

6-2011

# The Structure of Federal eGovernment: Using hyperlinks to analyze the .gov domain

Ryan Whalen

*Northwestern University, r-whalen@northwestern.edu*

Follow this and additional works at: [http://opensiuc.lib.siu.edu/pnconfs\\_2011](http://opensiuc.lib.siu.edu/pnconfs_2011)

---

## Recommended Citation

Whalen, Ryan, "The Structure of Federal eGovernment: Using hyperlinks to analyze the .gov domain" (2011). 2011. Paper 8.  
[http://opensiuc.lib.siu.edu/pnconfs\\_2011/8](http://opensiuc.lib.siu.edu/pnconfs_2011/8)

This Article is brought to you for free and open access by the Conference Proceedings at OpenSIUC. It has been accepted for inclusion in 2011 by an authorized administrator of OpenSIUC. For more information, please contact [opensiuc@lib.siu.edu](mailto:opensiuc@lib.siu.edu).

# **The Structure of Federal eGovernment: Using hyperlinks to analyze the .gov domain**

Ryan Whalen

Media, Technology and Society  
Northwestern University School of Communication  
[r-whalen@northwestern.edu](mailto:r-whalen@northwestern.edu)

Early research report; please contact author prior to citation

Prepared for the 4<sup>th</sup> Annual Political Networks Conference  
June 2011 University of Michigan, Ann Arbor, MI

## **Abstract**

This paper uses the hyperlink structure of federal web sites within the .gov domain to answer two research questions: to what degree does the online structure of the federal government mirror its offline hierarchy, and to what degree does the .gov web graph mirror the greater WWW graph. Findings of subgraph link analysis and Krackhardt's graph theoretical dimensions of hierarchy analysis demonstrate clear hierarchy within the .gov domain, but also suggest great discrepancies in the linking patterns of different government departments. Structural analysis suggests that the .gov web graph is indeed a fractal leaf of the greater WWW graph.

## **The Structure of Federal eGovernment**

While it may be tempting to think of *the government* as some sort of unitary entity working away in Washington, it would be a mistake. The American federal government is so vast, varied and complex that just comprehending on the most basic level how its hundreds of entities are related to one another is a challenging, if not unfeasible, task for most. Indeed, in his latest State of the Union speech President Obama suggested that this complicated structure was outdated, stating that: “We live and do business in the Information Age, but the last major reorganization of the government happened in the age of black-and-white TV. There are 12 different agencies that deal with exports. There are at least five different agencies that deal with housing policy” (Obama, 2011).

The federal government’s transition into the “information age” has brought it online where it has built a complex network of websites to provide information both to citizens and government employees. This move towards e-government has left us with digital traces of the government’s complex network of institutions, agencies and initiatives that we can use to help us better understand the structure of the American government and determine how the imperatives of online networked organization manifest in the context of a hierarchically structured institution like the government.

**Hyperlink Structure.** Hyperlinks can be used in many ways. In some cases they function as citations, showing readers where an author received her inspiration. In other cases they provide an interactive way to design and layout content, allowing viewers to click a hyperlinked image and view a larger version of the same. In many more cases they serve to link ideas or information, and thereby represent a semantic connection. The vast majority of inter-site .gov

links are of this last type. They guide visitors to information that the site authors and designers feel may be relevant. On a visit to [whitehouse.gov](http://whitehouse.gov) looking for information about the BP oil spill, a viewer may come across links to a number of agencies and initiatives relevant to the spill. These links represent a degree of similarity between the sites in question. When considered in aggregate, hyperlinks “reflect deep social and cultural structures” (Halavais, 2008) and can lay bare organizational relations that would not otherwise be evident.

Analyses of hyperlink structures have been used to examine diverse types of organizational relations. Many early academic studies of linking patterns looked at academic institutions themselves (e.g. Kim, Park, & Thelwall, 2006; M Thelwall, 2003; Wilkinson, Thelwall, & Li, 2003). These studies build on the tradition of bibliometrics, treating the hyperlink as a digital citation. This tradition of link mapping continues, and indeed has gone on to inspire at least one major project aimed at ranking institutions based to a large degree on hyperlink structures (“Ranking Web of World Universities,” 2011).

In her studies of policy networks, McNutt (Kathleen McNutt, 2006; K McNutt & Marchildon, 2009) argues that link structure can be taken as a proxy for real world structure and that when reliable data regarding page serves is unavailable, link structure can help establish a measure of organizational importance. Rogers (2010) echoes this argument for using the online to study the offline. He suggests that a digitally grounded approach can use online data to make claims about offline phenomena.

Certainly in the case of .gov linking, the links tell us *something* about organizational relations and importance. In many cases links will demonstrate the explicit hierarchical ordering of

governmental responsibilities. In other cases, they may shed light on how seemingly unrelated organizations are indeed “close” in an organizational structural context.

Link analysis has rarely been used to explore government structure, and when it has the scope of study has not included large portions of e-government domains. For instance, Petricek *et al* (2006) used intra-site link analysis of specific e-government websites to assess the quality of e-government design. Similarly, Li and Fu (2009) used link analysis of provincial government websites in China, correlating linking practices with agency efficiency. This study will move away from the micro-level analyses of government website linkage patterns to provide a more macro-level structural analysis of the .gov domain.

