


Seeing is believing: Designing visualizations for managing risk and compliance



R. K. E. Bellamy
T. Erickson
B. Fuller
W. A. Kellogg
R. Rosenbaum
J. C. Thomas
T. Vetting Wolf

This paper explores the design of visualizations that support mandated organizational compliance processes. We draw on the research literature to show how visualizations can operate as effective user interfaces for complex, distributed processes. We argue that visualizations can reduce the complexity of such processes, making them easier to manage, and can facilitate the communication and collaboration that are critical to supporting compliance. We describe the design and pilot deployment of a visualization that supports the IBM Sarbanes-Oxley Act compliance process, discussing design alternatives, the final design, its deployment, and lessons learned.

INTRODUCTION

Since the advent of the Sarbanes-Oxley Act (SOX), which established new or enhanced standards for corporate accountability, systems for tracking and managing compliance testing have become critically important to organizations. One of the challenges to designing such systems is managing the interface between people and the system's computational processes. By their very nature, compliance systems require that humans monitor processes and investigate and fix problems: There is no such thing as an "automatic" compliance system. The point of a compliance system is to ensure that humans are involved, because ultimately one or more persons must take personal and legal responsibility for compliance. Another challenge is the fact that, in an organization of any appreciable size, monitoring and other activities are likely to occur among a number of people in disparate locations. These challenges are compounded by the fact that human

performance differs from that of computers: People are prone to make a variety of errors of perception and calculation, including forgetting to enter information, neglecting to enter it in a timely fashion, and sometimes overlooking items. In short, compliance management systems are complex socio-technical entities whose smooth functioning involves a blend of technical, human, and social factors.

Our concern is with the design of socio-technical systems that span large organizations. In this paper, we report on the design and use of a shared, dynamic visualization of organization-wide pro-

©Copyright 2007 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the Journal reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to republish any other portion of the paper must be obtained from the Editor. 0018-8670/07/\$5.00 © 2007 IBM

cesses to aid tracking and managing compliance. We argue that the use of a visualization can amplify human cognition at an individual level, and that when it is visible to all participants, it will evoke social processes that can aid in managing compliance. At the same time, the shared nature of the visualization means that privacy becomes a key issue for ethical, cultural, and legal reasons, and so the design of the visualization involves making careful choices about the circumstances under which information is made visible to others.

MANAGING RISK AND COMPLIANCE

While there are a number of eminently practical reasons that companies should be concerned with managing their risks, the current interest in this domain is driven by recent legislation by the United States Congress. It is useful to understand this legal context, as it is the direct driver of the current push to devise better tools to support risk and compliance.

The Sarbanes-Oxley Act

SOX was signed into United States law on July 30, 2002, largely in response to a number of major corporate and accounting scandals. It establishes new or enhanced standards for corporate accountability. All publicly traded companies in the United States need to comply with this legislation. The scope of the law and the nature of the controls it specifies mean that compliance requires considerable work on the part of companies. In view of this, SOX is being put into effect on a staggered schedule, each section going into effect on a different date. The following are the main sections of the law:

- Section 302/906—Corporate Responsibility for Financial Reports
- Section 404—Management Assessment of Internal Controls
- Section 409—Real-Time Issuer Disclosures
- Section 802—Records Management

Section 302 of SOX is already in effect. To comply with Section 302, the chief financial officer (CFO) and chief executive officer (CEO) must personally certify the accuracy of financial statements and the efficacy of internal disclosure controls. Disclosure controls must be established and enforced at all levels within the company, with quarterly evaluation of the efficacy of controls by the company. All significant deficiencies, material weaknesses, and

acts of fraud must be disclosed to the audit committee. The company must establish and emphasize a culture of integrity, and the CEO and CFO must have confidence and trust in the people and process.

Section 404 for accelerated filers (companies with a public float greater than \$75 million) went into effect in November 2004. For nonaccelerated filers, the date has been extended to July 2007. It requires that management file annual reports on internal controls, and that these reports be attested to by external audit firms. All controls relating to financial reporting must be documented and tested for efficacy. All gaps and deficiencies of such controls must be reported. Companies must demonstrate an ability to monitor control compliance.

Sections 409 and 802 are not yet in effect. The ongoing implementation of Section 404 and the pending implementation of Sections 409 and 802 suggest that the way organizations manage their risk and compliance processes will be in flux for the near future, and that, in turn, there will be a significant need for approaches that facilitate an organization's ability to deal with these issues.

Before the advent of SOX, IBM already had a well-defined control process. This process was further developed to support SOX, and particularly SOX 404 reporting, the latter focusing on management oversight of internal controls, including quarterly evaluation of each control.

By the time we started working with the IBM controllers, all the controls had been identified and categorized by business process and by country. For example, one control—in the Accounts Receivable process category—requires that adjustments be checked to ensure that management has approved all adjustments that are not financial in nature; this control is tested in several countries, including Argentina, Australia, Austria, Belgium, and Brazil.

As one would expect, the process of testing controls is not automatic. Even if this were possible, it is not desirable, because SOX requires that individuals take personal responsibility for certifying the accuracy and efficiency of the control processes. Thus, people are an integral part of the system. Each process has an owner who monitors the set of controls in effect for that process (a process owner

may own several processes); each country has an owner who monitors the controls for that country; and each control has an owner who is responsible for seeing that the appropriate number of tests are run each quarter (or each testing period) and that defects are investigated, reported, and remediated. Process owners are responsible for coordinating the people who own their process controls and for determining that the defects are real, rather than a problem with the control itself.

Monthly reporting from the globally distributed owners for each control is supported by a form-based front end to a database, and all control information is stored in the database. A quarterly scorecard is generated that provides summary information about the control status, such as the number of unremediated defects since inception to date, the number of defects for the current quarter, and the number of samples for the current quarter. The scorecard is rolled up by business unit and business process.

Overall, IBM monitors on the order of hundreds of controls spanning dozens of countries (though not every control is relevant to every country). This generates a significant amount of data, and the problem of managing it is nontrivial, particularly because of the dynamic and distributed nature of the controls and control-related information.

