

GLOBAL ENERGY SYSTEMS NAVIGATOR: AN INTERACTIVE VISUALIZATION OF ENERGY AND CARBON FLOWS

Ljuba Miljkovic

UC Berkeley, School of Information
ljuba.miljkovic@gmail.com

Jeremy Huddleston

UC Berkeley, Department of Computer Science
jeremyhu@berkeley.edu

ABSTRACT

The Global Energy Systems Navigator is an interactive map of the flow of carbon and exergy through global human energy systems. It is an educational tool for anyone interested in better understanding the complex mechanisms by which we convert natural energy resources to end-uses. Our visualization differs from previous maps of this system in that it separates the presentation of data into a qualitative “network view” and a quantitative “details view”. The network view emphasizes the relationships between carriers, transformations, and accumulations of exergy and carbon in the system. It uses a unique combination of node depth and physics simulation techniques to layout the nodes. The details view illustrates the magnitude of exergy and/or carbon for a single system element in proportion to its inputs and outputs. From here, users can explore the network quantitatively. The details view offers linear and logarithmic scaling of the stacked bar graphs that represent the input and output elements. The Global Energy Systems Navigator is intended to serve a broad range of users, from researchers to the general public. User testing of our design and future design iterations is clearly needed to evaluate its efficacy.

Author Keywords

Energy, exergy, carbon, carriers, transformations.

INTRODUCTION

Climate change is one of the greatest global challenges of our time. No single technology or approach is sufficient to reduce the emission of greenhouse gases. Instead, a portfolio of technological, policy, and public outreach and education programs are needed. One such effort was taken on by the Global Climate and Energy Project (GCEP) at Stanford University, an organization that seeks new opportunities for research in climate change mitigation. Their most prominent public outreach effort was the development of maps tracking the flow of carbon and exergy through the human energy system. The data for these maps were compiled from dozens of peer reviewed journal articles, International Energy Agency and UN databases, and government reports [1].

Exergy is a thermodynamic term describing the useful component of energy, the fraction of energy that can be used to do work. For example, gasoline can be burned to run an automobile engine. Because much of its energy can be used to do useful work, gasoline is an exergetic fuel. By contrast, the heat produced by that engine is a form of energy, but has very little potential for work. Thus, it has very little exergy. Tracking exergy instead of energy, we can

more clearly identify opportunities to increase the efficiency of the human energy system since we are ultimately interested in sources and carriers of useful energy.

The exergy and carbon maps have been presented several times at GCEP symposia and have garnered significant interest from researchers, science educators, and the general public alike. They illustrate in great detail the complexity, interconnectedness, and general inefficiency of our global energy system. Understanding this system can offer scientists potential research opportunities for improving the efficiency of energy conversion processes. In addition, the maps can serve an educational role for the general public interested in where our energy comes from and how various sources contribute to greenhouse gas emissions.

The data that inform the maps are stored in a relational database and are of three types: carriers, transformations, and accumulations. Carriers are the mediums through which exergy and/or carbon flow through the system. For example, electricity and gasoline are carriers of exergy, while only gasoline is a carrier of carbon. The flow of exergy and carbon through carriers is measured in units of Joules/second and grams/second, respectively. Each carrier entry in the database contains relevant metadata such as a detailed description of its role in the global exergy system, as well as references to the data’s published source.

Transformations are processes by which exergy and carbon are passed from one carrier to another. For example, coal-fired power plants convert exergy in mined coal into electricity. The carbon contained in coal is released into the atmosphere, a carrier of carbon. At each transformation, a loss of exergy is incurred due to inherent inefficiencies of energy conversion. The total mass of carbon is conserved throughout the system. Each transformation entry in the database contains relevant metadata such as a detailed description of its role in the global energy system, the names of its input and output carriers, and transformation efficiency based on the aggregate of exergy inputs and outputs.

Accumulations are stores of exergy and/or carbon, measured in units of Joules and grams, respectively. A common example accumulation of exergy and carbon is the subterranean natural gas we extract for electricity and heat.

