Understanding Alliance Portfolios Using Visual Analytics

RAHUL C. BASOLE, TIMOTHY MAJOR, and ARJUN SRINIVASAN, College of Computing, Georgia Institute of Technology

In an increasingly global and competitive business landscape, firms must collaborate and partner with others to ensure survival, growth, and innovation. Understanding the evolutionary composition of a firm's relationship portfolio and the underlying formation strategy is a difficult task given the multidimensional, temporal, and geospatial nature of the data. In collaboration with senior executives, we iteratively determine core design requirements and then design and implement an interactive visualization system that enables decision makers to gain both systemic (macro) and detailed (micro) insights into a firm's alliance activities and discover patterns of multidimensional relationship formation. Our system provides both sequential and temporal representation modes, a rich set of additive cross-linked filters, the ability to stack multiple alliance portfolios, and a dynamically updated activity state model visualization to inform decision makers of past and likely future relationship moves. We illustrate our tool with examples of alliance activities of firms listed on the S&P 500. A controlled experiment and real-world evaluation with practitioners and researchers reveals significant evidence of the value of our visual analytic tool. Our design study contributes to design science by addressing a known problem (i.e., alliance portfolio analysis) with a novel solution (interactive, pixel-based multivariate visualization) and to the rapidly emerging area of data-driven visual decision support in corporate strategy contexts. We conclude with implications and future research opportunities.

CCS Concepts: • **Information systems** → **Decision support systems**; • **Human-centered computing** → **Visual analytics**; **Information visualization**; • **Social and professional topics** → **Computing and business**;

Additional Key Words and Phrases: Alliances, strategy decision support, visual analytics

ACM Reference format:
Rahul C. Basole, Timothy Major, and Arjun Srinivasan. 2017. Understanding Alliance Portfolios Using Visual Analytics. *ACM Trans. Manage. Inf. Syst.* 8, 4, Article 12 (August 2017), 21 pages.
https://doi.org/10.1145/3086308

1 INTRODUCTION

In an increasingly global and competitive business landscape, firms must collaborate and partner with other firms to ensure survival and growth (Gomes-Casseres 2015; Adner 2017). Corporate

This research was supported by the Tennenbaum Institute and the Institute for People & Technology at Georgia Tech, a Batten Fellowship provided by the Batten Institute at the University of Virginia Darden School of Business, and the Center for Global Enterprise.
Authors' addresses: R. C. Basole, School of Interactive Computing and Tennenbaum Institute, Georgia Institute of Technology, 85 Fifth Street NW, Atlanta, Georgia 30332, USA; email: basole@gatech.edu; T. Major and A. Srinivasan, School of Interactive Computing, Georgia Institute of Technology, 85 Fifth Street NW, Atlanta, Georgia 30332, USA; emails: {tmajor3, arjun010}@gatech.edu.

https://doi.org/10.1145/3086308

relationships enable a firm to create value that it cannot achieve on its own, provide access to new markets and technologies, and accelerate innovation (Gulati 1998). According to a recent study, nearly 50% of U.S. chief executive officers said they planned to enter a new strategic alliance with customers and suppliers as well as competitors, start-ups, and firms from different industries over the next year (PwC 2015).

While interfirm relationships, alliances, and strategic networks are extremely well studied topics, there still is a remarkably limited understanding on the evolutionary composition of a firm's alliance portfolio (Lavie 2017). An alliance portfolio can be defined as a firm's entire collection of direct partnerships (Das and Teng 2002). Decision makers learn about interfirm relationships one at a time, gaining only a partial perspective of a firm's interfirm activities and failing to see the big picture. This has serious business implications as decision makers may underestimate or entirely miss competitive strategies of other business ecosystem players (Basole 2014). In this study, we posit that if a firm's alliance portfolio is appropriately mapped, decision makers can identify relationship "traits" and "behaviors," enabling them to adjust their own strategies and even mimic successful formation strategies while avoiding bad ones. The ability to – timely, effectively, and interactively – understand current and future relationship formation strategies of partners and competitors is, thus, of strategic significance. To address this issue, we designed, developed, and evaluated an interactive, web-based alliance portfolio visualization system.

Our study makes important contributions to our understanding of visual decision support in a corporate strategy context, a topic situated at the intersection of information systems (IS) and strategic management. From a design science perspective (Hevner et al. 2004), our study describes the intricate journey of designing and implementing a visual exploration tool for temporal multivariate business data. Following (Goes 2014) classification of design science contributions to knowledge, our study focuses on a *known problem* (i.e., systematic understanding of alliance portfolios), with a *new solution* (i.e., interactive, temporal/sequential pixel-based visualization technique). Specifically, our system allows users to align alliance sequences, both temporally and sequentially. Second, the system enables "stacking" of multiple alliance portfolios to facilitate rapid visual comparison. Third, we allow users to dynamically filter all aspects of a firm's alliance portfolio, facilitating deep probing of the underlying data. Lastly, we computationally mine and visualize activity state transitions to allow users to understand possible strategic relationship moves. Collectively, our study thus contributes to the problem-solving paradigm advocated by leading design science scholars (Hevner et al. 2004; Peffers et al. 2007) and addresses the call for impactful design science research in the IS field (Gregor and Hevner 2013). From a practitioners perspective, our system allows decision makers to explore and compare alliance portfolios, gain insights into the multidimensional nature of partnership strategies, and make data-driven predictions about probable relationship formations and their timing. In doing so, we contribute to the nascent but rapidly emerging domain of visual enterprise analytics.

The remainder of the manuscript is organized corresponding to the design science schema advocated by Gregor and Hevner (2013). Section 2 presents related work. Section 3 describes our research method, including the underlying data, the design requirements, and the overall design process. Our design artifact is described in Section 4. Section 5 presents the evaluation. We conclude with a discussion and opportunities for future research in Section 6.

2 RELATED WORK

2.1 Alliance Portfolios

While prior work on alliances has predominantly focused on dyadic or network aspects, there has been an increased interest in understanding alliance portfolios (Hoffmann 2007). Alliance portfolios – defined as a firm's set of direct relationships – have been shown to be beneficial for risk reduction in uncertain business environments and enhanced knowledge exchange in knowledge-based industries (Ozcan and Eisenhardt 2009). Prominent topics include the (re)configuration, size, complexity, partner diversity, and geographic scope of alliance portfolios (Wassmer 2010). These issues are generally examined in relation to performance. Results that properly configured and managed alliance portfolios have significant positive effects on financial and innovation performance.

Most prior studies, however, tend to focus on a singular aspect (e.g., complexity or internationalization) rather than examining the various issues holistically. Similarly, they tend to focus on firms in individual industries, such as software (Lavie 2007), automobile (Jiang et al. 2010), or biotechnology (George et al. 2001), rather than cross industry comparisons. Possible explanations include the lack of comprehensive, integrated data and sophisticated analytical tools. However, with an increasing blurring of industry boundaries, an understanding of alliance strategies across the entire business ecosystems would be beneficial.

In this study, we overcome these issues by developing an interactive visual analytic tool that incorporates multiple empirical lenses of interest (e.g., size, scope, geography, industry). Specifically, our study addresses the problem of visualizing alliance portfolios, focusing on the temporal/sequential, multivariate, and geospatial aspects of alliance portfolios. In doing so, we augment decision makers' capability to comprehensively and systematically understand and compare alliance activities.

2.2 Data Visualization

Visualization of business data is an emerging area of interest to information systems research (Chen et al. 2012; Basole et al. 2015). Most prior work on alliances and interfirm relationship visualization has typically employed network representations (Basole et al. 2013). Fewer visualization studies have examined interfirm relationships from a longitudinal or compositional perspective. Given the temporal, multidimensional nature of alliance portfolios, we build on two streams of visualization research: (1) time and event data and (2) multi-attribute data.

2.2.1 Time and Event Data. Time and event-based data visualization has been a topic of great interest to the visualization community (Aigner et al. 2007, 2008, 2011). Prior work has focused on specific aspects of temporal data including cyclic time (e.g., SpiralGraph (Weber et al. 2001)), time points (e.g., TimeWheels (Tominski et al. 2004)), intervals (LifeLines (Plaisant et al. 1996)), and ordered (e.g., ThemeRiver (Havre et al. 2002)) versus branched times (e.g., PlanningLines (Aigner et al. 2005)). Our study extends this work by examining multi-attribute event data (e.g., alliance events). We draw inspiration from studies such as LifeFlow (Wongsuphasawat et al. 2011) and Outflow (Wongsuphasawat and Gotz 2012), which provide a visual overview of event sequences; LiveRAC (McLachlan et al. 2008), a visualization system that supports the analysis and side-by-side comparison of time-series data at multiple levels of detail; and EventTunnel (Suntinger et al. 2008), an interactive visualization of event streams for business process pattern analysis.