That structure may take a number of forms. It could mirror the hierarchical structure we have traditionally used to understand government. Indeed, Ravasz and Barabasi (Ravasz & Barabasi, 2003) demonstrate that the world wide web is itself hierarchically clustered, with high-degree nodes clustering together amongst a loosely linked mass of hierarchically-lower nodes. It is reasonable to suspect that the .gov domain displays this sort of hierarchical clustering. But without empirical study it is impossible to know whether or not the clustered modules conform to traditional offline institutional boundaries.

Our traditional ways of understanding federal government structure focus on the hierarchical organization dictated by the constitution and the legislation that empowers government entities. The organizational chart from the *U.S. Government Manual (The United States Government Manual, 2009, Appendix A)* reflects this traditional way of understanding government structure. In this type of rendering, relationships only explicitly exist in a vertical context, linking agencies and departments with their superiors and subordinates. However, it remains unclear whether

e-government structure mirrors this more traditional structural map. The logic of hyperlinks enables much flatter networks to emerge. Where jurisdictional or topic-area overlap or similarities occur, agencies can link laterally to others. But do they? This leads to this paper's first two research questions:

*RQ1a: To what extent does the web structure of the .gov domain mirror the offline hierarchy of the federal government and to what extent does it use hyperlinks to deviate from that structure?*

*RQ1b: How do different government agencies use hyperlinking to connect themselves to the federal government's web graph?*

The fact that the .gov domain reflects an organization that – in its offline form - is a complex, but hierarchically ordered, manmade network makes it especially interesting as an object of study. The web more generally has been studied as an example of manmade networks. It shares properties with networks from many domains. Like offline social relations, the world wide web is a “small world” where the number of links needed to navigate between any two pages is quite low (Adamic, 1999; Albert, Jeong, & Barabási, 1999). Both the size of websites – measured as number of pages/site – and the distribution of inlinks follow a power law. There are a few very large sites, and a few sites that receive a very large quantity of inlinks. One of the leading explanations for these phenomena, is the concept of “preferential attachment” wherein those with more inlinks come to be preferred, subsequently attracting even more inlinks (Barabási & Albert, 1999; Vazquez, 2003). In the context of hyperlinks, preferential attachment is often mediated by the way many people access online information. Using search engines – whose results tend to prefer high-degree sites – leads users to discover and subsequently link to sites with high-degree

(Chakrabarti, Frieze, & Vera, 2005). It remains unclear whether the preferential attachment phenomenon takes place within an institutionally bounded web subgraph like the .gov domain.

In a study examining subgraphs of the greater world wide web, Dill and colleagues (2002) demonstrated that the web's structure is fractal in that sub-structures mirror the structural traits of the super-structure. The Dill study included subgraphs based on content, geography and random selection, finding that all of the sub-graphs examined shared substantial structural traits with the web as a whole. These findings are interesting as they suggest the existence of unifying processes that underlie the generation of web subgraphs. However, none of the subgraphs studied used the web graphs of cohesive offline entities. The .gov domain is distinct from graphs bounded by geography, content or random selection. As opposed to the loose and informal organization inherent in something like a content-based subgraph, the .gov domain reflects a clear offline hierarchy. Empirically studying the .gov domain's structure can show us not only how the American government structures itself online but can also shed light on whether or not the processes leading to the near-universal web structural traits referred to above manifest in subgraphs that reflect an entity with clear offline structure. This leads to this paper's second research question:

*RQ2: Is the .gov domain structured as a fractal leaf of the greater WWW graph?*

In order to answer the questions posed above, this study maps and analyzes the hyperlinks of all federal .gov sites as described below.

## **Method**

To assemble the .gov hyperlink data, this study used LexiURL (Mike Thelwall, 2011) to query Yahoo's Site Explorer API. Using Yahoo's Site Explorer is a well-established method for link

analysis studies (see: Arakaki & Willett, 2009; Rathimala & Marthandan, 2010). It allows researchers to leverage a very large and detailed web graph, while avoiding many of the pitfalls of individually spidering the web space of interest. The queries used, formed to only return inlinks from the .gov domain while ignoring inlinks from the site itself, were constructed as follows:

```
linkdomain:www.whitehouse.gov –site:whitehouse.gov site:gov
```

This calls for all inlinks to whitehouse.gov and its subpages, but filters out within site links and only returns results from the .gov domain.

In order to limit results to federal sites I used a semi-manual spidering technique. The 85 government institutions and agencies listed in the *US Government Manual's* organizational chart seeded the search. The first wave of searches determined the inlink data for each of these organization's main websites. Following the first and subsequent waves, the results returned by LexiURL were parsed to identify all the new .gov sites encountered. I then manually coded each of these sites. Any site representing a federal agency, organization or initiative was coded as federal. All others, including state and municipal websites, were coded as non-federal. The federal sites were then used to generate another wave of searches. These waves continued until no new federal websites were discovered and the inlink data for each federal website had been determined.

