Network Analysis of Global Trade

Neel Guha, Andrew Han, and Alex Lin

I. Introduction

International trade consists of complex relationships between different countries, where changes in a single relationship could have repercussions on other countries and their relationships. Thus, modeling global trade networks is especially difficult as economic phenomena are not easily categorized as independent or dependent variables. Rather, the global economy is a dense network of interrelated and interdependent agents. As such, its underlying structure can be aptly conceptualized graphically as links between various economic players.

There is also a temporal component to the global economy. New economic players and resources emerge and disrupt existing trade networks in complicated ways. For instance, one such phenomena is the relationship between natural gas and crude oil. In the early 2010s, when the United States discovered hydraulic fracturing (fracking), an efficient way of extracting natural gas, the price of crude oil and the relative economic and political power of oil blocs like OPEC plummeted, leaving lasting impacts on world inflation, the energy market, and geopolitics.

Network analysis offers an attractive semantics with which we can begin to quantitatively model such geopolitical and economic phenomena. Network based techniques provide us with new tools that shed light on questions such as the importance of certain countries (centrality) or the relationships between different countries (community detection).

In this paper, we investigate how the global trade network evolved between 1996 and 2005. This time period captures major economic developments ranging from the entrance of China into the WTO to the ”opening up” of the Indian economy (among others). These events - and the effects of globalization - have been qualitatively studied and are believed to have a significant impact on the global trade network. We hope to be able to quantitatively communicate these effects and developments. We specifically focus on how new trading relationships formed during this time and the communities that emerged.

II. Literature Review

Zhu et al. offer guidelines to explore the temporal element to economic network [2]. They explore communities that existed before and after China's entrance to the

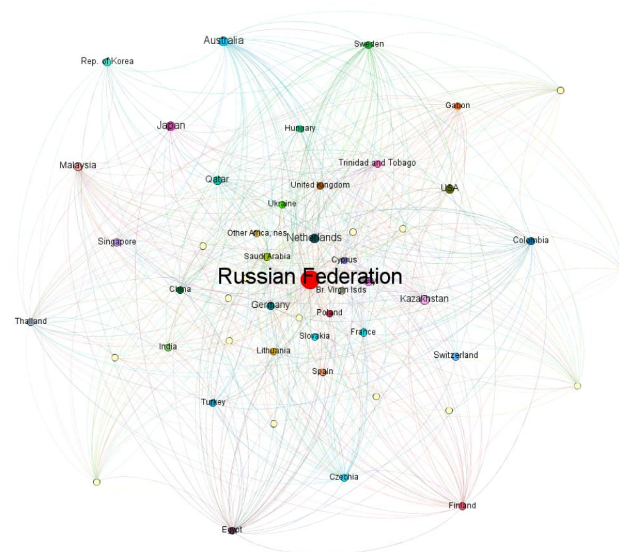

Fig. 1. Global fuels trade, 2005

World Trade Organization in 2001. Communities are a particularly useful feature of network structure; they are based on the idea that a random graph is not expected to have a community structure and are a measure of groups of nodes that have many (or strong) edges between them and weak edges between groups. A notable insight from the work found that the addition of such a large player as a new node was in part the reason why a prominent community (the Asia-Oceania community) collapsed as China developed trading partnerships with Germany and the U.S. The authors utilized a modularity optimization method introduced by Newman and Girvan for community detection. Noting that such a metric did not account for the influence of individual nodes in the network, the authors also utilized metrics developed by De Leo V et al. for community core detection and Cerina et al. for the global strength of a node in a community.

[3] expands on previous research by applying the cascade shock model to real world trade graphs. One of their underlying assumptions was that declining exports result in reduced growth, which itself results in a reduction of imports. This further propagates through the

economic network and trade flows, causing a reduction in exports and imports. Specifically, their approach updates responses to shocks by aggregating received impulses, emitting a proportional impulse based on a linear response function. Nevertheless, the authors provide a great deal of creative groundwork. The authors focus primarily on generating a cascading-shock update model to measure the influence of perturbations from a single node to other nodes.