IMPORTANCE OF VISUALIZATION

This section addresses why visualization is a promising approach to compliance management. This discussion draws on work in the field of social computing (e.g., Erickson and Kellogg¹ and Olson and Olson²).

The term *visualization* might refer to a wide range of representations, from static depictions of data to full interactive applications. We primarily discuss interactive visualizations, though some or even many of the cognitive, motivational, and social benefits of visualizations that we discuss could be realized as static displays. However, in most cases the benefits are enhanced when the user has the ability to filter, interact with, and manipulate the data. In the section “Lessons Learned,” we contrast dashboards (a static representation and not always of visualized data) with interactive visualizations. For an interesting example of an interactive visualization, see Reference 3.

Visualization at the individual level

There is a large body of research literature on information visualization and the cognitive effects of perceiving and conceptualizing information; summarizing it is beyond the scope of this paper.

■ A visualization should be thought of as a user interface to a control task, not as a report or a report component ■

Instead, we draw on the categorization and explanation of the benefits of visualization by Card et al., particularly how visualizations can increase people’s visual capabilities and amplify cognition.⁴ They outline six ways in which visualizations can provide benefits: (1) by increasing memory and processing resources available to the user, (2) by reducing the search for information, (3) by enhancing the detection of patterns, (4) by enabling perceptual inference, (5) by using perceptual attention mechanisms for monitoring, and (6) by encoding information in a malleable medium.

Cognitive psychologists have built up an empirical picture of these benefits over the last 30 years. For example, visualizations can increase working memory and the cognitive processing resources available to a user because some visual elements can be processed in parallel by the (relatively fast) human perceptual system;⁵ some work that would otherwise be performed by the slower and limited-capacity cognitive system can be offloaded to the perceptual system;⁶ and keeping pertinent details “in the world” rather than in the user’s head can increase the amount of working memory that can be devoted to problem solving.⁷

Visualizations can reduce the effort needed to search for information because they are dense, portraying a large amount of data in a small space.⁸ They can enhance pattern detection through aggregation and abstraction of data⁹ (when this is combined with the ability to focus down to details, it can be particularly effective). Visualizations support perceptual inference, making some problems obvious⁶ and allowing a large number of elements to be monitored. Finally, when visualizations are interactive and malleable, they can be directed by users to examine particular areas of interest.

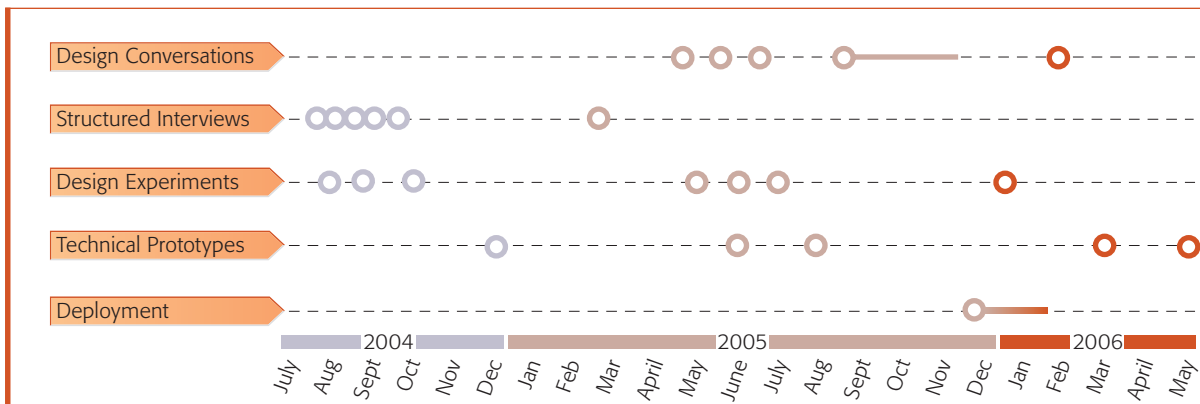


Figure 1
A schematic representation of the methods used in the design of the IBM SOX compliance visualization

Visualization for groups

Visualization is also a powerful tool when used by groups of people to collaborate, particularly across distances. Erickson and Kellogg¹ use the concept of *social translucence* to explain why this is so: Visualizing people and their activities leads to awareness and mutual accountability. Socially translucent systems create common resources by which people can more effectively coordinate their behavior. Such systems also support the emergence of social dynamics—such as peer pressure, imitation, the creation of norms, accountability, and other phenomena—that help to motivate collective effort toward common goals. In a series of studies of a socially translucent group-chat tool called Babble, Erickson and colleagues documented several effects on group process, including encouraging informal expression, minimizing the social overhead required for information exchange, and allowing remote members to keep in touch and quickly reestablish context.^{10,11}

Similar social dynamics can emerge from information visualizations. Erickson et al. describe the *task proxy*, a visualization of the state of a task by individuals throughout an organization.¹² In this representation, information about the task is in the foreground, so that its overall state can be discerned at a glance, but the elements of the task state reflect whether individuals have completed their bit of the task or not. Exposing such a visualization to the people in the organization generates social dynamics, such as peer pressure (“I’d better get this done; I’m almost the last one”), requests for assistance (“I see Tessa is already done; I’ll ask her if the update

messed up her system”), or offers of assistance (“No one in John’s group has started on this; I’d better let them know”). The visualization provides a mirror to the organization of its own behavior and common ground for task participants to discuss what is happening and to coordinate task completion.

Halverson and colleagues have taken a similar tack in creating a customizable interactive visualization for bug-tracking data in software development (see Reference 13 for a screen capture of the actual visualization). This visualization supports software developers and project managers in managing both technical and social issues. It does this because the visualization is compact, yet allows users to monitor for problems throughout an extremely large data set (over 10,000 bugs), and when shared among all the development team members, this visualization provides a basis for negotiating the priority of various bug fixes.

DESIGN PROCESS

The process of designing and deploying the risk and compliance visualization spanned almost two years. In this section, we describe the basic methods brought to bear on the process and the participatory approach we used to develop and deploy the design.

