

## Dashboard Design for Rich and Rapid Monitoring

Stephen Few

*Visual Business Intelligence Newsletter*

November 2006

This article is the fourth in a five-part series that features the winning solutions to the 2006 Data Visualization Competition that I judged for the Business Intelligence Network ([www.B-EYE-NETWORK.com](http://www.B-EYE-NETWORK.com)). The fourth scenario in the competition asked participants to create a dashboard for airline executives.

Here's the scenario as it was described to participants:

*You are a consultant who has been hired by a U.S. commercial airline to design a dashboard for its executives. The information that the executive team wants to monitor has been identified and now it's your job to create the dashboard's visual design. You must try to display all of this information in some manner on a single screen such that the executives will be able to quickly identify anything that needs their attention and then have the means to discern enough about the situation to decide if they can ignore it for now or must perhaps take some action. It is up to you to determine the appropriate manner, level of detail, and means to display each piece of information.*

I supplied participants with an Excel spreadsheet that included a large collection of actual and target performance measures for the last 12 months (November 1, 2004 through the current date of November 15, 2005) and a few other items without history, including the current top 10 and bottom 10 routes based on cancellations and delays, as well as the reasons for cancellations and delays by percentage.

While reviewing the solutions, in addition to clarity of communication and ease of use, I looked for dashboard designs that provided the following:

- A single-screen display, which wouldn't force people to move from screen to screen or scroll around on a single screen to see all of the information
- Display media, such as graphs, which communicate performance information directly, clearly, accurately, rapidly, and with enough context to be meaningful
- Visual emphasis on the information that most requires the viewer's attention
- An arrangement of the information that is logical and supports the way airline executives would likely want to view it
- Greater focus on what is currently happening than on what happened in the past
- An aesthetically pleasing display (that is, easy on the eyes)

It is much harder to design a dashboard than an individual graph. This is because a dashboard combines a large collection of information, often disparate, which can easily

become cluttered. Great care and skill is required to design a dashboard that communicates richly and rapidly.

## Solutions that Fell Short

Before showing you the winning solution, I'm going to reverse my course from the previous articles in this series and begin by looking at a few of the solutions that fell short in ways that are worth noting. It's worthwhile to know what to avoid. None of the solutions shown in this article are bad, but we'll gain valuable insights by examining them closely, looking for ways to improve them, including the winning solution.

This first solution (Figure 1) exhibits several good design decisions, including an appropriate use of dot plots (the bottom two graphs), but several aspects of the design could be improved. Before reading my critique below, however, take a minute to examine this dashboard on your own to determine what works and what doesn't.

