CHIMIT '07: 2007 Conference on Human Interfaces to the Management of Information Technology

Cambridge, MA
March 30–31, 2007

Summarized by Alva Couch (couch@cs.tufts.edu) and Marc Chiarini (marc.chiarini@tufts.edu)

Just what do system administration and ethnography have in common? One might think that ethnography is about studying cultures, but in a broader sense, ethnography means "writing about and describing humans and human behavior." For better or worse, system administration involves many human-to-human and human-to-machine interactions. Ethnographers have been studying the behavior of system administrators and other technical professionals, and they have come to some surprising conclusions about the nature of our profession, and how and why we interact socially with humans and technically with machines and networks.

CHIMIT '07 brought together human factors experts, system administrators, and researchers to study how the practice of system administration is affected by human factors such as software interfaces and social context. Papers analyzed how system administrators work, analyzed failures of human process, proposed novel interfaces to administrative data, and suggested improvements in process, software, and structure of support organizations. Unlike the typical LISA paper, which describes "how" best to perform a task or perhaps the options available for addressing a problem, papers at CHIMIT concentrated upon "what system administrators really do" and "why things are the way they are" from a human perspective. The conference was sponsored by ACM, in cooperation with USENIX.

OPENING PLENARY

■ *Information Technology in the Wild*

Stephen Barley, Stanford University

Stephen Barley has been engaged in more than 20 years of ethnographic studies of technical professionals and how they fit into organizations. "Barley's law" is that "what you get is never quite what you plan." First-order effects of employing a technology are often followed by second-order effects that are primarily sociological in nature. Technologies alter work practices, which change or create social divisions, which become demographic divisions (adopted as part of personal identity) over time.

Barley spun an elaborate tale illustrating this principle, starting with the nineteenth-century industrial revolution. The divide between professionals and technical support staff is—in many contexts—a sociological second-order effect of the introduction of technology. Barley's definition of a technician is someone who "creates a bridge between the real and the representation." The accepted myth is that professionals (such as those responsible for interpreting X-rays and business data) can do "anything that technicians can do." Barley's ethnographic studies show that there is little overlap between "professional" and "technical" staff expertise sets and neither is well-equipped in the skills and knowledge necessary to do the job of the other.

Stories of technological change teem with examples of unintended sociological consequences. Telecommuting was first promoted as a way to reduce pollution. Companies could "comply" with the Clean Air Act by setting up telecommuting programs, even though these programs statistically account for the activities of less than 1% of the workforce. This led to a large number of telecommuting programs that few employees used, along with marketing ploys to attract business from the almost insignificant population of telecommuters. Barley and others collected over 3000 articles on telecommuting over many years and demonstrated—through ethnographic analysis—how the story evolved over time.

Another social factor is the rise of the independent consultant. The consultant is perceived as a free agent who takes long vacations and pursues a leisurely existence. Ethnographic studies show that consultants work harder, and for longer hours, than their full-time counterparts, and they often ignore long-term financial planning, which is built into most employee compensation plans. Roughly one-half of all consultants surveyed have no retirement savings plan at all. Conversations after the talk revealed an innovative strategy: There are consultant organizations that function like "employers of record," allowing consultants to "appear" to be fully employed and accrue benefits through the organization without answering to the organization.

Barley asserts that positive change in organizational structure and sociology seldom comes from inside an organization; there must be a "contravening force" to accomplish change. The nineteenth-century evolution of the manager created a situation analogous to that of user versus system administrator and was only countered by the evolution of contravening labor unions and organizations. The information revolution has created a divide that separates "professional" and "technical" staff in much the way that the industrial revolution separated "management" from "labor," and Barley's assertion is that only contravening forces similar to guilds or labor unions will create a balance between "professional" and "technical" needs.

PAPER SESSION I: STUDIES OF SYSTEMS MANAGEMENT

- ***Design Guidelines for System Administration Tools Developed Through Ethnographic Field Studies***

  Eben M. Haber and John Bailey, IBM Almaden Research Center

System administration tools have several weaknesses that can be discovered by ethnographic methods and that have not been suggested by system administrators themselves. Using direct observations of practice as a guide, the authors identify several improvements that could be made in system administration and configuration management tools, including support for planning and rehearsal, support for long-running change operations with limited change windows, and non-blocking user interfaces that do not force the system administrator to wait at the terminal for long operations to complete. Other suggested enhancements include progress indicators for long operations as well as execution time prediction.