Once the final web graph was assembled, each site was coded by its agency affiliation. The coding scheme included the executive, legislature, each executive-level department, inter-agency, and independent agency sites (see Table 1).

Table 1
Subgraph categories
Executive (EXC)
Legislative (LEG)
Inter-Agency (INT)
Independent (IND)
Housing and Urban Development (HUD)
Interior (DOI)
Justice (DOJ)
Treasury (TD)
Energy (DOE)
Commerce (DOC)
Transportation (DOT)
Defense (DOD)
Agriculture (USDA)
Homeland Security (DHS)
Education (ED)
State (DOS)
Labor (DOL)
Veteran's Affairs (VA)

With each site coded by agency, I generated graphs of each agency's sub-network and calculated intra and inter-agency linkages. Subsequently, structural and linking patterns were assessed using Krackhardt's graph theoretical dimensions of hierarchy (GTD) and with t-tests of log-normalized mean links.

**Results.**

The complete network contains 1077 nodes – each representing one unique federal government website. The graph contains 37 700 weighted edges representing a total of 492 495 links. Departments vary greatly in their numbers of both sites and links (see Table 2). Budget information for each cabinet-level department from the proposed 2012 budget

Table 2			
Subgraphs			
	% of nodes	Out Links	Budget
DHS	2.04	8185	43.2
DOC	4.18	13920	8.8
DOD	2.23	11158	553
DOE	4.46	35460	29.5
DOI	5.48	35626	12
DOJ	5.39	21314	28.2
DOL	1.39	3072	12.8
DOT	2.6	10398	13.4
ED	1.86	7314	77.4
EXC	3.99	5998	
HHS	9.38	87449	79.9
HUD	1.02	2394	42
IND	21.63	100470	
INT	24.23	63096	
LEG	1.3	36006	
DOS	1.67	21382	47
TD	4.55	5123	14
USDA	2.04	19019	23.9
VA	0.56	5111	61.85

(*Fiscal Year 2012 Budget of the U.S. Government*, 2011) shows no clear correlation between a department's budget and the number of sites or links it creates.

Research Question 1a asks to what extent the structure of the .gov domain mirrors the offline hierarchy of the federal government. To address this question we can examine hierarchy within the .gov domain. Because hyperlinks are low-cost relations that do not necessarily conform to offline hierarchy, various agency web graphs could deviate from the offline hierarchy of centralized control and one-way relations. Krackhardt (1994) has identified four structural signatures indicative of hierarchy and proposed index calculations to describe them. The four elements of Krackhardt's graph theoretic dimensions of hierarchy (GTD) are:

*Connectedness:* The degree to which a graph is maximally connected. A completely connected graph will have an index score of 1, a graph of entirely isolated nodes will have an index score of 0.

*Hierarchy:* The extent to which ties are reciprocated within the graph. An index of 1 demonstrates no reciprocation while an index of 0 demonstrates a 100% reciprocation rate.

*Efficiency:* The extent to which nodes are ordered in a tree-like structure with each node only having one source of in-links. Calculated by determining the difference between the actual number of links minus the maximum possible number of links. An index score of 1 demonstrates that a graph is joined with just enough links to join each node, while an index score of 0 demonstrates that every possible link exists.

*Least upper boundedness:* The degree to which each pair of nodes are downstream from a third node. This measurement is meant to quantify unity of command in organizations. In the context of a webgraph like the .gov domain it demonstrates to what extent pairs of sites are connected via more centralized hubs. An index score of 1 demonstrates that all pairs of actors share common

upstream hubs (as with a k-out star type formation), while an index score of 0 is indicative of an in-star type formation, where node pairs do not share upstream hubs.

Table 3				
Krackhardt's GTD				
	Connectedness	Hierarchy	Efficiency	LUB
Full Gov	1	0.1103	0.9516	0.9975
IND	1	0.2218	0.9386	0.9824
DHS	1	0	0.7476	1
DOC	0.9131	0.2564	0.842	0.993
DOE	1	0.0417	0.6716	1
DOI	1	0.0984	0.8294	0.9982
DOJ	1	0.1639	0.8786	0.9935
DOL	1	0.25	0.9011	0.989
DOS	1	0.4733	0.6691	0.9779
DOT	1	0.1402	0.8803	1
ED	1	0.4762	0.9667	0.9167
EXC	1	0.3792	0.9762	0.9857
HHS	1	0.1372	0.7854	0.9762
HUD	1	0.5714	0.9778	0.8667
INTER	0.9835	0.3851	0.9612	0.9658
LEG	1	0.1429	0.6795	1
TD	1	0.598	0.9136	0.969
USDA	1	0.4009	0.9333	0.9857
VA	1	0	0.6667	1
Mean	0.994558	0.2551	0.851068	0.978805

We see in Table 3 that, except for hierarchy, the various .gov subgraphs, and the .gov domain more generally display very high GTD index scores. This suggests that, apart from moderate levels of reciprocation, the .gov domain is structured as a hierarchical tree with centralized sites surrounded by tiers of subordinate sites.