[4] builds a classifier to detect when crises occur (specifically natural disasters) by analyzing the resulting effects on the structure of the global trade network. They identify the effect crises have on various network measures (like random walk betweenness centrality) and various subgraphs that commonly occur in the wake of a crisis. Drawing on the results of the work of Milo et al., the paper also demonstrates the applicability and value of triad significance profiles in analyzing the global trade network. Milo et al. present a uniform way to study the structural similarities between graphs even when they differ significantly in connectivity and structure. They introduce a way to calculate the Significance Profile (SP) of a network by comparing the frequencies of every subgraph in the network of interest to equivalent subgraphs in a collection of randomly generated networks (of similar degree). This allows us to determine the relative importance of certain subgraphs in the network by calculating when a specific subgraphs frequency exceeds random expectation. The original paper shows how triad significance profiles can be combined with other network measures (random walk betweenness centrality, link density, etc) to predict crises.

III. Data Collection

A. Data Source

The United Nations Comtrade Database [1], which offers global trade data on an annual basis as far back as 1962. Using the Comtrade API, we obtained bilateral trade data for a number of commodities. In accordance with the standard HS Classification for commodities [5], we chose to analyze three broad categories of industrial commodities: $fuels$, $plastics$, and $rubbers$. Thus far, we have retrieved trade data for those commodities for the following years:

- Fuels, from 1996 to 2015
- Plastics, from 1996 to 2005
- Rubbers, from 1996 to 2005

B. Programmatic Collection

The Comtrade Database API places a number of restrictions on how much data we can retrieve at once. Since they limit us to 100 API calls an hour and only allow retrieval of up to five years of trade data at a time, we wrote a script of the following form to call their API in regular intervals.

```
commodities = [fuels, plastic, rubber]
countries = [all countries]
for commodity in commodities:
    for country in countries:
        for five_year_range in years:
            make_api_call()
sleep(1 hour)
```

We then dumped the results to a number of `JSON` files (such as `fuels_1996-2005.json`), which we could later re-load into, in this case, ten NetworkX and/or SNAP graphs (one for each year).

IV. Approach

We now discuss how we represent the trade network, and the array of metrics we use to explore its evolution over time.

A. Trade Network Representation

We can consider the global trade network as a directed graph where nodes correspond to countries and edges correspond to an import/export relationship for a certain good. An edge from A to B represents country A exporting a certain good to country B (or B importing from A). The weight of this edge is proportional to the magnitude of the trade. For the sake of simplicity we assumed an undirected graph in our analysis. In this graph, the magnitude of each edge is the sum of the trade in either direction. Note that this is not the "net trade" but the raw sum. If country A exported 50 units to country B and country B exported 10 units to country A, we'd calculate a weight of $50 + 10 = 60$. We use Networkx [6] and SNAP [7] to model our network.

B. Node and Edge Distributions

We first explore the manner in which new edges are created between nodes in the network. We're interested in exploring any patterns that dictate how these edges are formed. Prior work has suggested that the global trade network tends to follow a preferential attachment model – we hope to explore its applicability in this case. We explore this by primarily examining the degree distributions of newly created edges. This can offer us insight into the way the trade network is expanding. For example, if we observe the creation of edges between high degree nodes, then we'd be able to draw some conclusions on whether global trade is becoming more concentrated/centralized.

C. Centrality

Measures of graph centrality (betweenness, degree, etc.) offer us the ability to capture the importance of certain nodes to the overall network. We leverage these metrics to identify the most important nodes in our network. By looking at how the importance of certain nodes change over time, we can evaluate how the graph evolves.

Additionally, the Girvan-Newman technique for community detection (which we discuss more in the following section) uses such metrics as an input, iteratively removing edges with the highest centrality. The vanilla version of the algorithm uses betweenness centrality to determine a "most valuable edge", but we believe that it may not capture qualitative characteristics of what is generally considered to be "important" nodes and edges in economic networks.