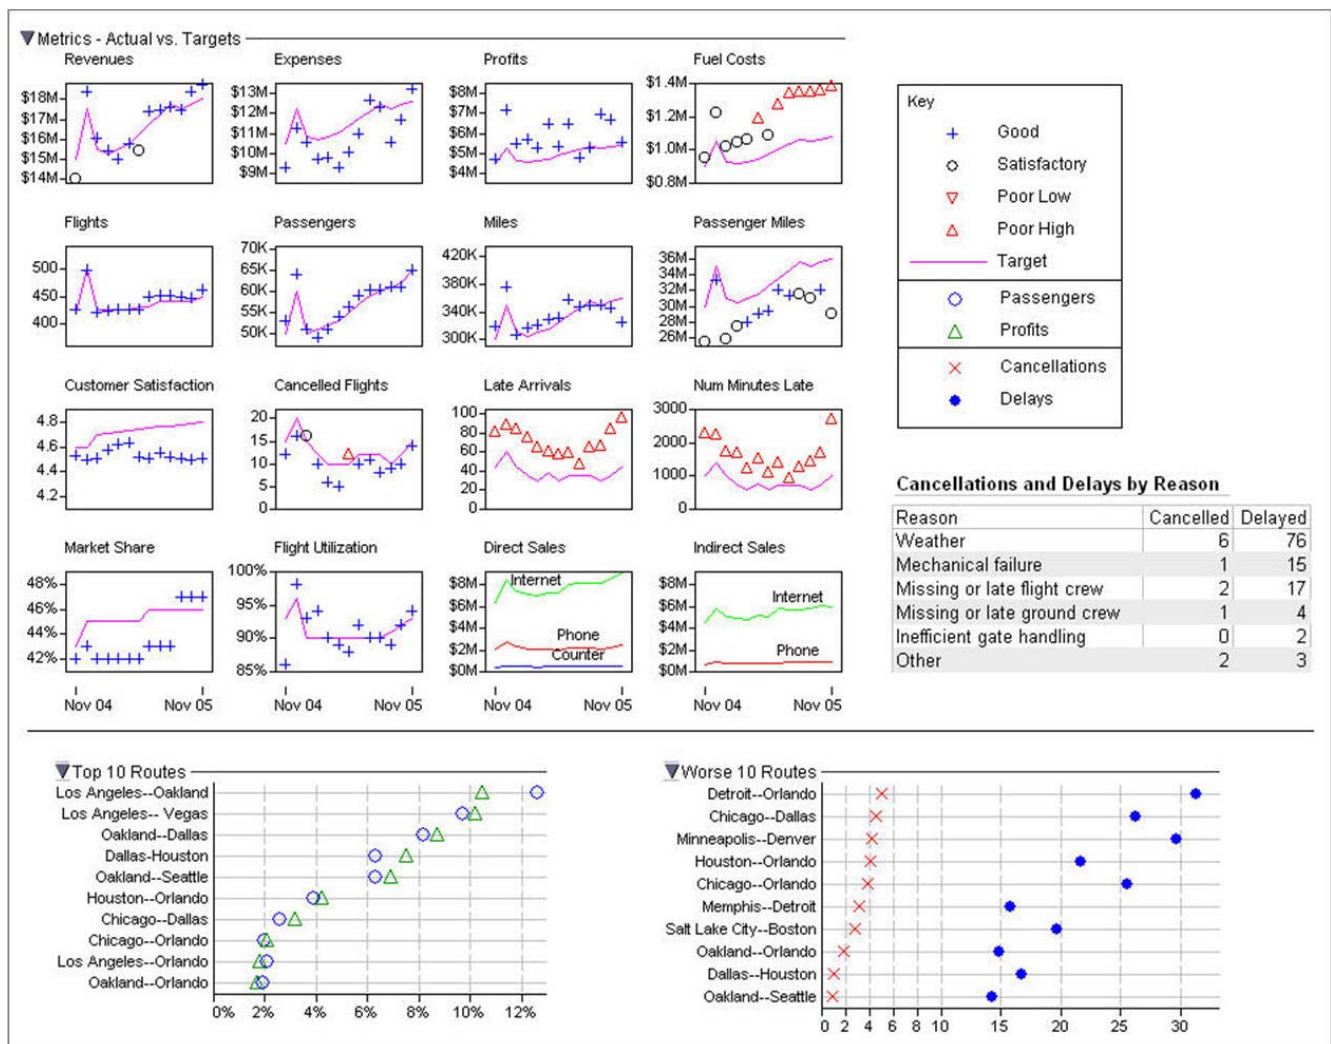


Figure 1

Here's my list of the items that fall short in some way:

- Given that their purpose is to help us monitor what's going on, dashboards ought to focus on current measures. Historical information may also be included to put what's going on in the context of past performance, but the past is less important than what's going on now. The primary measures in this dashboard (the top left section named "Metrics – Actual vs. Targets"), however, are displayed in a way that features present and past performance equally. Readers must work too hard to sift what's going on right now from the sands of history.
- The purpose of including historical data in the metrics section is to show the predominant trends and patterns that have led up to the current measures, yet this is not easy to see based on the discrete data points (circles, triangles, and plus signs). Line graphs would have made it much easier to see the pattern of change through time.
- I believe that an attempt was made to organize the major measures that appear in the upper left-hand section, because they do not appear in the same order that I provided them in the spreadsheet. Organizing this information is important, but nothing was done to visually reinforce this arrangement. Nothing was done to group measures that belong together.
- The monthly values representing actuals in the metrics section are larger than they need to be (that is, the large circles, triangles, and plus signs), resulting in visual clutter.
- The last two graphs in the metrics section are not separate primary metrics, but a breakdown of the revenue metric by sales channel. As such, revenue by sales channel should either be more closely associated with the revenues metric (the upper left graph) or separated into another section of the dashboard to distinguish them from the primary metrics.
- The legend takes up too much space and is much too visually dominant. The legend merely plays a support role in the dashboard, so it should not take up much space or demand much attention.
- By grouping the legends for all the graphs into a single section, many of the legend items are located far from the graphs that they explain, making it much harder than necessary, when examining graphs, to find out what the various symbols mean (for example, the green triangles that appear in the bottom left-hand graph).
- Nothing has been done to adequately draw the viewer's eyes to those items of information that require attention. You must search for the red triangles, which represent poor performance values, but only the last value in each graph (the one for the current month) is of primary interest. You are distracted by the red triangles that represent poor performance in the past when looking for those in the present.
- The borders around the small line graphs are too prominent, resulting in visual clutter that could easily be avoided by light borders.
- All the graphical symbols used to encode values in the graphs (circles, upward pointing triangles, downward pointing triangles, plus signs, and X's) lack fill color except one, the blue circle that represents flight delays, which calls attention to flight delays as uniquely different from the other values, which is not the case.
- The decision to display "Cancellation and Delays by Reason" as a text table rather than graphically seems arbitrary. A graphical display would present this information in a way that could be scanned more quickly, supporting easier comparisons between the reasons.

- The gray background color that appears in alternate rows of the “Cancellations and Delays by Reason” table is unnecessary. This technique can be used to help viewers’ eyes track easily across rows in a table, but isn’t needed in this small table. For wider tables, when it is useful to assist viewers’ eyes in tracking across a single row, if bands of fill color are used, they should be just visible enough to do the job and no more. Otherwise, they contribute to visual clutter.

Let’s move on to another solution (Figure 2). Again, take some time to critique it on your own before reading my comments.

