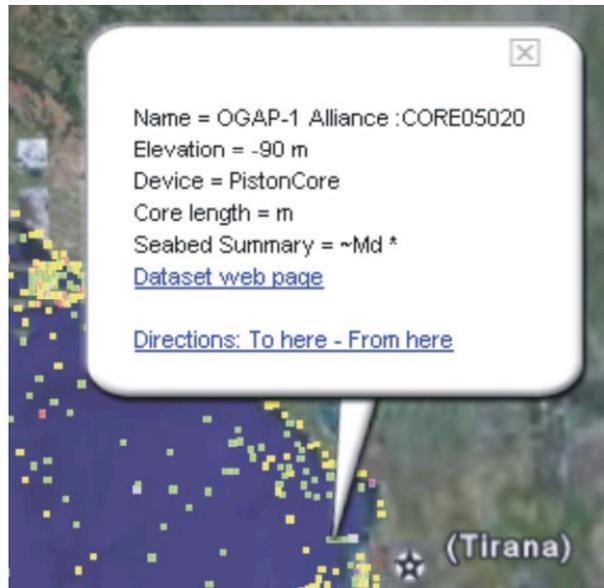


Figure 13 Close up of prototype data bubble for a particular drill hole. Displayed in the bubble are the name, elevation, drilling method and other data related to the hole.

Preliminary work by Chris Jenkins of the Institute of Arctic and Alpine Research (INSTAAR), at the University of Colorado at Boulder, have taken the data accessible by dbSEABED (Jenkins 2005) and integrated the data points into the Google Earth application (Figure 11, 12). The second goal of the CoreNavigator component is to take the set of cores of interest (COI) and create an input file to load into the Corelyzer application.

3.2. Corelyzer, the Central Environment

Corelyzer, is the central application in the CoreWall Work Suite and the author's contribution to the CWS. The aim of the application is to be a central environment to browse, annotate and analyze geological cores. Once an individual has acquired core

section images and numerical logs, Corelyzer can be used. This application fits in multiple locations in the generalized core description and analysis workflow, but is designed to be part of any ICD/VCD work and to create second order data. Again the goal of Corelyzer is to meet the previously stated points for improving core workflow: display scalability, data scalability, interactivity, the ability to annotate, the ability to extend capabilities, the ability to access data from various places and to support multiple operating system platforms.

Designed to be scalable with respect to displays and data, Corelyzer uses OpenGL to hardware accelerate rendering of core images and data, in order to keep the application interactive. Java is the front-end language of the application, used to make Graphical User Interfaces (GUIs) and a plug-in system. Using Java and OpenGL, Corelyzer has successfully run on multiple operating systems. The user-interface is designed to be intuitive with a focus on a simple set of interaction schemes. Plug-ins allow outside developers to create modules for the application to suite their respective needs. This also allows new features to be created, tested and incorporated, such as to access data from various databases. The following sections will describe, in detail, how Corelyzer satisfies the goals to improve workflow.

3.2.1. Corelyzer Architecture

A major focus of Corelyzer is to be able to display vast amounts of numerical datasets generated from measurements of core sections juxtaposed with high-resolution digital line scan images whole-round or split surface of the core sections.

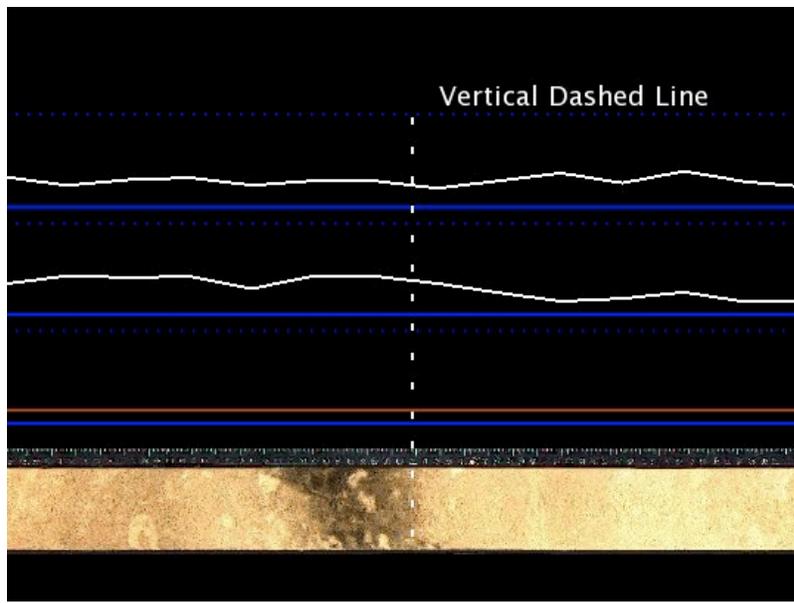


Figure 14 A screenshot of Corelyzer. We see a core section and two graphs plotted alongside the image. Here a user has selected a graph and a vertical dashed line appears to help visually align the plots with the respective location on the image.

The architecture of Corelyzer involves mixing Java and native C/C++ at different levels of the application (Figure 15). This design choice was made to rapidly develop user interfaces, keep time taken to perform rendering with OpenGL to a minimum and be able to make use of the full memory heap given by the user's operating system for a process. Java bindings for OpenGL (JOGL) is an open-source development to allow Java programmers to using natively implemented OpenGL functions. Corelyzer uses this technology to create OpenGL contexts, but has separate native methods implemented to effectively manage data, images, other visual objects for rendering. These native methods are amalgamated into the Corelyzer's Scene Graph library.

The use of Java also helps in creating an effective plug-in system. Plug-ins are modules created outside the Corelyzer source code, which can be made to not only

process incoming data, but also create new user interfaces to create events and manage user input.

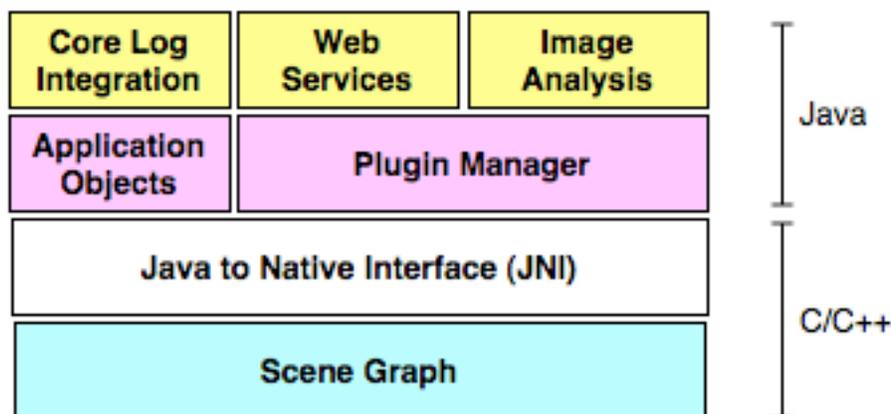


Figure 15 Corelyzer architectural layers. The Scene Graph layer is in C/C++ for rendering efficiency. The user interface and plug-in management are done in Java for rapid development. Example plug-ins are placed on top of both the plugin manager and application objects because they can make use of any public object in Corelyzer.

Development of Corelyzer has focused on supporting multiple platforms, and has been successfully built for Microsoft Windows, Mac OS X, and several Linux distributions.

3.2.2. How Corelyzer Meets Display Scalability

Unlike studies of the use of multiple monitors with multiple windows and applications being run (Hutchings et. al. 2004), Corelyzer is a single application that

spans across multiple monitors. Multi-screen display systems have become a central focus for some institutions because those systems can facilitate information richness in applications, and they provide more interesting problems to solve. While other work has focused on using clusters to create and display visualizations (Jeong et. al, 2005) (Krishnaprasad 2004), Corelyzer focuses on scaling to a multi-screen display driven by a single computer. The main reason for this is to simplify infrastructure needs so that a scientist can begin performing analysis right away. However, future developments such as the LambdaTable (Krumbholz et. al, 2005) can provide an environment to promote collaboration and build community interactions using such a system.