Research question 1b asks

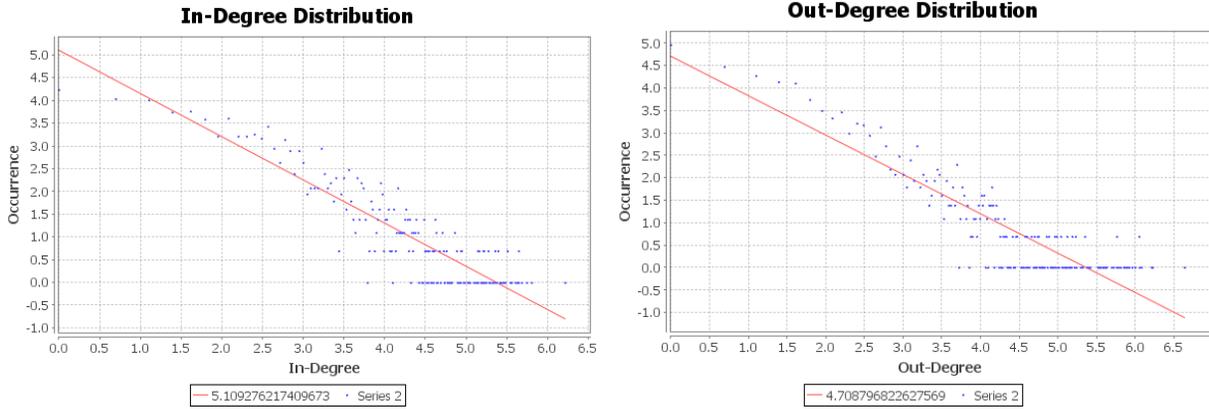
how departments connect themselves to one another and to the greater .gov web graph. By examining the self-linking and out-linking practices of the various departments we can see that some of the silos we saw in the traditional representation of government relations manifest in the .gov domain, while other agencies are more apt to form lateral links. Table 4 shows each subgraph's linking practices segmented by out-links, self-links, mean out-links (outlinks/n-1, where n=number of subgraphs), and the ratio of self to mean-out links.

Table 4				
Linking Practices				
	Out	Self	Mean Out	Self:Out
DHS	6570	1615	365	0.25
DOC	11859	2061	658.8333	0.17
DOD	2008	9150	111.5556	4.56
DOE	23025	12435	1279.167	0.54
DOI	25045	10581	1391.389	0.42
DOJ	7856	13458	436.4444	1.71
DOL	1305	1767	72.5	1.35
DOT	8417	1981	467.6111	0.24
ED	7087	227	393.7222	0.03
EXC	5514	484	306.3333	0.09
HHS	46394	41055	2577.444	0.88
HUD	2249	145	124.9444	0.06
IND	41388	59082	2299.333	1.43
INT	49102	13994	2727.889	0.28
LEG	22640	13366	1257.778	0.59
DOS	19552	1830	1086.222	0.09
TD	3933	1190	218.5	0.30
USDA	18206	813	1011.444	0.04
VA	4952	159	275.1111	0.03
Mean	16163.26	9757.526	897.9591	0.6

We can see that, for the most part, subgraphs show a preponderance of self versus out links. A t-test comparing the two means (log-normalized to address the large size discrepancy between subgraph sizes) demonstrates that this self-linking preference is statistically significant ( $t(36)=3.37, p=.0018$ ). Recall that self links in this context are defined as links to other websites

affiliated with the subgraph in question, not links within a given website. We can consider the Self:Out ratio above as a sort of silo measurement, where higher values (i.e. the Department of Defense's 4.56) demonstrate a much more silo-like web graph and lower values (i.e. Veteran's Affairs's .03) show a propensity to form lateral links across the greater .gov graph.

**Fractal graphs.** Research question 2 asks whether the .gov domain is structured as a fractal leaf of the greater WWW graph. One of the most well documented structural traits of the WWW is the power law distribution of in and out links (Adamic, 1999; Albert, et al., 1999). Examining both the in and out link distributions of federal .gov sites (Figure 1) clearly shows similar power law distributions. Furthermore, the magnitude of the power law demonstrated is a conservative estimate as links to and from non-.gov sites are not taken into account in the present data.



**Figure 1**

Like the greater WWW graph, the .gov domain is a “small world” structure with a diameter of only 5 and an average path length of 2.41. The average clustering coefficient – which we can interpret as a measure of closeness – is a relatively high 0.571.

The above suggests that RQ2 can be answered in the affirmative. The .gov graph is, for the most part, structured similarly to the greater WWW graph. The logic of hyperlinked structures does indeed extend to those formed by organizations with prior non-hyperlinked structures.

**Discussion**

We can use the linking data in Table 4 to subdivide the various subgraphs into three categories:

networked, moderately-siloed and highly-siloed.

Table 5 gives a breakdown of these categories.

Networked subgraphs have more than 10 out links to every self link, moderately siloed subgraphs have between about 5 and 1 outlinks for every self

Table 5		
Linking Categories		
<b>Networked</b>	<b>Moderate Silo</b>	<b>High Silo</b>
ED	DOC	DOL
VA	DOT	IND
USDA	DHS	DOJ
HUD	INT	DOD
EXC	TD	
DOS	DOI	
	DOE	
	LEG	
	HHS	

link, and highly siloed subgraphs are more likely to form self links than they are out links.