2.2.2 Multi-Attribute Data. A similarly rich area of visualization research involves multi-attribute data. The Value and Relation (VaR) system (Yang et al. 2004), for instance, leverages dense pixel displays to visualize large sets of data. By focusing on the sequence of events, an idea central to alliance portfolios, Fischer et al. (2012) are able to condense the representation to view large time frames. Tominski et al. (2012) suggest an alternate approach by stacking multi-attribute data and integrating time through appropriate ordering of bands. One particularly relevant approach to visualizing alliance portfolios, and similar to pixel displays, is to use a matrix-based attribute approach due to the clear and flexible metaphor of matrices (Alsallakh et al. 2014). Examples include ConSets (Kim et al. 2007), which orders sets and elements to rows and columns, respectively, OnSet

(Sadana et al. 2014), in which elements are explicitly assigned to each cell in the matrix, Frequency Grids (Micallef et al. 2012), where glyphs are placed into each cell, or WebCANVAS (Cadez et al. 2000), where rows depict an entire sequence, and each cell within a row is color coded by category.

3 METHODOLOGY

3.1 Data

We use Thomson Reuters SDC Platinum[1] (henceforth, SDC) for our study. SDC is considered one of the most comprehensive data sources for the study of alliances and interfirm relationships and is widely used by researchers and practitioners (Schilling 2009). SDC contains detailed information on nearly one million global alliances, joint-ventures, and acquisitions in over 1,000 industry segments from 1970 to the present. The underlying data is curated from U.S. Securities and Exchange Commission (SEC) filings, trade publications, wires, and news sources, validated by an extensive, expert staff, and updated monthly. Each alliance is described by numerous characteristics, including dates (announced, signed, terminated, extended), relationship type (e.g., strategic, marketing, research & development, licensing, supply, tech transfer), corresponding market segments (based on Standard Industry Classification (SIC) codes), and a textual synopsis of the relationship content and terms.

Following discussions with corporate decision makers, we focused our study on relationships formed by all firms in the Global S&P 100 and IT firms in the S&P 500 since its inception. We chose this study context as the ICT space is perhaps one of the most dynamic ecosystems of recent times, with new actors emerging frequently and the need to form partnerships to survive is well documented. Moreover the two S&P lists provide a set of well-established global companies with rich and extensive alliance portfolios. In total, we sequenced the alliance portfolio of 155 enterprises for our study. One hundred one unique firms come from the Global S&P 100 (because companies drop in/out over time); 54 come from the S&P 500. We followed this by extracting all alliances for these companies formed between 1990–2015. We chose this time period as it includes three very dynamic industry life cycles, including the dot-com boom (1995–2000), the millennium cycle (2000–03), and the financial crisis (2007–08). Our final dataset consisted of 11,597 relationships (70% strategic, 28% R&D, 21% manufacturing, 4% supply, 16% licensing, 24% marketing, and 14% tech transfer) with 5,858 unique partner companies (49% U.S.-based). Twelve percent of relationships involved multiple partners; 38.9% were cross-region and 47.7% were cross-country. Partner companies came from 523 different market segments (at four-digit SIC code level).

3.2 Design Requirements

While SDC is a rich data source, its current user interface has major limitations. First, it is a standalone application and data is accessed and explored through a cumbersome query-based approach. Second, results are presented through massive list-based reports, with each row depicting a single relationship, that have to be exported as messy comma separated files or highly unstructured text files. This output format makes systemic understanding of a firm's alliance portfolio immensely challenging. Third, given this output format, comparison across companies is virtually impossible without some significant data wrangling.

The design and development of our system was motivated by these limitations and guided by a two-phase field study with corporate strategists and strategy researchers. In doing so, our aim was to follow a broader approach to understanding and eliciting requirements (Burnay 2016). In the first phase, we conducted eight 60-minute face-to-face interviews with corporate strategy experts

[1] http://tmsnrt.rs/1XFvhUV.

(five practitioners and three scholars), all actively pursuing competitive intelligence analysis tasks. All experts had extensive work experience (>15+ years). The objective was to understand primary data sources, the broader context of their analysis tasks, the tools they are currently using, as well as their critical and desired functionalities for a visual analytic system. We analyzed and coded the responses to identify key design requirements evident from the interviews. We asked participants to use a think-aloud protocol (Nielsen 1994); their responses were digitally recorded and transcribed. Using a topic frequency analysis we identified common concepts, words, themes, and requirements. Responses were coded as "must have" to "nice to have" along data, functionality, and analysis dimensions. The researchers clustered these responses after discussion and resolved any ambiguities. In the second phase, these requirements were sent back by email to the same eight experts for validation, prioritization, and refinement. The final list consisted of six core requirements that drove our overall design:

— **R1: Provide coordinated macro ("big picture") and micro ("nuts and bolts") perspectives of alliance portfolios.** One of the fundamental tasks of competitive intelligence is to understand the overall composition of a firm's alliance portfolio as well as individual relationship details. Analysts want to see both a holistic picture as well as individual activities. As one practitioner put it, "the big picture allows us to see patterns, but we need to be able to drill down to specific relationships to understand their content." The visualization tool, thus, should allow an analyst to understand patterns of relationship activities as a whole, while a detailed view should provide the ability to view individual relationships.
— **R2: Provide a longitudinal (temporal) perspective of alliance formation patterns.** Firms form relationships over time. It is therefore not surprising that every expert participant emphasized the need to see a longitudinal, temporal perspective of a firm's relationship activities. Two interesting side issues emerged from this design requirement. First, our experts did not prefer traditional line charts to depict the evolution of events. While line charts showed the cumulative growth of events relationships there was a strong consensus that they did not adequately capture the multidimensional nature of relationships and "missed the essence of what makes a relationship." Secondly, most experts (5/8) operated using two different temporal modes: a true-time and a sequential mode. Most existing tools used the former mode to understand when events occurred. However, they felt it masked their ability to quickly see competitive moves/actions and instead wanted to see a sequential perspective of the data implemented. If provided, they also preferred the ability to quickly switch between the two temporal modes.
— **R3: Enable rapid exploration of the multidimensional nature of an alliance portfolio (type, scope, geography, market, complexity).** All relationships are multidimensional in nature, formed for different purposes, with different types of partners across different markets and countries. All experts expressed the need to rapidly explore the multidimensional characteristics of relationships. Specifically, they wanted to understand the composition, the scope, and the complexity of the relationship activities. Based on our field study, current tools require users to specify the characteristics of relationships of interest through a query; results are then returned in tabular format. This, however, can take valuable time and experts urged us to think of ways to accelerate the process, perhaps through point-and-click designs.
— **R4: Facilitate high-level, quick comparison of alliance portfolios.** A secondary but also important task of competitive intelligence is to compare multiple alliance portfolios. One expert argued that she often wants to "simply [...] visually inspect the portfolio of two firms and gain a quick [pre-attentive] understanding of commonalities and differences."

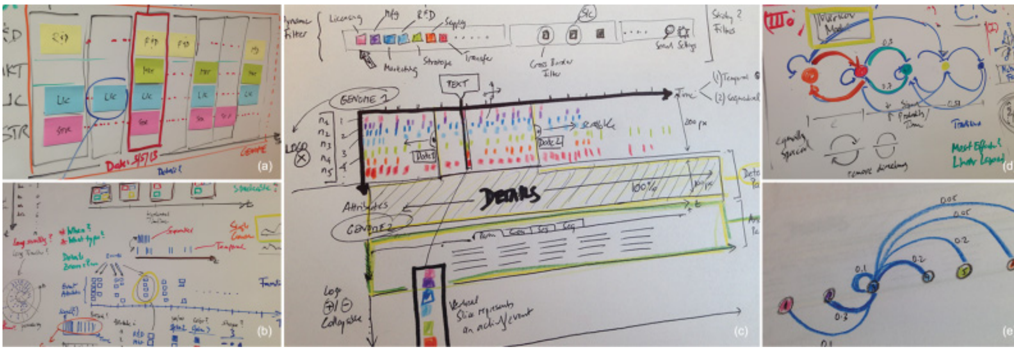

Fig. 1. Following guidelines by McKenna et al. (2014), the development of our visualization involved an iterative process of design ideation, sketching, prototyping, and refinement between members of the research team and three prototypical corporate users. As corporate analysts are often not familiar or clear with potential visualization alternatives, we used dry-erase boards (a,b, and d), color paper sketches (c,e), and post-it notes (a) to brainstorm potential visualization ideas.

