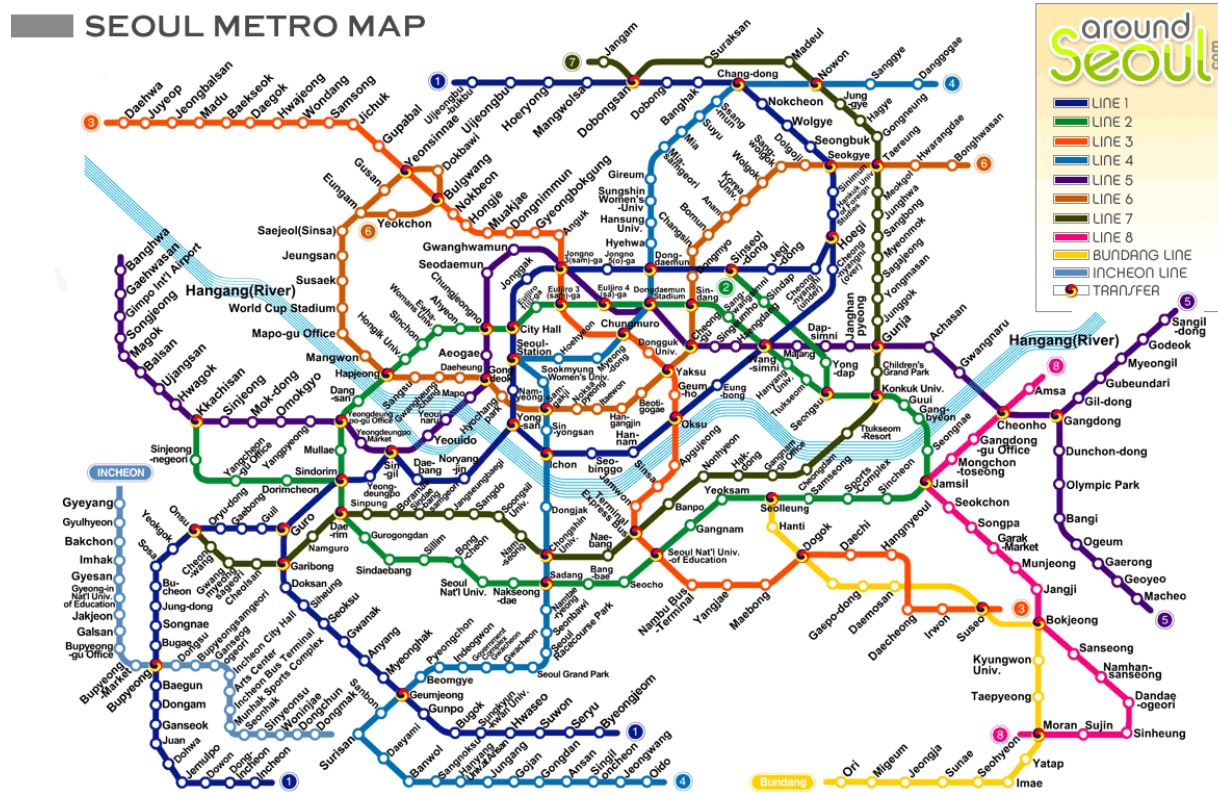


An everyday example of architecture documentation: Subway maps

I've always been interested in how you write down software designs so that others can use them. If I'm completely honest I will admit that I find that question more interesting than how you come up with the designs themselves (although I believe the act of writing them down is inextricably intertwined with coming up with them). As Len Bass, Felix Bachmann, David Garlan, James Ivers, Reed Little, Paulo Merson, Robert Nord, Judy Stafford and I work on finishing the second edition of *Documenting Software Architectures: Views and Beyond*, architecture documentation is on my mind more than usual these days.

I'm in Seoul, Korea, as I write this. Yesterday I rode the Seoul subway across the city to do some shopping, and I used the subway maps in the stations and on the trains to guide my journey. As I sat on the train as it made its way from station to station, I stared at the system map inside each car:



It occurred to me, as it has in the past, that subway maps are very good everyday examples of *architecture documentation*. A subway system architecture isn't a software architecture, for sure, but it's certainly an architecture. I started wondering what parts of our *Views and Beyond* approach to

documentation I could recognize in a subway map and the associated pieces of information the public can find in a subway station.

Stakeholders

My first piece of advice to any architect producing a piece of documentation is “Know your stakeholders.” Who are the stakeholders for a subway map like the one above?

For my thought exercise, the stakeholder was someone who wanted to go from Point A (my hotel) in Seoul to Point B (Namdaemun Market, where I wanted to shop) by riding the subway. In other words, me.

I was a *rider*, surely a class of stakeholder. Unlike most of my fellow riders that day, I was a rider with only a rudimentary knowledge of Seoul – that is, I was a *visitor*, not a *resident*. Non-rider stakeholders for a map like this might include owners of businesses near subway stations, who want to tell their customers how to reach them. Or drivers of cars who want to pick someone up or drop someone off at a subway stop.