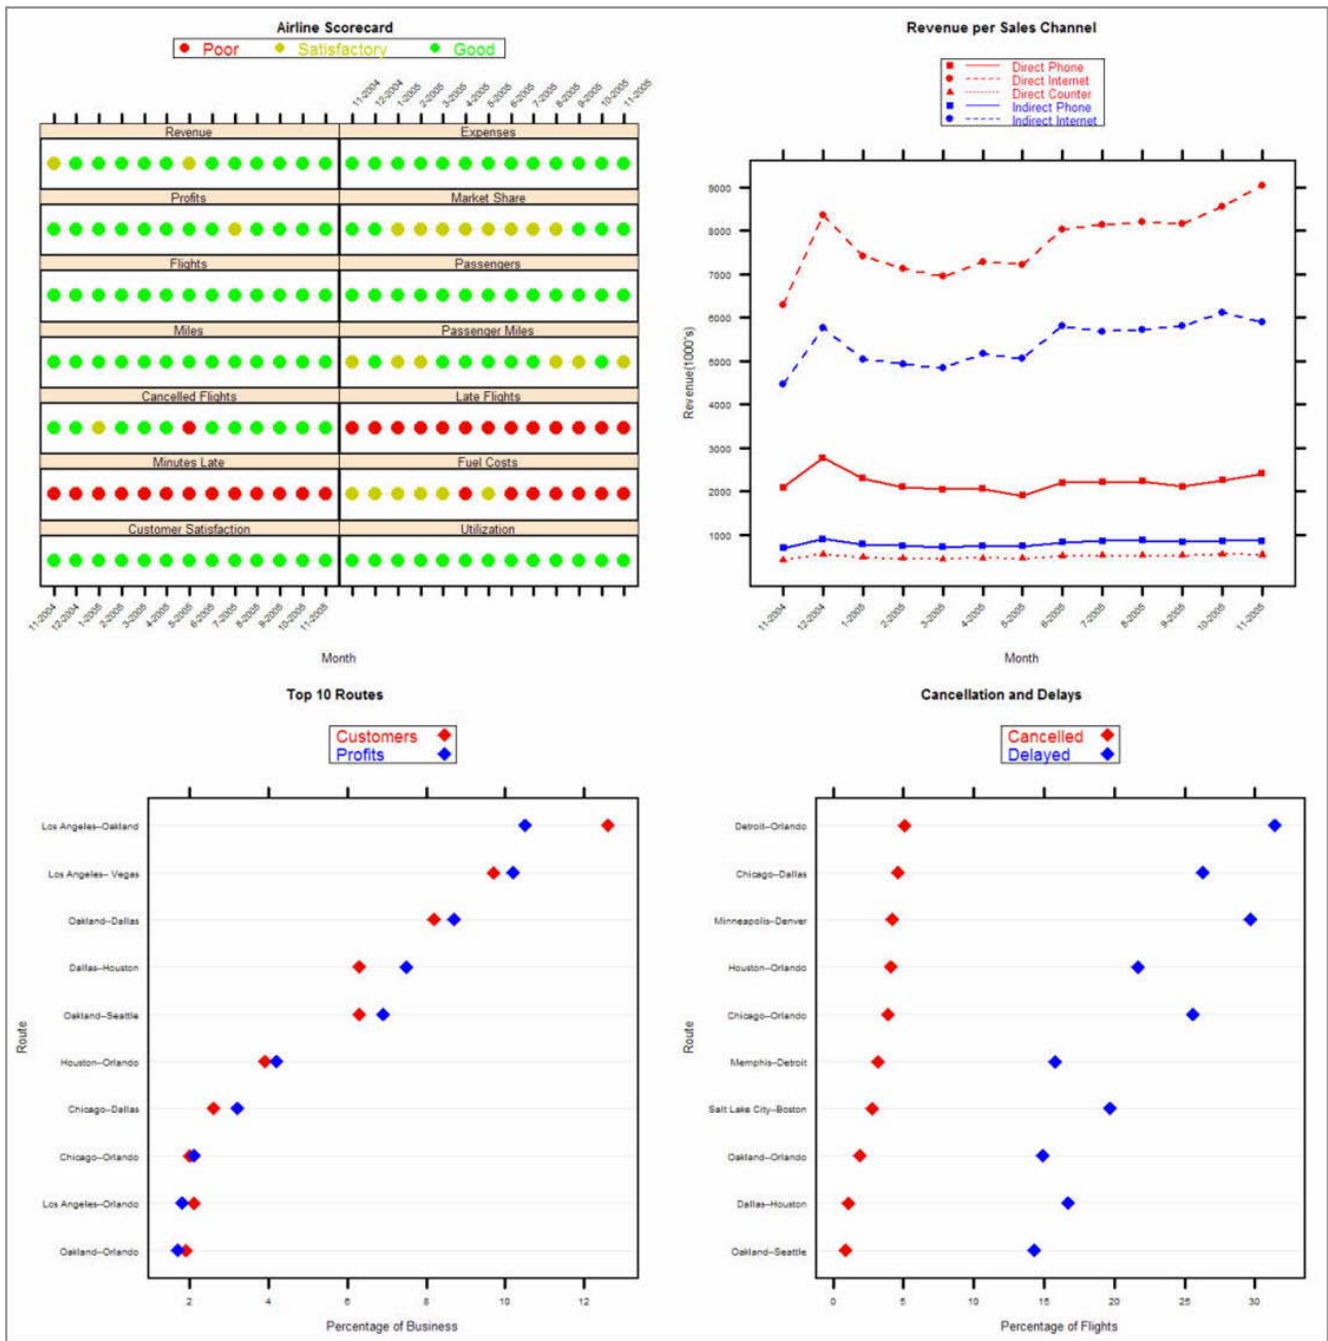


Figure 2

One of the design features that I worked quite well in this dashboard was the use of distinct hues (blue and red) in the “Revenue per Sales Channel” graph to distinguish direct and indirect sales. I also liked the clean look of the dashboard, but as we’ll see, this look was accomplished by eliminating useful information.

Some of the same problems that I pointed out regarding Figure 1 apply to this and other solutions as well, but I won’t bother to repeat anything that has already been mentioned. Here’s a list of the new problems that I found in this solution:

- No attempt was made to organize the primary measures that appear in the upper left section of the dashboard. They appear in the same order in which I provided them in the spreadsheet. To be fair, given no opportunity to interview airline executives to determine how these measures should be arranged, participants in the competition were on their own to determine a useful arrangement, and the order in which the measures were provided might have seemed adequate. In some of the other solutions, however, we’ll see attempts to organize these measures that have merit.
- The way the primary measures were displayed doesn’t provide enough information about historical trends and patterns. By using a heatmap approach, with a colored circle per month to indicate good (green), satisfactory (yellow), or poor (red) performance, we miss the rich view of change through time that a simple line, as in a line graph or sparkline, could provide.
- Also in the primary measures display, the red and green colors that were used to encode good and poor performance could not be discriminated by most people who are color blind. If red and green must be used for this purpose, this problem could be corrected by making the green color much less saturated or lighter than the red, which would make them appear differently even to those who are color blind.
- I have one more observation regarding the display of primary measures. Why do the month labels appear on the bottom for the graphs on the left and on the top for those on the right? This forces viewers to alter the way they read measures on the left versus those on the right and wastes space by requiring room for month labels to appear both on the bottom and top.
- The use of fully saturated versions of blue and red to encode values in multiple graphs creates a visual connection between values of the same color, but this connection doesn’t exist in the data. For example, the color red represents poor performance, direct sales, customers, and cancelled flights—all very different meanings. It is better to use less saturated (and thus less visually dominant), neutral colors, such as grays, reserving more noticeable colors for times when you want viewers to associate the information they encode wherever they appear.