D. Community Detection

Prior work has suggested that the Girvan-Newman algorithm has had success in community detection. The algorithm proceeds as follows:

1) Calculate the betweenness centrality for all pairs of nodes and assign betweenness scores to all edges
2) Remove the edge with the highest betweenness score
3) Recalculate the betweenness centrality between all nodes and the betweenness scores for all edges
4) Repeat steps 2 and 3 until there are no more edges left in the graph

The algorithm's intuition is that edges between communities have high betweenness scores. By eliminating them, we're can partition the graph and identify subcommunities.

We experimented with various inputs to Girvan-Newman. We used the following methods to assign scores to edges:

1) Largest Trade: We score each edge by the magnitude of the trade between the two countries. Thus, at each iteration we remove the edge with the highest weight (trade magnitude). Intuitively, though this metric does not capture global centrality in the way that betweenness centrality does, it does serve to rank edges by relative value, at least in a local sense.

2) Degree Weighted Trade: We score each edge by it's trade magnitude weighted against the degrees of the edge's source and destination nodes. Formally, for an edge E,

$$E_{\text{score}} = \log(E_{\text{weight}}) * deg(E_{\text{source}}) * deg(E_{\text{destination}})$$

We add this modification to the above method in order to also take into account the relative importance of its neighbors. Intuitively, a high trade weight between two "important" nodes should be more important.

3) Weighted Edge Betweenness Centrality: We calculate the edge betweenness centrality and weight it by the proportion of the trade value. This allows us to discriminate more from less important edges (as defined by how the economic value of the trade). Formally for an edge E with a betweenness centrality score of w_{bc},

$$E_{\text{score}} = E_{\text{weight}} * w_{bc}$$

4) Weighted Edge Betweenness Centrality by Degree: We score edges using edge betweenness centrality as above, but also weight them by the node degrees. Formally for an edge E,

$$E_{\text{score}} = E_{\text{weight}} * w_{bc} * deg(E_{\text{source}}) * deg(E_{\text{destination}}$$

where tv is the weight or magnitude of the edge and b_e is the betweenness centrality of the edge.

E. Community Volatility

We examine the consistency of communities in our graph over the 1996-2005 time period. While we should expect the communities remain relatively similar from year to year, changes in global trade patterns could disrupt certain communities and give rise to new ones. By examining the persistence (or lack thereof) of collections of specific communities, we can infer a lot about global economic trends.

Specifically, we formalize this as a country specific *purity score*. Given a country n_i and two times t_j and t_k we can calculate the purity score $p_i(j,k)$ as

$$p_i(j,k) = \frac{C(n_i,j) \cap C(n_i,k)}{C(n_i,j) \cup C(n_i,k)}$$

where $C(n_i,j)$ is the set of countries in the same community as n_i at time j. For a given community, we can then also calculate a community purity score as the average purity of all countries in that community.

One of the challenges of evaluating our methods is that there's no ground truth for communities in the global trade network. We therefore attempt to evaluate our methods based off of the following dimensions

- **Consistency**: Though we expect communities to change over longer periods of time, from year to year they should remain largely similar. We can use the purity scores mentioned above to quantify the consistency of communities. Stronger community detection measures should be resistant to yearly trade fluctuations and consistently produce the same communities.
- **Qualitative**: Additionally, we can evaluate our community detection methods from a qualitative perspective. We should intuitively expect countries