Figure 1 provides a high-level overview of the types and distribution of methods involved, and their relationship to the deployment. The key points illustrated by the figure are that we employed a number of different methods to advance the design process and that the different methods were often pursued in parallel (although there is a tendency for

those in the upper portion of the figure to be used earlier than those in the lower portion of the figure). While it is not uncommon for design to be portrayed as a linear process, in our experience, such interleaved parallel activity is the usual case, especially when the design process involves working with multiple stakeholders in various organizations, as this one did (see Reference 14). Some activities, such as design experiments, can be started immediately by members of the core team, whereas others, such as structured interviews, may require lead time to set up (arranging to interview the CEO of a major financial institution can require considerable political work as well as waiting for an opening in his or her schedule), and still others, such as deployment, may be constrained by organizational processes.

Methods used

We began by familiarizing ourselves with the design territory. In addition to reading background material and literature, we interviewed people involved in compliance management. In parallel, we began a series of design experiments, which helped us develop a sense of the possibilities afforded by visualization techniques. We then implemented an early technical prototype and moved fluidly between design experiments, technical prototype development, and design conversations with our users and stakeholders. This work eventually culminated in a deployment, although we continued our other activities.

Structured interviews

We used structured interviews early in the design process to learn how IBM executives were thinking about risk and compliance. These interviews helped us understand the risk and compliance domain in general and the concerns of executives (who, under SOX 302, must certify the compliance process) in particular. The executives were primarily concerned with achieving a standard and unified approach to controls testing and reporting processes used in different parts of the organization. They viewed unification as essential if they were going to be able to monitor, manage, and track the controls testing data for the current quarter.

Design experiments

Learning by creating visualizations was an important part of the design process. Some visualizations took the form of quick sketches to explore an idea,

such as exploring how to present a high-level view of an organization's compliance. Others were detailed explorations of a particular design feature; for example, we created four alternative designs to explore how to represent the state of compliance of a business process, and we did a series of sketches to investigate how to structure the compliance representation to best reveal interesting patterns in the data. We used storyboarding techniques to explore the details of interactions through a particular use of the visualization. For example, we created a storyboard to demonstrate to IBM controllers how they would use a compliance visualization to identify and focus on a particular defective control that needed attention.

Design conversations

We held detailed conversations with end users and stakeholders about specific design problems. The conversations ranged from e-mail exchanges to semiformal meetings, and topics ranged from general issues to discussions focused on a specific design issue, storyboard, or iteration of a design solution. For example, after meeting with controllers to learn about their process and discuss the design problem, we exchanged e-mails in which we asked specific questions about their process, such as how it was represented and encoded in the controls database, how they used the controls database to support their process, and their scheme for numbering controls.

Technical prototypes

To ensure that our designs were technically feasible, we built working prototypes of some that we were considering deploying. Many were user-interface prototypes to explore the interactions we were proposing. We also conducted studies of how to structure the Extensible Markup Language (XML) from which the visualizations are generated. Finally, we used our technical prototypes to investigate performance and communication between the underlying analytic components and databases that we were intending to use.

Participatory approach

Our purpose was to create visualizations that would support real-world risk and compliance processes. As a consequence, it was critical to deploy our prototypes to the actual people who would be using the real system and to provide access to the actual data.

Doing this introduced a number of difficulties. The real nature of the work and the demands of the organizational contexts in which it took place often shaped what we did and how we had to do it. For example, because of the quarterly financial cycle, our prototype had to be ready to deploy for the fourth quarter, even though it was not as fully featured as we would have liked. Running a prototype against the actual SOX database was another challenge. The technology staff who hosted the SOX database was unwilling to let us deploy technology on their working server, and so we had to run against a replicated copy of the database. This limited our ability to create a visualization that reflected up-to-the-minute information in the database, as we could replicate only twice a day without affecting database performance. Although we were reassured by the executive controllers that the update frequency was sufficient, it proved to be insufficient for other stakeholders in the controls process.

Furthermore, we wanted the controllers to feel that they owned this system. That was one reason that we involved them early, engaging them in interviews and in conversations about various design experiments and technical prototypes. In addition, we decided that we should make use of their process for deploying new technology. Thus, although we hosted the prototype on our server, they created the training materials for it, announced it, and distributed it by means of a link from the corporate controls Web site.

WORKING VISUALIZATION

Our goal was to create an interactive picture to capture and summarize compliance data. We wanted users to see at a glance the processes and controls that needed further investigation. A good visualization should reveal patterns in the data that are important for understanding the current state of the compliance process and the actions that need to be taken. For example, the visualization should make it easy to see common controls tested in different countries that have unremediated defects or to see controls that repeatedly have remediated defects each testing period. It should then be possible through interaction with the visualization to obtain the information needed for necessary action on the part of the user, such as contacting the person responsible for a control.

Rationale

We anticipated that a visualization would support the monitoring of control status and the analysis of patterns of defects. In particular, a visualization should enable stakeholders in the control process to answer the question, “How is our organization doing?” Rather than the typical dashboard approach of providing a high-level summary to answer this question, we thought that it was important to provide a picture so that people could actually “see” how the organization was doing. This was accomplished by two tactics: using small colored shapes to represent the data of interest and grouping the data closely together so that visual comparisons were facilitated. Thus, for example, grouping controls by business process would allow people to monitor a particular process and to compare that process with others.

These tactics were critical to support monitoring and pattern recognition throughout this large data set. Pattern recognition is a central benefit of visual processing, and an appropriate visualization can support this process, changing the task from a cognitive one to a perceptual one.⁶ A visualization that contains a visual representation of each control and groups these controls along dimensions that are important for the control process should make it easier to spot patterns in the control data. Patterns that could be shown in a compliance visualization include processes with defective controls, controls with defects that have the same control process but are run in different countries, and repeated control defect-remediation cycles.