Onward now to one more solution that fell short in several ways (Figure 3). This solution will probably strike you immediately as more visually interesting, mostly due to its use of color, but take the time to consider whether this visual interest translates to an effective display.

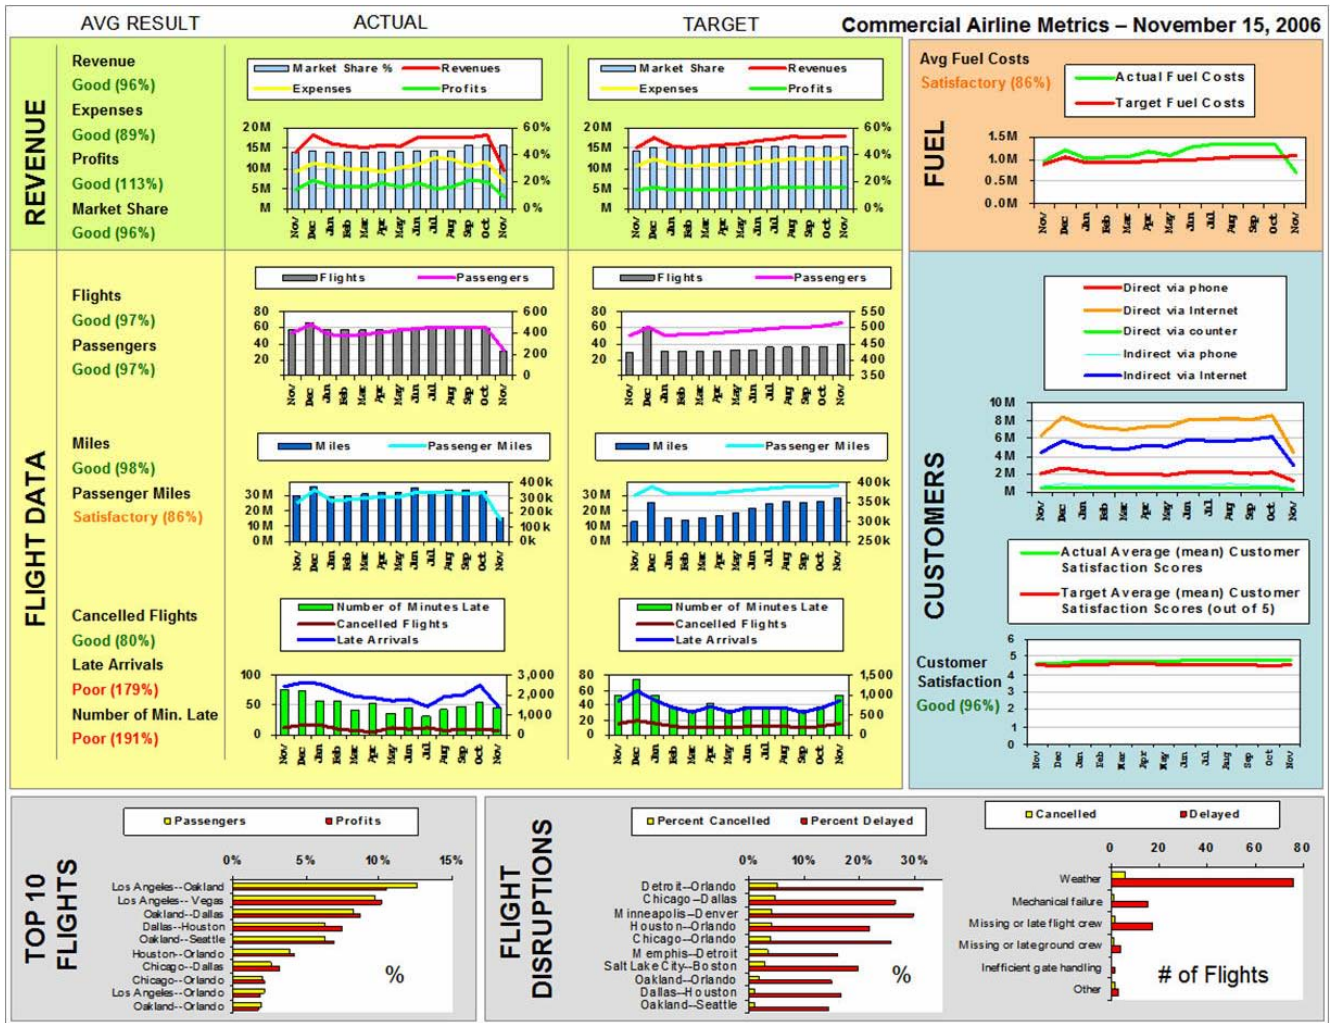


Figure 3

One of the first things that I noticed about this dashboard was its attempt to organize the information into meaningful groups and to use visual means to clearly distinguish them. This is helpful in that it probably comes closer to supporting the mental model that airline executives maintain of the business. I find the use of distinct hues to visually separate the sections a bit overdone, however. If other more subtle means of separated these groups had been used, color could have been used to draw viewers' eyes to items that need their attention. As it is, our eyes are not drawn to particular items based on their importance or need for attention.