Separating the subgraphs as in Table 5 shows us that those agencies with more public-service orientations tend to create more lateral links across the .gov domain. Networked departments like Education, Veteran's Affairs and the USDA all have substantial public information responsibilities and all also cooperate with other agencies on many inter-agency initiatives. The Department of Health and Human Service's moderate-siloing is somewhat surprising as – of all the cabinet-level departments – it is one of the most prolific suppliers of information pertinent to individual citizens. However, much of that information takes the form of public information campaigns that inhabit their own website within the HHS subgraph. These tend to be well integrated into the HHS subgraph, but have few external links, leading it to have a relatively high siloing measure.

Those subgraphs scoring higher in their siloing measure run fewer inter-agency initiatives and generally are less public-information oriented than the others. The independent agency subgraph was one of the most siloed. While they are a heterogeneous bunch, the independent agencies tend to be quite highly specialized organizations, and thus generate few links to other government subgraphs. Similarly, the departments of Justice and Defense tend to not engage in many interagency initiatives and thus form a great many more links within their own subgraphs than laterally across the greater .gov graph.

The .gov domain's structure as a fractal leaf of the greater WWW graph strengthens prior findings of a universal logic to hyperlinked structures. To the best of my knowledge no previous study had confirmed that the web graph of an institutionally bounded offline organization would conform to the structural traits of the WWW. We now know that - at least in the instance of the U.S. federal government - that is the case.

**Limitations.** Any study of a hyperlink network is limited by the fact that there is no ideal method for extracting hyperlink data from the WWW. There are advantages and disadvantages to each method, whether one uses off the shelf software, programs his own crawler, or uses third-party data like Yahoo's web graph. This study used the Yahoo web graph, leveraging its crawling expertise, but also being subject to its black box nature.

Another weakness comes from the nature of studying a specific top-level domain. Examining only .gov sites and their linking relationships precludes us from understanding the greater context. This study provides no insight into how the .gov domain relates to the greater WWW. It also means that the network statistics described above must be understood in context. For instance, any measurement of degree or centrality only takes into account other .gov sites. The in degree of many .gov sites is much higher than what is presented here if one takes into account .com, .org and other in linking sites.

**Future Work.** While this work provides insight into how e-government is structured, it is only a beginning. Future work using both hyperlinks and other elements of relationship – like data sharing - could help provide a more nuanced view of government structure. Data from and about the individuals responsible for building and maintaining .gov websites would give a much richer understanding of background e-government processes. In continuing this work I hope to run ERGM/p\* analyses of the various subgraphs to better allow for comparison between the different agencies and levels of government.

**Conclusion.** The hyperlink graph of the federal .gov domain helps nuance our understanding of government structure. It shows us that while the web graphs of various levels of government and cabinet-level agencies are hierarchically ordered, they integrate with one another in very

different ways. Some remain relatively bounded within their traditional organizational silos while others are much more well networked with one another. In addition, the fractal nature of the .gov domain further confirms the universal structure of the world wide web.

## References

- Adamic, L. (1999). The Small World Web. In S. Abiteboul & A.-M. Vercoustre (Eds.), *Research and Advanced Technology for Digital Libraries* (Vol. 1696, pp. 852-852): Springer Berlin / Heidelberg.
- Albert, R., Jeong, H., & Barabási, A. (1999). Diameter of the World-Wide Web.
- Arakaki, M., & Willett, P. (2009). Webometric analysis of departments of librarianship and information science: a follow-up study. *Journal of information science*, 35(2), 143.
- Barabási, A., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509.
- Chakrabarti, S., Frieze, A., & Vera, J. (2005). *The influence of search engines on preferential attachment*.
- Dill, S., Kumar, R., McCurley, K., Rajagopalan, S., Sivakumar, D., & Tomkins, A. (2002). Self-similarity in the web. *ACM Transactions on Internet Technology (TOIT)*, 2(3), 205-223.
- . *Fiscal Year 2012 Budget of the U.S. Government*. (2011). Washington, D.C.: Government Printing Office Retrieved from <http://www.whitehouse.gov/sites/default/files/omb/budget/fy2012/assets/budget.pdf>.
- Halavais, A. (2008). The Hyperlink as Organizing Principle. In J. Turow & L. Tsui (Eds.), *The Hyperlinked Society: Questioning connections in the digital age*. Ann Arbor: University of Michigan Press.
- Kim, H., Park, H., & Thelwall, M. (2006). Comparing Academic Hyperlink Structures with Journal Publishing in Korea. *Science Communication*, 27(4), 540.
- Krackhardt, D. (1994). Graph Theoretical Dimensions of Informal Organizations. In M. J. Prietula (Ed.), *Computational Organization Theory*. Hillsdale NJ: Lawrence Erlbaum Associates.