What other stakeholders come to mind? I don’t mean stakeholders like the people who dug the tunnels or measure the electric current in the third rail; they have other documentation for that. I mean stakeholders who are members of the subway system’s general public.

What kind of view?

What kind of architecture view is being shown by that map?

To answer this, let’s first ask what the elements and relations are. I think the elements of the map are *stations*, and *tracks* that connect the stations to each other. The relation shown by the map is “is connected to,” establishing an adjacency relationship between pairs of stations. The map also shows a logical grouping of stations that are transitively connected to each other; the groups are called *routes*, and each is denoted with its own color.

The routes are not superfluous. At stations where two routes meet, we would not know what adjacent station a departing train would next visit without knowing its route. A red-route train won’t suddenly go visit a blue-route station just because it’s adjacent.

What are the topological restrictions? Two tracks can’t occupy the same space (although they can be very close to each other). A route must be contiguous – no gaps. Every station has to be connected to at least one track, and every track it’s connected to is connected to a different station. Tracks don’t branch. Routes are allowed to loop, but they don’t have to. Routes don’t branch.

There are behavior restrictions as well. A train can only visit a station adjacent to the last station it entered that’s on the same route. (In other words, trains only travel on tracks and don’t magically hop from station to station arbitrarily.) If a route loops, we could restrict things so that trains don’t ever

backtrack; I don't know if any subway systems allow that or not. (If a route doesn't loop, but rather has two endpoints, then trains must backtrack.)

After establishing these things, I think the route map is a component-and-connector view of the subway system. It shows connections between stations that are used when the system is in operation – that is, at “runtime.”

What would a module view of the subway system look like? I think it might represent the abstract layout for a subway station, which might then be instantiated in several places (and in several ways)¹.

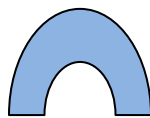
What would an allocation view look like? An allocation view maps elements in “our” architecture (here, the subway system's stations and routes) to elements outside that architecture. One allocation view might be an allocation of groups of train cars (i.e., trains) to routes, but this isn't a piece of documentation that we would normally find as a member of the public visiting the subway.

Context

In *Views and Beyond*, we advise architects to add a context diagram to each view they produce. A context diagram shows how the architecture shown in the view relates to its environment. We suggest that there are times when you might want to combine the primary presentation with the context diagram for clarity and conciseness.

Context turns out to play a surprisingly important role in subway maps.

By consulting my travel guide to Seoul, I knew that my hotel was four blocks east of the Yeoksam station, and the market I wanted to visit was near the Hyeonhwajong station. Good luck finding either of these just by scanning the map I showed earlier. But my travel guide also told me that the Yeoksam station was on the Number 2 line – that's the green one – and I knew the hotel was south of the river. In fact, a map in the guide showed that the hotel was south of where the river made a shape like this:

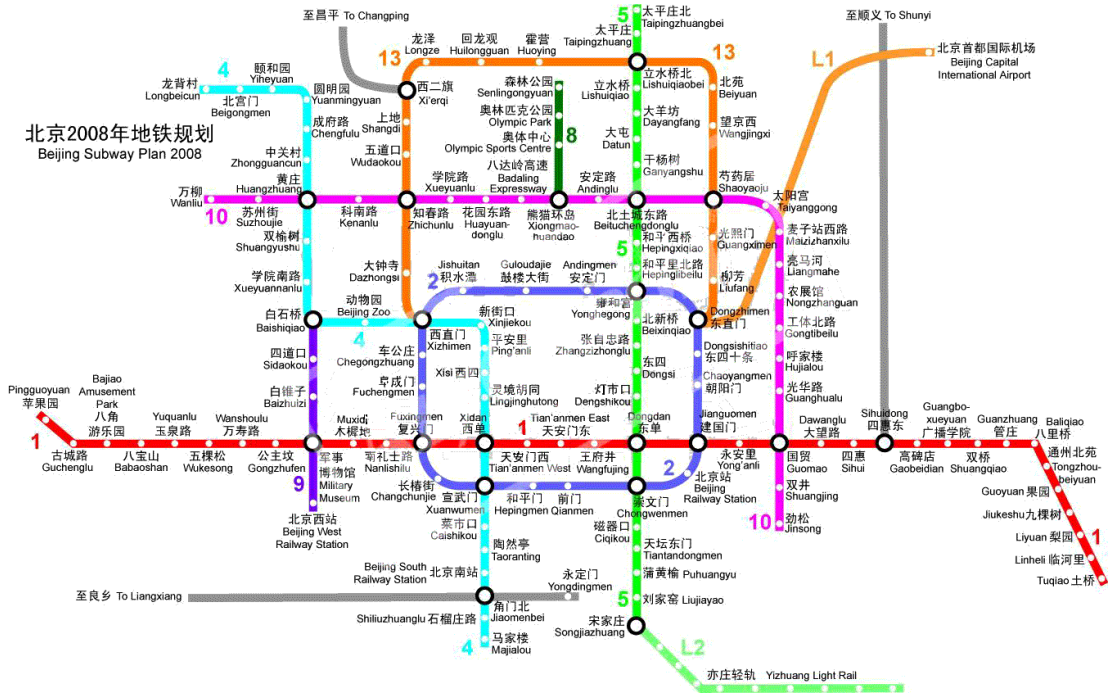


Now have a look. The Yeoksam station is pretty easy to find once you know these two bits of additional information. My travel guide had similar information for the market.