Prior work has shown that when using multiple screens, developers must be aware that some visualizations benefit from removing the appearance of disjointed visuals due to screen border seams, termed “mullions.” Two types of visual disruptions created by seams have been defined (Mackinlay et. al. 2004). One is based on the space used as a container of graphical objects, the other used "as a metric field to position objects meaningfully with respect to quantitative axes." Container space disruption mitigation involves giving the appearance that a graphical object physically continues, undeterred by the mullions (Figure 16). Given that the predominant visual objects in Corelyzer are the high-resolution images, container based disruption mitigation has been implemented. This reduces the perceptual impact that seams create (Mackinley 2004). The same method for mitigating visual disruptions has also been implemented in the Scalable Adaptive Graphics Environment (SAGE) (Jeong et. al. 2005). In order to provide a simple way of customizing the display configuration, so that Corelyzer can automatically calculate the container based disruption mitigation, Corelyzer employs a user interface

that allows a user to customize the display layout (Figure 17), for example, to use Corelyzer display data spanning two 30-inch displays (Figure 18).

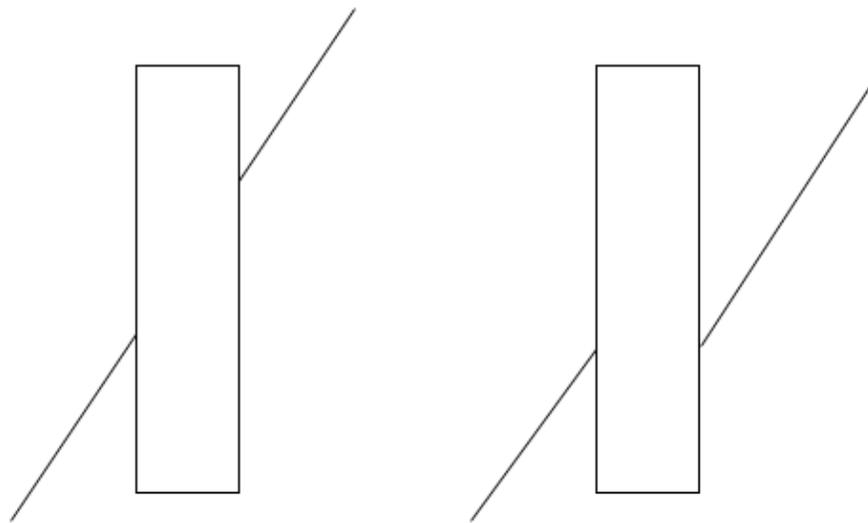


Figure 16 Two versions of a line crossing a vertical seam. One is drawn with container space disruption mitigation (left), and the other without (right). The version on the left reduces the perceptual impact the version on the right creates (Mackinley 2004).

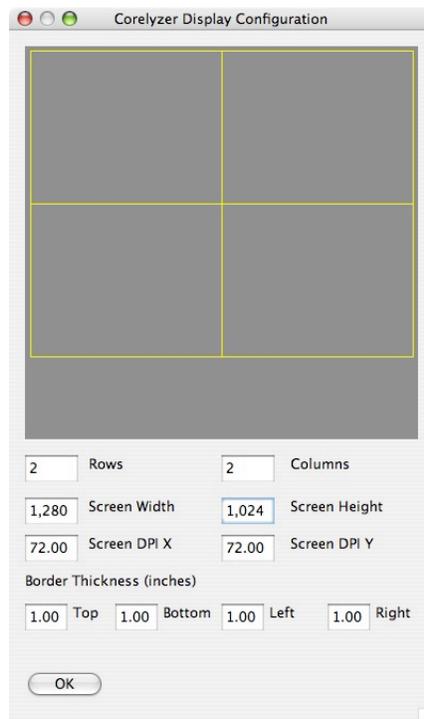


Figure 17 Corelyzer display configuration dialog. Used to describe the screen layout, screen dimensions, and other characteristics to help mitigate container based disruption and provide life-size views of core sections.



Figure 18 Display scalability of Corelyzer shown here running on two 30" display monitors, each at 2560 x1600 pixels, and Corlyzer running on a laptop.

3.2.3. How Corelyzer Meets Image Data Scalability

Hardware accelerated graphics is commonplace with modern desktop and laptop computers. Using hardware accelerated graphics allows Corelyzer to take advantage of the large texture memory that exist on graphics cards (approximately 64 MB to 256 MB at the time of this writing). Using textures allows less data to be unnecessarily transferred over the system bus when rendering an image over multiple frames. Aimed to lower disk usage to view larger amounts of data, an incomplete mip-map pyramid is used for level-of-detail control. Corelyzer uses a quad-tree structure at each level of detail to make use of an application-level Least Recently Used (LRU) texture paging system to reduce transfer of texture blocks from disk to main memory, and the texture paging from main memory to video memory that is built into the OpenGL drivers. The use of the Scene Graph layer of the Corelyzer architecture (Figure 15) was used not only to reduce the overhead of moving from Java to C/C++ for every OpenGL call, but also to be able to more flexibly use the full heap size given by the user's operating system. As explained in section 2.2, the JVM typically does not make full use of a heap available to an application.

After defining the overall architecture, decisions were made on how to keep the rendering moderately efficient. In particular, what texture block sizes needed to be used for the LRU texture paging system. Experimentation was done to determine the optimal texture block size to use. Four dimensions were compared fairly: 64 x 64, 128 x 128, 256 x 256 and 512 x 512. The scenario consisted of 21.8 meters of core imagery. Each image had a height of 1000 pixels and the horizontal resolutions of the images were 254 DPI, which effectively gives 218,000 pixels horizontally and 623 MB of data. The experiment was constrained to display the images without any scaling aside from scaling the images

to fit to the depth scale along the bottom of the visualization. Figure 19 shows that as the square block dimension grew the time taken to render a frame dropped overall.

With the knowledge that the S3TC/DXT texture compression OpenGL extension has been specified in November 2001 (SGI 2001) and implemented on every major graphics chipset since then, the experiment was re-run using 512 x 512 compressed textures.