Configuration management is an especially hard problem which deserves special attention and tool capabilities. Observation of system administrators indicated that enhanced collaboration tools for quick exchange and comparison of information are needed and that communication and trust between administrators is often a major bottleneck. To illustrate this, a video shows a chat and phone session between two administrators in which trust and disbelief play a major part in impeding troubleshooting. One administrator cannot believe that another is correct in asserting that the base configuration of a device has changed. It is the communication between the two administrators that is the bottleneck, rather than the technology.

- ***Deciding When to Trust Automation in a Policy-Based City Management Game: Policity***

  Kenya Freeman Oduor, IBM Software Group; Christopher S. Campbell, IBM Almaden Research Center

Trust is a major factor in accepting "self-managing" systems as part of IT infrastructure. It is difficult to observe the social bases of system administrator trust directly, but one can reason from analogy to other kinds of expert systems. In this paper, the authors utilized relationships between players and expert "assistants" in a computer simulation game as analogous to relationships between system administrators and autonomic control systems. The chosen game was Policity, a computer simulation in which players try to run a large city by making infrastructure decisions.

Researchers embedded reliable and unreliable "assistants" in the game, asked students to play the game, and observed their reactions and evolving trust relationships. Some students were given access to high-reliability assistants, whereas others' assistants functioned at a lower level of reliability and sometimes gave poor advice. Observation of actions showed that unreliable assistants were eventually ignored by the players, and reliable assistants were also ignored, but to a lesser degree. One conclusion was that trust is not a simple matter of reliability; it also requires some exposure of the underlying decision process. This study suggests that, to be trusted, an autonomic management solution should reveal some of the logic behind its decisions and/or statistics on the accuracy of previous decisions.

- ***Towards an Understanding of Decision Complexity in IT Configuration***

  Bin Lin, Northwestern University; Aaron B. Brown and Joseph L. Hellerstein, IBM T.J. Watson Research Center

Low-level configuration procedures for machines are error-prone, are complex, and lead to inconsistencies in service, but many administrators still perform much—if not all—configuration by hand. To evaluate whether automation is an appropriate substitute for manual changes, it is appropriate to seek metrics for the complexity of the human part of the process. The authors identify three kinds of complexity:

- Execution complexity: How complex and error-prone is the execution of commands?
- Parameter complexity: How difficult is it to give commands appropriate parameters?
- Memory complexity: How much must one remember in order to configure systems effectively?

To understand how these kinds of complexity could affect nonexpert system administration, the authors appeal to an analogous system: route planning on a map.

Human subjects were asked to solve a variety of route-planning problems and their attitudes and actions were observed. Initial results showed wide variations in human performance. A factor analysis showed that task completion time varied most as a result of the amount of guidance given and constraints imposed. An analysis based upon only these factors yielded some surprising results.

It is a popular belief that showing guidance will shorten the time taken to complete a task. In this study, showing guidance information lowered the user-perceived difficulty of route planning but did *not* improve their performance in using the tool. In fact, task times were similar in the two cases, whereas error rates were somewhat lower when guidance information was *not* shown. This very counterintuitive result suggests that guidance information may be an impediment rather than an aid when attempting to assist nonexpert system administrators in completing complex tasks.

PAPER SESSION II: DESIGNING FOR IT MANAGEMENT

- ***Cube Management System: A Tangible Interface for Monitoring Large Scale Systems***

  Elliot Jaffe, Aviva Dayan, and Amnon Dekel, Hebrew University of Jerusalem

In large installations, it is especially difficult to analyze and exchange logging information for technical problems in a way that is easily exchanged with others. The Cube Management System (CMS) is a log analysis system based upon an innovative tangible interface. Physical "cubes" are arranged in a grid in an interface in which cubes can be physically picked up and moved around. Each cube is an active computing device whose location in the grid corresponds to the identity of a processing element such as a rack, blade, or disk. The status of the device to which a cube corresponds is indicated by three LEDs: red for trouble, yellow for marginal operation, and green for proper function. Each cube is labeled (via an alphanumeric display) with the component to which it corresponds, and it also contains current log information and problem descriptions. Cubes can be moved to any physically accessible location in the IT infrastructure and also passed among system administrators in order to share information. Prototypes have been implemented both with physical cubes containing standalone computers and with virtual cubes in a 3D graphical interface.