Carriers, transformations, and accumulations are linked in a network of exergy and carbon flows. The system generally flows from primary resources to end uses in an alternating sequence of carriers and transformations. For example, a primary resource such as sunlight is converted to an intermediate carrier such as electricity by solar cells (a transformation), which is then used in end-use transformations that serve human needs such as lighting or cooking. However, because carbon is conserved in the system and some intermediate exergy carriers are used in transformation of primary resources, exergy and carbon do loop back through the system. Exergy and carbon from accumulations flow directly to carriers, without an intervening transformation step.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2008, April 5–10, 2008, Florence, Italy.

Copyright 2008 ACM 978-1-60558-011-1/08/04...\$5.00

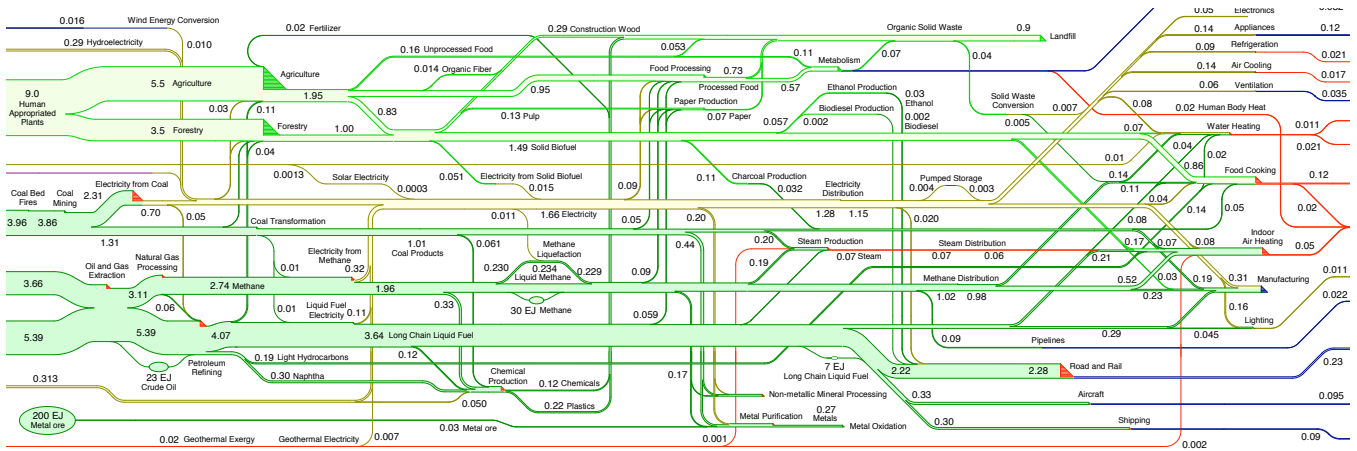


Figure 1. An section of the exergy map showing intermediate exergy-containing carriers and transformations.