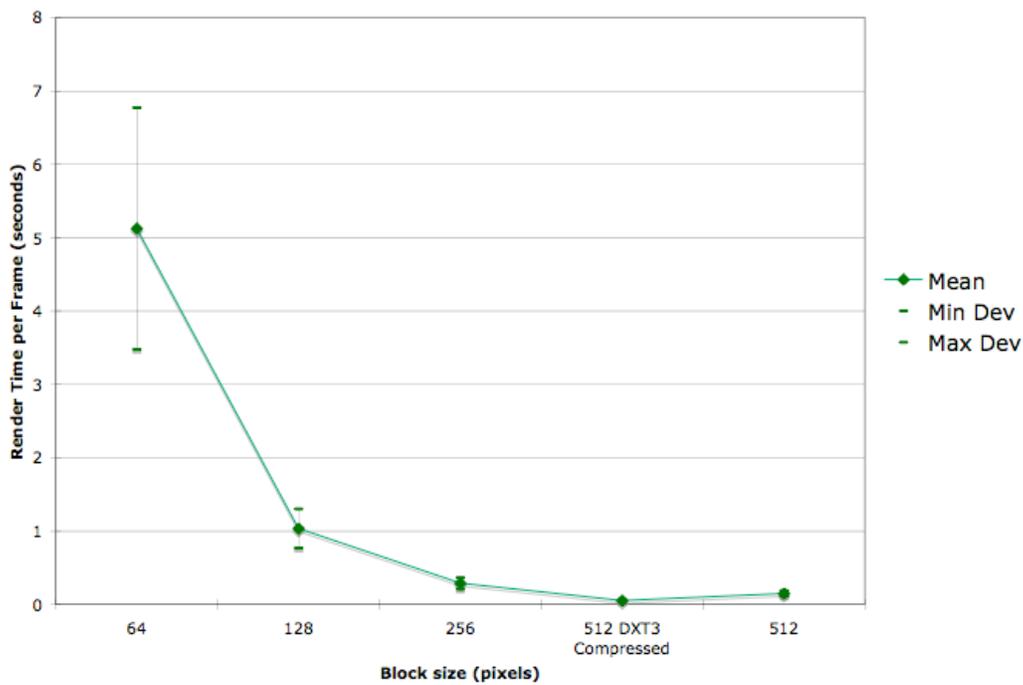


Figure 19 Average Time to Render with Min and Max Deviation

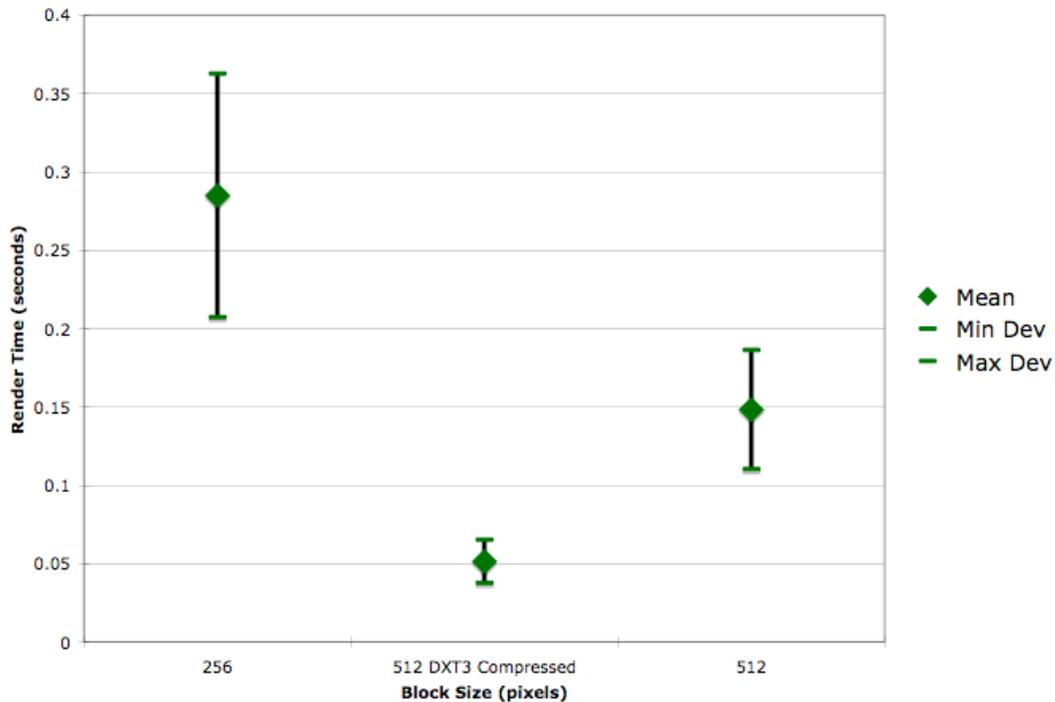


Figure 20 Detailed look at Average Time to Render with Min and Max Deviation, for block sizes 256, 512 and 512 with compression

The experiment showed that not only did average rendering time decrease, but so did the minimum and maximum deviations. From Figure 20, we surmise that an average frame-rate of Corelyzer when simply browsing is approximately 20 frames-per-second (FPS), if we use compressed texture blocks of dimensions 512 x 512 pixels.

3.2.4. Annotations

Being able to annotate a core is an important feature because this is the key method for creating ICDs. Corelyzer ties annotations directly to a depth relative to the top of a core section. This information, tied with the sub-surface depth of a core section can give the depth of the annotation with respect to the top of the borehole.

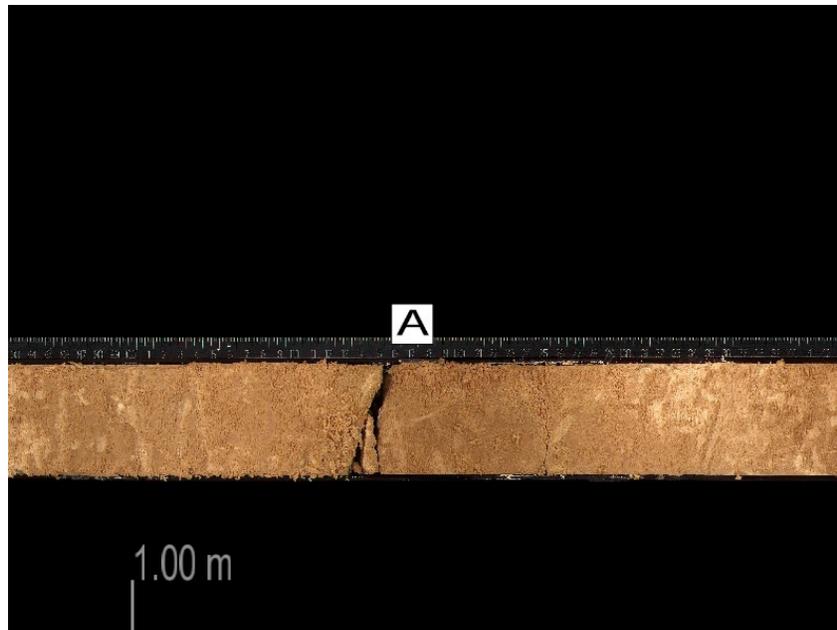


Figure 21 An annotation marker placed along side a core section. Given the visual feature on the core section image, the annotation is probably related to the gap in the core.



Figure 22 An annotation describing a feature visible in the previous figure. Here we see that the user has categorized it under geophysics (lower left corner). At the top of the window are prior entries made, and at the bottom new entries can be created for insertion into the discussion.

Above in Figure 21 and Figure 22, we have a view of an annotation marker that indicates to a user that an annotation has been placed along that depth in the section, and what type of annotation it is. Different markers indicate different annotation types, and in this case we have an annotation that refers specifically to a point at that depth. Once a user selects

the annotation they can continue a thread of thought and discussion with themselves or others (Figure 22), simply by adding a new entry into the annotation. This is similar to the anchoring of textual annotations on web documents (Brush et. al. 2001). Whereas work by Olsen et. al., in 2004, had annotations with a more free form construct based on the use of virtual pens. While extremely flexible, storing handwritten notes can become illegible and Olsen's work was not tested with multiple users. Thus a decision was made to use HTML 1.1 documents to store annotations as text and images. Now individuals can create annotations with references to either images captured through a microscope, text or provide links to other related work in order to effectively describe a core.

This method of creating annotations combined with the extensibility of Corelyzer, discussed in the following section, is used to allow synchronous and asynchronous discussions to take place as cores are recovered by the ANDRILL MIS project. With the same extensibility, Corelyzer can be modified to used structured annotations (i.e. annotations that follow a syntax) so that they can be parsed, validated against a dictionary, and searchable.

3.2.5. User Interaction

Corelyzer has two main interaction schemes. The first is the traditional menu based user interface that users of many systems are accustomed to (i.e. menu bars with drop down menu items). The second interaction scheme is used to navigate through the visualization area. Using a scroll-wheel mouse, or laptop trackpad, individuals can pan, zoom and interact with objects in the visualization windows. Both left and right mouse button events are context sensitive and take into account modifier keys. For instance,

holding down the left mouse button on an unoccupied space in the virtual world will automatically place the application in a click-and-drag style of panning. Where as left mouse button presses over a core section will allow the user to slide the core section image horizontally along the designated horizontal track that the image belongs to.