I found that interesting. That perfectly wonderful architecture picture was almost useless to me until I had two pieces of *context* with which to relate it: The location of my points of interest relative to a station, and the location of both relative to a feature of the world outside the subway system (in this case, a river).

¹ Thanks to David Garlan for this observation.

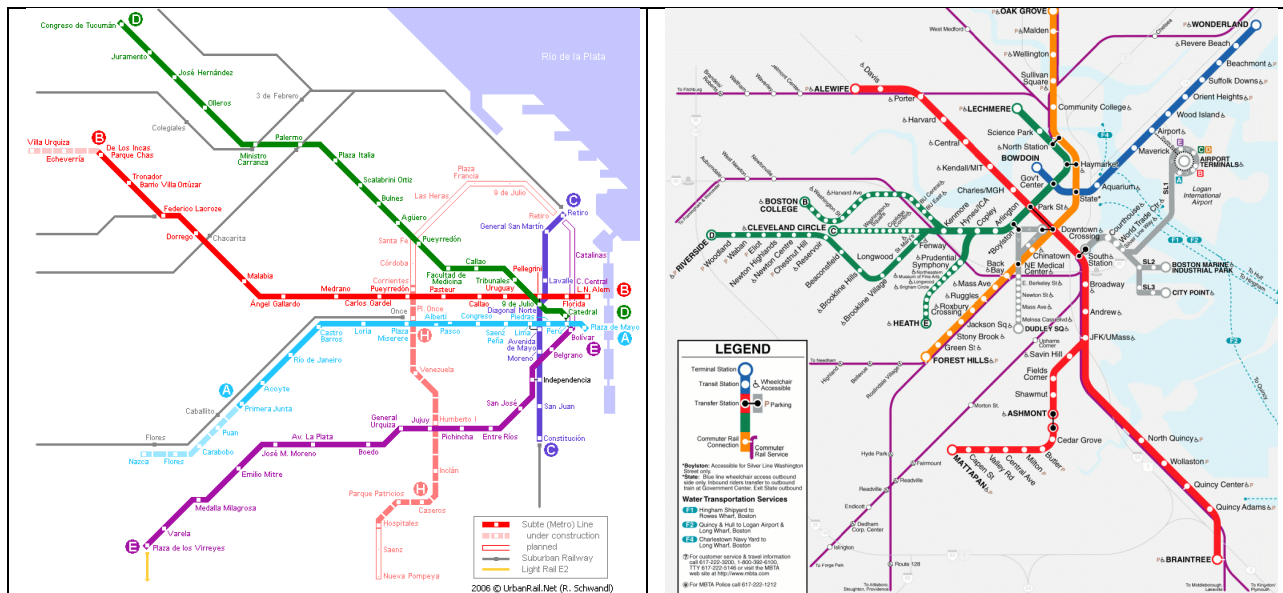
And Beijing^{ix}:



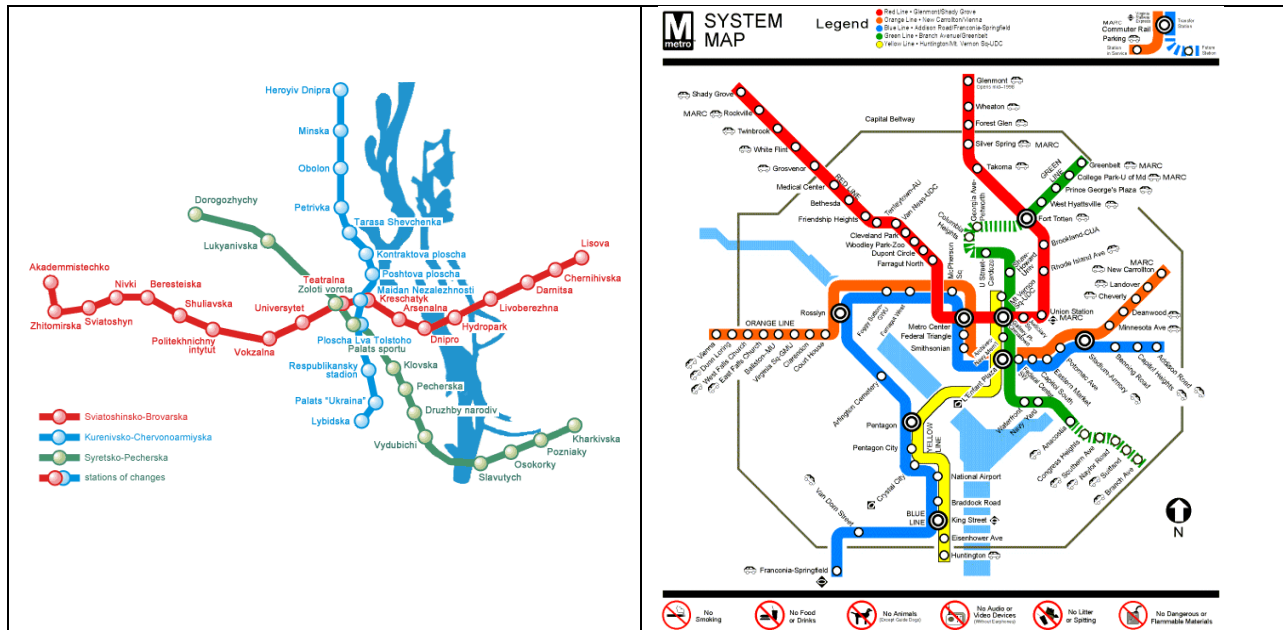
Maps with a little context: Water, water everywhere