Trouble isolation and delegation of responsibility in CMS are accomplished via a "zooming" mechanism that changes cube identities in the grid to represent different kinds of components at different functional levels. A system administrator places a cube of interest into a special tray on a grid, which has the effect of changing the grid so that it represents components at another level. Thus one can very quickly zoom down from, for example, a rack with a problem, to the blades in the rack with problems, to the disks of each blade that have problems. The mechanism preserves context by allowing one to set aside a cube at one level and return to it after having drilled down to a problem spot. By shuffling cubes on the work surface and handing them to technicians responsible for problem resolution, one can quickly locate, identify, and resolve problems at varying scales. An additional advantage is "ambient information transfer": A colored light allows multiple technicians to be subtly aware of the health of monitored systems without taxing other important cognitive functions.

- ***Activity-based Management of IT Service Delivery***

  John Bailey, Eser Kandogan, Eben Haber, and Paul P. Maglio, IBM Almaden Research Center

Outsourcing has been thought to be a panacea for a variety of information technology problems, but it is difficult to transition from insourced to outsourced IT solutions. The ethnographic study of the outsourcing process in a particular organization suggests that the key to a smooth transition is a "transition manager" who coordinates design and deployment staff. Current transition managers oversee the transition process via a spreadsheet of task states and/or existing project management or ticketing software.

The transition manager actually requires five kinds of information:

- People: Staffing needs for each task
- Resources: Requirements for hardware, connectivity, etc.
- Calendar: Schedule and expectations for task completion
- Tools: Integrated email mechanisms for "context passing" from manager to staff and vice versa
- Process: An ITIL diagram of the transition process and its state

The paper proposes that a "big board/shared workspace" for the transition manager would greatly improve the efficiency of the transition process, integrate data sources, and make outsourcing practical in more cases. The workspace would need to extend traditional tasks with ad hoc steps, support team collaboration, and support reuse of ad hoc activity knowledge (including mining and analysis of such). An added benefit would be improved cost estimates of service delivery activities. The big board is not without obstacles, of course: How does one assure the efficient and accurate mapping of activity data? What are the requirements of the data mining and analysis components? How can the "Rockwell Effect"—in which staff feel their every move is being recorded and analyzed—be mitigated? Will filters, "perspectives," and alerts be sufficient to help the transition manager deal with information overload?

- ***IT Ecosystems: Evolved Complexity and Unintelligent Design***

  Jim L. Lentz and Terry M. Bleizeffer, IBM Software Group

Most IT infrastructures are not designed from scratch; rather, they evolve over time as a result of external and unforeseen forces, leading to final designs with significant flaws. The results of this evolution are often inferior to designing a new solution from the top down. Evolving via small changes can lead to ad hoc employee organization, poorly interlocked and defined responsibilities, overly diverse IT environments, and overly "personalized" IT practices.

A different set of design issues is encountered when the problem is approached from the top down. "Unintelligent design" practices include:

- Concentrating upon a small subset of the overall solution
- Making simplifications that make the product less useful
- Relying upon GUI design alone without considering process
- Blaming architects for complexity of the ecosystem
- Believing that all customer requirements must be addressed exactly as described

Getting control of complexity in an evolved ecosystem is difficult. The authors suggest strategic encapsulation of complexity. Complex parts of process are limited to those that are mission-critical, and domain experts are trained to manage those processes via specialization, much as the general practitioner often refers a patient to a specialist. In this way, the final design preserves the business process and places expertise where it matters, in the most complex parts of the process.