3.2.6. Corelyzer and Extensibility

After an NSF workshop focused on CoreWall took place in May 2006, community response was positive to a proposal to create an extensible environment. What was shown was a Corelyzer plug-in that used web-services developed at CHRONOS to access IODP-TAMU's JANUS database. The plug-in allowed users to simply query for existing image data, and directly retrieve the images from their respective URL's. This was an example of how to bring in new features to the Corelyzer system without having to make major modifications to the base application source code.

The original purpose of the plug-in system was to create an architecture that would allow the creation of user interfaces in order to bring similar data from heterogeneous databases together, effectively making Corelyzer a central work environment. This architecture was developed to allow individuals to easily install new modules and could be used to access a central database during an expedition, such as LIMS, or merge data from other databases such as LIMS, the Petrological Database of the Ocean Floor (PetDB) (Lehnert 2005), or the Integrated Data Management for Sediment Geochemistry (SedDB) (Lehnert 1999). Eventually, after the CoreWall Workshop several new capabilities were defined such as the ability to count clasts (rocks formed from pre-existing rock fragments) automatically (ANDRILL 2006, May), incorporating image analysis tools, integration of down-hole experiment logs, various

methods of describing cores and the ability to create hardcopy documents of annotations and descriptions made.

3.3. CoreCLIP, a Corelyzer Plug-in to Analyze Down-hole Logs

A common request for Corelyzer at the CoreWall Workshop has been to integrate data that is gathered separately from the physical core sections. In particular, scientists want to be able to integrate down-hole geophysical logs, or data gathered by sending instruments to measure properties along the walls of a drill-hole. The logs are typically gathered as sets of numerical values associated with given properties and registered to given depths within a hole. A common activity with these numerical data is to correlate the data from multiple adjacent holes together to create a single composite log that defines the entire interval from where the cores were taken to ensure complete recovery. Vertical offsets in drilling in adjacent holes are used to compensate for possibly missing data in any one hole, and to account for expansion of cores in response to their recovery to the Earth surface conditions and as they equilibrate to room temperatures. The process of correlating between logs can be fairly subjective and is often based on matching patterns of plots along graphs that are common to the multiple logs. The Core Log Integration Platform (CLIP) is comprised of the applications SPLICER (Figure 23) and SAGAN (deMenocal et. al.1993) and development to update and port the applications has resulted a single application called NCLIP (Kamp 2006). These programs have been designed to splice, stretch and shrink the graphical logs to correlate with the composite logs.

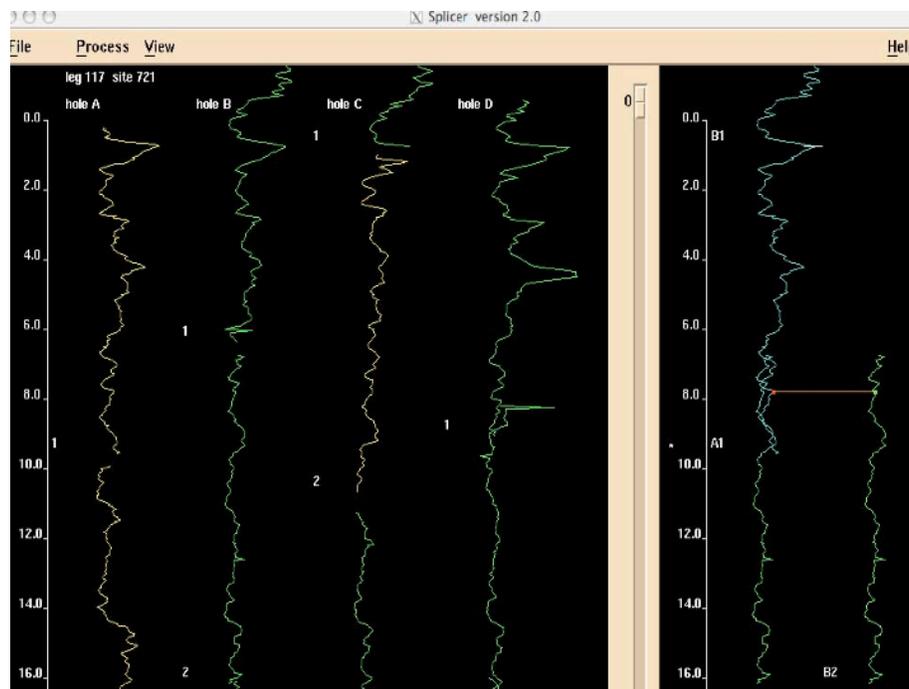


Figure 23 SPLICER with source logs (left) used to create a composite log (right)

Scientists feel that the process of having individual numerical logs that can be correlated and used to form composites would benefit from having the core section images resampled to create a composite image that matches the composite log. To do this CoreCLIP, short for the Corelyzer Core Log Integration Plug-in, will be tasked with this development. As a first step, the resulting output from NCLIP composite building can be taken and brought into Corelyzer to create the composite images with the composite logs displayed alongside. Due to Corelyzer's flexibility, a second step will be the development to integrate the NCLIP interface more directly into CoreCLIP. Development of CoreCLIP will be done in collaboration with Insight Access Group Partners, Ltd. (IAGP), after the NCLIP application has stabilized.

3.4. Workflow Database

The Workflow Database is designed to be in use along side the Corelyzer application, which would be during ICD/VCD work and creation of second order data. During syn-drilling operations, synchronization of descriptions and interpretations of data can be a useful collaborative tool. Once all syn-drilling operations are complete and scientists continue analyzing cores at various geographic locations, it is useful to be able to continue collaborating in the same manner. There still exists the concern over how to restrict or control access to moratorium data. Following are points of interest for the Workflow Database to allow collaboration and help existing workflow.

3.4.1. Synchronized Sessions for Collaboration

Collaboration is key for any scientific drilling operation, as scientists need help from lab technicians, drill operators, curators, and other P.I.s. As a convenience to using the CoreWall Work Suite, and the Corelyzer application in particular, the Workflow Database should have a mode of holding a shared, synchronized Corelyzer session for the duration of a drilling operation, or some user defined time. This entails:

- Broadcasting data availability
- Ensure moratorium policy
- Sharing annotations

Geoscientists who study geological cores are effectively empiricists, or experimental scientists who require tangible objects. Birnholtz et. al. describe empiricists

as people who do not like to share data that required hard-work to acquire, but in the case of scientific drilling, sharing some data can be key to the success of an expedition otherwise the ICD process will be compromised due to lack of detail. However, actual numerical values have to be protected to prevent individuals from violating the data moratorium. While collocated individuals working together can often form a sense of trust (Rocco et. al. 2001), policies must still be enforced to protect interested parties.

As noted before, bigger operations need to hold meetings every day on order to keep everyone synchronized and aware of what operations have transpired, what findings scientists are gathering from recovered cores and what drilling decisions have been made. On top of needing to know when data becomes available, there needs to be a mechanism to support sharing annotations to communicate, whether synchronously like most instant messaging applications, or asynchronously (Weng et. al. 2004) in order to effectively meet the communication needs. The same architecture should be usable for post-drilling collaboration as well.

4. Detailed Case Study: ANTArctic DRILLing

The McMurdo Ice Shelf project is a part of the ANDRILL Program, which will start in October 2006 and drilling ending December 2006. In order to improve communication and improve quality of work, ANDRILL will be making using of Corelyzer and a prototype Workflow Database system.