- Li, C., & Fu, X. (2009). Link Analysis on Government Website Based-on Factor. In W. Liu, X. Luo, F. Wang & J. Lei (Eds.), *Web Information Systems and Mining* (Vol. 5854, pp. 108-115): Springer Berlin / Heidelberg.
- McNutt, K. (2006). Research Note: Do Virtual Policy Networks Matter? Tracing Network Structure Online. *Canadian Journal of Political Science/Revue canadienne de science politique*, 39(02), 391-405. doi: doi:10.1017/S0008423906060161
- McNutt, K., & Marchildon, G. (2009). Think Tanks and the Web: Measuring Visibility and Influence. *Canadian Public Policy*, 35(2), 219-236.
- Obama, B. (2011). *The State of the Union*. Washington, DC.
- Petricek, V., Escher, T., Cox, I. J., & Margetts, H. (2006). *The web structure of e-government - developing a methodology for quantitative evaluation*. Paper presented at the Proceedings of the 15th international conference on World Wide Web, Edinburgh, Scotland.
- . Ranking Web of World Universities. (2011) Retrieved Feb. 4, 2011, from [www.webometrics.info](http://www.webometrics.info)
- Rathimala, K., & Marthandan, G. (2010). Exploring hyperlink structure of electronic commerce websites: a Webometric study. *International Journal of Electronic Business*, 8(4), 391-404.
- Ravasz, E., & Barabasi, A. L. (2003). Hierarchical organization in complex networks. *Physical Review E*, 67(2), 26112.
- Rogers, R. (2010). Internet Research: The Question of Method - A Keynote Address from the Youtube and the 2008 Election Cycle in the United States Conference. *Journal of Information Technology & Politics*, 7, 20.

Thelwall, M. (2003). What is this link doing here? Beginning a fine-grained process of identifying reasons for academic hyperlink creation. *Information research*, 8(3), 8-3.

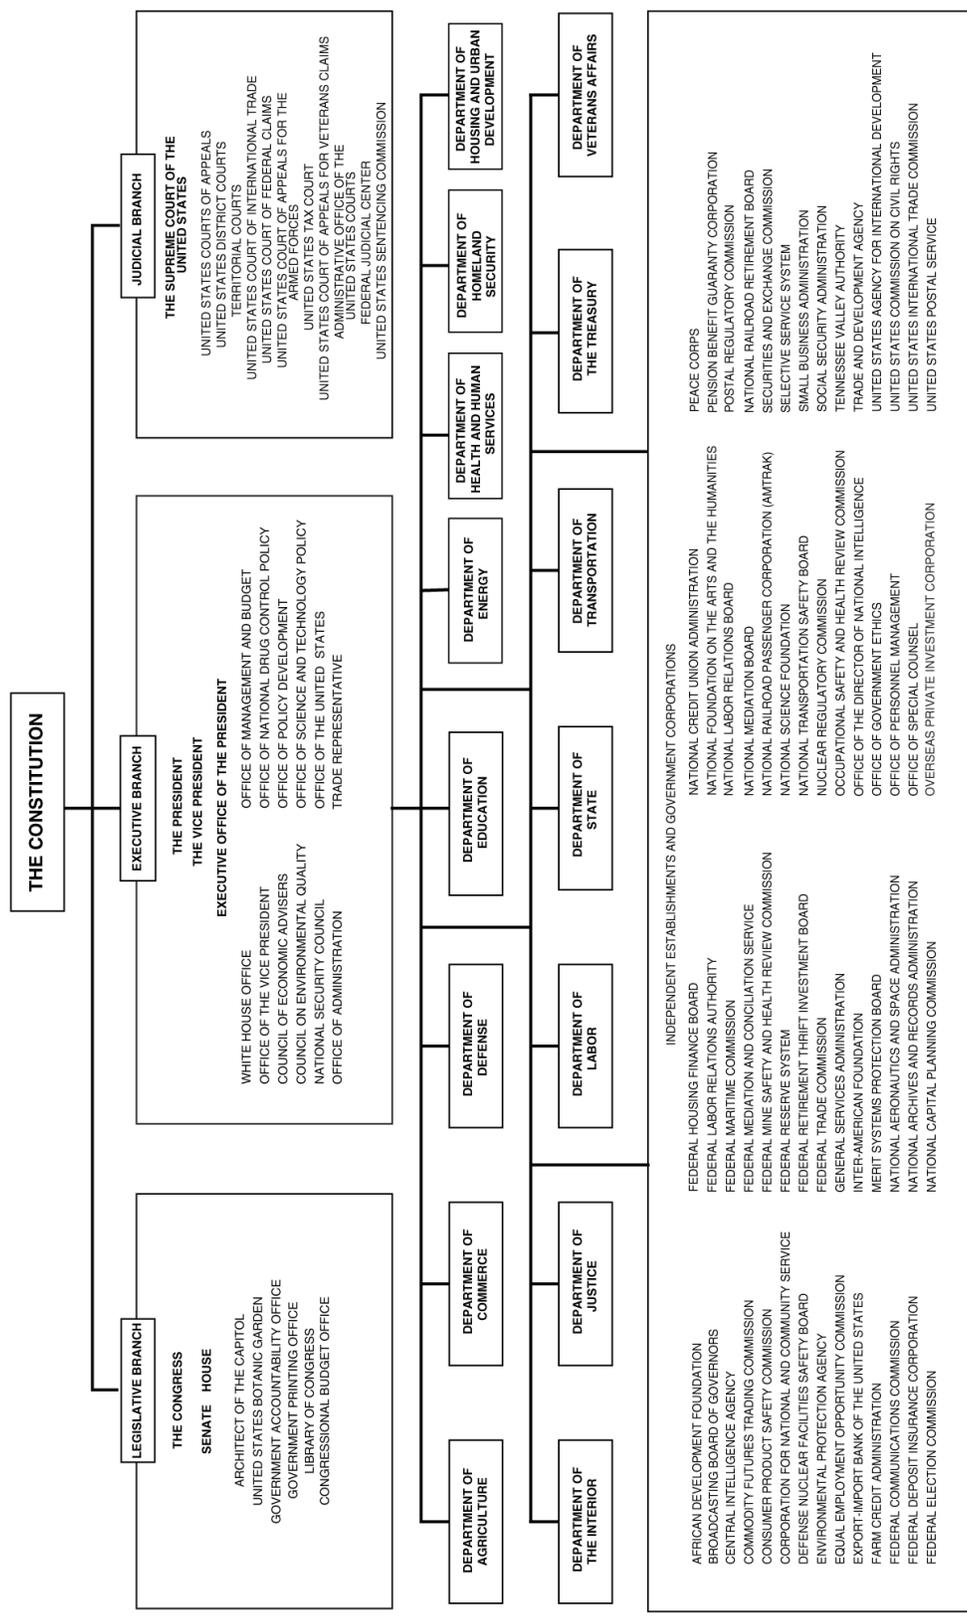
Thelwall, M. (2011). LexiURL [software], from <http://lexiurl.wlv.ac.uk/>