PANEL

- ***Is Automation a Panacea for Management of Information Technology?***

  Michael Beck, Emerging Technologies Group; David N. Blank-Edelman, Northeastern University; Tom Limoncelli, Google; Paul P. Maglio, IBM Research

  Moderator: Alva Couch, Tufts University

Is automation a panacea for management of information technology? Yes and no. The panel was in some sense in agreement on this answer: Yes, for the easily automated processes; no, for the inherently human ones. Human processes include dealing with policies and workflow, including account management. The practical approach of capturing scripts and replaying them works in many cases, and autonomics promise to automate all but troubleshooting in the long run.

DINNER INVITED TALK

- ***Interfaces for Everyone***

  Rob Kolstad, Ph.D.

Interfaces—including complex ones—are everywhere. The talk began by a look at the number of buttons and dials and indicators that the average person has to manipulate in order to prepare for work on a given day, including the alarm clock, shower, toaster (with labels in German), shoelaces (including hundreds of ways to tie laces), and automobiles. This was followed by a detailed tour of computer interfaces through history, starting with the abacus and ending with the Wii. Throughout history and a regular day, we deal with fairly complex interfaces that require training and context. History shows, as well, that interface developments occur in spurts with long intervening pauses. The last interface development occurred over 15 years ago.

A common way of designing interfaces is to refer to "design patterns," but Rob believes that many of these have become "design ruts." These include the incessant insistence upon use of a mouse when a keyboard will do; the lack of respect for users' prior training in typing; and the idea that command-line interfaces are "too dangerous" or "provably less effective than CLIs." The subsequent discussion included the question, "Should people have a driver's license in order to utilize a CLI?" There was no clear answer.

PAPER SESSION III: TECHNOLOGY AND USE

- ***Network-Centricity: Hindered by Hierarchical Anchors***

  Steve Abrams and Gloria Mark, University of California, Irvine

Ethnographic study of the interactions of a failed design team exposes several social "anchors" that impede progress when moving from hierarchical ("stovepipe") to network-centric business infrastructure. A division of a large company wished to institute a division-wide calendaring system and combine data from several existing site-specific calendaring systems into one coherent calendar. A design team composed of members from three sites was given six weeks to design a new calendaring solution to be adopted by all sites. Ethnographic researchers observed all communications and meetings and analyzed the results of the interactions.

In the first meeting, two design philosophies emerged that kept the sites from communicating and compromising. The "bottom-up" camp was concerned with pleasing the boss and integrating existing technologies, while the "top-down" camp was concerned with surveying user needs and selecting a new and distinct division-wide solution. Four kinds of unresolvable differences and inflexibility arose:

- External authority: An authority figure requires one kind of solution.
- Technology: Our current design only supports one kind of solution.
- Asymmetric communications: Leaders are unresponsive to communications from other site representatives.
- Work commitments: Other job requirements interfere with deliverables.

These "anchors" were repeated throughout three weeks of meetings, in different guises, and finally led to team disbandment. The key lesson is that any successful transition process must necessarily avoid these patterns of interaction.

- ***Managing Technology Use and Learning in Nonprofit Community Organizations: Methodological Challenges and Opportunities***

  Cecelia Merkel, Umer Farooq, Lu Xiao, Craig Ganoe, Mary Beth Rosson, and John M. Carroll, Pennsylvania State University

This study concerns how to develop sustainable information technology in small-scale nonprofit organizations. Small nonprofit organizations often suffer from ad hoc planning processes, a lack of risk analyses, and migratory volunteer workforces, leading to impromptu and "opportunistic" IT development strategies: Make changes while

someone is available to make them. This often leads to unreliable or even unusable IT services.

Researchers utilized the "participatory action research" model in which they are both partners and observers of technology development practices. In several case studies, they attempted to partner, suggest improvements, and then "fade out" of the picture to leave a more sustainable infrastructure.

The case studies illustrate the usefulness and limitations of technology partnering. A food bank providing 12,000 bags of groceries to 800 households was hampered by the lack of a technology plan and needed help in transitioning from ad hoc to structured IT planning. A watershed council was stalled on a Web site design issue; partners helped them define audiences of the Web site and tune it to the organizational mission. A historical society suffered from an ad hoc approach to technology in which long periods of stasis were interrupted by occasional grants; partnering failed to remedy lack of interest in moving toward a more stable solution. Lessons learned include that partnering can help nonprofits engage in planning processes, leverage domain expertise in designing or improving IT infrastructure, and work toward acceptance of the need for change.