Although it is important to get an overview of the entire data set, it is also essential to be able to restrict the view in order to see more detail. Thus, filters are important so that more specific questions can be asked. Time and place are always important and allow for such questions as: “How is the United Kingdom doing?” and “How did we do last month?” Similarly, specific information about a process or control is also necessary to follow up on observed problematic patterns. Thus, clicking on different parts of the visualization should reveal more information.

Deployed prototype

Figure 2 is a visualization for the IBM SOX compliance process. (It is for illustration purposes only; the data is fictitious.) At the top is a gray

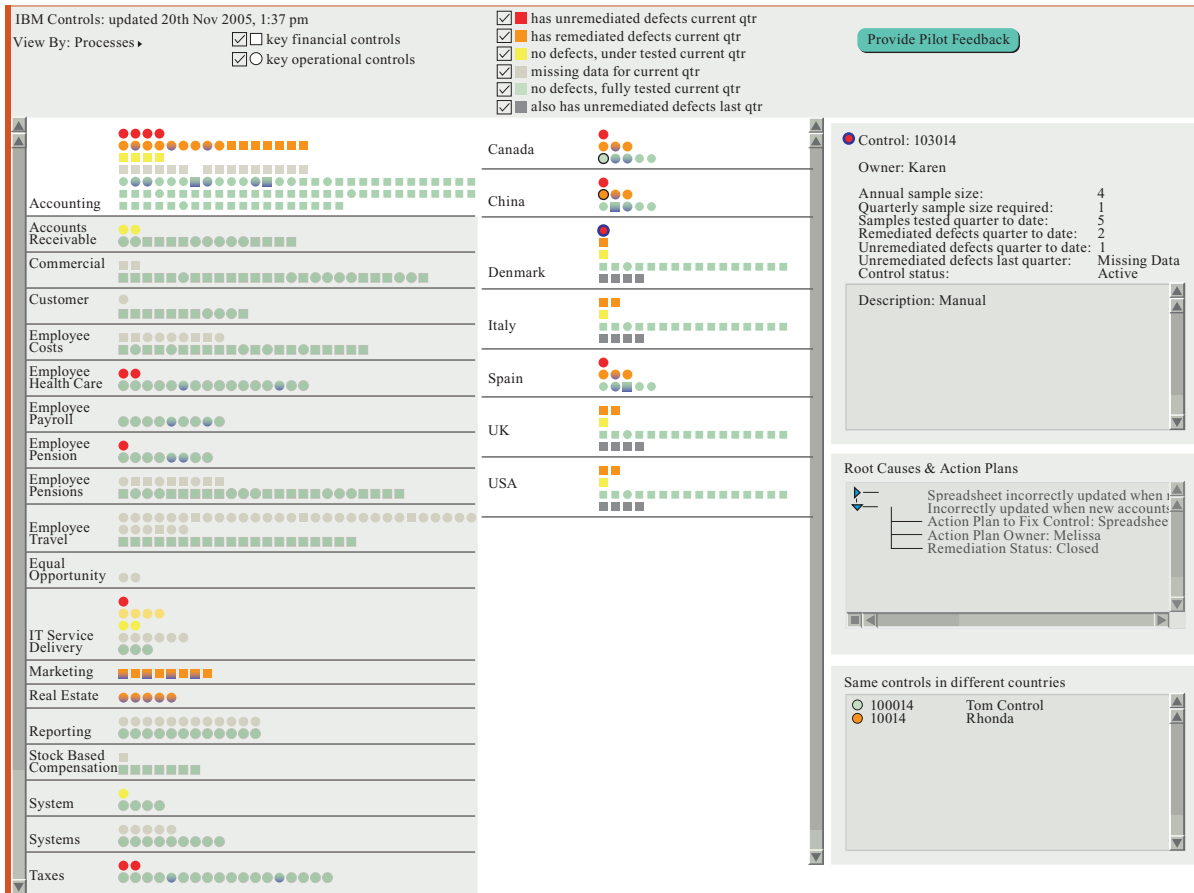


Figure 2
 The deployed SOX visualization (the data is fictitious)

header that contains the legend, a “View By” drop-down menu, and check boxes that allow the user to configure the control types shown in the visualization. Below the header, the screen is divided into three columns. The first column (left) contains a scrollable list of business processes ordered alphabetically. (In actuality, there are approximately 150 business processes.) Each colored shape represents a control for that process. Shading indicates whether the control had a defect the previous quarter, and its color indicates its status. The second column (middle) provides a more detailed view of a single process: It shows the controls for a selected process broken down by country. The third column (right) provides a view of the details of a single control. Thus, the visualization shows, from left to right, a view of all processes, a view of all controls of a single process broken down by country, and the details of a single control within a process.

An executive viewing the processes in the first column might wish to explore the topmost process, Accounting, noting that it has four magenta controls (i.e., controls with defects), and a number of other controls that are under-tested (yellow) or have no data at all (gray). Selecting this process would provide an expanded view of it in the middle column (as shown), with its controls grouped by country. Thus, the executive could quickly see that—while the controls in Italy, the United Kingdom and the United States are in fairly good shape—Canada, China, Denmark, and Spain give more reason for concern.

When the executive clicks on a particular control in the middle column, the details for that control appear in the right column. These details include a description of the control and the root causes and action plans for any defects. The control that is selected in the figure is in Denmark and is indicated

by a heavy blue border around the shape. Its details are shown in the right column. Some other controls in the middle column (in Canada and China) are also highlighted by a thin black border. These represent the same control, but run in different countries. These controls are also listed at the bottom of the right column.

From the drop-down menu on the left, the user can group the controls in the first column either by the business process, as shown in the figure, or by country. When the left-column view is by country, the middle column becomes a breakdown of the controls for the country by process. Using the check boxes, the user can turn the display of controls of a certain status or type on and off. Thus, for example, the user might choose not to show controls that are fully tested, and when he or she unchecks that option, the green squares are not shown in the visualization.

The end result is an interactive visualization that shows aggregate status at one instant with the ability to selectively explore particular elements in greater detail. A consequence of this approach is that users will have to negotiate a learning curve to interact with the visualization effectively. Without some training and experience, they may find it difficult to understand the significance of what is shown, to navigate the data, or to best configure the visualization to serve particular needs.