. *The United States Government Manual*. (2009). Washington, D.C.: Government Printing Office.

Vazquez, A. (2003). Growing network with local rules: Preferential attachment, clustering hierarchy, and degree correlations. *Physical Review E*, 67(5), 56104.

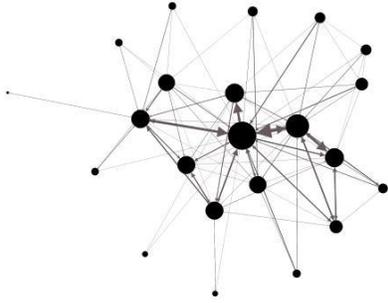
Wilkinson, D., Thelwall, M., & Li, X. (2003). Exploiting hyperlinks to study academic Web use. *Social Science Computer Review*, 21(3), 340.

# Appendix A

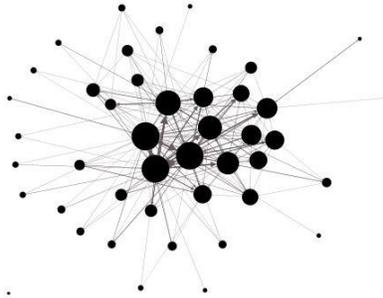


## Appendix B

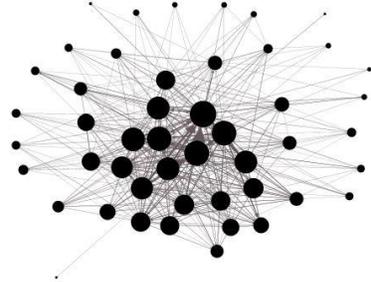
DHS



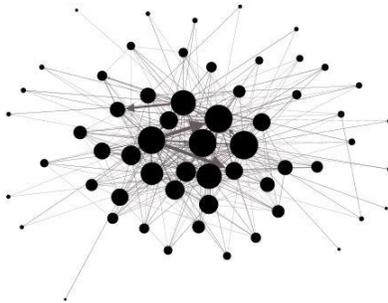
Commerce



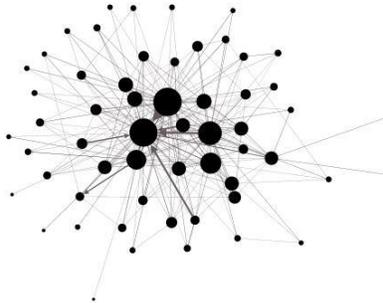
Energy



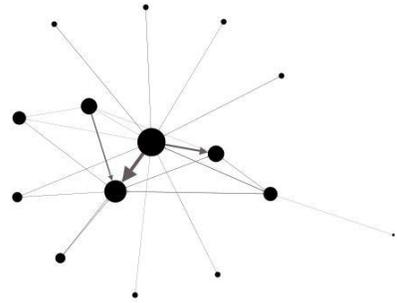
Interior



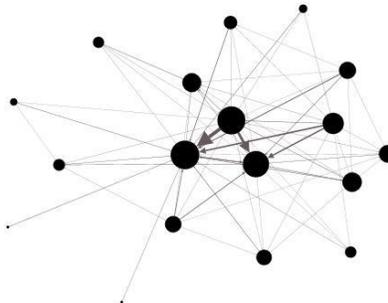