Another expert was interested in the ability to compare a small set of portfolios at once, and then rearrange or remove portfolios from the analysis as desired.

— **R5: Provide capabilities to understand likely paths and timing of alliances.** Perhaps the most intriguing desirable design requirement was the ability to integrate analytic capabilities into the visual representation. The primary focus was to provide data-driven evidence into competitive moves. As one expert put it "it would be great to harness all this data for predictive capacities. I would like to know what alliance a firm may pursue next and when."
— **R6: Provide a simple, intuitive, and familiar interface.** While all experts were familiar with traditional financial and business analytic tools (e.g., Google Finance, Tableau, Bloomberg, R), they also emphasized that they were not real power users of interactive visualization tools. None of the experts preferred command line tools. The key design requirements they suggested are that the tool had to be simple, intuitive, and familiar in order to be readily adopted by other prototypical users. They also mentioned that they were particularly comfortable with click-and-point interfaces, but stressed that simplicity was key over having complex functionalities.

3.3 Design Process

The development of our visualization involved an iterative process of design ideation, sketching, prototyping, and refinement between members of the research team and three corporate users. As decision makers are often not familiar or clear with potential visualization alternatives, we used dry-erase boards, color paper sketches, and post-it notes to brainstorm potential visualization ideas (see Figure 1).

The core challenge of our research was to design an intuitive and easy to understand representation that appropriately captured the multidimensional and temporal nature of alliance activities and was sufficiently compact that it could be "stacked" for comparing multiple portfolios (R1, R5, R6). We drew much of our inspiration from design suggested in prior work, including CloudLines (Krstajić et al. 2011), event sequence simplifications (Monroe et al. 2013), rectangular views (Nocke et al. 2003), PatternFinder (Fails et al. 2006), and Pixel Bar Charts (Keim et al. 2002). The research team came up with several glyph design alternatives that were identified as particularly effective in prior research on multidimensional visualization (Borgo et al. 2013). It was evident from the

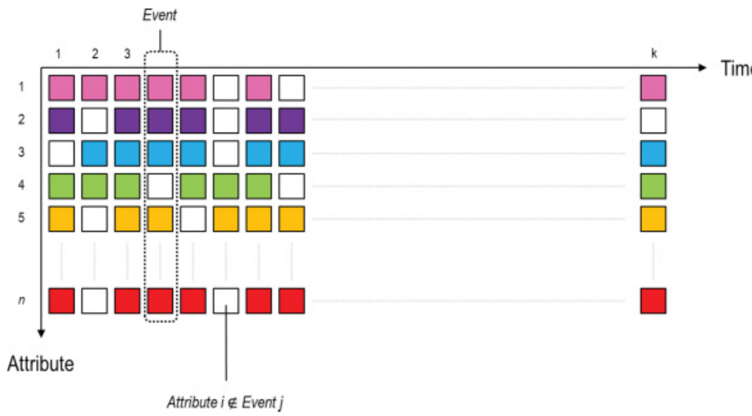

Fig. 2. Conceptual illustration of an alliance portfolio visualization. Alliance events are arranged horizontally (time axis), while event attributes are arranged vertically (attribute axis). The grid approach enables fast identification of patterns and outliers within a portfolio. Binary attributes are color-encoded. If attribute is not present, it is left blank.

discussion with our users that the temporal aspect of the alliance portfolio was perhaps the most critical design bottleneck; the visual representation had to enable users to quickly understand and compare the temporal and sequential order of events (R2). This requirement thus constrained our design alternatives. Prior research has advocated several alternative methods for timeline design (e.g., 2D and 3D spirals, wrapping). However, following some informal tests we conducted, we observed that decision makers felt less comfortable with any of the more eccentric design methods and advocated us to consider a linear timeline approach due to general familiarity. We followed their suggestion and ultimately used a traditional horizontal axis to denote time as it not only best reflected the cognitive model of prototypical users but also utilized screenspace more effectively when stacking multiple portfolios.

A subsequent design issue then included how to best encode the multidimensional aspects of a relationship activity (R3). Given the many different combinations but individual importance of alliance types, a simultaneous depiction of all dimensions was deemed critical. Prior research suggested many different multidimensional visualization techniques (e.g., parallel coordinates, starplots, small multiples) The key tradeoff in choosing an appropriate encoding was that it had to accommodate a large number of dimensions (if needed), provide some level of compactness to allow for reasonable stacking of alliance portfolios, and effectively handle the horizontal time aspect. The research team went through several different design iterations before settling on a vertically stacked pixel-like approach (see Figure 2). Based on the feedback, this glyph (i.e., visual design) met all the characteristics of compactness, stackability, temporal orientation, and readability/interpretability corresponding to our six elicited design requirements.

Given that alliance types are binary in nature (yes/no), we opted to choose a bright, categorical color scheme. Colors were chosen from the Google Charts color scheme. We also used a dark visualization background as it mimicked existing financial analysis tools (Dang and Lemieux 2013), allowed attribute colors to stand out, and created a vibrant interface.

4 SYSTEM DESCRIPTION

Our alliance portfolio visualization tool consists of two main regions (shown in Figure 3). The **top region** (Figure 3(a)) consists of a rich global filter bar, which allows users to (1) set date ranges,

Fig. 3. Our system consists of two regions. The top region (a) contains the global filters. The bottom region (b) acts as the area for stacking alliance portfolios. Each portfolio has three panes. The overview pane (c) provides the overall portfolio. Details are provided in the details pane (d). Derived or computed information from the portfolio is presented in (e). Results are organized in pages (f) for readability.

(2) choose the type of display mode (sequential or temporal), (3) select relationship types (licensing, marketing, manufacturing, R&D, strategic, supply, and tech transfer), (4) select cross-border (region or country), (5) cross-SIC (2,3, and 4-digit level), and (6) multi-partner relationships. These filters were selected and placed in this order per user feedback. The filter bar is pinned to the top of the browser window, facilitating access to it without having to scroll even after adding multiple portfolios.

The **bottom region** serves at the main visualization area where one (or multiple) alliance portfolio can be stacked. Portfolios can be added through a drop-down list of corporate logos with incremental search (filter-as-you-find) capability. We considered several different ways of adding portfolios to the visualization area, including simple search and drag-and-drop, but found the incremental search/drop-down list combination to be most effective as the total number of alliance portfolios increased. We chose to use both corporate logos and text descriptions to leverage corporate brand recognition and facilitate user recall. Logos were obtained from brandsoftheworld.com. The number in line, right-aligned to the label/corporate logo indicates the size of the alliance portfolio.

4.1 Overview Panel

The overview (or context) panel (Figure 3(c)) displays the entire alliance portfolio. There are two modes of display: sequential (default) and temporal. The sequential display mode orders events from left-to-right without an explicit position of time; the temporal display mode time-aligns the events. A horizontal axis above the portfolio provides a reference to the sequence number or date of the alliance, corresponding to the selected display mode. The numbers to the left of the portfolio indicate the total number of alliance types in the portfolio. Hovering over a relationship type number emphasizes that type while fading all others, enabling a user to explore the pattern of a specific attribute. Users can also hover over alliance entries in the portfolio and a tooltip appears with information on the firm(s) involved, the date the event occurred, the underlying alliance types, as well as a synopsis text of the relationship (see Figure 6(b)).

4.2 Focus Panel

Alliance portfolio details and analyses are provided through two on-demand panels. The focus panel (Figure 3(d)) appears by either clicking-and-dragging a selection window (rubberbanding) over the portfolio or by clicking on the details nub below the logo. In early design versions, we contemplated placing the focus panel above the portfolio as it is commonly done in other financial/ business focus+context visualizations. However, this design conflicted with the visual flow of stacking portfolios and, based on user feedback, was the less preferred location for the detail pane. In doing so, our design follows the well-established information visualization mantra "overview first, details on demand" (Shneiderman 1996) and also our participants' request to provide "simplicity first, complexity on demand." If the user opts to use a rubberbanding approach to select parts of the portfolio, a movable, stretchable selection window appears. This allows users to quickly explore and browse the entire alliance portfolio and get dynamic details. The start and end values are shown in tooltips to give the user feedback on the selection range.