Here are a few other aspects of this solution that undermine its effectiveness:

- Separating the displays of actuals from their targets into separate graphs made them much too hard to compare.
- All of the graphs in the revenue and flight data sections use a combination of bars and lines to encode the values, but the choice of bars or lines was arbitrary. When displaying time-series data, bars should be used to feature individual values and support the comparison of one value with another, and lines should be used to feature trends and patterns of change through time. Looking at the upper left graph as an example, do we want to see revenue, expenses, and profits as they change through

time, but only focus on individual market share values? It would have worked better to use lines for all of the items, and to perhaps use bars of the same colors for the current month only, which would give the current measures greater prominence.

- All three of the graphs in the gray sections along the bottom would have been easier to read and less cluttered if displayed as dot plots rather than bar graphs, similar to the dot plots that appear in the previous two solutions.
- Vertically-oriented text is hard to read. Although it is not a big problem that the headings for the various sections of this dashboard run vertically, finding a way to orient them horizontally would have been worth the effort.

## Better Solutions

Before getting to the winning solution, let's look at two more entries that were quite good, but were slightly edged out of the winning spot. Take a look at Figure 4 to see if you can spot ways that it surpassed the previous solutions in effectiveness, as well as ways it could be improved.

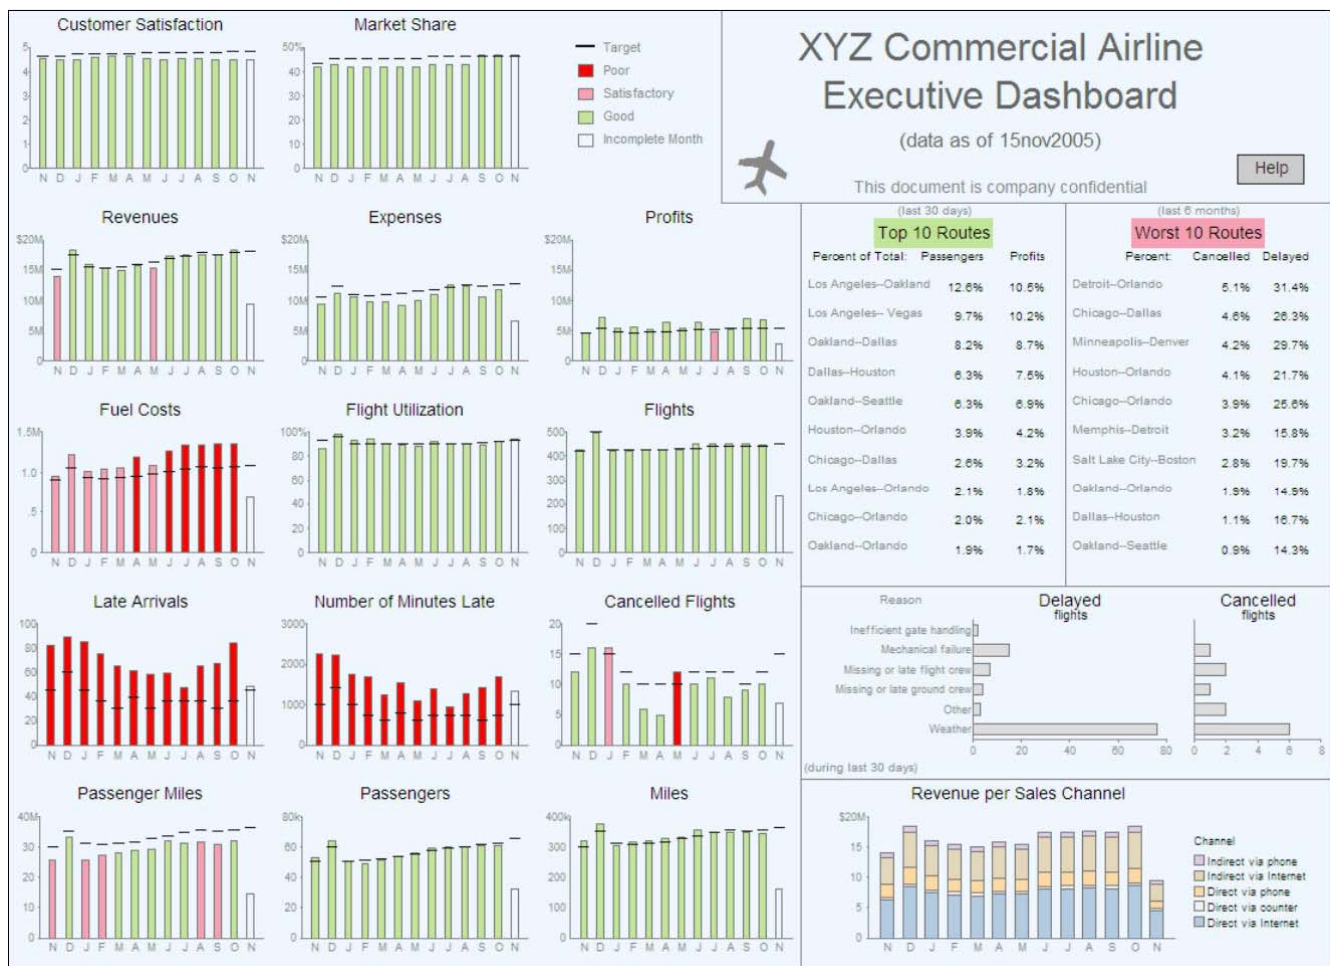


Figure 4: Submitted by Robert Allison of SAS Institute

Robert Allison submitted the winning dashboard solution for the 2005 Data Visualization Competition that I judged for *DM Review* magazine. Once again, he has demonstrated his dashboard design prowess. Here are a few aspects of Robert's design that I particularly like:

- The restrained use of color for most data throughout the dashboard allowed the use of bright red to make items that need attention stand out clearly.
- Colors were used meaningfully, without confusion. Notice that the light green, light red, and bright red colors that were used to represent poor, satisfactory, and good performance were not used anywhere to mean something different.
- Choosing to display the top 10 and worst 10 routes as text tables probably makes sense, because of how this information is used. Most likely, airline executives don't need to compare the values associated with the top 10 and worst 10 routes, but simply need an ordered list to alert them to these routes for possible action.
- Displaying the reasons for delays and cancellations in separate but related graphs allowed them to be better scaled. Placing them in the same graph as we saw in Figure 3 caused the bars representing cancellations to be very short and therefore much harder to read and compare.
- An attempt was made to organize the primary measures such that related measures appear on the same row.
- A simple means to access information about the dashboard, when necessary, was provided in the form of a help button.