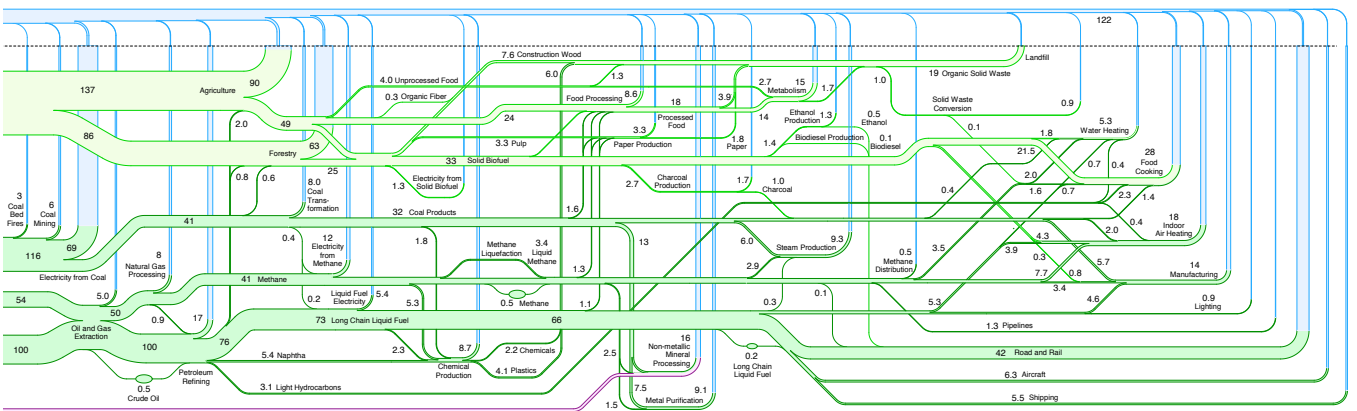


Figure 2. An section of the carbon map (1/4 of total) showing intermediate and end-use carbon-containing carriers and transformations.

Exergy and Carbon Map

GCEP originally developed a large (7' x 4') visualization of these data in the form of two large sankey diagrams: one for carbon and the other for exergy (see Figures 1 and 2). The horizontal bars represent carriers of exergy or carbon. Thicker bars carry proportionally more exergy or carbon. Transformations occur at the intersection of two or more carriers. Small triangles represent the loss of exergy associated with a given transformation.

This approach has the benefit of showing the entire system at once. It also clearly expresses the complexity and interconnectedness of the network as well as the relative magnitudes of exergy/carbon flows and transformation losses.

However, sankey diagrams used this way have several serious drawbacks. First, the maps of carbon and exergy must be drawn separately, making it difficult to co-locate carriers and transformations which contain both exergy and carbon. Second, the maps must be very large in order to be legible. The potential reach of such maps is severely limited if they can only be effectively presented as a large poster. Third, a static map offers no way to explore the data

in greater depth. In this case, users might want to view the citation for an exergy carrier of interest or find out exactly what “long-chain liquid fuels” are. The current map offers no facility for deeper exploration or metadata. Finally, the maps’ layouts are hand-drawn and largely arbitrary. Updating values and sizes of the carriers and transformations is therefore very labor intensive. Because the layouts were designed primarily to reduce overlap of the carrier bars, it is very difficult to aggregate all inputs and outputs of any given carrier. A simple question like “what are all the uses (outputs) of electricity?” is very difficult to answer with the current map.

Interactive Visualization

The interactive visualization developed here, called the Global Energy Systems Navigator, aims to address all the shortcomings of the static sankey diagrams while retaining the expressiveness of the original maps. To this end, our visualization layout is dynamically generated by a computer. It allows “drill-down” for additional information about a given carrier, transformation, or accumulation such as a citation, detailed description of its role in the energy system, and magnitude of carbon or exergy. By limiting the scope of information visible at any given time, our

visualization can be made accessible on the web for standard personal computers. This was accomplished by dividing the visualization into two separate views: a “network view” with limited quantitative data that focuses on conveying the interconnectedness of carriers and transformations; and a “details view” which visually presents quantitative data for a given node and its inputs and outputs.

APPROACH

The Global Energy Systems Navigator was built using Adobe Flex 3 [2] and the Flare [3] visualization library. These technologies were chosen because of their strong graphics and visualization capabilities. Also, they run through the ubiquitous Adobe Flash Player [4] in modern web browsers.

Network View

The network view presents the data broadly in a classic node-and-edge visualization. Carriers, transformations, and accumulations are presented as nodes of various shapes and colors with thin edges connecting them. This view emphasizes the interconnectedness of the data and allows users to explore paths through the energy and carbon systems. The nodes are grouped so that primary resources such as sunlight and coal appear to the left, intermediate carriers and transformations such as electricity and gasoline appear in the middle, and final end-uses such as lighting and vehicle propulsion appear to the right. This basic layout conveys the directional nature of that data and mirrors the flow illustrated in the original maps.

Our layout algorithm has three primary goals:

1. Present a dominantly “left to right” flow
2. Minimize the number of edge crossings
3. Limit node clustering and eliminate node overlap

Our approach uses multiple passes through the data combined with a physics simulation to generate a reasonable layout that achieves these goals.