4.3 Analysis Panel

The lower panel (Figure 3(e)) contains the analysis results derived from the corresponding portfolio. It contains five interconnected tabs (partners, market segments, countries, regions, and sequence), which act as cross-linked filters that interact with the alliance portfolio visualization. The tab categories reflect the key analyses areas users typically wish to explore when understanding the portfolio, as identified during our design requirements phase. We initially considered placing tabs across the panel top, but user feedback suggested to place them on the left for ease of navigation and readability.

4.3.1 Partners. The **Partners** tab shows the top partners of the focal enterprise in a list-based results format. Results are shown as company name and number of occurrences in the selected part (date range) of the portfolio. The default sort order is by occurrence count, however, it can also be sorted alphabetically. We added this option as users sometimes wanted to quickly find an entry that was not a top entry. If the analysis list is longer, results are provided in subsequent pages. We chose to organize results this way to limit the total information shown at once and to mimic representation of familiar search tools. In Figure 3, for instance, we see the top partner companies for Hewlett Packard based on the selected portion of the portfolio.

4.3.2 Segments. The **Segments** tab contains information of the key market segments involved in the portfolio. We use four-digit SIC codes and their descriptions (taken from osha.gov) to label market segments. The default sort order is by number of occurrences, however, it can also be sorted alphabetically.

4.3.3 Countries and Regions. Given the increasingly global nature of business, understanding of countries and regions in an alliance portfolio are important. We thus provide two geographic analyses: the **Countries** tab lists the partner's headquarter country; the **Regions** tab groups countries into major geographic regions based on World Bank classifications. Both lists are default sorted by occurrence count and are sortable alphabetically as well.

4.3.4 Sequence. The final tab facilitates **Sequence** analysis and includes an activity state visualization of a firm's probability and timing of alliance activities. The inclusion of this tab was motivated by the feedback we received from experts who wanted to understand a firm's strategic moves and their timing.

The sequence of alliances can be modeled as a simplified Markov process model (see Appendix A.1. for details). We represent transitions as an arc diagram. The nodes depict the activity

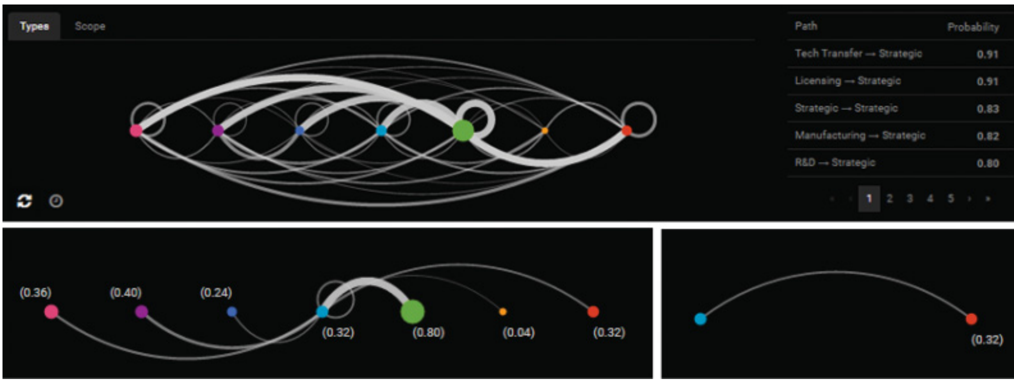

Fig. 4. Three different activity state visualizations are shown. Users can select to either show transitions between alliance type or scope. In these examples, we only depict alliance types. Edges can represent the probability or the average time. The user can switch between these encodings using the icons in the lower half of the pane. A sortable details list complements the data encoded in the visualization.

states, while the edges depict either the likelihood (percentage) or the timing (days) of a transition between two states. Directionality of edges is encoded in clockwise fashion (i.e., edges in the upper half of the arc diagram indicate transitions from left to right and edges on the bottom represent transitions from right to left). Thickness of the edge correlates to the probability of going from one state to another or the average number of days taken to transition between two states. The transition probabilities and transition duration are computed based on the active selection window in the alliance portfolio. This allows an exploration of how a firm's relationship forming activity has changed over different time periods. Hovering over a node highlights all in- and out-bound edges and corresponding edge values are shown. Single clicking a node highlights its connections and freezes the selection. Clicking a node connected to the selected node highlights and freezes the selection to show only the transition between that pair of nodes. Double-clicking on the canvas around the arc diagram clears all selections. The arc diagram is linked to the alliance portfolio, hence, any selection of nodes or edges is added as a filter.

The sequence tab (Figure 4) is organized into three sections. At the top, there are two tabs that allow users to switch between the two state categories (Type and Scope). At the bottom left are icons that allow users to switch the transition encoding between probability (likelihood) and average days (timing). The right panel provides the values encoded by the arc diagram as a sortable list.

4.4 Interactions and Filters

Our system integrates several key interaction techniques. It allows for focus+context, selection through rubberbanding, dynamic filters and previews, rearrangement of portfolios, details-on-demand, and multi-selection. Moreover we use animations to help transition from sequential to temporal modes. All regions and tabs within a portfolio are dynamically linked, facilitating brushing and linking. Portfolios can be removed by clicking on the [x] icon next to the corporate logo. Portfolios can also be rearranged by moving them above or below another portfolio.

All tab entries in the analysis pane, including the Markov visualization, can be used to further filter the alliance portfolio. Filters are additive and are carried over to other tabs. We chose to use this approach to allow users to easily drill down and explore the data deeply. We allow multi-selections as it allows users to pick as many filters within and across tabs. If a filter "selection" is made in any of the tabs, a filter icon appears on the corresponding tab. The selection can be removed/deselected by clicking on it again; if multiple selections are made in a tab, the entire tab can also be cleared by clicking on the remove filter icon.

4.5 Implementation

Our system is a web-based application implemented in HTML and JavaScript. The alliance data is loaded from JavaScript Object Notation (JSON) files, parsed, and stored on the client. In doing so, the application is able to provide dynamic filtering without the delay that would accompany interaction with the web server. The AngularJS framework is used to structure the web application using a model-view-controller paradigm. A combination of HTML/Scalable Vector Graphics (SVG) elements provide the interface. The alliance portfolio and corresponding Markov visualizations are rendered using Data-Driven Documents (D3). The PourOver JavaScript library is used to provide fast filtering of the data model.

5 EVALUATION

We conducted a three-stage user evaluation study to assess the potential usability and usefulness of our system for the analysis of alliance portfolios (Plaisant 2004; Venable et al. 2012). We also sought to assess fulfillment of the design requirements described earlier. In the first stage, we asked participants to answer specific questions about alliance portfolios and conducted an open-ended analysis in a realistic scenario. We followed this up with additional feedback about potential and desired capabilities. In the second stage, we asked participants to freely use the system and provide feedback about the usability and utility. We opted for a predominantly qualitative user study, complemented by a short survey, as we believed that broad and deep insights would be best attained through observation in combination with an open interview format (Lam et al. 2012). We intentionally did not perform a formal usability evaluation or utility assessment. Instead, we sought to learn potential usability and learnability issues, observe what aspects of the system would and would not be used, determine what capabilities provided the most value, and identify gaps and missing capabilities (Prat et al. 2015). Moreover, we used different samples and metrics corresponding to the focus of the particular stage with the intent to understand cross-usability. This approach has been shown to be appropriate for a project such as ours (Nunamaker Jr et al. 1990; Shneiderman and Plaisant 2006).

5.1 Stage 1: Task-Oriented User Study

5.1.1 Participants. Five corporate strategists participated in the task-oriented user study. All participants had extensive executive experience (25+ years) in technology, consulting, and finance. All had external facing roles and were thus intimately familiar with the importance of alliance. None of them had seen or tried our system prior to the study. They did not receive any compensation for participating in the study, but offered unlimited access afterward.