■ ***Supporting Expertise Awareness: Finding Out What Others Know***

Christian Dörner and Volkmar Pipek, University of Siegen; Markus Won, University of Bonn

This paper describes an attempt to utilize activity log information to characterize expertise profiles of members of a "freelance network" of more than 200 technical professionals. Rather than using a "yellow-pages" approach in which each member lists his or her domains of expertise, the authors engaged in a three-year study in which they attempted to infer members' domains of expertise from logs of member activities such as browsing Web sites and interacting with newsgroups. Members who posted to particular newsgroups and/or read expertise-centric news sites or documents were deemed to have more expertise, based upon the time intervals in which they sustained such activity.

The expertise filtering is accomplished by an environment called eXact, which takes as input a graphical representation of the filtering to be done. The filtering is done in several stages, including collection of the statistics, filtering for privacy, and generation of hypotheses that match gathered data. These hypotheses are collected in a "knowledge garden" and used as additional information in building teams and training new members. The result is a set of hypotheses that—in practice—represent "extra information" but are not considered definitive as an expertise measure.

Two major obstacles were encountered during this study. First, there was no resolution to the "eternal struggle" between privacy and the need to characterize individuals. Second, it would be possible to "fake" expertise data by generating meaningless log entries, though the authors assert that there would be strong negative social ramifications of such behavior if discovered, including ostracism.

PAPER SESSION IV: USABILITY AND SECURITY

■ ***Looking for Trouble: Understanding End-User Security Management***

Joshua B. Gross and Mary Beth Rosson, Pennsylvania State University

In this ethnographic study, twelve users handling sensitive data were surveyed about their attitudes toward security. What do users know, how do they manage it, and whom do they perceive as responsible? The end result of this work is to define "personas" that can be utilized in designing better security policies. A "persona" is a personality type, defining likelihood of specific behaviors, unlike a "role," which instead describes responsibilities. Informally, three kinds of personas arose in the study: "vigilant," "reliant," and "careful." A vigilant person takes personal responsibility for security, a reliant person places the responsibility for security upon others, and a careful person relies upon incomplete knowledge and is "careful" within that context.

Eleven of the twelve users exhibited significant lack of knowledge about security issues. One user reported "I feel lost." There is serious concern over security; one user reported "I would never work in this field again if I were involved in a major incident." But there is a disconnect between "concern" and "responsibility." Only one of twelve users took personal responsibility for security; the others placed responsibility on technical solutions or staff. This lack of knowledge extends to physical security; several participants had confidential files open on their desks during the interview.

The authors conclude by listing several methods whereby the knowledge of users about security might be improved. These include activities that build trust in institutional policy, use of information escrow (in which a trusted human agent relays all sensitive information), and adding visual and audio cues in software to increase trust and social presence. Knowledge of the personas of users can be utilized in designing better risk-reduction measures.

■ ***User Help Techniques for Usable Security***

Almut Herzog and Nahid Shahmehri, Linköping University

Security is not a user's primary task; understanding current security mechanisms involves specialized knowledge that almost no users possess. A typical security pop-up, such as "Do you wish to trust this IP address?" usually translates in the mind of a user into "Do you want to get your work done or not?." Thus these pop-ups, in the context of an average user, provide no protection and are merely an an-

noyance. To address this, one must express to the user—clearly and concisely—the nature of the decision to be made and its impact. The authors consider the role of online help in aiding the user to make intelligent security decisions.

There are several forms of online help, including documentation, context-sensitive help, assistants, wizards, staging, and social navigation. Staging refers to the process of training the user in steps toward more advanced security aims (e.g., training the user on one new security idea per day). Social navigation refers to the practice of showing the statistics for each decision for other users on the network; this is controversial because the majority may in fact choose incorrectly. Several capability matrices show the strengths and weaknesses of each approach but in the end all are less than reliable and offer no substitute for built-in security mechanisms that cannot be overridden. The authors indicate a "slight preference" for wizard-based mechanisms in cases where changes are made infrequently.