The network graph being visualized possesses two qualities that make it difficult to layout with pre-built techniques. The first is multiple root nodes: initial inputs of exergy and carbon into the system. The earth is an open system that receives inputs of exergy from the sun and other sources. The second is loops in the system where exergy or carbon cycle from end-uses to primary resource. This occurs in the case of electricity when it is used to run a coal-fired power plant which itself generates electricity. Typical tree-drawing algorithms cannot handle multiple root nodes or loops, so we set out to design our own.

Layout Algorithm - Phase 1

In the first phase of the layout algorithm, the data network is traversed and the depth of each node is recorded. Because our graph contains loops, we do not store true “depth” but rather a minimum and maximum depth. The minimum depth of a node is defined as the minimum number of edges one must traverse from an input to reach that node. A node’s maximum depth is the maximum

number of edges one can traverse from an input to reach that node without looping.

The minimum and maximum depths for all input nodes is zero. The graph is then explored in a breadth-first manner by maintaining a queue of untraversed edges. The first time a node is visited, its minimum and maximum depths are set to one greater than the source of the edge traversed to reach it. If a node has already been visited, then these values are compared to the node’s current depth and updated.

Since our graph is partitioned into primary resource, intermediate, and final end-use groups, this depth information is stored as depth within each group. If an edge crosses from one group to the other, the minimum and maximum depths of the node in the new group is set to zero instead of being incremented.

Due to their direct connection to carriers, accumulations are considered to have the same depth values as their associated carrier.

Layout Algorithm - Phase 2

During the algorithm’s second phase, the depth data generated in the first phase is used to determine reasonable node placement along the horizontal axis; placement along the vertical axis is handled in phase three.

The second phase begins by initializing the horizontal position of a node as a weighted average of its minimum and maximum depths calculated in phase 1: 75% of the maximum depth and 25% of the minimum depth. Next, the horizontal positions of the nodes are scaled to values between 0 and 1 within each group (primary resource, intermediate, and final end-use).

After this initialization, three passes are run, repositioning the nodes based on their locations and the locations of connected nodes in the previous pass. Recalculated positions use a weighted average of the node’s previous position and the positions derived from each of that node’s edges. If the node we are positioning is the source of the edge, then the weighted contribution from that edge is “slightly” to the left of the edge’s destination. Similarly, if the node is the target of that edge, the weighted contribution is “slightly” to the right of the edge’s source. In our implementation, “slightly” is just a desired offset that we set to 0.20. For edges that cross groups, the contribution is either 0 or 1 and weighted less than other edges.

After repositioning, we rescale and possibly translate the node positions to place them between 0 and 1 to start the next pass or move on to phase three.

Layout Algorithm - Phase 3

In the third phase, a physics particle simulation is run on all of the nodes to determine their final placement.

The nodes’ positions are first initialized using the results from phase 2. The horizontal component of the primary resource group is scaled to fit between 0 and 1/3. The horizontal component of the intermediate group is scaled and translated to fit between 1/3 and 2/3. Similarly the horizontal component of the final end-use group is scaled

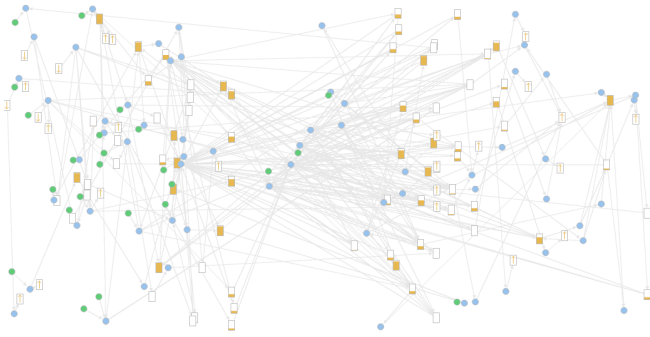


Figure 3. Network view layout based on weighted average of minimum and maximum depths in the network.