DEPLOYMENT

The deployment took place between December 2005 and February 2006. The deployed version was written in Squeak¹⁵ (an open-source version of Smalltalk) and required downloading a plug-in. The availability of the compliance visualization was announced on the front page of the IBM corporate controls Web site, and there was a link (with a “NEW” bubble icon) to a page describing the visualization and the pilot deployment. The visualization was updated twice daily with a current snapshot of control data.

Because we ran the pilot near the close of a financial quarter—a very busy period for our users—the pilot involved only 20 participants. The global controls coordinators, both those responsible for particular processes worldwide and those who were responsible for particular countries, were the main participants. They are responsible for monitoring

the controls for a specific process or set of processes, or a specific country or set of countries.

We used a number of methods to collect feedback. First, a “Provide Pilot Feedback” button on the visualization (see Figure 2) was used to point the participants to a wiki where they were able to provide written feedback. Second, we interviewed seven of the controls-process and country owners who participated in the pilot deployment. Third, we conducted an online survey that gathered input from seven different controls-process and country owners, asking them about their role in the controls process, their experiences using the controls database and scorecard, and their feedback on the visualization. Finally, we talked to the controls process executive team—the group that provided input during our participatory design process—about the deployed design.

LESSONS LEARNED

In this section we report on the lessons learned to date, based on the nearly two years of work described herein. This is a work in progress and is directed toward a very particular situation. Nevertheless, we believe that these observations can provide value to others pursuing similar efforts.

Visualizations as user interfaces to control tasks

A visualization should be thought of as a user interface to a control task, not as a report or a report component. The visualization was designed to support the controllers, and in particular, the team of controllers who were overseeing the SOX controls process. It was this team that participated in our design process, and consequently, their needs and influences shaped what we produced. When we were doing our design work, they were looking for a reporting solution, that is, a view that would substitute for the quarterly scorecard which they used to summarize their data. Because their main interest was in seeing compliance status across all their processes, we chose not to support the input of control data in the prototype, but, we learned this did not serve the needs of other controllers. Several controllers who did not participate in the design process reported that while they liked the visualization that provided an overview and allowed them to see the controls status and compare global controls across countries, they still needed access to the database. As one informant said, “Any time I see

something that doesn't look right, I need to do something."

From our post-pilot interviews and surveys of the participants, we learned that the controls testing, monitoring, and reporting tasks are complex. There are many subtasks and items that need to be tracked. Examples mentioned by our interviewees include: keeping track of how many tests need to be run and at what intervals; ensuring that defects are due to particular errors, as opposed to systemic issues with that control; and tracking the e-mail exchanges about control remediations and action plans. They also mentioned other activities that occur in support of compliance management: e-mail exchanges, phone meetings, surveillance of the CFO dashboard, forays into the controls database to find out detailed information about a particular control, and filling out forms to add information to the controls database. This suggests that controllers are already using many tools and doing many small subtasks as part of their overall responsibilities. This, in turn, necessitates continual switching between tasks and tools and leads to overload and frustration. Thus, although our visualization was quite attractive to those whose primary job was overseeing the state of compliance, it was "just one more tool" to those who were involved in managing details of the process. Nevertheless, we believe that an interactive visualization of the sort we designed has promise as a single online place that controllers can go to perform compliance-related tasks. In addition to providing input to the database, it should provide an integrated view and allow access to all the functions that controllers require. It would be even better if the visualization also provided integrated support to track the state of all the components of subtasks, even those initiated in a separate tool, such as e-mail; it would be useful if the visualization could allow controllers to see the e-mail trail for a particular control.

To create a visualization that provides this additional functionality is a considerably larger and more complex design task. The complexity is not so much due to the work of integrating the additional functionality as it is to the logistic and political issues raised by deploying an application that changes and updates the data in the actual corporate SOX database. Such changes would require that our work become an integral part of the complex and lengthy software development process for technol-

ogy that is created and deployed for IBM internal business support. The decision to begin our design work with a visualization that could be pointed at a cordoned-off replica of the real data was not just due

■ The patterns that are important for users to manage are at the level of the controls ■

to the interests of those controllers involved in our design process; it was also a politically feasible way to integrate a research system into the core practices of a working organization.

Visualizations as landscapes

Visualizations should show a landscape, not an isolated summary view. A common type of solution to manage risk and compliance is the executive dashboard.^{16,17} We started our design work by exploring such dashboards. They typically provide a series of portlets that show the summary status of different aspects of a business. The Hyperion dashboard¹⁸ is one of the best dashboards in the industry. The executive user can view the summaries reported in these portlets to get an overview of the current status of the organization and can hone in on any problem areas. Typically, a user can click on an item and be provided with more detailed information. From our interviews with executives early in the design process and from our own experiences designing such dashboards, we learned that such summary visualizations do not provide the executive with an adequate integrated picture of the status of the organization. While the dashboard provides iconic abstractions representing the data of the individual controls, the picture of the organization that it provides is at too high a level. The patterns that are important for users to manage are at the level of the controls. For example, executives told us that they needed to be able to see which controls are failing in multiple countries, or which countries have multiple controls with open action plans. Most important, the executives told us that they needed to see the specific defective controls, as these were the ones of concern. One commented: "Here's the real issue. Let me get underneath. I look at this report and I don't know whether I'm in trouble. The fact that I have an unremediated defect is not a show stopper." Reporting the data

abstracted away from the control details means that the executive must inspect each of these separate reports, determine the details, and integrate these details into a coherent mental picture. This is a lot of work.

Also, depending on the metrics used to create the visualization, the summary status can be misleading. Understanding the true status of an organization's compliance requires "seeing" the details underlying the summary. That summaries are misleading was evident in a number of our design explorations in which we attempted to use a single object to represent a whole body of processes. The downfall of this idea was that one was unable to see how many processes comprised that object. Often even the smallest of trouble spots needed immediate attention, but these small trouble spots were overpowered and lost in the status of the majority. In short, to monitor compliance, averages or other summary statistics are often not what is needed. One needs a display that brings the problem areas to the foreground and permits their immediate identification, even if they are proportionately small.