4.1. ANTarctic DRILLing Scenario

The scenario of Corelyzer usage is fairly simple. The goal of the project is to use Corelyzer as be the application that integrates and visualizes all of the data coming out of the expedition during drilling, and to provide an environment to promote discussion.

ANDRILL will be operating with teams that work 12-hour shifts on a 24 hour, 7 days a week schedule, and hold daily meetings to discuss operations and analyze data. For instance, the science team hopes to analyze and interpret the results obtained visible from the split-core and whole-core images. Scientists would like to be automatically notified when new data is available to view. Unlike “Screen Crayons” (Olsen et. al. 2004), which focused on single users, annotations from multiple users must be managed.

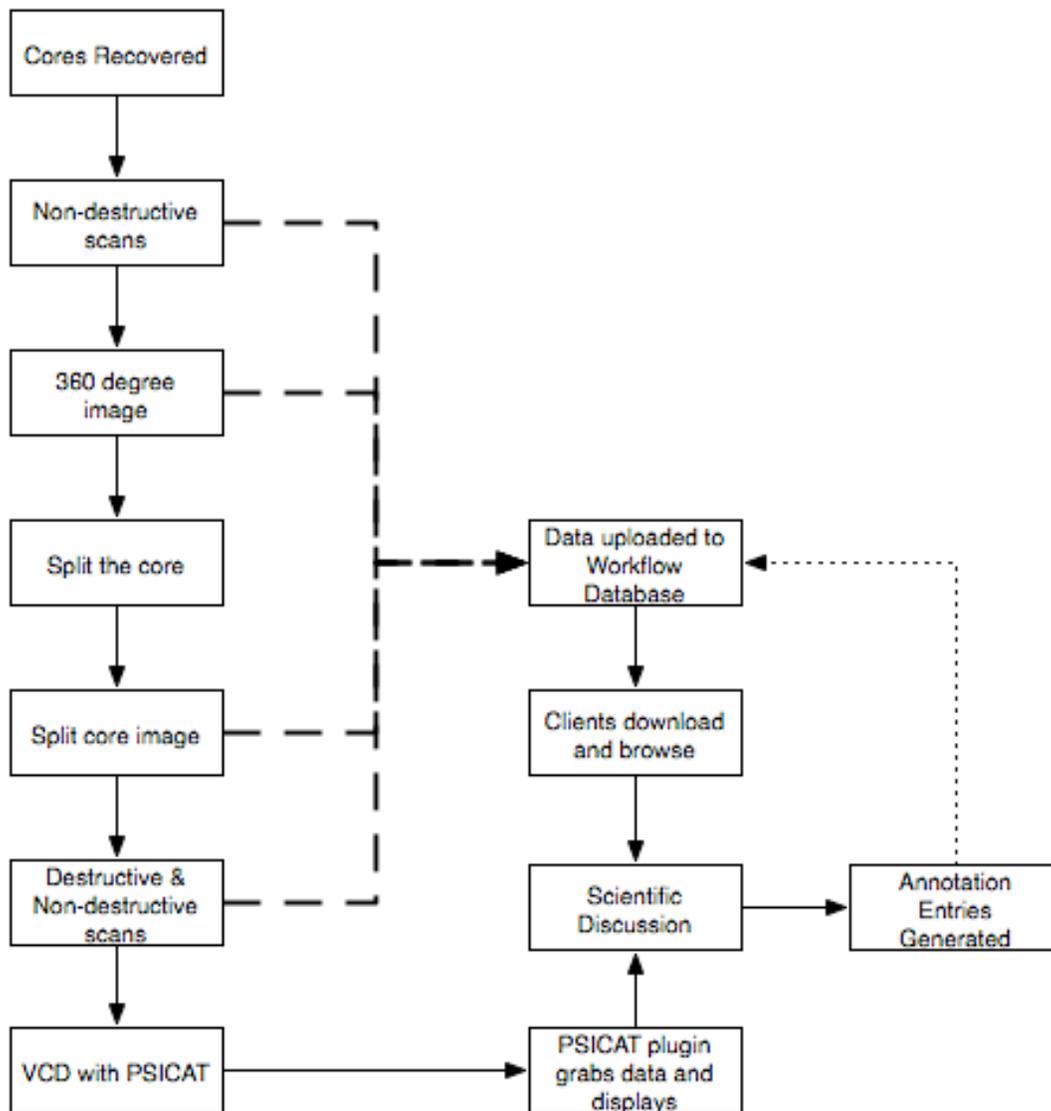


Figure 24 A diagram detailing which points of the core workflow involve interaction with Corelyzer and the prototype Workflow Database. Incoming data from images and sensors (dashed lines) and annotations from discussion (dotted line) are inserted into the database, to be viewed in Corelyzer when a client plug-in is used to download and browse existing data. PSICAT is used for Visual Core Description (VCD) and integrated into Corelyzer via a plug-in, to enhance scientific discussion.

In Figure 24, we see at which points in the generalized core workflow (Figure 11) do Corelyzer and the prototype Workflow Database (WDB) integrate to enhance the workflow. The dashed lines in Figure 24 indicate at which stages do the curator and administrators use a client to push first order data to the WDB, which then notifies other users of the existence of such data. Using a plug-in developed specifically for the ANDRILL scenario (ANDRILL Session Plug-in), scientists will be able to easily pull in the visuals necessary to complete their work. As scientists discuss data and interpretations, annotation entries can be generated and submitted to the WDB. The following sections discuss the workflow integration in more detail, data management issues, how Corelyzer extensibility helps to integrate third party applications into a more complete workflow, and how user feedback can be used to validate the work.

4.2. Physical Integration into the ANDRILL Workflow

Each individual scientist may run Corelyzer on his or her own laptop, hardware requirements permitting. There will also be two communal Corelyzer-dedicated workstations at McMurdo station. Sedimentologists will use the first system in order to support the VCD process by providing a means to view the high-resolution imagery alongside the physical cores. The second system will be used in a common space to provide individuals with a readily available means of contributing to discussions in the shared, synchronized Corelyzer session. Brignull et. al., have performed studies on communally shared interactive systems. The results from their study said that a shared system should:

- Be able to fit with the artifacts used in the space
- Provide physical and digital flexibility with respect to arrangement
- Be designed so a community can absorb the applications into existing activities
- Provide an intuitive set of display-based interactions

Corelyzer can already provide for all four requirements that follows from their study, point for point. Corelyzer runs on any computer, and computers are common artifacts in any space. Physical and digital flexibility are provided for by the display scalability, and by user interfaces to help describe the display layouts (Figure 17). Point three was discussed in detail in previous sections concerning core workflow. The final point is fulfilled based on the OpenGL windows with simple interactions such as panning, zooming, context based left and right clicks of the mouse (see User Interactions).

4.3. Image and Numerical Log Management

First order data is defined here as data that has been gathered by sensors (MSCL), scanners (digital line-scan camera for imagery) and Initial Core Descriptions (VCD / ICD). ANDRILL, unlike the other drilling operations mentioned, is less likely to be acquiring soft sediment cores. This means that they will be able to acquire whole-round images of cores more often, as displayed in the ANDRILL workflow (Figure 6), and be able to more accurately describe the orientation of a core section to help with understanding of structural geology and geophysics. It should be understood that sometimes the whole-round images do not span the length of an entire core, especially if

there is too much water and the core disaggregates. ANDRILL will be drilling a single hole but creating two sets of images or more per core section.

The initial set of non-destructive scans and images are taken when the core section is whole, and this is done at the drill site. Once transported to McMurdo Station, the WDB is updated with these section names, whole-core images of particular depth intervals, and non-destructive whole-core measurements. As the core workflow continues the WDB will be updated with scans from split-cores, data from discrete analyses, and creation of Visual Core Descriptions (VCD) will be made using PSICAT. The integration of PSICAT will be discussed in the following section, but first the management of the images and numerical logs must be detailed.