5.1.2 Tasks and Procedures. Each participant's session was conducted individually. The tutorial phase included an overview of the system, the underlying data, and key functionalities. A member of the research team performed a step-by-step walk through of the interface components and provided several illustrative examples. All participants were asked to use their own computers to gain hands-on experience. At the end of the tutorial, participants were encouraged to ask clarifying questions. The tutorial phase lasted approximately 20 minutes.

In the second phase (practice) of the study, we created a set of five questions. These questions tested the participants' knowledge of basic functionalities as well as probed their ability to explore the data. For example, we asked them, "How does the alliance portfolio of Company X compare to

Table 1. Evaluation Results (Study 1; n = 5)

Task	Accuracy	Time
Name one of Intel's R&D partner companies and describe the alliance.	5/5	2.3
How do the alliance portfolios of Intel and Qualcomm compare?	4/5	3.4
What regions do Intel's software (SW) partners come from?	5/5	2.1
What time period has seen the largest number of tech transfer alliances for Microsoft?	5/5	1.6
What is the most likely alliance type that follows an R&D alliance for IBM?	4/5	1.8

that of Company Y?" We had the participants attempt to answer each question on their own. When they had difficulties or did not know how to proceed, we then assisted and showed how to determine the answer of the question. We also reminded the participants about particular capabilities when it appeared that they had forgotten about them. We decided to use specific questions, rather than open-ended exploration, because we felt that this type of directed analysis would better ensure that participants encountered most aspects of the interface. Furthermore, if participants could not answer these small, direct questions, then open-ended analysis likely would suffer as well. This phase took between 20–45 minutes.

In the third phase (evaluation), we gave all participants a set of five questions to answer and an open-ended scenario (see Table 1). The questions were similar in nature to the practice phase. We did not direct the participants to any functionality. The motivation for this phase was to understand if participants had learned about the system and the functionalities. A research member logged whether the participant was able to answer each question correctly or partially, and how long it took. The open-ended question was to analyze the alliance portfolios of three large technology companies. We asked participants to "think out loud" and stop when they felt they had come to an insight. At the end of the study, we asked participants to fill out a short Likert-style value-based evaluation survey about their impressions of the system, the utility, and usefulness, as well as if the tool generated confidence in the data, provided the essence of the data, and generated value to the end user. The final phase took 30 minutes, resulting in a complete session between 70 to 110 minutes.

5.1.3 Results. The task-oriented user study produced both encouraging results and helpful feedback about the design of our tool. The overall assessment from the users was that our system would be very useful in understanding, analyzing, and comparing alliance portfolios. All participants found that the UI was simple, intuitive, and easy to use. While not explicitly assessed, four participants commented on the ease of learning the functionalities of the tool. The comments also suggested that our system would be particularly useful for supporting partnership initiatives. All participants appreciated the ability to change the temporal mode of the visualization, the ease with which to add, remove, and arrange portfolios, and the ability to gain both macro and micro perspectives. Perhaps the most discussed functionality of the tool was the activity state transition visualization. One participant noted "that this may be the crystal ball strategists are looking for." At the same time, however, two participants noted that they would have liked to have seen data of smaller companies, given the growing prominence of start-ups. Several other interesting suggestions were made, including the ability to save an image of a portfolio for a report, the ability to annotate a portfolio, and the ability to copy-and-paste the alliance text. While these functionalities were not explicitly built into our system, we considered them excellent suggestions for future

Table 2. Evaluation Results (Study 2; 5-point Likert Scale, n = 16)

Assessment	μ	σ^2
The visualization was easy to use.	4.61	0.33
The visualization was easy to learn.	4.83	0.16
The visualization enabled me to discover insights about a firm's alliance activities.	4.61	0.52
The visualization enabled me to ask insightful questions about a firm's alliance activities.	4.42	0.45
The visualization helped me generate knowledge about a firm's alliance activities.	4.27	0.34
The visualization conveyed an overall essence (or take-away sense) of a firm's alliance activities.	4.56	0.28
The visualization helped me generate confidence about a firm's alliance activities.	4.12	0.71

versions. An evaluation of the task results revealed that all participants were able to successfully answer nearly all questions (see Table 1).

5.2 Stage 2: Web-Based, Value-Driven User Study

5.2.1 Participants. The second user study was conducted entirely online. The purpose of this study was to reach a wider audience and to deploy the system in the wild. We identified 50 corporate strategists and strategy scholars with an expressed interest in alliances and interfirm relationships. All participants had significant work experience (15+ years). None of the participants had seen or tried our system prior to the study. They also did not receive any compensation for participating in the study. Sixteen participants accepted the invitation and used the tool, for a response rate of 32%. Thirteen invitees declined to participate in the study due to unavailability; one invitee declined due to a lack of interest; 20 invitees did not respond.

5.2.2 Tasks and Procedures. Personal invitations were sent to each prospective participant by email along with a detailed tutorial, a link to a password-protected version of our system, and a link to a web-based survey. A friendly reminder was sent midweek to encourage participation. Participants were not asked to complete any specific tasks, but rather to explore the tool and any of the portfolios freely. The system and survey were administered for one week. The survey included questions about ease of use and learning, insights gained, confidence, and data essence following a value-based evaluation approach (Stasko 2014). Participants were also asked to provide feedback on strengths and weaknesses as well as other comments.

5.2.3 Results. Similar to the first evaluation study, the overall assessment of our system was very positive (see Table 2). All participants used the tool extensively. We analyzed log-data of the various sessions and found that users spent on average 334 minutes with the tool (median 261 minutes). They most frequently accessed the Partners tab (100% users) followed by the activity state visualization (90%). All users switched between the temporal and sequential display modes and used rubberbanding to get insight into portfolio details. Interestingly, participants opened predominantly a single alliance portfolio. A possible interpretation of this is that they may have focused their analysis on a single firm or did not find the desired portfolio. Across the board, our tool received positive remarks. One participant, for instance, noted that it was particularly easy to compare "the overall size of alliance activities" and pursue "deep probing of different alliance types" using the rich filter panel. We thus felt confident that our system had practical merit. Moreover, all participants felt that the system provided a basis for interesting follow-on questions, including the "potential correlation of activity sequences with operational, financial, and innovation performance measures."

5.3 Stage 3: Real-World Setting

5.3.1 Participants. While the first two user studies measured both novelty effects and demand characteristics, an ideal design science study should also provide insights into how users work with the visualization in a real-world setting. To gain insight into the true utility, we deployed our system into the work environment of ICT market analysts. Through the research team's personal network, we identified 11 market analysts located in North America, Europe, and Asia. Three analysts accepted the invitation, for a response rate of 27%. All others declined due to technical conflicts of integrating the system into their work processes. The three respondents had only moderate overall work experience (average of 9 years), however, all had spent their entire career analyzing the ICT industry and, thus, were very familiar with the domain. More importantly, all three analysts had similar levels of visual literacy and extensive prior knowledge with analytical tools. None of the three analysts had seen or tried the system before and received no compensation for participating in the study.

5.3.2 Protocol. Each participant received a link to a password-protected login page. We also provided a simple Skype-based tutorial. Based on our prior suggestions, we asked the three participants to freely explore and use the system over a period of one month. In case any issues or concerns came up, we asked participants to directly reach out to the research team. We periodically checked on the analysts. At the end of the study period, participants were asked to provide detailed written and verbal feedback on their experience with the system.

5.3.3 Observations. The three participants used the system on average twice a week per month for at least 1 hour each day. All participants logged into the system at the beginning of the day when new tasks or projects were assigned. One participant, for instance, noted that she used the system exclusively when she was asked to quickly research a company in relation to a major event or announcement, such as a product release or new market entry. She commented that it "helped her accelerate [her] learning about a company and its idiosyncratic partnership behaviors." Another participant noted that the ability to explore historical data was similarly important. ("Understanding past behavior of competitors, particularly during financial downturns, allows me to gauge where they may go next"). Over the duration of the study, however, participants used the system more on-demand, in particular, for "knowledge gathering" (Participant 3) or "insight confirmation" (Participant 2) tasks.

Our system appeared to be particularly valuable when comparing portfolios. Participant 3, for instance, commented, "to my surprise, Google has a much shorter alliance portfolio than Apple, but when considering the last 10 years, Google had nearly 4 times (34) as many strategic alliances than Apple (9), including an interesting supply chain alliance with an eye glass manufacturer. Hello, wearables?!"