For example, in one of our initial designs, individual controls were not shown; rather a block representing the business process was shown, and the color of that block was to be determined either by averaging the status across controls or by showing the color of the control with the "worst" status. To see the status of individual controls required a mouse-over on a particular process. Even using the "worst" control status to determine the color of the process block did not give the controllers sufficient information to determine which process needed their immediate attention. Because it was important to know how many controls for a process had defects, using the average control status to determine color was not helpful either, as there might be the same number of controls with defects in a large block of magenta as in a smaller block of magenta. A heat-map visualization, by its very nature, draws attention to large blocks of magenta, whereas, for controllers, blocks of both sizes need their attention.

In response to this, we created a representation that depicted the actual controls along with the use of color to show status. With this, the controllers can visually inspect the representation and compare the relative size of the blocks of color to determine the extent to which a particular business process is in

compliance. The feedback we collected suggests that this representation appealed to the controllers and was readily understood. One drawback, however, is that several controllers did complain that the control shape was too small in size, and the difference between the shapes was hard to discern. This is the downside of a landscape approach, which puts all the information on one screen.

Although our initial analysis of dashboards had led us to discount summary statistics, feedback from the deployment caused us to reconsider this position and arrive at a more nuanced understanding. Controllers told us that summary statistics do provide a useful high-level orientation. One remarked, "A graphical view would be useful in reporting if I could capture a view and put it in a status report. Sometimes you get lost in the spreadsheet numbers, where a graph can make a point a whole lot better."

Because all controls could not be seen on one screen, users wanted to know the number of controls of different status, for example, how many controls had unremediated defects. Although summaries are a good way of reporting what happened during the quarter, they are not a good approach for day-to-day management of compliance. It does not help a controller to know that four controls have defects; the controller needs to know the specific controls, who owns them, and what is being done to remediate those defects. Further, knowing that there are four controls with defects this quarter and there were four controls with defects the previous quarter does not let the controller know whether these are the same four defects, and so on. Thus, while summary statistics are insufficient on their own, they are not without value. The best solution is a combination: summary statistics that provide orienting information and a landscape visualization that provides an overview of processes and controls and a way to obtain detailed information from an integrated context.

Visualizations as part of an evolutionary process

Visualizations must be able to evolve as the process for an individual or a group evolves or as the overall compliance process evolves. While the controls database is defined to support the generic controls process, individual processes may have exceptions. For example we were told by one of the process controllers that his process is not broken down in

terms of countries (as is defined by the generic process), but in terms of brands. To use the controls database, the unique structure of this process was mapped to the database structure. However, because this was done by creating an individual process for each brand, comparisons of controls across brands could not be made. Such a comparison would be useful, because the same control runs in different brands, just as for other processes the same control runs in different countries. A second example described by an interviewee is a process in which some of the controls are tested on a schedule that is different from the quarterly schedule set up in the database. This means that for some of the controls, there will be missing data one quarter, because that control is only tested twice a year. Having the visualization show “Missing Data” for that control is misleading, as the data is not missing; it is just not expected to be there. The information about how often each control should be tested is not kept in the database, but in a separate spreadsheet.

Not surprisingly, exactly what information controllers need to see depends on what they need to do. For example, executives who have responsibility for overseeing the whole control process need to have a view of all the processes and be able to see the proportion of controls with defects. They would also like to see how important a defect is in relation to financial risk. In contrast, controllers who are overseeing a specific country or process need to see just those controls for which they are responsible and the underlying details. It is these detailed views that are important to them, as they need to make sure that appropriate actions are being taken by control owners to remediate the control. Although we did not talk to any control testers, presumably they would want to see how their controls were doing, ensure that their status was accurately reflected in the database, and that no mistakes had been made in data entry.

In addition, for a particular controller or group of controllers, what they need to see and do depends on where they are in a reporting period. The work is cyclic, and the tasks change across the testing and reporting cycles. For example, controllers told us that during the quarter when they are monitoring the status of controls and checking on any remediations in progress, they need to be able to focus on the data for a particular control. At the end of the quarter, when they are reporting the results of the

testing, they need to see overviews and summary information. Being able to associate the kinds of views they are looking at with a personal or organizational time line would add significant value to such a visualization.

We had assumed that the overall compliance process was defined, thus defining the visualization needed to show the patterns in the compliance data. However, we learned through our experience working with the controllers that the compliance

■ Visualizations must be able to evolve as the process for an individual or a group evolves or as the overall compliance process evolves ■

process is continually being refined. As controllers go through the process and make use of the data, their understanding of how a compliance process might work evolves, their notion of what data can be collected changes, and they use the data to answer new kinds of questions. For example, when we interviewed process owners, they requested that trends be added to the visualization so that they could see the behavior of controls over time. They had recently started to experiment with viewing data this way and integrating trend analyses into their quarterly reports. When probed further, we learned that as the controls process had now been running for some time, there was sufficient historical data to make trend analysis worthwhile. We found this point interesting, as we had probed the importance of a historical view of the data in some of our original design sketches. However, because these design sketches had not received particularly positive reactions at the time they were presented, we had placed such views lower in our design features list.

The evolution of the controls process leads to new requirements for the visualization. This suggests that—just as the visualization needs to allow individuals and groups to create custom views to support their unique variants of the process—the overall visualization needs to be malleable as the generic controls process evolves. For example, a customizable visualization would maintain the

same visual elements, but allow these elements to be associated with different organizational structures. It would also allow visual features, such as color, to be associated with different control attributes.