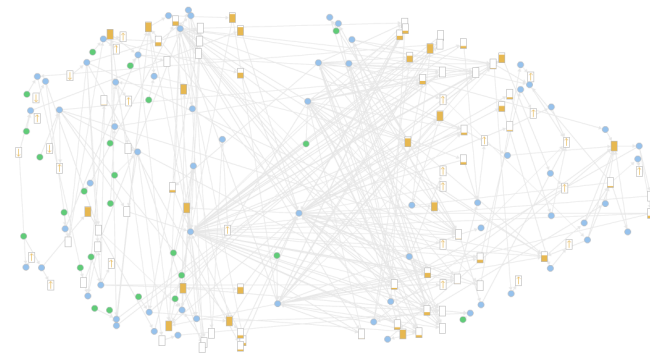


Figure 4. Network view layout with supplemental physics simulations.

and translated to fit between $2/3$ and 1 . The vertical positions of each node is initialized to a random value between 0 and 1 . The result of this initial placement can be seen in Figure 3.

During each step of the particle simulation, the total force on a particular node is the sum of three types of forces:

1. An exponential boundary force
2. An inverse-squared inter-particle repulsion force
3. A spring force along edges

The exponential boundary force is used as a “bumper” near the edge of the region that should contain the node. This force falls off very rapidly when a particle is inside the region and increases rapidly as a particle approaches the boundary. The inverse-squared inter-particle repulsion force is used to prevent the nodes from overlapping. If any two nodes are close together, this force will separate them in the subsequent steps of the simulation. The spring force along the edges “linearizes” the graph, resulting in connected nodes with similar vertical positions and subsequently shorter edges.

In order to help the simulation converge to an acceptable state, we set a relatively low efficiency for the simulation, 50% per step. This allows oscillations to quickly dissipate and results in position changes being much more in line

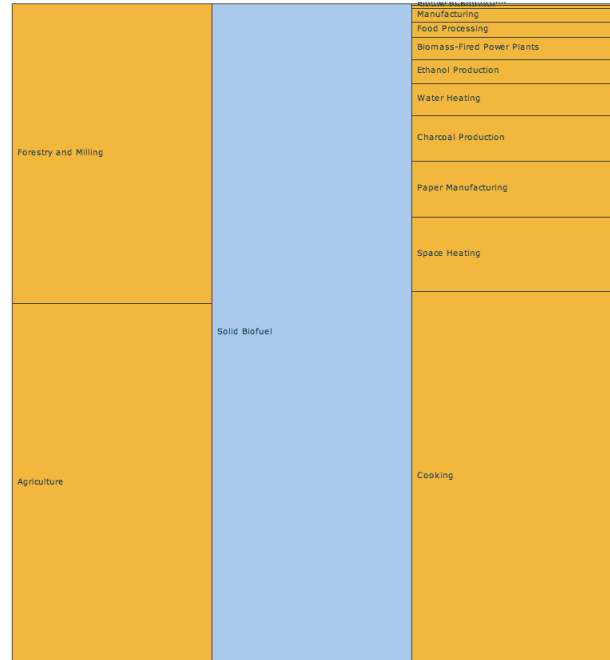


Figure 5. The details view visualizing a carrier. The quantity of exergy flowing from the input (on the left) and to the output transformations (on the right) is indicated by the heights of the stacked elements.

with the applied force of the current step. The result of 20 steps through the simulation can be seen in Figure 4.

Details View

The details view presents a quantitative view of a single carrier or transformation’s carbon or exergy magnitude with its various inputs and outputs stacked on either side of it (see Figure 5). This view is presented when a user double-clicks on a node in the network view, indicating a desire to learn more about the details of this element. From here the user can navigate the network by clicking on any visible carrier or transformation. The selected element’s inputs and outputs are aggregated dynamically and presented as before, on either side of the new central element. Changes in element positions are animated to help users follow the changes in the visualization.

Linear and Logarithmic Scaling

This work introduces a novel logarithmic approach to visualizing stacked bar graphs. Simple bar graphs with logarithmic scales are used to visualize sets of data spanning many orders of magnitude. Traditional stacked bar graphs are often used to visualize components that make up a whole. In our visualization, we combine these ideas to present a logarithmic stacked bar graph for visualizing inputs and outputs spanning many orders of magnitude.

This is accomplished by two levels of stack scaling. The first scales the height of the a stacks relative to the central stack. All stacks will have the same height except when a transformation is the central element in the details view. The output stack in this case will be lower because of the requisite loss of exergy in every transformation. In the case

of very inefficient transformations, it is useful to scale the height of the output stack logarithmically so its constituent elements can be seen.