Now, let's consider a few ways that this dashboard might be improved:

- The blue background is a little too dark, creating too little contrast between it and the gray text for easy reading.
- In all the bar graphs on the left, which display the primary measures, the current month includes the least amount of information, because color has not been applied to it, despite the fact that it is the most important month. I believe that the proper emphasis on the current month would be shown if past history were encoded as lines for actual and target values without color coding for good, satisfactory, and poor, and only the current month were shown as a bar with color coding.
- The light green and light red colors, because they have similar levels of intensity, might not be distinguishable by those who are color blind.
- Using stacked bars in the "Revenue per Sales Channel" graph makes it hard to see the trends for any one channel. A line graph with a line for each sales channel would have displayed the changes through time more clearly and made it easier to compare the channels to one another.



Let's look at one more good solution (Figure 5) before getting to the winner.

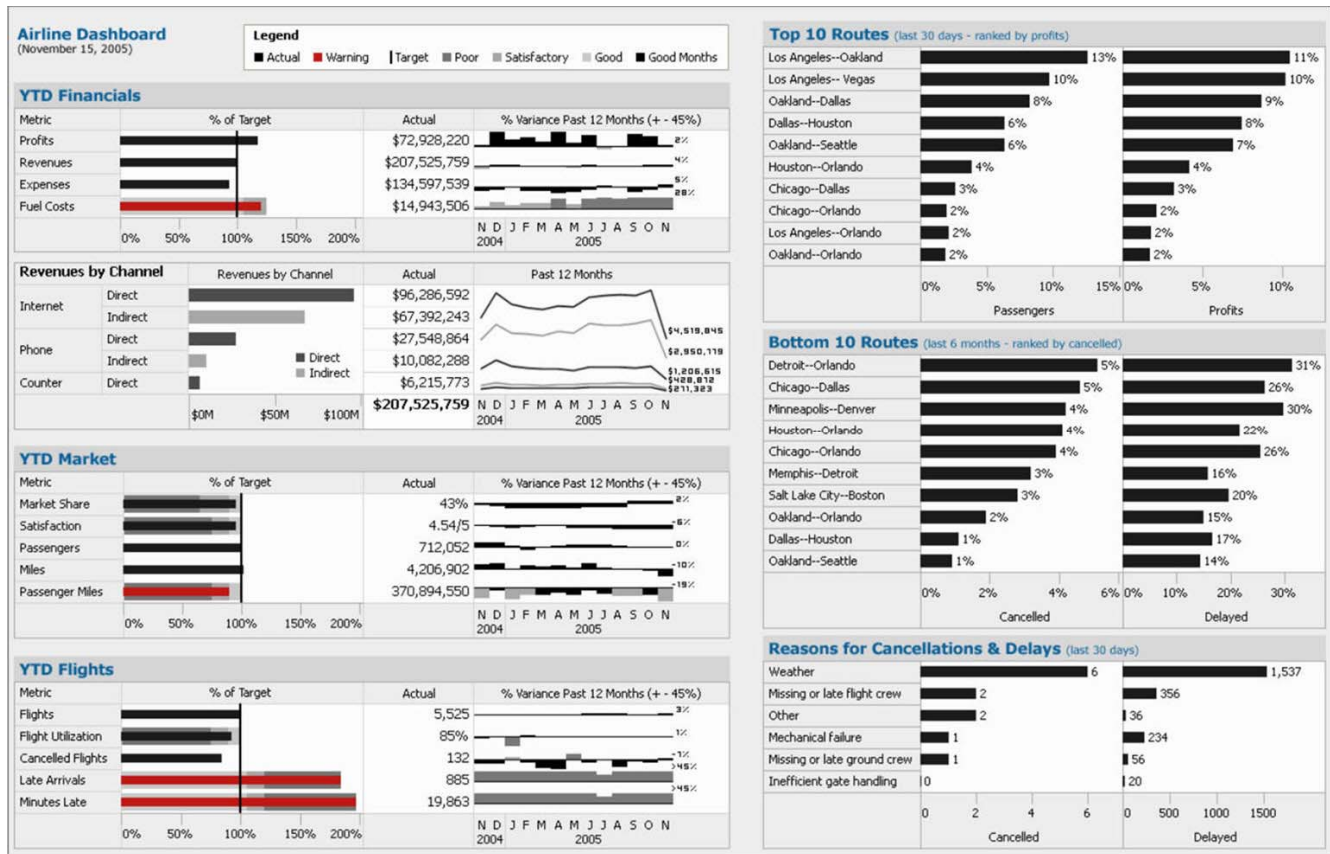


Figure 5: Submitted by Jock Mackinlay of Tableau Software

Jock Mackinlay of Tableau Software, who submitted this entry, was the overall winner of the DM Review magazine's 2005 competition. In this year's competition, Jock submitted the winning solution to scenario #3, featured in October, and tied with another participant for the winning solution for scenario #5, which I'll feature next month. Here are a few of the strengths of Jock's dashboard solution:

- This dashboard much more directly measures performance than the others so far. Notice how the primary measures that are shown in the left half of the dashboard are displayed as percentage of target, making performance crystal clear and rapidly recognizable.
- The measures were nicely organized into groups of related information.
- Color was used effectively to highlight what needs the most attention. The only time a non-gray color appears—red in this case—is when attention is required.
- The white background in the graphs provides nice contrast for viewers to easily read the data.
- Overall, the visual design of this dashboard is very clean, well balanced, and aesthetically pleasing.