Both images and numerical logs fall under the data moratorium. Images are affectively owned by all science team members, and can be viewed freely. Therefore it is trivial to host an image using a web server and allow the ANDRILL Session Plug-in (ASP) to copy the image locally and process it for viewing. However, all science team members do not own numerical logs from MSCL devices. Individual attributes, such as magnetic susceptibility, are crucial parts of research for different principal investigators. As such, individuals should not be able to download the actual numerical values of the core logs, but the science teams as a whole are able to view simple plots of the logs. Figure 25 describes the process an individual must undergo in order to view a graphical plot of numerical logs, or to access a data file of the actual logs in order to perform an analysis. By the time an individual is at decision point number three, the individual must be interacting with a database administrator. This inter-personal interaction is designed

specifically to prevent moratorium policies from being violated unless special permission was given.

For Corelyzer to display the graphical plots, the WDB transfers the numerical logs as per core section tables. The data is stored in main memory by the Corelyzer plug-in, in main memory (i.e. no copy of the data exists in non-volatile memory). Once a user exits the application, there will be no existing copy of the core logs. The only way an individual would be able to violate the data moratorium and access the data would be if they develop a plug-in to do so. This is highly unlikely to occur as the intended audience is not well versed in computer science or programming, and time spent in Antarctica is devoted to other tasks. One critical factor missing in this phase of development is communication with individuals “off the ice.” ANDRILL scientists communicate with the globally distributed science team members via low-bandwidth satellite. It is not feasible to directly serve data to the off-ice individuals directly from Antarctica unless data were flown to Christchurch, New Zealand on harddrives and placed on a server to access the Internet. Due to time constraints, this problem was not solved for the 2006 MIS drilling project. Fortunately, this system can be used post-drilling in the same manner, in order to continue any shared discussions.

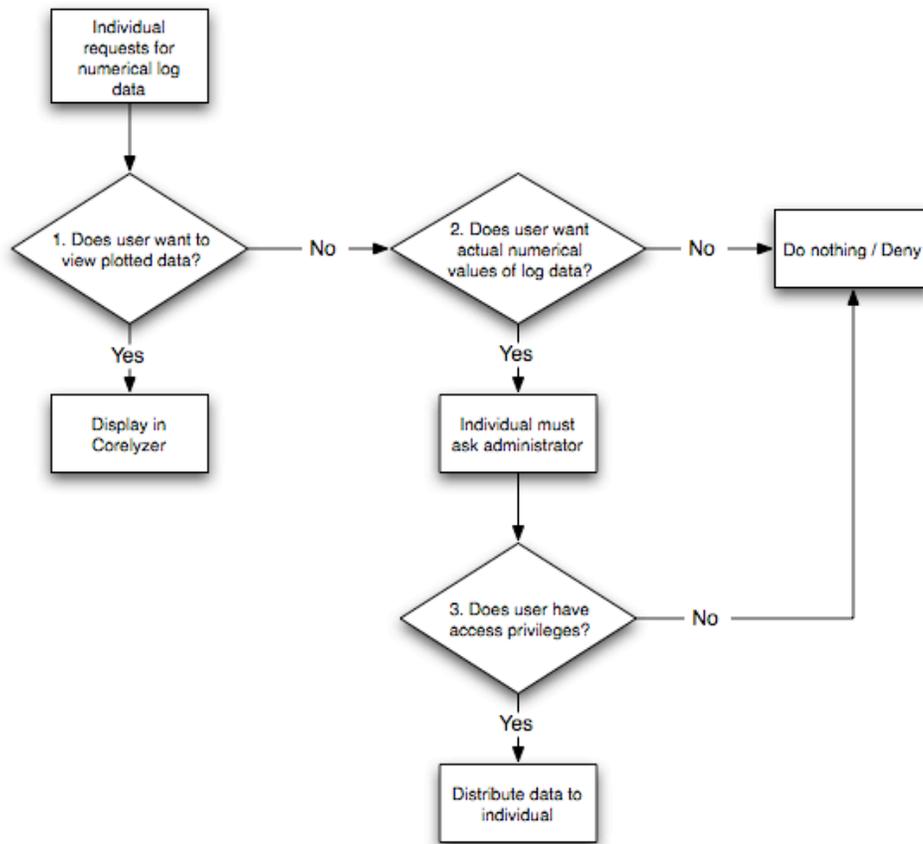


Figure 25 Numerical log access flowchart. Viewing data is less difficult than acquiring actual values, in order to preserve the data moratorium.

4.4. Corelyzer and PSICAT Integration

Visual Core Descriptions (VCD) for the upcoming ANDRILL MIS project will be done with the PSICAT application. In order to produce a more complete integration of data, the resulting VCDs should be displayed in some meaningful manner within Corelyzer. To fulfill this need, CHRONOS has begun development of the PSICAT Corelyzer Plug-in (PCP). The PCP uses VCD files created by PSICAT as input, and

performs the necessary rendering of the VCD to display in Corelyzer. The plug-in is capable of doing this using an internal construct in Corelyzer that allows a plug-in to freely draw anything that may be deemed meaningful to the plugin, within a requested rectangular area. The rectangular areas provided have been termed Free Draw Areas (FDA).

4.5. Community Response to Corelyzer

In order to validate that the technology created for the purpose of improving core workflow, the Corelyzer application, prototype Workflow Database and ANDRILL Session Plug-in were demonstrated to the ANDRILL community. In particular, an ANDRILL P.I., David Harwood, and staff scientist, Richard Levy, experimented with the applications hands-on. Both individuals are key members of the ANDRILL Science Management Office (ANDRILL-SMO) and responded with the following positive comments.

"I'm very impressed ... Prior drilling projects, like the Cape Roberts project, suffered from the ability to have all the science team see the data, the core and have it readily accessible to them, ... to maximize their scientific output. (Corelyzer) is a superb, easy to use template for the integration of a lot of data. I can see potentials for improving it as we continue to march along into the future."

- David Harwood, Ph. D., ANDRILL P.I.

“One of the things I've been impressed with ... is the (ANDRILL session plug-in) application that's been developed for the Corelyzer. It's pretty user friendly, pretty easy to pull in core images, get the information we want to view onto the screens. For our (science team members) who don't want to spend a lot of time learning new software, I think the plug-in is simple enough to allow us to effectively use the program. It's going to be interesting to see how it functions, but I think we're in good shape.”

- Richard Levy, Ph. D., ANDRILL Staff Scientist

At the MIS pre-drill meeting, held in early September 2006, scientists that would describe the cores in the sedimentology room performed a simulation run, using PSICAT for Visual Core Description and Corelyzer to help aid in viewing the cores. The simulation did not have physical property logs to go along with the cores because the cores were borrowed, but the response to Corelyzer was very positive. One scientist said that with the physical properties positioned alongside the core images, Corelyzer would be a valuable tool in interpreting cores.

5. Future Work

The CoreWall Work Suite development is not complete. Development for more effective rendering of datasets need to be incorporated, and development of different plug-ins must be made. Following necessary technical modifications, user studies support by the Internal Review Board (IRB) must be made to more fully evaluate the usability and effectiveness of the CWS as a whole.

Several technical improvements must be made to the Corelyzer system. First, a modification to the rendering system should be made to provide a more smooth transition when zooming and panning. Google Earth does this by rendering the lowest resolution of a particular grid block and gradually filling in with higher resolution images as new frames are created. Google Map does this by simply not rendering any new visuals until a

user has ceased moving the viewpoint. Either of these two methods, or a combination thereof, like Magic Carpet (Svistula 2006) will work. Secondly a system to export a current user Corelyzer session is needed to create publishable documents. For instance, creation of ICD barrel-sheets containing all annotations of a core section. A prototype output design is currently in the works, and is pending community response before implementation begins.