Visualizations as support for communication and coordination

A single picture created from a “visual language” that maps to the domain elements can support communication and coordination among people

■ We learned through our experience working with the controllers that the compliance process is continually being refined ■

who are focused on different parts of the task and concerned with different granularities of analysis. Controllers work in teams and need to coordinate with respect to the current status of the controls. They reported that the scorecard is insufficient for coordination as it does not provide detail down to the control level. It is at the level of the individual control that problems typically need to be addressed, and communication among a control owner, country owner, and process owners about individual controls are typical. Controllers do coordinate by referencing the database to see a description of the defect. One controller pointed out that this description was missing from our visualization (a feature we needed to add immediately). Controllers also use the database for initial information on what actions have been taken regarding a particular control. Although controllers use the database in these ways, they spend a large amount of time communicating by e-mail and telephone. In particular, controllers whose responsibility it is to monitor processes or countries typically act as *alerters*, that is, they send an e-mail to the appropriate people when they need to attend to an issue at a particular time. Also, groups working on one process tend to have regularly scheduled status meetings to ensure that everyone is aware of any change to controls and of the progress made on problematic issues.

Our design participants felt that the visualization could support their coordination tasks and make

their alerting role easier. Unfortunately, due to limitations of the deployment (too few users and too short a deployment time), we were unable to verify this. However, we believe that the visualization could be extended to better support alerting activity with the addition of alert tracking and timing features. For example, allowing the controller to specify alerts of the form: “Send an alert to Marge Johnson if control 10 is not remediated by June 24, 2006.” Further evidence that such an alerting feature is important came from some process owners who stated that the visualization might reduce the number of e-mails being exchanged among those managing the overall SOX compliance process and the process owners. At the same time, they also raised the concern that such a visualization might lead to micromanagement, because it increased the desire to ensure that there would be no magenta on the display each quarter; that is, there would not be any controls with unremediated defects. They felt this would not be a good thing, as often there was a reason for a control to have an unremediated defect, and as long as that control was being managed appropriately, there was no need to bring additional pressure to bear on process owners (which would, in turn, require additional e-mails to explain what was happening). As one participant commented, “I don’t need anyone looking over my shoulder. ‘Missing’ could be misinterpreted by someone unfamiliar with the process. [...] I don’t want someone jumping to a conclusion. I don’t have that concern with the scorecard because no one is looking. [But] I like this visualization a lot better than the scorecard. I like this.” One way to manage this issue, explored by Erickson et al.,¹² might be to restrict the ability of users to see and communicate with owners of controls for which they are not directly responsible. Another possibility would be to refine the set of states for controls to include a classification for “unremediated but being effectively managed” controls.

The need for personalization and customized views is problematic from a social computing standpoint. For a visualization to enable coordination of work, it needs to provide a shared view of the information. This allows people to point to something on the screen and have a shared context for conversation. Once people start customizing views, the shared context is lost. As one informant put it, “I wouldn’t go completely away from the canned views. They are valuable, especially where reporting, [it] is

useful that everyone is looking at it the same way.” For this reason, we would argue that customized views need to be shareable. Furthermore, there needs to be a single visual language used so that the concepts and objects important for SOX controllers have a single and consistent visual referent. For example, in our visualization, the visual language represents an operational control as a small circle and a defect as a magenta color. These shapes and colors can be reused in different custom views, so long as the mapping between the visual language and the domain is maintained. Thus, a process name written in magenta text, for instance, would probably indicate to all people familiar with the visual language that the process had a defect.

CONCLUSION

Visualizations can play an important role in facilitating compliance processes throughout an organization. Reducing the complexity of the process by turning cognitive operations into perceptual ones is key to this simplification. Providing a picture that can be used as a common point of reference for conversation and communication is another crucial attribute of visualizations. This is particularly important when compliance is a global process involving multiple languages and cultures. While a single picture is important, individuals who have different roles and responsibilities within the process have different tasks they need to perform, and thus must be able to create custom views of the whole. These customizations can be achieved through filtering and support for multiple organizations. Custom views and filters must maintain the same visual language and be shareable so that they support effective communication among stakeholders about the status of the compliance process.

Organizational processes are dynamic; they evolve in response to organizational pressures and learning. Visualizations that support such processes must be malleable if they are to continue to be used as the process evolves. This requires support for customizing so that those most familiar with the changing process, typically the end users (the controllers in the case of the SOX visualization described here), can transform their visualization as the process changes.

In this paper, we have focused mainly on a specific example of a design process and the resulting system for an internal risk and compliance application. However, we believe that many aspects of both

the design process that we describe and elements of the solution hold as well for a wide variety of similar application areas. In particular, there are numerous other regulations that require organizations to assess and address risks and to ensure that they are in compliance. In the United States, such areas include tax laws, equal opportunity laws, Occupational Safety and Health Administration standards, accessibility standards, environmental standards, privacy, and ensuring the physical security of people, plants, assets, and data. Although these laws, standards, and regulations differ from country to country, the basic issues of awareness, comprehension, and compliance are universal. In each case, there is a need for cooperation and collaboration, the issues are complex, and a well-designed visualization can help. In all of these areas, a great deal of change has occurred as business and the regulatory environment have co-evolved, and there is every reason to believe that change will continue. Thus, fostering an organizational culture that is conducive to competence in these areas and providing a compelling interactive and visual way to approach issues of compliance can be expected to be more effective and less confusing than simply issuing employees an ever-changing set of rules and regulations.

ACKNOWLEDGMENTS

We would like to thank our management for the support of this work. We would also like to thank the IBM controller team for participating in our design process and for their enthusiastic support of this work.

CITED REFERENCES

1. T. Erickson and W. A. Kellogg, “Social Translucence: An Approach to Designing Systems That Support Social Processes,” *ACM Transactions on Computer-Human Interaction* 7, No. 1, 59–83 (2000).
2. G. Olson and J. Olson, “Distance Matters,” *Human-Computer Interaction* 15, No. 2/3, 139–178 (2000).
3. The Baby Name Wizard’s NameVoyager, iVillage, <http://babynamewizard.com/namevoyager/Inv0105.html>.
4. S. K. Card, J. D. Mackinlay, and B. Shneiderman, *Readings in Information Visualization: Using Vision to Think*, Morgan Kaufmann, San Francisco, CA (1999).
5. A. Treisman, “Preattentive Processing in Vision,” *Computer Vision, Graphics, and Image Processing*, 31, No. 2, 156–177 (1985).
6. J. H. Larkin and H. A. Simon, “Why a Diagram is (Sometimes) Worth Ten Thousand Words,” *Cognitive Science* 11, No. 1, 65–99 (1987).