All participants used filtering, stacking, and temporal/sequential modes extensively to compare and contrast portfolios. However, one participant commented that "while I can understand and visually compare two companies, it would've been ideal if I could have overlayed two portfolios to see the commonality and/or differences." Temporal and time-shifted comparisons were also voiced as desirable future functionalities.

While not explicitly stated by all participants, it was clear that our system also had some "real-time" data limitations that they would like to have addressed in future releases. "There's no question that I spend the majority of my time examining large established companies. But at the same time, I need to be knowledgeable about what's happening next, particularly with smaller emerging companies. It would be great if you could port data from Crunchbase or CB Insights into the system." These comments emphasize that to be truly integrated into daily workflows, more automated data pipes must be established.

Despite some of these limitations, all three participants expressed their delight to work with our system. As one participant in particular noted: "I work with lots of analytical and visual tools on a daily basis. However, most of them are just spreadsheets or simple graphs. I very much welcomed the visually appealing and easy to learn interface. The ability to interact with and explore different facets of the data dynamically was awesome" (Participant 2). "It was kind of cool to play with the tool and gain insights quickly. My colleagues often peeked in over my shoulder and clearly wanted to use it, too."

6 CONCLUDING REMARKS

This study presented our journey in designing and implementing an interactive visualization system for understanding alliance portfolios. High-level design requirements included rapid comparison and deep, multi-faceted data probing of alliance portfolios as well as analysis of alliance formation strategies. Our results show that our system enables decision makers to gain both systemic (macro) and detailed (micro) insights into a firm's alliance activities.

More broadly, our study reveals the complexity of designing and developing visual analytic tools for enterprise decision-making contexts. First, our design process confirms that despite using an iterative approach, artifact functionalities can be ineffective and have potentially unintended consequences. Our tool, for instance, provided a Markov model visualization for alliance formation probabilities. While effective initially, once decision makers used this functionality, they immediately requested more complex sequence mining capabilities. Our design process, thus, revealed that while initial functionalities satisfy initial design requirements, new design expectations rapidly emerge once the artifact is used and the capabilities are understood. This observation is in line with the idea that design is really a continuous search process (Walls et al. 1992; Markus et al. 2002; Peffers et al. 2007). Designers of artifacts should, thus, not implement a rigid design but accommodate extensibility and scalability to emerging functionalities, such as going from a single alliance activity to more complex sequences of alliance formations. In doing so, seamless and rapid generativity and deeper usage of the design artifact can be achieved.

Another unexpected consequence of interactive visual analytic artifacts, and related to feature extensibility and scalability, is the insatiable appetite for additional data. Decision makers expect that tools can either bring in data or export the data into other tools they use. In our context, decision makers wanted to compare alliance formation patterns to financial and innovation performance data. Our system was not designed to accommodate this. Future design of visual analytic artifacts should, thus, at the minimum, allow exporting of relevant data/analysis for use in other tools or, in an ideal case, have the ability to onboard new data easily.

Visual enterprise analytics is unquestionably a nascent, but rapidly emerging domain of relevance for IS, computing, and strategy researchers. Our design study reveals some of the interesting opportunities and challenges in designing effective systems. Our hope is that by exposing them, we encourage others to pursue similarly exciting studies.

APPENDIX

A.1 Computation of Transition Probabilities and Duration

We modeled each firm's sequence of alliances as a stochastic Markov process (Ross 1996). Following design requirements, we consider two types of "states:"

— Alliance types (e.g., Strategic, Licensing, Marketing, Manufacturing, Supply, R&D, Tech Transfer)[2]
— Alliance scope (i.e., the number of alliance types)

As analysts are interested in transition probabilities (in percentage) and timing (average number of days), the edges between states can be selected to depict either the likelihood or the timing of moving from state i to state j. The transition probabilities and lengths are computed based on the selection window (by default it is the entire portfolio), enabling an exploration over different time periods.

For a given date range $[a, b]$, let $p(t_1 \rightarrow t_2)$ indicate the probability of transitioning from an alliance containing a relationship of type t_1 to an alliance containing a relationship of type t_2. Then, $p(t_1 \rightarrow t_2) = \Phi_{t_1 \rightarrow t_2}/\Phi_{t_1}$, where $\Phi_{t_1 \rightarrow t_2}$ is the number of transitions from an alliance $\subseteq t_1$ to an alliance $\subseteq t_2$, and Φ_{t_1} indicates the total number of transitions from an alliance $\subseteq t_1$ to any other state in the network during $[a, b]$.

For a given date range $[a, b]$, let $d(t_1 \rightarrow t_2)$ indicate the average number of days taken for a transition between an alliance containing relationship type t_1 to an alliance containing a relationship of type t_2. Then, $d(t_1 \rightarrow t_2) = \sum \Delta^1_{t_1 \rightarrow t_2}, \Delta^2_{t_1 \rightarrow t_2}, \ldots/\Phi_{t_1 \rightarrow t_2}$, where ${\Delta_{t_1 \rightarrow t_2}}^1$ is the time taken (in days) for the first transition from an alliance $\subseteq t_1$ to an alliance $\subseteq t_2$, ${\Delta_{t_1 \rightarrow t_2}}^2$ is the time taken for the second transition from an alliance $\subseteq t_1$ to an alliance $\subseteq t_2$ and so on; $\Phi_{t_1 \rightarrow t_2}$ is the total number of transitions from an alliance $\subseteq t_1$ to an alliance $\subseteq t_2$ in the network during $[a, b]$.

Similarly, we calculate the transition probabilities and elapsed time between alliances of different scopes as follows. For a given date range $[a, b]$, let $p(s_1 \rightarrow s_2)$ indicate the probability of transitioning from an alliance of scope s_1 to an alliance of scope s_2. Then, $p(s_1 \rightarrow s_2) = \Phi_{s_1 \rightarrow s_2}/\Phi_{s_1}$, where $\Phi_{s_1 \rightarrow s_2}$ is the number of transitions from an alliance $\subseteq s_1$ to an alliance $\subseteq s_2$ and Φ_{s_1} indicates the total number of transitions from an alliance $\subseteq s_1$ to any other state in the network during $[a, b]$.

For a given date range $[a, b]$, let $d(s_1 \rightarrow s_2)$ indicate the average number of days taken for a transition between an alliance of scope s_1 to an alliance of scope s_2. Then, $d(s_1 \rightarrow s_2) = \sum \Delta^1_{s_1 \rightarrow s_2}, \Delta^2_{s_1 \rightarrow s_2}, \ldots/\Phi_{s_1 \rightarrow s_2}$, where ${\Delta_{s_1 \rightarrow s_2}}^1$ is the time taken(in days) for the first transition from an alliance $\subseteq s_1$ to an alliance $\subseteq s_2$, and ${\Delta_{s_1 \rightarrow s_2}}^2$ is the time taken for the second transition from an alliance $\subseteq s_1$ to an alliance $\subseteq s_2$. Lastly, $\Phi_{s_1 \rightarrow s_2}$ is the total number of transitions from an alliance $\subseteq s_1$ to an alliance $\subseteq s_2$ in the network during $[a, b]$.

A.2 Sample Alliance Portfolios of S&P 500 Companies

A complete set of alliance portfolios can be found at [URL not included due to the review process]. Figure 6 presents a select set of alliance portfolios.

A.3 Sample Usage Scenario

Consider the case of a corporate analyst who is interested in understanding the evolving relationship portfolio of three incumbent technology companies (Hewlett-Packard, IBM, and Microsoft), which, due to a changing technology landscape and emergence of new entrants, have faced significant pressure to transform their product and innovation strategies.

The analyst begins by selecting the three companies from the drop-down logo list. The three alliance portfolios are stacked on top of each other. The analyst notices that IBM has the largest

[2]Note: If we considered an exhaustive Markov model, there would be 127 unique states (as the order in which types appear do not matter). However, based on analyst feedback, the primary concern is to determine whether an alliance "contains" a relationship type. We, thus, have only seven states.

Fig. 5. Sample alliance portfolios.

portfolio in contrast to Microsoft and HP. She also notices that all three firms tend to predominantly form strategic alliances (green), that R&D alliances (light blue) were formed early on, and that there has been an increase in tech transfer alliances (red) more recently. Curious about what R&D alliances IBM has formed recently, the analyst rubberbands over a portion of IBM's portfolio to get a more detailed view. The details pane appears below the portfolio. She then selects the R&D filter. Hovering over the latest R&D event she reads that IBM formed an agreement with NEC in Sept. 2008 on semiconductor technology.