Searching for subway maps on the internet, I discovered that many of them put in a very basic piece of geographical context information: Rivers or other major bodies of water in or near the city. Usually the bodies of water are as stylized (straight-line edges, 45-degree angles) as the subway lines themselves.

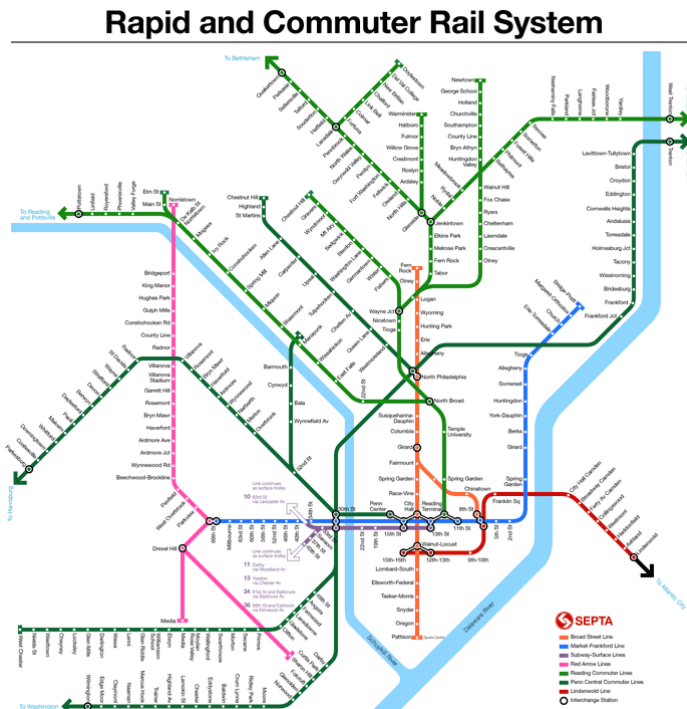
For example, here are subway maps for Buenos Aires^x and Boston^{xi}:



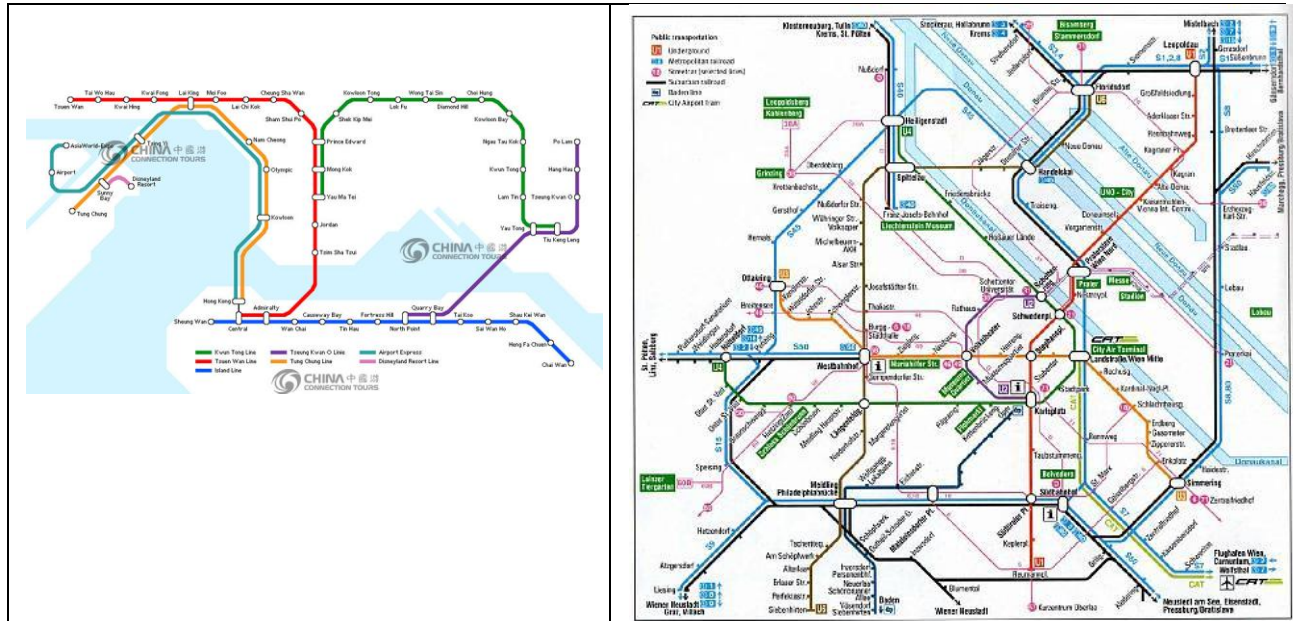
Kiev^{xii} and Washington^{xiii}:



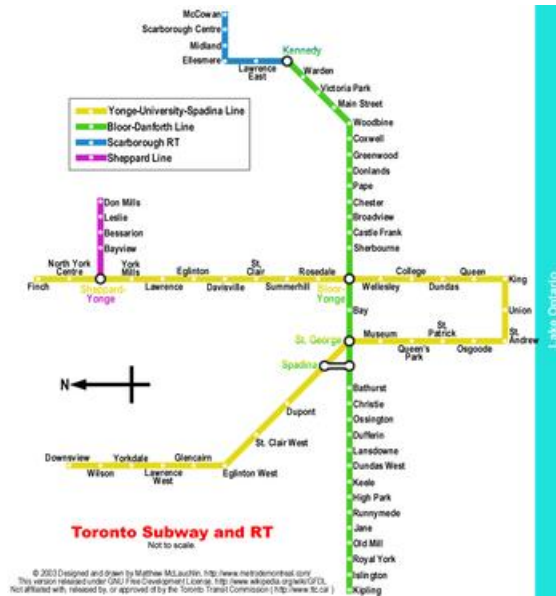
Philadelphia^{xiv}:



Hong Kong^{xv} and Vienna^{xvi}:



But showing water doesn't always help. Here's a subway map of Toronto^{xvii} that shows a major body of water (Lake Ontario) as a geographical reference, but rendered in such a bland way as to be useless. This map says every station is west of the lake. Right. Pretty sure I knew that already. Thanks a bunch.



The next one^{xx} shows a great deal of geographical information along with a few more surface landmarks, such as “Universal Studio,” and what I think is a major highway running vertically along the upper right side of the map (disappointingly, that symbol isn’t explained in the key, so points off for that):



Which of these three maps of the Osaka subway system do you prefer? I think the answer depends on how you’re using the map at the time. If I want to find a subway station, I prefer the third map. If I’m already on the train and know my destination, I prefer the first map.

This tracks intuitively with the purpose of context. If I’m looking for a station, I’m *outside* the system and I need some geographic references. If I’m on a train then I’m *inside* the system, and my attention can remain focused inside the system.

So, with apologies to Bucharest, Milan, Kyoto, Busan, Sapporo, Moscow, Mexico City, and Beijing, I confess I see the usefulness of your context-free maps. If I’m riding on the train and know my destination, that’s all I need. But I had to use some other information – usually a street map showing the subway stations – before I could do that.

A case study of context and controversy: New York City subway maps

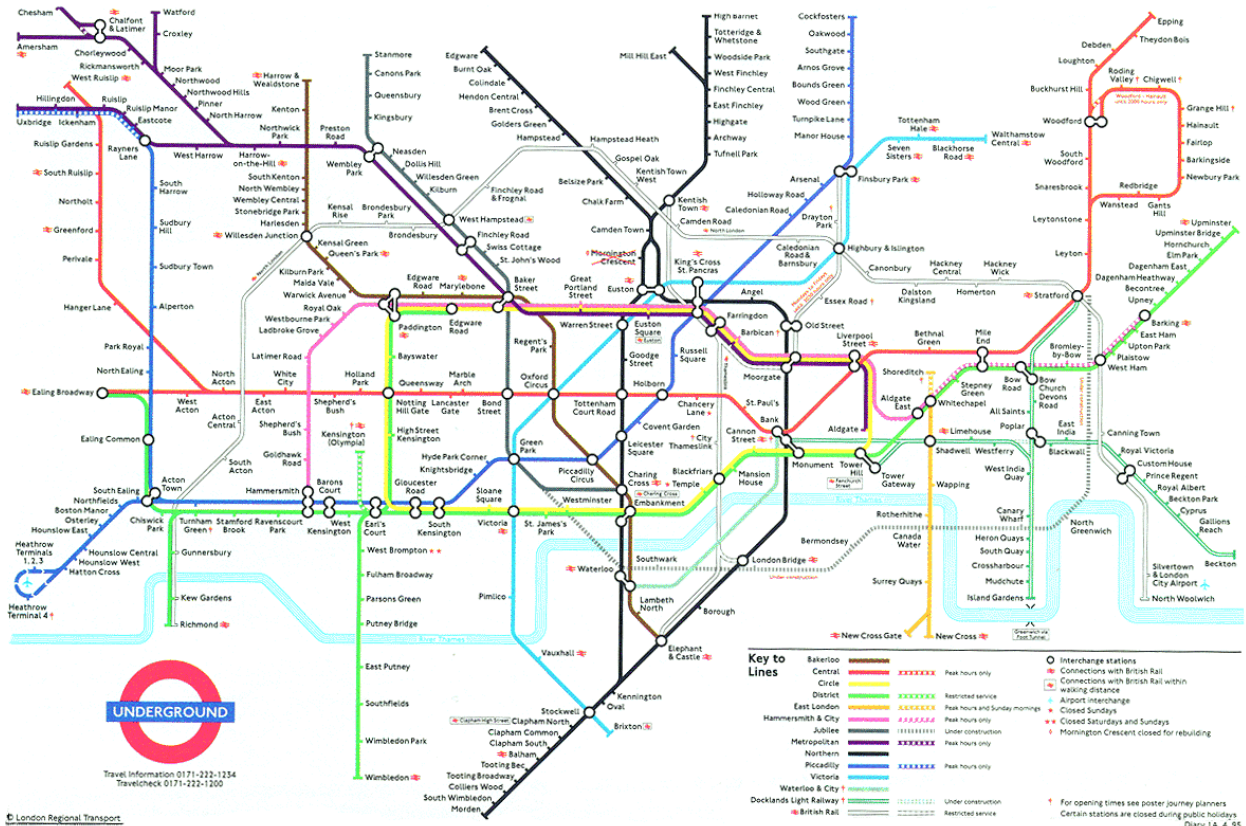
The New York subway *map*, not to mention the subway itself, is so a part of New York culture that it has attained iconic status. In 1972, the city adopted a map^{xxi} designed (and I do mean *designed*) by Massimo Vignelli³:



This abstract style for subway maps is traced back to the London Underground map that took pretty much its present form in 1933⁴. Here's the modern version of that map^{xxii}:

³ Massimo Vignelli is a major figure in design and architecture. You can read his bio at <http://www.vignelli.com/awards/massimo.pdf>. You can see a clip of him talking about his 1972 map at <http://www.helveticafilm.com/vignellimap.html>.

⁴ This is a map style that has been emulated by subway map makers around the world: Parallel lines, 45- or 90-degree turns, stylized bodies of water, drab backgrounds, and a minimum of geographical information. Did you notice that *every other map in this paper* (except for the third Osaka map) follows this same abstract layout with only minor variations? This visual style is so well-known that you can easily find “maps” in this style of the human digestive system, a landscape of internet technology companies, a snapshot of the publishing industry, and a genealogy of internet search engine optimization strategies. This style has become so ubiquitous that it's difficult to imagine any other. But subway maps used to consist of regular geographic or street maps of cities, with the stations and actual track locations superimposed.



For context, it shows the Thames and nothing more.

Here's how *The New York Times* explained the minimal context information in the 1972 map:

"In 1972, Mr. Vignelli designed a completely new schematic map... that showed New York's subway routes as rich, contrasting stripes of color, marching in lock step across a white background, and turning only at 45- or 90-degree angles. In contrast to the brilliance of the subway routes, aboveground New York was almost invisible: the outlines of the boroughs were stubby and squared-off; the parks were gray boxes; and the water was tan. The map defiantly ignored the city's geography: the Broadway line was shown crossing the Eighth Avenue line at 42nd Street (they actually cross at Columbus Circle); Bowling Green appeared above Rector Street (it's below); and Central Park was a small square rather than a tall rectangle. "Of course I know Central Park is rectangular and not square," Mr. Vignelli said... "Of course I know the park is green, and not gray. Who cares? You want to go from Point A to Point B, period. The only thing you are interested in is the spaghetti." As it turned out, New Yorkers were interested in more than the spaghetti. Almost as soon as Mr. Vignelli's map arrived at stations, people started complaining about its failure to describe the city's geography. Tourists were getting off the

subway at the bottom of Central Park and trying to stroll to the top, for example, expecting a 30-minute walk.⁵ (New York Times, September 3, 2006^{xxiii})

Here, Vignelli is arguing that people don't need the context. People disagreed. In 1979, New York City adopted a new subway map that

“showed more geographical information than any previous New York subway map. It was the first since the 1930's to reproduce the street grid; it gave neighborhood names and pointed out major landmarks. Parks, of course, were green, and the water was blue. Subsequent maps have imported even more geography, and added more information.” (ibid)

The stakeholders, valuing usefulness over stylistic purity, won out. Context (in the sense of an architectural context diagram) matters.

View packets

In the *Views and Beyond* approach, we suggest breaking up large and complex pictures up into a series of smaller pictures, and then documenting each one of those. We call these view packets; a view packet is the smallest bundle of architecture documentation you'd give to a single stakeholder.