7. D. A. Norman, *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*, Addison-Wesley Publishing Company, Reading, MA (1993).
8. E. R. Tufte, *The Visual Display of Quantitative Information*, Graphics Press, Cheshire, CT (1986).
9. S. K. Card, G. G. Robertson, and J. D. Mackinlay, "The Information Visualizer, an Information Workspace," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New Orleans, LA (1991), pp. 181-186.
10. E. Bradner, W. A. Kellogg, and T. Erickson, "The Adoption and Use of BABBLE: A Field Study of Chat in the Workplace," *Proceedings of the Sixth European Conference on Computer Supported Cooperative Work*, Copenhagen, Denmark (1999), pp. 139-158.
11. T. Erickson, D. N. Smith, W. A. Kellogg, M. Laff, J. T. Richards, and E. Bradner, "Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of Babble," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Pittsburgh, PA (1999), pp. 72-79.
12. T. Erickson, W. Huang, C. Danis, and W. A. Kellogg, "A Social Proxy for Distributed Tasks: Design and Evaluation of a Working Prototype," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Vienna, Austria (2004), pp. 559-566.
13. C. A. Halverson, J. B. Ellis, C. Danis, and W. A. Kellogg, "Designing Task Visualizations to Support the Coordination of Work in Software Development," *Proceedings of the ACM SIGCHI Conference on Computer-Supported Cooperative Work*, Banff, Alberta, Canada (2006), <http://jellis.org/research/jbe-cscw2006.pdf>.
14. T. V. Wolf, J. A. Rode, J. Sussman, and W. A. Kellogg, "Dispelling 'Design' as the Black Art of CHI," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Montréal, Québec, Canada (2006), pp. 521-530.
15. Squeak, <http://www.squeak.org>.
16. R. Brath and M. Peters, "Dashboard Design: Why Design is Important," *DM Direct Newsletter* (October 15, 2004), http://www.dmreview.com/editorial/newsletter_article.cfm?nl=dmdirect&articleId=1011285&issue=20081.
17. A. Marcus, "Dashboards in Your Future," *Interactions* **13**, No. 1, 48-60 (2006).
18. The Hyperion Compliance Management Dashboard, Hyperion Solutions Corporation, http://www.hyperion.com/products/business_intelligence/dashboards/compliance_dashboard.cfm.

Accepted for publication October 16, 2006.

Published online March 22, 2007.

Rachel K. E. Bellamy

IBM Research Division, Thomas J. Watson Research Center, 19 Skyline Drive, Hawthorne, New York 10532 (rachel@us.ibm.com). Dr. Bellamy is a research staff member. She has been practicing user-centered and participatory design for the last 15 years. She has a Ph.D. degree in cognitive psychology from the University of Cambridge. Dr. Bellamy's research focuses on the design and programming process.

Thomas Erickson

IBM Research Division, Thomas J. Watson Research Center, 19 Skyline Drive, Hawthorne, New York 10532 (snowfall@us.ibm.com). Mr. Erickson is a research staff member whose work involves studying and designing systems

that support computer-mediated communication and collaboration. His current interests include social software design, computer support for interactions among strangers, and collective intelligence systems. He has an M.A. degree in cognitive psychology from the University of California at San Diego. Mr. Erickson is a member of the editorial boards of the *Journal of Computer Mediated Communication* and the *Journal of Computer Supported Cooperative Work*.

Brian Fuller

IBM Corporation, 1 New Orchard Road, Armonk, New York 10504. Mr. Fuller works for IBM business controls and is responsible for IT strategy and tools that support the operations and consolidation of control metrics. He has a background in IBM financial and planning systems, with a strong focus on IT controls. He has a B.A. degree in management of information systems from Western Connecticut State University. Mr. Fuller is a member of the Institute of Internal Auditors.

Wendy A. Kellogg

IBM Research Division, Thomas J. Watson Research Center, 19 Skyline Drive, Hawthorne, New York 10532 (wkellogg@us.ibm.com). Dr. Kellogg is manager of social computing at the IBM Thomas J. Watson Research Center, where her work focuses on computer-mediated collaboration systems for supporting work in organizations. She has a Ph.D. degree in cognitive psychology from the University of Oregon. Dr. Kellogg is a member of the IEEE and the Association for Computing Machinery (ACM) and a member of the ACM Queue editorial board. She was elected an ACM Fellow in 2002.

Rhonda Rosenbaum

IBM Research Division, Thomas J. Watson Research Center, 19 Skyline Drive, Hawthorne, New York 10532 (rhonda1@us.ibm.com). Ms. Rosenbaum is a software engineer. Currently her work focuses on data management, processing, and restructuring of data for the SOX visualization project. Ms. Rosenbaum has an M.S. degree in computer science from Polytechnic University.

John C. Thomas

IBM Research Division, Thomas J. Watson Research Center, 19 Skyline Drive, Hawthorne, New York 10532 (www.truthtable.com). Dr. Thomas is a research staff member working on understanding, measuring, and reducing psychological complexity. Recent work includes developing a socio-technical pattern language, the business uses of stories and storytelling, and designing the user experience for a dynamic learning environment. Dr. Thomas has a Ph.D. degree in experimental psychology from the University of Michigan and is a licensed psychologist in New York State.

Tracee Vetting Wolf

IBM Corporation (tlwolf@us.ibm.com). Ms. Vetting Wolf was a designer at the IBM Thomas J. Watson Research Center. Her research interests included understanding how to support social interactions and interpersonal collaboration gracefully within online social settings. Ms. Vetting Wolf recently joined Lotus as a product designer for real-time collaboration products. She has a B.S. degree in graphic design and a Master of Architecture degree from the University of Minnesota. ■