To put these events in a longitudinal perspective, she switches from the sequential to the temporal view. The activities in the portfolio rearrange through smooth transitions to the appropriate temporal position. She deselects R&D to see all alliance types. Hovering over the various alliance filters reveals that all three companies formed their last supply alliance in 2003 or prior. Similarly, hovering over manufacturing reveals that all three companies did not form any manufacturing alliances after 2010, suggesting a shift from a product-oriented organization to a service-oriented organization.

Given that so few alliances were formed in recent years, the analyst shifts her focus on a recent timeframe of significant activities by using the date filter (2005–2009). She notices that both IBM (37) and Microsoft (38) have approximately the same number of tech transfer alliances. She removes HP from her investigation to focus on just those two companies. Next, she expands the

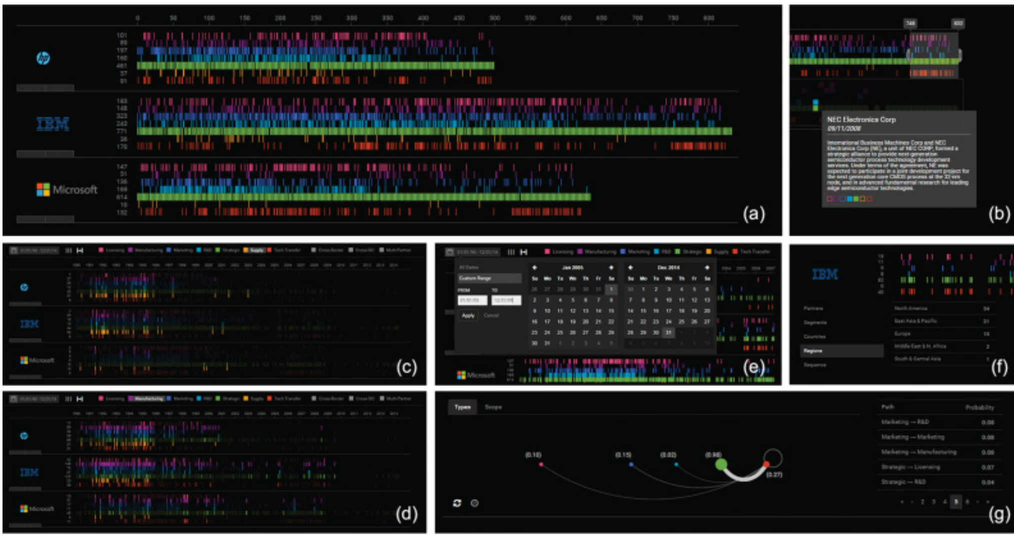

Fig. 6. Series of screenshots depicting a sample use-case scenario. (a) An analyst adds three alliance portfolios (HP, IBM, and Microsoft). (b) Hovering over an alliance reveals details. Hovering over filters highlights those alliances meeting the criteria and fades out others, in this case, (c) supply and (d) manufacturing. (e) User filters by date range. (f) Regions tab reveals geographic strategies of a firm. (g) Markov model visualization highlights most likely next strategic moves, in this case, moves after a tech transfer alliance.

analysis pane for both IBM and Microsoft and clicks on the regions tabs. She notes that IBM has focused its effort outside North America, primarily on the East Asia and Pacific (18) region, while Microsoft focused more on Europe (23), suggesting a diverse global partnership strategy. She then switches to the Sequence tab for both firms. While the two Markov visualizations appear similar, she is interested in what actions each firm pursues after a Tech Transfer alliance. Hovering over the Tech Transfer node, she finds that IBM's most likely subsequent alliance is strategic (97%) followed by another tech transfer (49%) or licensing (22%). Only 8% of subsequent alliances result in R&D; switching to the time view in the Markov visualization, she sees that an R&D alliance is formed 16 days after a tech transfer alliance. She then switches to Microsoft. Hovering over the Tech Transfer node, she finds that Microsoft's most likely subsequent relationship to a tech transfer alliance is also strategic (97%) but significantly lower for tech transfer (26%), licensing (11%), and R&D (3%). Comparing the two sets of probabilities, she concludes that IBM and Microsoft, despite having similar numbers of tech transfer alliances, use widely different alliance formation strategies during this timeframe.

ACKNOWLEDGMENTS

The authors would like to thank Peter Evans, John Stasko, Hyunwoo Park, and Raul Chao for valuable feedback on earlier versions of the manuscript.

REFERENCES

Ron Adner. 2017. Ecosystem as structure an actionable construct for strategy. *Journal of Management* 43, 1 (2017), 39–58.

Wolfgang Aigner, Silvia Miksch, Wolfgang Müller, Heidrun Schumann, and Christian Tominski. 2007. Visualizing time-oriented data: A systematic view. *Computers & Graphics* 31, 3 (2007), 401–409.

Wolfgang Aigner, Silvia Miksch, Wayne Muller, Heidrun Schumann, and Christian Tominski. 2008. Visual methods for analyzing time-oriented data. *IEEE Transactions on Visualization and Computer Graphics* 14, 1 (2008), 47–60.

Wolfgang Aigner, Silvia Miksch, Heidrun Schumann, and Christian Tominski. 2011. *Visualization of Time-oriented Data*. Springer Science & Business Media.

Wolfgang Aigner, Silvia Miksch, Bettina Thurnher, and Stefan Biffl. 2005. PlanningLines: Novel glyphs for representing temporal uncertainties and their evaluation. In *Proceedings of the 9th International Conference on Information Visualisation*. IEEE, 457–463.

Bilal Alsallakh, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, and Peter Rodgers. 2014. Visualizing sets and set-typed data: State-of-the-art and future challenges. In *Proceedings of the Eurographics Conference on Visualization (EuroVis)–State of The Art Reports*. 1–21.

Rahul C. Basole. 2014. Visual business ecosystem intelligence: Lessons from the field. *IEEE Computer Graphics and Applications*5 (2014), 26–34.

Rahul C. Basole, Trustin Clear, Mengdie Hu, Harshit Mehrotra, and John Stasko. 2013. Understanding interfirm relationships in business ecosystems with interactive visualization. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2526–2535.

Rahul C. Basole, Martha G. Russell, Jukka Huhtamäki, Neil Rubens, Kaisa Still, and Hyunwoo Park. 2015. Understanding business ecosystem dynamics: A data-driven approach. *ACM Transactions on Management Information Systems (TMIS)* 6, 2 (2015), 6.

Rita Borgo, Johannes Kehrer, David H. Chung, Eamonn Maguire, Robert S. Laramee, Helwig Hauser, Matthew Ward, and Min Chen. 2013. Glyph-based visualization: Foundations, design guidelines, techniques and applications. *Eurographics State of the Art Reports* (2013), 39–63.

Corentin Burnay. 2016. Are stakeholders the only source of information for requirements engineers? Toward a taxonomy of elicitation information sources. *ACM Transactions on Management Information Systems (TMIS)* 7, 3 (2016), 8.

Igor Cadez, David Heckerman, Christopher Meek, Padhraic Smyth, and Steven White. 2000. Visualization of navigation patterns on a web site using model-based clustering. In *Proceedings of the 6th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 280–284.

Hsinchun Chen, Roger H. L. Chiang, and Veda C. Storey. 2012. Business intelligence and analytics: From big data to big impact. *MIS Quarterly* 36, 4 (2012), 1165–1188.

Thomas Luan Dang and Victoria L. Lemieux. 2013. A Functional Framework for Evaluating Financial Visualization Products. In *Financial Analysis and Risk Management*. Springer, 115–153.

Tushar K. Das and Bing-Sheng Teng. 2002. Alliance constellations: A social exchange perspective. *Academy of Management Review* 27, 3 (2002), 445–456.

Jerry Alan Fails, Amy Karlson, Layla Shahamat, and Ben Shneiderman. 2006. A visual interface for multivariate temporal data: Finding patterns of events across multiple histories. In *Proceedings of the 2006 IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 167–174.

Fabian Fischer, Florian Mansmann, and Daniel A. Keim. 2012. Real-time visual analytics for event data streams. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing*. ACM, 801–806.

Gerard George, Shaker A. Zahra, Kathleen K. Wheatley, and Raihan Khan. 2001. The effects of alliance portfolio characteristics and absorptive capacity on performance: A study of biotechnology firms. *Journal of High Technology Management Research* 12, 2 (2001), 205–226.