The second level of scaling occurs within a stack, on the elements themselves. To scale these elements linearly, we first calculate the ratio of their exergy or carbon values to the stack total. This ratio is multiplied by height of the stack to determine the height of the individual components. To scale these components using a logarithmic scale, the ratio of the log of their exergy or carbon values is used instead of the ratio of the values themselves. Thus, an element that appears twice as high as another in a given stack contains 10 times the exergy or carbon.

Navigation Bar

A navigation bar gives users access to a list of all carriers, transformations, and accumulations in the system at once. It enables real-time text searching by name and filtering of the elements by type (carrier, transformation, and accumulations) or position in the map (primary resource, intermediate, and final end-use). The navigation bar is visible in both network and details views.

RESULTS AND DISCUSSION

By visualizing the exergy and carbon data in two separate network and details views, we were able to retain the benefits of the large sankey diagrams while addressing our design goals of automatically laying out the elements, exploring the data at various levels of detail, and presenting the visualization on the web in a size viewable on standard displays.

Network View

We used a combination of weighted averages of minimum and maximum node depth and physics simulations to achieve an effective layout of the nodes in the network view. A node depth approach alone produces the left-to-right flow of nodes that we desire (see Figure 3). However, this approach only affects node positions along the horizontal axis. As a result, nodes may overlap or cluster. When the physics simulation passes are applied, node clustering is reduced as is edge overlap (see Figure 4).

Interaction

The network view presents the entire energy and carbon system through interconnected carriers, transformations, and accumulations. The appearance of and interaction with these nodes helps users develop an understanding of this complex system.

Carrier, transformation, and accumulation nodes are easily distinguishable. Carriers and accumulations are represented as blue and green dots, respectively. Transformations are encoded as small bar graphs whose height reflects the efficiency of that transformation. At a glance, users can spot transformations of high or low efficiency.

When a user hovers the cursor over any node, its name and the names of its inputs and outputs are shown. Also, the connecting edges are highlighted in red (for inputs) and green (for outputs). Arrowheads at the end of the edges also



Figure 6. Navigation bar with node icons and input/output tags.

encode connection directionality. By merely hovering the mouse over various nodes, users can explore the network, and easily see which nodes connect to each other.

After highlighting a node and its connections, clicking it locks that state in place. The user is now free to navigate a path through the system by clicking on one of the connected nodes. If the user follows the green edges, she will be traveling “forward” through the system in the direction of energy end-uses. Following the red edges will move her “backwards” towards primary energy resources. Clicking on the white space between nodes will clear this locked state. Double-clicking on any node, will take the user to the details view, where she can learn more about the quantities of exergy and carbon contained there.

Details View

The details view provides a highly quantitative visualization of a selected element. It accomplishes this by only displaying a single carrier or transformation with its inputs and outputs.

Interaction

When a carrier is the central element in the details view, input transformations illustrate how that carrier is produced. For example, electricity is produced by coal-fired power plants, natural gas-fired power plants, nuclear power plants, etc. Output transformations indicate how that carrier is used in the energy system. In the case of electricity, there are many output transformations such as heating, lighting, manufacturing, etc. The height of each input and output bar is proportional to the total exergy flux through the central carrier. By selecting a logarithmic scaling of the bars, smaller contributors to the whole can be more easily identified. Toggling the exergy/carbon button redraws the stack for carbon flux instead of exergy.

The interaction for when a transformation is the central element in the details view is very similar. Input carrier

bars to the left illustrate which carriers the central transformation acts on. For example, cooking is performed with solid biofuels, propane, natural gas, etc. The outputs of this transformation are carriers containing carbon and any remaining exergy. In the case of cooking, carbon is released to the atmosphere, and exergy is left over in warm air and radiant heat. Because all transformations incur a loss of exergy, its output stack of elements will be lower than its input stack.

Using a non-linear scaling of for stacked bar graph elements is a powerful method of visualizing values that span many orders of magnitude and sum up to a whole. Such a scaling technique is not limited to logarithmic scales. It could be possible to use an exponential scale within one stack and a logarithmic scale within another if needed.