The ways that this dashboard could be improved are minor—nitpicking in most cases—but even minor improvements can translate into significant benefits for those who use our dashboards. Here’s my list:

- The three sets of “% variance past 12 months” bar graphs, at first glance seem a little difficult to read, because they are squeezed into such little vertical space, but actually, they do what they need to do fairly well. They are only meant to give a quick sense of the historical values that have led to the current values. I include this in the list of potential improvements only because I believe that by using a line to encode this information the historical trend of variance to target would be slightly easier to see.
- The graphs on the right, which display the top and bottom routes, work well, but it is questionable whether this information requires graphical display. Do these routes need to be compared to one another or simply listed in order? Even if the graphical display is useful, the fact that these two graphs are larger than all the others gives them prominence that is probably undeserved.
- The lines that encode revenues per sales channel for the past 12 months are not labeled. You can tell which line is which channel because the lines appear from top to bottom in the same order as the bars in the “Revenues by Channel” graph to the left, but will this always be the case? Because the sales channels are grouped as direct and indirect within the larger categories of Internet, phone, and counter sales, if the value of one of the channels surpasses another, it might not be possible to preserve this arrangement and still order the sales channels from top to bottom in the same order as the lines. This problem could be solved by labeling the sales channels “Direct Internet,” “Indirect Internet,” “Direct Phone,” etc., rather than the current hierarchical grouping, which would allow the channels be ordered by sales amount without difficulty, no matter what their values.
- The last problem is perhaps the biggest of this meager list. By including the current month, which is incomplete (data has only been collected through November 15<sup>th</sup>), in the “Revenues per Channel” line graph, the downward slope of the lines in November suggests a severe decline in sales, which is not the case. The current month must be displayed differently from previous months to correct this erroneous suggestion.

## The Winning Solution

It’s finally time to take a look at the winning solutions. Andreas Flockermann of [BonaVista Systems](#) (a new business intelligence software company located in Germany) submitted the winning solution for this scenario, which he created using Excel, with the help of his company’s add-in product named [MicroCharts](#).

Take a moment to look at Andreas’ solution (Figure 6) to determine for yourself how well he succeeded in communicating this large collection of information in a small amount of space for effective monitoring.