Paulo B. Goes. 2014. Editor's comments: Design science research in top information systems journals. *MIS Quarterly* 38, 1 (2014), iii–viii.

Benjamin Gomes-Casseres. 2015. *Remix Strategy: The Three Laws of Business Combinations*. Harvard Business Press.

Shirley Gregor and Alan R. Hevner. 2013. Positioning and presenting design science research for maximum impact. *MIS Quarterly* 37, 2 (2013), 337–355.

Ranjay Gulati. 1998. Alliances and networks. *Strategic Management Journal* (1998), 293–317.

Susan Havre, Elizabeth Hetzler, Paul Whitney, and Lucy Nowell. 2002. Themeriver: Visualizing thematic changes in large document collections. *IEEE Transactions on Visualization and Computer Graphics* 8, 1 (2002), 9–20.

Alan Hevner, Salvatore T. March, Jinsoo Park, and Sudha Ram. 2004. Design science in information systems research. *MIS Quarterly* 28, 1 (2004), 75–105.

Werner H. Hoffmann. 2007. Strategies for managing a portfolio of alliances. *Strategic Management Journal* 28, 8 (2007), 827–856.

Ruihua Joy Jiang, Qingjiu Tom Tao, and Michael D. Santoro. 2010. Alliance portfolio diversity and firm performance. *Strategic Management Journal* 31, 10 (2010), 1136–1144.

Daniel A. Keim, Ming C. Hao, Umesh Dayal, and Meichun Hsu. 2002. Pixel bar charts: A visualization technique for very large multi-attribute data sets. *Information Visualization* 1, 1 (2002), 20–34.

Bohyoung Kim, Bongshin Lee, and Jinwook Seo. 2007. Visualizing set concordance with permutation matrices and fan diagrams. *Interacting with Computers* 19, 5–6 (2007), 630–643.

Miloš Krstajić, Enrico Bertini, and Daniel A. Keim. 2011. Cloudlines: Compact display of event episodes in multiple time-series. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2432–2439.

Heidi Lam, Enrico Bertini, Petra Isenberg, Catherine Plaisant, and Sheelagh Carpendale. 2012. Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Visualization and Computer Graphics* 18, 9 (2012), 1520–1536.

Dovev Lavie. 2007. Alliance portfolios and firm performance: A study of value creation and appropriation in the US software industry. *Strategic Management Journal* 28, 12 (2007), 1187–1212.

D. Lavie. 2017. Alliance capability. *The Palgrave Encyclopedia of Strategic Management*, Mie Augier and David J. Teece (Eds.). Palgrave Macmillan, UK.

M. Lynne Markus, Ann Majchrzak, and Les Gasser. 2002. A design theory for systems that support emergent knowledge processes. *MIS Quarterly* (2002), 179–212.

Sean McKenna, David C. Mazur, James Agutter, and Michael Meyer. 2014. Design activity framework for visualization design. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2191–2200.

Peter McLachlan, Tamara Munzner, Eleftherios Koutsofios, and Stephen North. 2008. LiveRAC: Interactive visual exploration of system management time-series data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* ACM, 1483–1492.

Luana Micallef, Pierre Dragicevic, and Jean-Daniel Fekete. 2012. Assessing the effect of visualizations on Bayesian reasoning through crowdsourcing. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2536–2545.

Megan Monroe, Rongjian Lan, Hanseung Lee, Catherine Plaisant, and Ben Shneiderman. 2013. Temporal event sequence simplification. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2227–2236.

Jakob Nielsen. 1994. *Usability Engineering*. Elsevier.

Thomas Nocke, Heidrum Schumann, Uwe Böhm, and Michael Flechsig. 2003. Information visualization supporting modelling and evaluation tasks for climate models. In *Proceedings of the 2003 Winter Simulation Conference*, Vol. 1. IEEE, 763–771.

Jay F. Nunamaker Jr, Minder Chen, and Titus D. M. Purdin. 1990. Systems development in information systems research. *Journal of Management Information Systems* 7, 3 (1990), 89–106.

Pinar Ozcan and Kathleen M. Eisenhardt. 2009. Origin of alliance portfolios: Entrepreneurs, network strategies, and firm performance. *Academy of Management Journal* 52, 2 (2009), 246–279.

Ken Peffers, Tuure Tuunanen, Marcus A. Rothenberger, and Samir Chatterjee. 2007. A design science research methodology for information systems research. *Journal of Management Information Systems* 24, 3 (2007), 45–77.

Catherine Plaisant. 2004. The challenge of information visualization evaluation. In *Proceedings of the Working Conference on Advanced Visual Interfaces.* ACM, 109–116.

Catherine Plaisant, Brett Milash, Anne Rose, Seth Widoff, and Ben Shneiderman. 1996. LifeLines: Visualizing personal histories. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* ACM, 221–227.

Nicolas Prat, Isabelle Comyn-Wattiau, and Jacky Akoka. 2015. A taxonomy of evaluation methods for information systems artifacts. *Journal of Management Information Systems* 32, 3 (2015), 229–267.

PwC. 2015. The 2015 CEO Survey by PwC. Retrieved from http://www.pwc.com/gx/en/ceo-survey/.

Sheldon M. Ross. 1996. *Stochastic Processes.* Vol. 2. John Wiley & Sons New York.

Ramik Sadana, Timothy Major, Alistair Dove, and John Stasko. 2014. OnSet: A visualization technique for large-scale binary set data. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1993–2002.

Melissa A. Schilling. 2009. Understanding the alliance data. *Strategic Management Journal* 30, 3 (2009), 233–260.

Ben Shneiderman. 1996. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the 1996 IEEE Symposium on Visual Languages.* IEEE, 336–343.

Ben Shneiderman and Catherine Plaisant. 2006. Strategies for evaluating information visualization tools: multi-dimensional in-depth long-term case studies. In *Proceedings of the 2006 AVI Workshop on BEyond Time and Errors: Novel Evaluation Methods for Information Visualization.* ACM, 1–7.

John Stasko. 2014. Value-driven evaluation of visualizations. In *Proceedings of the 5th Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization.* ACM, 46–53.

Martin Suntinger, Hannes Obweger, Josef Schiefer, and M. Eduard Gröller. 2008. The event tunnel: Interactive visualization of complex event streams for business process pattern analysis. In *Proceedings of the IEEE Pacific Visualization Symposium (PacificVIS'08).* IEEE, 111–118.

Christian Tominski, James Abello, and Heidrun Schumann. 2004. Axes-based visualizations with radial layouts. In *Proceedings of the 2004 ACM Symposium on Applied Computing.* ACM, 1242–1247.

Christian Tominski, Heidrun Schumann, Gennady Andrienko, and Natalia Andrienko. 2012. Stacking-based visualization of trajectory attribute data. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2565–2574.

John Venable, Jan Pries-Heje, and Richard Baskerville. 2012. A comprehensive framework for evaluation in design science research. In *Proceedings of the International Conference on Design Science Research in Information Systems.* Springer, 423–438.

Joseph G. Walls, George R. Widmeyer, and Omar A. El Sawy. 1992. Building an information system design theory for vigilant EIS. *Information Systems Research* 3, 1 (1992), 36–59.

U. Wassmer. 2010. Alliance portfolios: A review and research agenda. *Journal of Management* 36, 1 (2010), 141–171.

Marc Weber, Marc Alexa, and Wolfgang Müller. 2001. Visualizing time-series on spirals. In *Infovis*. IEEE, 7.

Krist Wongsuphasawat and David Gotz. 2012. Exploring flow, factors, and outcomes of temporal event sequences with the outflow visualization. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2659–2668.

Krist Wongsuphasawat, John Alexis Guerra Gómez, Catherine Plaisant, Taowei David Wang, Meirav Taieb-Maimon, and Ben Shneiderman. 2011. LifeFlow: Visualizing an overview of event sequences. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1747–1756.

Jing Yang, Anilkumar Patro, Huang Shiping, Nishant Mehta, Matthew O. Ward, and Elke A. Rundensteiner. 2004. Value and relation display for interactive exploration of high dimensional datasets. In *Proceedings of the 2004 IEEE Symposium on Information Visualization (INFOVIS'04)*. IEEE, 73–80.

Received August 2016; revised February 2017; accepted April 2017