As new technical improvements are made to the base system code, more plug-ins need to be developed to meet community needs. Plug-ins can fulfill a variety of needs. The first plug-in would be to support image analysis using ImageJ (National Institutes of Health), a tool with its own variety of plug-ins to perform texture analysis, particle removing, angle measurements, automatic gamma correction, etc. Secondly, a final implementation of CoreCLIP to provide integration of data gathered from a borehole, MSCL data and core section images. CoreCLIP would allow scientists to create a composite log and a composite image representing the scientist's vision of what the complete record preserved in the ground would appear to be. Third, individual scientists could benefit if a plug-in was developed to allow users to view sample requests, for example, which requests for physical properties experiments were made, and to be able to make additional requests from Corelyzer. This type of plug-in would help curators to avoid having to expose repeatedly those core sections that have been sealed for curation, and it would help scientists to identify whether new cores must be procured to perform any necessary analysis. Fourth, a series of plug-ins should be made to support education and outreach efforts of the various scientific drilling operations. The National Science Foundation cares deeply about educating the world and developing future generations of

scientists of high caliber. Education and outreach activities include both the informal (e.g. museums) and formal (e.g. high school, college and universities) types. Informal education can make use of Corelyzer if a plug-in is used to create real-time animations to tell the story of a core and how it records aspects of the Earth's history. Formal education can benefit from plug-ins that allow students to browse existing public preliminary reports and easily acquire and display the data to create their own reports, and perform analyses of their own.

As Corelyzer, and the accompanied tools that comprise the CWS are more actively used, studies should be performed to evaluate the system as a whole. Evaluations will be critical in improving the system's interface, reliability, and flexibility for customization. Other studies can be performed to help further understand the domain of Computer Supported Collaborative Work (CSCW). Several recent publications have dealt with collocated and distributed work groups. It has been determined that being collocated or not usually does not affect performance (Bos et. al. 2004) (Bos et. al. 2006), but it has been speculated that if everyone in a group is co-located then outside points of view may be lost that can enrich the work (Bos et. al. 2006). What also seems to form with mixed co-located and distributed members of a group is two separate "in-groups" (Bos et. al. 2004). The two groups are formed because co-located members practically ignore distributed members, forcing distributed members to form a group of their own. The experiment that determined this was done in a simulated environment where the goal was to score high in a game. Does this situation arise in a real-world environment such as ANDRILL's, where ANDRILL not only has individuals in Antarctica, but outside as well? This is one of the questions to ask when solving the problem of synchronizing

communication for the 2007 ANDRILL drilling operations. By this time, perhaps CSCW organization theory (Barley et. al. 2004) will exist that can be tested in a scientific-drilling environment.

6. Conclusion

Geological cores are cylindrical objects containing natural records of environmental change that are retrieved for the sake of understanding the past climates of the Earth. Whether retrieving cores from the world's oceans, lakes, ice sheets or continents, the workflow of studying geological cores follow a common path. However, because of the cultural differences between various scientific drilling projects, creating a system to work with all projects can be difficult. This is compounded by the fact that certain data fall under moratorium, requiring special access, and that the amount of data that exists and is generated is beyond the scope of existing applications. The contribution of this thesis is:

1. A generalized workflow model that merges the workflow of three scientific drilling organizations. A clear definition of the workflow not only aids in the development of effective software but also enables future IT specialists to quickly understand the fundamental requirements of scientific drilling projects. This workflow could potentially assist researchers who study and develop workflow technologies (Lei et. al. 1997).
2. A mapping of technology, implemented as the Corelyzer application, to this generalized model that is useful to anyone interested in developing technological solutions for the scientific drilling community.

3. A validation of the efficacy of Corelyzer on a detailed case study- namely ANDRILL
4. A discovery of new problems that need to be solved, new adjustments that are needed to the generalized workflow model based on findings from the case study. With a proposed solution comprised of multiple applications, including Corelyzer, to create the CoreWall Work Suite.

While development of the whole CoreWall Work Suite is not complete, the ANDRILL case study has provided a tangible example and positive support that the current strategy is proceeding in the right direction. More work is needed to complete the system and meet other needs of the community. As the needs are met, more questions will arise relative to understanding workflow models and Computer Supported Collaborative Work.

APPENDIX A: Geology Terms Used

clast – rock formed from pre-existing rock fragments

colorimetry – measurement of the intensity of color

geophysics – the physics of the earth

lacustrine – of, relating to, or associated with lakes.

limnology – the study of the biological, chemical, and physical features of lakes and other bodies of fresh water.

lithify – transform (a sediment or other material) into stone

lithology – the general physical characteristics of a rock or rocks in a particular area

magnetic susceptibility – the fact of being influenced by magnetic force; measurement used to determine metal content

sediment – particulate matter that is carried by water or wind and deposited to the surface of the land or bottom of a body of water, and may in time become consolidated into rock

spectrometry – measurement of the entire range of wavelengths of electromagnetic radiation

REFERENCES

- Co-Chiefs, Staff Scientists, Shipboard Scientific Party, Ocean Drilling Program: Leg 210 Preliminary Report, Drilling the Newfoundland Half of the Newfoundland–Iberia Transect: The First Conjugate Margin Drilling in a Nonvolcanic Rift, November 4, 2003, http://www-odp.tamu.edu/publications/prelim/210_prel/210toc.html
- Co-Chiefs, Staff Scientists, Shipboard Scientific Party, Ocean Drilling Program: Leg 207 Preliminary Report, Demerara Rise: Equatorial Cretaceous and Paleogene Paleooceanographic Transect, Western Atlantic, April 28, 2003. http://www-odp.tamu.edu/publications/prelim/207_prel/207toc.html
- Cohen, Andrew S., *Paleolimnology : The History and Evolution of Lake Systems*. Cary, NC, USA. Oxford University Press, Incorporated, 2003.
- ANTarctic Drilling Homepage, <http://www.andrill.org/>
- Ocean Drilling Program: Leg 202 Preliminary Report, Southeast Pacific Paleooceanographic Transects , http://www-odp.tamu.edu/publications/prelim/202_prel/202toc.html
- IODP MSP Data Portal: Expedition 302, Arctic Coring Expedition (ACEX), http://iodp.wdc-mare.org/front_content.php?idcat=115
- PSICAT Homepage, http://portal.chronos.org:80/gridsphere/gridsphere?cid=tools_psicat&JavaScript=enabled
- CHRONOS's Paleontological-Stratigraphic Interval Construction and Analysis Tool (PSICAT), Corewall Workshop, Washington D.C., May 8-10, 2006, <http://www.ev1.uic.edu/cavern/corewall/pubs/workshop/poster-corewall-legal.png>
- Sun Microsystems, Java J2SE 5.0 Garbage Collector Ergonomics, <http://java.sun.com/j2se/1.5.0/docs/guide/vm/gc-ergonomics.html>
- Google, Google Earth Enterprise, http://earth.google.com/earth_enterprise.html
- Google, Google Earth KML Documentation, <http://earth.google.com/kml/>
- Jenkins C., DBSeabed Homepage, <http://instaar.colorado.edu/~jenkinsc/dbseabed/>
- OpenGL S3TC/DXT Texture Compression Extension Specification, http://oss.sgi.com/projects/ogl-sample/registry/EXT/texture_compression_s3tc.txt
- GLAD - Global Lakes Drilling vessel, http://www.dosecc.org/html/body_glad800.html