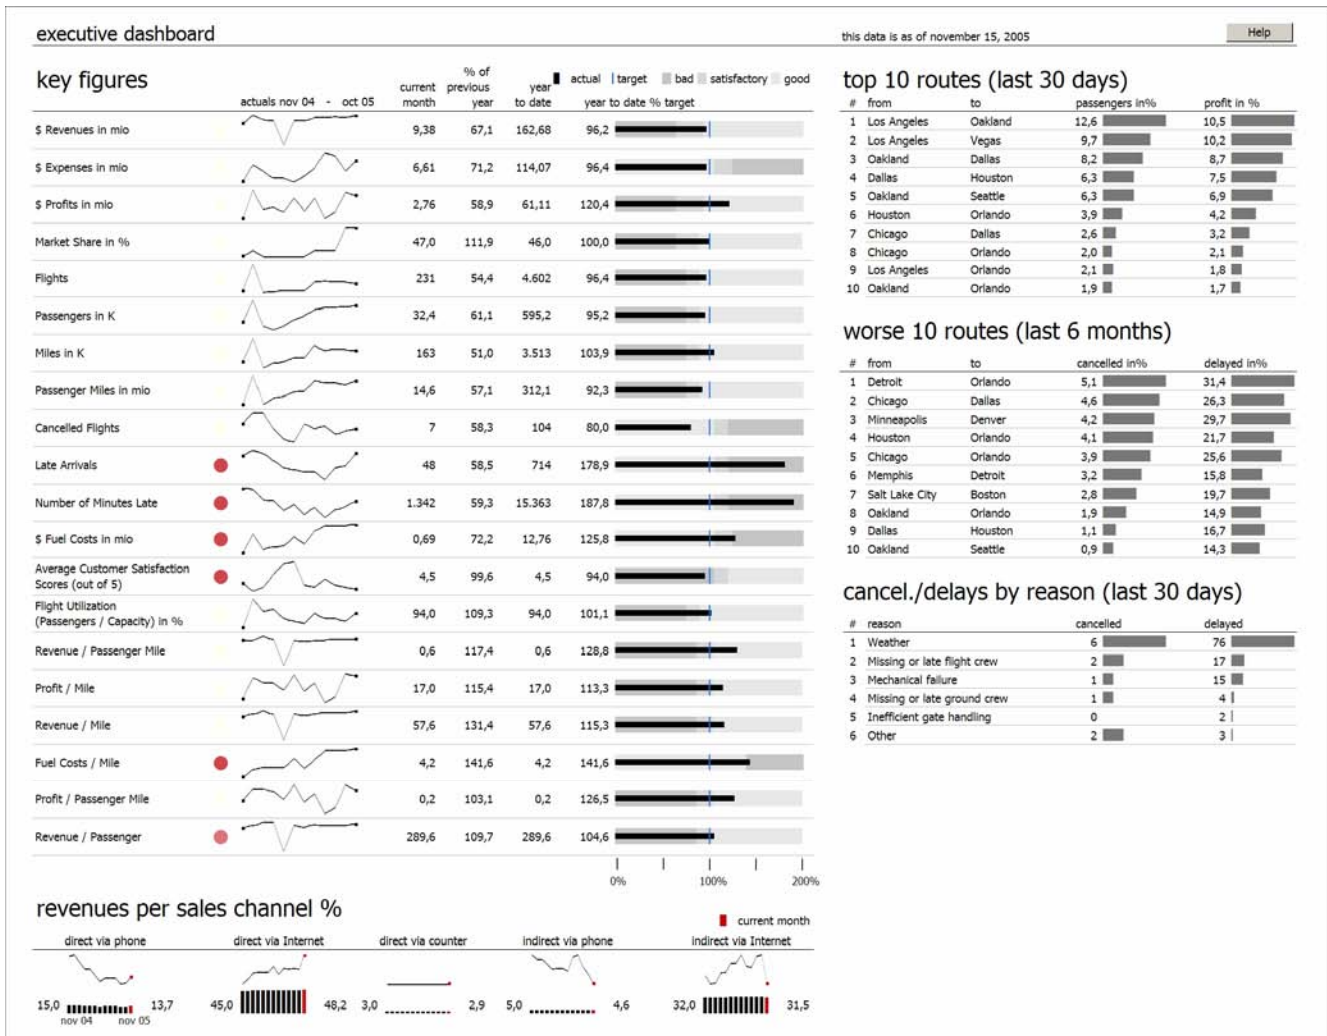


Figure 6: Andreas Flockermann of BonaVista Systems' winning solution.

Andreas' solution might not be perfect (we'll look at how it could be improved in a moment), but it succeeds in a few ways that placed it slightly ahead of the last two solutions that we examined. Here's a list of its design achievements, including its unique achievements highlighted in **black**:

- No space was wasted by overemphasizing unimportant content, such as the title of the dashboard.
- It is very easy to know where to look, based on the red icons, because this is the only non-gray-scale color that appears on the dashboard (except for the final values in each of the "revenues per sales channel %" graphs at the bottom, which I'll address in a moment).
- Precise values have been included for the measures in the form of text in addition to their graphical representation. This makes it possible for viewers to quickly scan the graphics for the basic information that they need about performance, including patterns, trends, and relationships, but to also see precise values that are important to them, presented in a way that doesn't distract them when scanning the graphics.
- The appropriate amount of emphasis was placed on the current measures compared to historical measures. The use of sparklines gives the historical information less

emphasis than the current information, displayed as bullet graphs, yet the sparklines adequately display the historical trend and patterns that provide the necessary context.

- The top ten routes and the worst routes have been displayed primarily as a table, but with the addition of bars to graphically display the relative magnitudes of the values in a simple, visually subdued manner that doesn't overly emphasize the importance of this part of the dashboard.
- The simplicity of this design, with almost no unnecessary visual content, provides an exceptionally clean presentation of the information, causing it to stand out and tell its story with a clear voice.
- The other solutions that we've examined displayed revenues by sales channel in a single graph, which makes comparisons between the channels easy, but makes it hard to see the patterns of change through time for the sales channels with the lowest values. For example, notice how flat the indirect phone sales channel line looks in Figure 1, compared to the sparkline that represents it in Figure 6. There is a great deal of change that cannot be seen in a graph that is scaled to include much greater values, such as those for the direct internet sales channel. By using a combination of sparklines to show change through time and the small bar graphs with a common scale to allow comparisons between the five channels, Andreas was able to support both needs effectively.
- A few additional measures, expressed as ratios, such as revenue per mile, were included, which I did not provide in the spreadsheet. Although we cannot say for sure that airline executives would find these measures important enough to monitor, the fact that they were included indicates the level of thought about the measures that dashboard designers should always pursue.

Despite its exceptional qualities, even the winning solution can be improved in minor ways, such as:

- Similar to some of the other solutions, the primary measures in this dashboard have not been organized into meaningful groups.
- The "revenues by sales channel" graphs would communicate useful ranking information if the sales channels were sorted from highest to lowest revenues.
- The power of the red icons to attract attention was slightly undermined by the use of red to also indicate the current month's values in the "revenues by sales channels" graphs. Making the current month's values look different from the rest gives them useful emphasis, but a color other than red would have worked better. For example, using gray lines and bars for past values and a black point on the sparklines and black bar on the bar graphs for the current month would have worked just as well, without compromising the eye-catching quality of the red icons.
- Although the presence of the white space in the bottom right corner of the dashboard is not necessarily a problem, the available space might have somehow been used to enlarge other parts of the display, such as the "revenues by sales channel" graphs.

I hope you've found these dashboard solutions and my critiques informative. Most dashboards are poorly designed, often to the point of absurdity. Hopefully, articles such as this, which feature effective visual design, will provide useful examples of how much better dashboards can be. Next month, I'll feature the winning solution to the final scenario of the

Business Intelligence Network's 2006 Data Visualization Competition. You won't want to miss this, because in the final scenario, participants were given the opportunity to strut their stuff by submitting their best data visualization work of any kind, without restriction.

---

## About the Author

Stephen Few has worked for over 20 years as an IT innovator, consultant, and teacher. Today, as Principal of the consultancy Perceptual Edge, Stephen focuses on data visualization for analyzing and communicating quantitative business information. He provides training and consulting services, writes the monthly *Visual Business Intelligence Newsletter*, speaks frequently at conferences, and teaches in the MBA program at the University of California, Berkeley. He is the author of two books: *Show Me the Numbers: Designing Tables and Graphs to Enlighten* and *Information Dashboard Design: The Effective Visual Communication of Data*. You can learn more about Stephen's work and access an entire library of articles at [www.perceptualedge.com](http://www.perceptualedge.com). Between articles, you can read Stephen's thoughts on the industry in his blog.