Justice



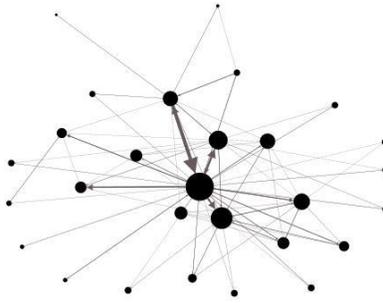
Labor



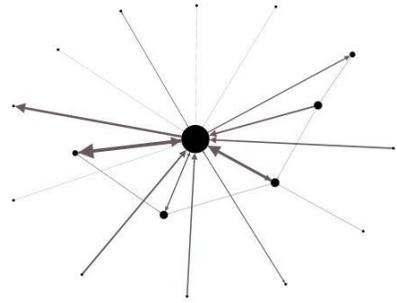
State



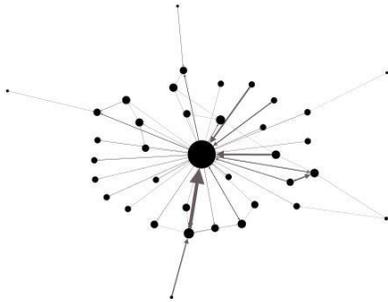
Transportation



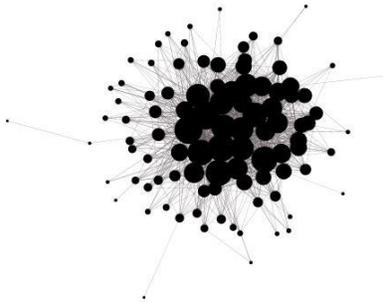
Education



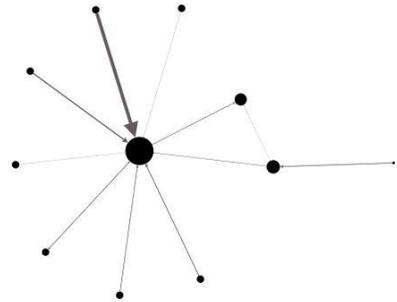
Executive



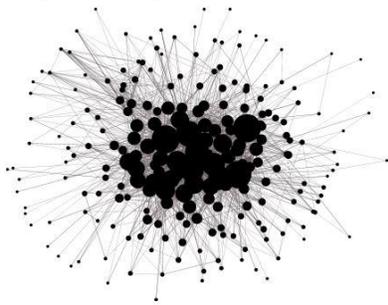
HHS



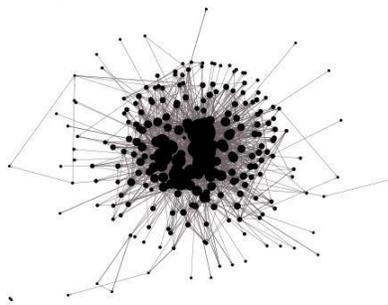
Housing & Urban Development



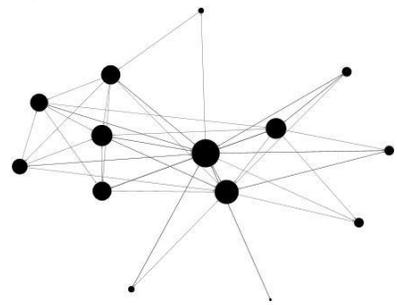
Independent Agency



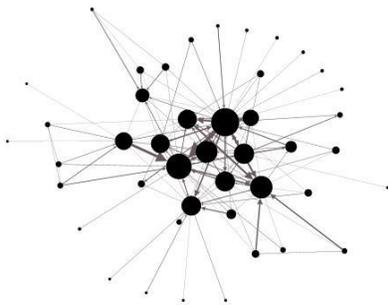
Interagency



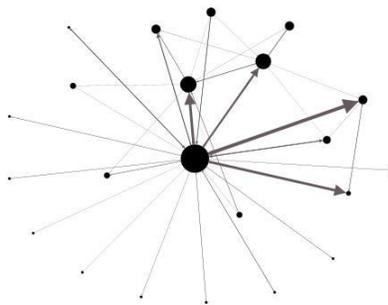
Legislative



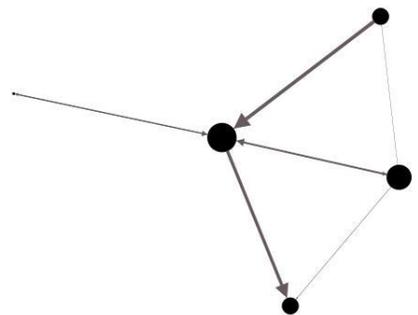
Treasury



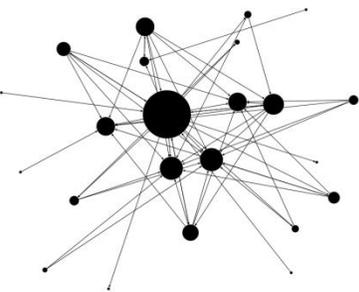
USDA



Veteran's Affairs



Defense



Created by Vizuals, <http://vizuals.com/visuals/>