Are there any view packets in subway systems? Yes. As I was riding the Seoul subway, it occurred to me that a view packet was on the wall of each car. I couldn't find an image of one, so I'll do a little coarse surgery on the big Seoul subway picture to show you what I mean:



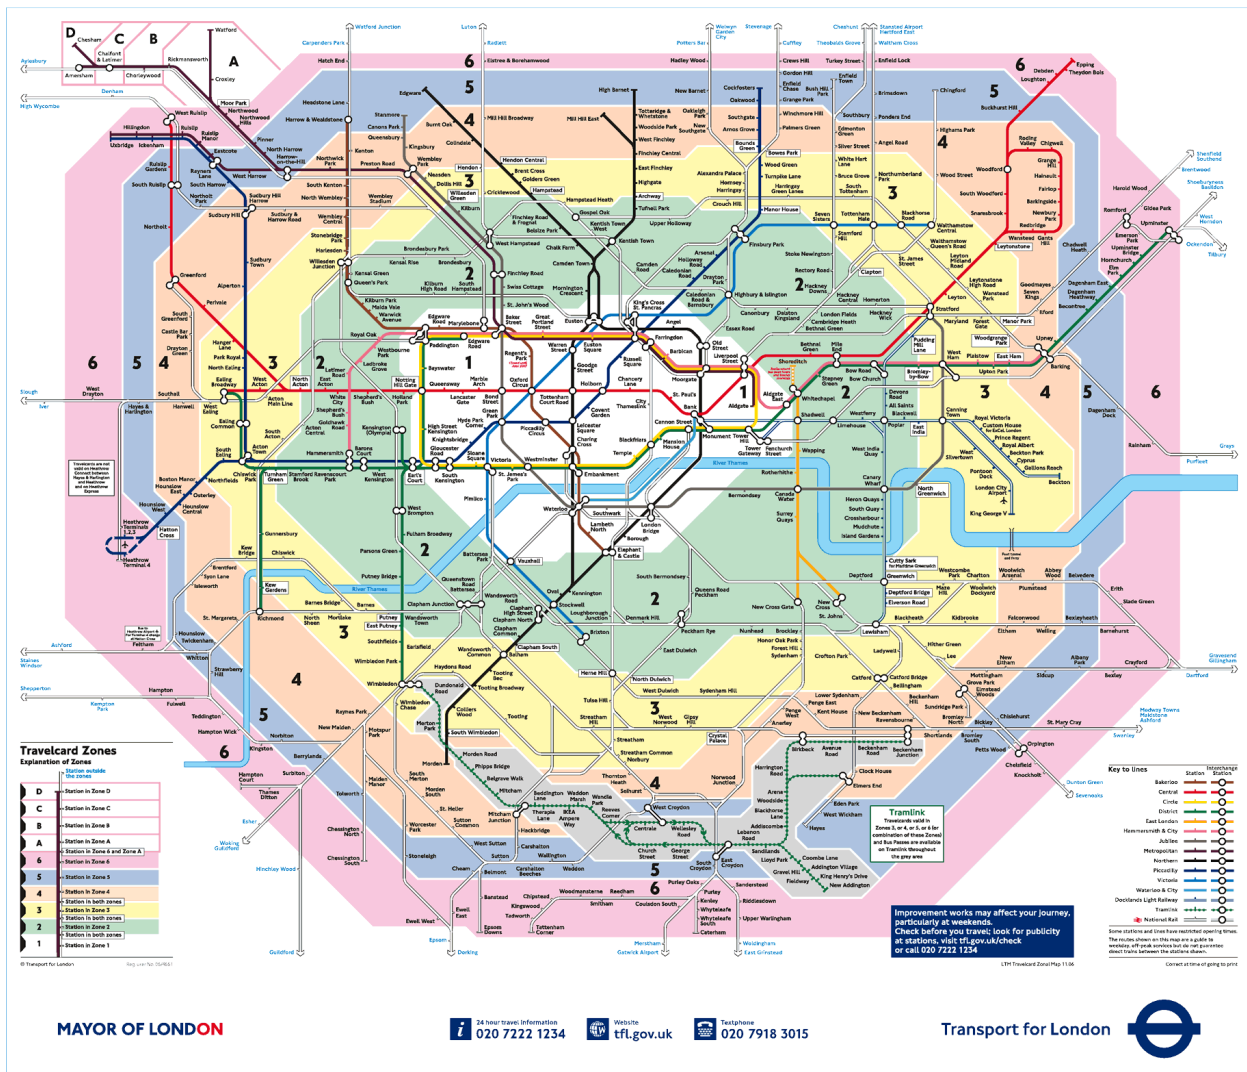
Each train car showed only a part of the map, the route that the train was running. Who are the stakeholders for this “view packet”? The riders on Route 1, of course. You can use this to count the number of stops left before you need to disembark.

⁵ This sentence is saying that a particular kind of *stakeholder* found the documentation wanting for a particular use.

Overlays

In Views and Beyond, an overlay is a picture showing information from two or more otherwise-separate views. I found two examples on-line of subway maps that I believe represent overlays; I'll show one below.

Many subway systems have zones that determine the fare; same-zone rides are cheaper than cross-zone rides, for instance. To determine your fare, you have to know the zone(s) of your origin and destination stations. How would you document that? Some subway systems do it with a table – look up your station, and there's the zone it's in. But since your system map shows all the stations anyway, it might be more convenient to overlay the system map with the zone information. Here's a zoned map^{xxiv} of the London transport network that I believe is an overlay:



A separate “zone view” would consist of stations and zones (tracks don’t matter). Our “station and track view” shows, well, stations and tracks. This picture above both – a classic overlay.