- DOSECC – Drilling, Observing and Sampling of the Earths Continental Crust,
<http://www.dosecc.org/index.html>
- Lenhart K., PetDB – Petrological Database of the Ocean Floor, <http://www.petdb.org/>
- Lenhart K., SedDB – Integrated Data Management for Sediment Geochemistry,
<http://www.seddb.org/>
- Levy, R., ANTArctic DRILLing and the CoreWall Suite, CoreWall Workshop,
Washington D.C., May 2006,
http://www.evl.uic.edu/cavern/corewall/pubs/workshop/ANDRILL_CW_DC.pdf
- Group A Session 2 Breakout Group Presentation, CoreWall Workshop, Washington
D.C., May 2006,
<http://www.evl.uic.edu/cavern/corewall/pubs/workshop/GrpASess2-3.pdf>
- deMenocal P., Esmay A., SPLICER and SAGAN,
<http://www.ldeo.columbia.edu/BRG/ODP/ODP/CLIP/clip.html>
- Kamp W., NCLIP, <http://www.iagp.net/NClip/>
- Mackinlay, J.D.; Heer, J. "Wideband displays: mitigating multiple monitor seams," ACM
International Conference on Human Factors in Computing Systems (CHI 2004);
2004 April 24-29; Vienna, Austria. NY: ACM; 2004; 1521-1524.
- Jeong, B., Jagodic, R., Renambot, L., Singh, R., Johnson, A., Leigh, J. Scalable Graphics
Architecture for High-Resolution Displays, Proceedings of IEEE Information
Visualization Workshop 2005, Minneapolis, MN, October, 23, 2005
- Brush, A. J., Barger, D., Gupta, A., and Cadiz, J. J. 2001. Robust annotation
positioning in digital documents. In Proceedings of the SIGCHI Conference on
Human Factors in Computing Systems (Seattle, Washington, United States). CHI
'01. ACM Press, New York, NY, 285-292.
- Olsen Jr., D; Taufer, T; Fails, J.A.; "ScreenCrayons: Annotating Anything", ACM UIST
2004 Conference Proceedings, Santa Fe, New Mexico, 24-27 October 2004, pp
165 - 174.
- Hutchings, D.R., Smith, G., Meyers, B., Czerwinski, M., Robertson, G., "Display Space
Usage and Window Management Operation Comparisons between Single
Monitor and Multiple Monitor Users," AVI 2004 Conference Proceedings,
Gallipoli (LE), Italy, 25-28 May 2004, pp 32 - 39.

- Rocco, E., Finholt, T.A., Hofer, E.C., & Herbsleb, J.D. (2001, April). "Out of sight, short of trust," Presentation at the Founding Conference of the European Academy of Management. Barcelona, Spain.
- Weng, C., Gennari, J., "Asynchronous Collaborative Writing through Annotations," Proceedings of CSCW 2004, 6-10, November, 2004, Chicago, Illinois, USA, pp 578 – 581
- Lei, K. and Singh, M.. A Comparison of Workflow Metamodels, Proceedings of the ER-97 Workshop on Behavioral Modeling and Design Transformations: Issues and Opportunities in Conceptual Modeling, Los Angeles, November 1997
- Dimitri Svistula, Magic Carpet, <http://www.evl.uic.edu/cavern/mc/index.html>
- National Institutes of Health, ImageJ: Image Processing and Analysis in Java, <http://rsb.info.nih.gov/ij/>
- Bos, N., Olson, J., Nan, N., Shami, N., Hoch, S., Johnston, E., "'Collocation Blindness' in Partially Distributed Groups: Is There a Downside to Being Collocated?" , Proceedings of CHI 2006, 22-27, April, 2006, Montreal, Quebec, Canada, pp1313 – 1321
- Brignull, H., Izadi, S., Fitzpatrick, G., Rogers, Y., Rodden, T., "The Introduction of a Shared Interactive Surface into a Communal Space," Proceedings of CSCW 2004, Chicago, IL, USA, 6-10 November, 2004, pp 49-58
- Tollinger, I., McCurdy, M., Vera, A., Tollinger, P., "Collaborative Knowledge Management Supporting Mars Mission Scientists," Proceedings of CSCW 2004, Chicago, IL, USA, 6-10 November, 2004, pp 29-38
- Lampe, C., Johnston, E., "Follow the (Slash) dot: Effects of Feedback on New Members in an Online Community," Proceedings of GROUP 2005, Sanibel Island, Florida, USA, 6-9 November, 2005, pp 11-20
- Birnholtz, J., Bietz, M., "Data at Work: Supporting Sharing in Science and Engineering," Proceedings of 2003 International ACM SIGGROUP Conference on Supporting Group Work, Sanibel Island, Florida, USA, 9-12 November, 2003, pp. 339 – 348
- Bos, N., Shami, N., Olson, J., Cheshin, A., Ning, N., "In-group/Out-group Effects in Distributed Teams: An Experimental Simulation," Proceedings of CSCW 2004, Chicago, IL, USA, 6-10 November, 2004, pp. 429-436
- Barley, S., Kiesler, S., Kraut, R., Dutton, W., Resnick, P., Yates, J., "Does CSCW Need Organization Theory? Proceedings of CSCW 2004, 6-10, November, 2004, Chicago, Illinois, USA, pp 122-124

Krishnaprasad, N., Vishwanath, V., Venkataraman, S., Rao, A., Renambot, L., Leigh, J., Johnson, A., Juxtaview – a Tool for Interactive Visualization of large Imagery on Scalable Tile Displays. Proceedings of IEEE Cluster 2004,, 20-23, September, 2004, San Diego, CA, USA, pp 411 – 420

Krumbholz, C., Leigh, J., Johnson, A., Renambot, L., Kooima, R., Lambda Table: High Resolution Tiled Display Table for Interacting with Large Visualizations. Proceedings of 5th Workshop on Advanced Collaborative Environments, Redmond, Washington, 2005

VITA

- Name: Arun Gangadhar Gudur Rao
- Education: M.S., Computer Science, University of Illinois at Chicago, 2006
B.S., Computer Science, University of Illinois at Chicago, 2003
- Research Experience: Electronic Visualization Laboratory, Department of Computer Science, University of Illinois at Chicago, 2003 – 2006
- Publications
- Rao, A., Chen, Y., Lee, S., Leigh, J., Johnson, A., Renambot, L. Corelyzer: Scalable Geologic Core Visualization using OSX, Java and OpenGL, Apple's Worldwide Developers Conference 2006, 07/08/2006 - 08/11/2006
- Rao, A., Rack, F., Kamp, B., Fils, D., Ito, E., Morin, P., Higgins, S., Leigh, J., Johnson, A., Renambot, L. CoreWall: A Scalable Interactive Tool for Visual Core Description, Data Visualization, and Stratigraphic Correlation, Eos Trans. AGU,86(52), Fall Meet. Suppl., San Francisco, CA, 12/05/2005 - 12/05/2005
- Rao, A., Kamp, B., Noren, A., Ito, E., Schnurrenberger, D., Morin, P., Leigh, J., Rack, F., Johnson, A., Renambot, L. CoreWall: A Visualization Environment for the Analysis of Lake and Ocean Cores, Geological Society of America Abstracts with Program, Denver, Colorado, 12/12/2004 - 12/12/2004
- Renambot, L., Rao, A., Singh, R., Jeong, B., Krishnaprasad, N., Vishwanath, V., Chandrasekhar, V., Schwarz, N., Spale, A., Zhang, C., Goldman, G., Leigh, J., Johnson, A. SAGE: the Scalable Adaptive Graphics Environment, Proceedings of WACE 2004, 09/23/2004 - 09/24/2004
- Krishnaprasad, N., Vishwanath, V., Venkataraman, S., Rao, A., Renambot, L., Leigh, J., Johnson, A., JuxtaView – a Tool for Interactive Visualization of Large Imagery on Scalable Tiled Displays, Proceedings of IEEE Cluster 2004, San Diego, CA, 09/20/2004 - 09/23/2004