However, logarithmically scaled bar graphs remove a user's ability to additively compare stack components. The sum of two elements in a stack cannot be compared with a third. Similarly, in logarithmic mode, the scaling of the input stack will likely be different than the scaling of the output stack, so it would not be trivial to compare components across stacks.

Navigation Bar

In either the network or details view, the navigation bar provides access to all the elements in the energy and carbon system (see Figure 6). Elements can be searched using real-time text queries and filtered by element type and position in the network using the provided toggle buttons. Each element list item contains an icon to the left of its name, which is identical to its node in the network view. Carriers and accumulations are blue and green dots, respectively. Transformations are orange efficiency bars. To the right of the name is a tag indicating the number of that element's inputs and outputs.

Interaction

By design, the interaction of the navigation bar is very similar to that of the network view. This provides a consistent user experience across two major visual components of the application. Hovering over a carrier or transformation listed in the data grid highlights that node and its connections in the network view. Clicking on a list item locks the highlighted state, while double-clicking takes the user to the details view. Also, hovering over the element icon displays a tooltip with quantitative exergy and carbon data for that element.

FUTURE WORK

There are many areas in which the Global Energy Systems Navigator can be improved. Increasing data density while retaining ease-of-use is a high priority for all future work.

Improving the network layout algorithm

Currently, the layout algorithm is not guaranteed to converge. The algorithm produces acceptable results after a few iterations, but there is no guarantee that the layout is converging to a stable state. Eventually, we would like to use the layout to animate in real-time when nodes are

filtered with the carbon/exergy toggle button or to move less relevant nodes away when a node and its connections are highlighted. Further work is needed to build a convergent layout algorithm.

Additionally, each step through our physics simulation is $O(n^2)$. With some effort, it could be replaced with an $O(n \log(n))$ approach or even possibly an $O(n)$ approach, but optimization of the physics will not be a priority until a convergent layout algorithm is achieved.

Multiple data encoding in the network view

Though the network view is designed to present the data broadly, we believe more data could be encoded in the shapes, sizes, and even relative positions of the nodes without causing distraction.

Edge clustering in network view

The network layout could be improved by applying algorithms to cluster edges along common paths through the system [5]. By forming edge bundles, we could greatly reduce the number of edge crossings in the network view.

Path history

These exergy and carbon data could be further explored by recording and comparing the efficiency properties of various paths through the network. Users could, for example, compare the carbon footprints of cooking with solid biofuels vs. electricity produced from solar power. To do this, however, we would need more detailed information about specific carrier-transformation combinations that currently do not exist in the database.

User testing

As a tool intended for broad audiences, usability studies would help determine whether our design addresses the needs of our users.

CONCLUSION

In creating the Global Energy Systems Navigator, our goal was to improve the sankey diagram maps of carbon and exergy flows developed at GCEP by overcoming their size and static layout limitations while increasing the interactivity and readability on standard computer displays. We constructed a web application featuring a layered visualization of the exergy and carbon networks which separates the presentation of the data into a qualitative "network view" and a quantitative "details view". The network view emphasizes the relationships between carriers, transformations, and accumulations of exergy and carbon in the system while the details view illustrates the magnitude of exergy and/or carbon for a single system element in proportion to its inputs and outputs. These views were strengthened by a unique combination of node layout techniques in the network view and a novel logarithmic stacked bar graphing technique in the details view. Though there remains much work before this application is ready to be presented to the public, it demonstrates the great potential of a small, layered visualization for presenting complex global energy and carbon systems.

ACKNOWLEDGMENTS

We would like to thank the Global Climate and Energy Project at Stanford University for collecting the data that drive this visualization and for providing user feedback on iterative designs.

REFERENCES

1. Herman, W. Quantifying global energy resources. *Energy* 31, 12 (2006), 1685-1702.
2. Adobe Flex 3. www.adobe.com/products/flex
3. Flare. flare.prefuse.org
4. Adobe Flash Player. www.adobe.com/products/fashplayer
5. Cui, W. et al. Geometry-Based Edge Clustering for Graph Visualization. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1277-1284.