Documenting Behavior

In Views and Beyond, we prescribe documenting behavior as part of a view. What would the behavioral documentation for a subway system look like? I think it would look something like this schedule board (the example is from Kyoto^{xxv}):



Put all of these together – that is, all of the schedules for all of the trains – along with the maps showing architectural structure, and you can do all the kinds of analysis you’d expect for a subway system: track usage, collision avoidance, passenger throughput, simulation and modeling of (for instance) performance if a train breaks down, and so forth.

Oh, and notice the “view packet” – a partial “stations and tracks” view – underneath the arrival/departure data for that route. It shows the layout of the route, plus routes that intersect.

Schedules for individual trains reaffirm the behavior restrictions we mentioned earlier; for example, trains follow their assigned routes and go from station to station in the order determined by the track layout.

Interfaces

What about interfaces? There are all sorts of interfaces present in subway systems: Trains’ wheels interfacing with tracks, trains with platforms (so people can step and off safely), people with train cars

(via doors and seats and hand-hold loops), and many others. But is there subway system documentation aimed at the general public that explains an interface?

I think so. A key interface between a rider and the system is the protocol involving money. Pay some money, get to ride. Here's a key part of that interface (this example is from the Shanghai system^{xxvi}); the instructions on the screen constitute its documentation:



Other parts of this same interface include whatever you get out of the machine (a fare card? a token?) and what you do with it to enter the system (scan it? drop it in a slot?)

Another key interface is the physical interface between people and the subway stations. Enter the station and follow the posted directions – the documentation – to get to your train.

Summary

Subway maps are component-and-connector diagrams of subway systems. How much context to show on the map is a surprisingly important issue. The documentation concepts of behavior, interfaces, and view packets all show up in the everyday information the public is shown to help them utilize these systems.

I hope you've enjoyed this little thought experiment. It has made me more confident that the basic Views and Beyond approach is a good fit for architectures other than the purely software kind.

If you disagree with some of my conclusions or want to extend the Views and Beyond analogy in ways I didn't think of, write and let me know or post a comment.

Happy commuting,



Postscript: Hierarchical components?

David Garlan asked a wonderful question about this analogy: Are there hierarchical components in a subway map? There are, of course, in "regular" maps. National maps are elaborated by state or provincial maps. Those are elaborated by city maps. Some city maps are elaborated by maps that show, for example, details of the city center.

Can you find an example in subway maps?

ⁱ <http://www.aroundseoul.com/images/seoulSubwayMap.gif>

ⁱⁱ <http://www.aboutromania.com>

ⁱⁱⁱ http://www.dsi.unimi.it/img/map_metro.png

^{iv} <http://www.traveljapan.com.au/images/kyoto/kyoto%2520subway%2520map.gif>

^v <http://www.lookatkorea.com>

^{vi} http://upload.wikimedia.org/wikipedia/commons/thumb/c/c4/Sapporo_Subway_Map_pl.svg/800px-Sapporo_Subway_Map_pl.svg.png

^{vii} http://www.aisee.com/graph_of_the_month/metro.htm

^{viii} <http://mexico-on-line.com/mexico-city-maps/mexico-city-metro-subway-map.gif>

^{ix} <http://www.beijingchina.net.cn/transportation/beijing-subway-map.jpg>

^x www.bairesgetaway.com/SITE/AREA/SUBWAY_MAP.html

^{xi} subway.umka.org/.../washington-square.html

^{xii} www.pixopark.com/kyev_subway_map.php

^{xiii} www.nidcr.nih.gov/.../Maps/MapMetro.htm

^{xiv} bellehellama.wordpress.com/.../

^{xv} : www.china-tour.cn/.../HongKong-Subway-Map.htm

^{xvi}

<http://www.vickywoodard.com/album/Vienna%20Aug%202007/Trip%20Maps/slides/vienna%20subway%20map.html>

^{xvii} www.fact-archive.com/.../Toronto_Subway_and_RT

^{xviii} www.urbanrail.net/as/osak/osaka.htm

^{xix} <http://www.bento.com/pix/subway-osaka2.gif>

^{xx} : www.bento.com/kansai/platform-osaka.html

^{xxi} http://www.mitchmagee.com/blog/wp-content/uploads/2008/01/system_1972.jpg

^{xxii} <http://www.barclayweb.com/DESTINAT/unitedki/ENGLAND/undgr.gif>

^{xxiii} <http://www.nytimes.com/2006/09/03/nyregion/thecity/03maps.html>

^{xxiv} <http://www.uktravelbureau.com/images/london-underground-tube-map-revised.gif>

^{xxv}

http://2.bp.blogspot.com/_qvZ1h1bgBRo/SrAo1Lb_23I/AAAAAAAAAil/YqTx4GxJOiE/s400/kyoto_subway_time_sch_edule.JPG

^{xxvi} gochina.about.com/.../ss/Use_Subway_2.htm

Copyright 2009 Carnegie Mellon University.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN “AS-IS” BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and “No Warranty” statements are included with all reproductions and derivative works.

External use. This document may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

For information about SEI reports, please visit the publications section of our website (<http://www.sei.cmu.edu/publications/>).