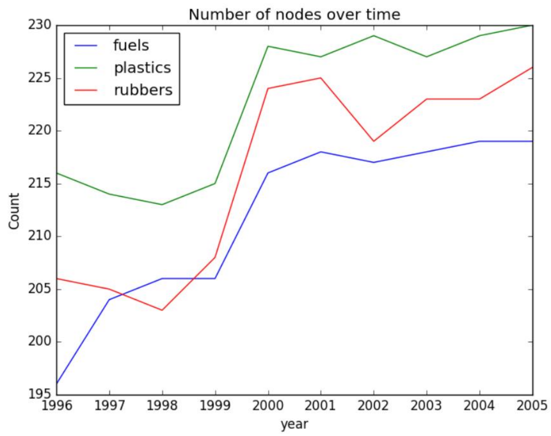

Fig. 2. Number of nodes in each of the trade graphs over time

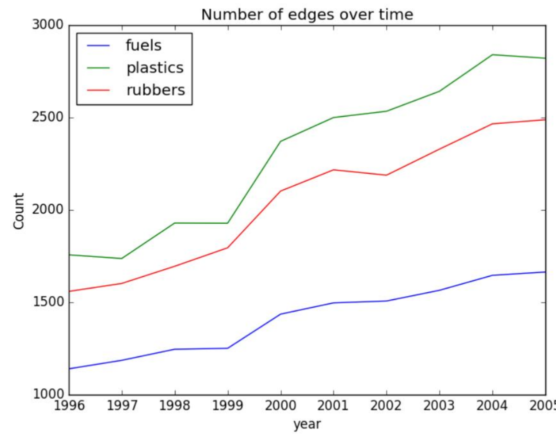

Fig. 3. Number of edges in each of the trade graphs over time

within the same region or with similar trading patterns to be grouped together.

We calculate several different versions of the purity score:

1) Year-over-year Purity (YOYP): For every pair of successive years (1996 and 1997, 1997 and 1998, etc.) we calculate purity score for each country. We then report the mean purity score across all countries over all pairs of years as well as the mean purity score for each country across all years. This allows us to identify which countries are the most "stable" in their communities.

2) 1996-2005 Purity Score: We calculate the mean purity score across all countries for the difference in communities between 1996 and 2005. This quantifies the "volatility" of the communities and allows us to identify which communities have changed the most during this timeframe.

V. RESULTS

A. Global Graph Changes

Before delving into the regional effects of node/edge additions, we analyzed how the graph evolved on a global scale over time. Based on figures 2 and 3, we see that between 1996 and 2005 the number of edges and nodes increased somewhat steadily – a likely result of more countries joining the global trade economy. This would confirm the general hypothesis in mainstream economic and political discourse that world trade has become more globalized, with more countries interacting and trading with one another than ever before. This is also supports by the densification power law, which dictates that the number of edges in a graph grows faster than the number of nodes.

We should also observe the distribution of these additional nodes/edges. Figures 6 and 7 are the histogram for the degree distribution and the degree rank for the fuel trade network respectively. We observe that uniformly, the degrees of nodes in the graph appear to be increasing.

B. New Nodes and Edges

One of the subtopics of interest to us in this project was the process through which new nodes/edges "joined" the global trade network. Figure 4 demonstrates that new edges in the fuel trade network are more likely to occur between existing nodes (as opposed to new nodes).

The global trade network has often been described as following a preferential attachment model, where the degree distribution follows the power law. We can examine the applicability of this claim by analyzing the degree distribution of the source and destination nodes of all edges added between 1997 and 2005. For a given edge e_i between nodes n_0 and n_1, we define the

$$\text{source} = \min(\text{degree}(n_0), \text{degree}(n_1))$$

$$\text{destination} = \max(\text{degree}(n_0), \text{degree}(n_1))$$

The source node corresponds to the less connected node in the edge and the destination node corresponds to the more connected node. In a preferential attachment model, we should expect that edges are more likely to be formed between low and high degree nodes. Figure 5 is the degree distribution for all source and destination nodes across all edges between 1997 and 2006. The distribution of source degrees is in line with what we expect and there is a strong skew towards low degree nodes. However, the destination distribution appears to be somewhat uniform. Table I contains the most common destination countries. Nearly all the countries are strong regional traders. This appears to suggest that the trade network follows a "partial" preferential attachment

TABLE I
MOST FREQUENT DESTINATION COUNTRIES BY EDGE COUNT

Top Destination Countries	Number of Edges
Thailand	88
Colombia	85
Czech Republic	75
Egypt	75
Australia	74
Malaysia	72
Russian Federation	71
United Rep of Tanzania	71
Sweden	70
Trinidad and Tobago	69

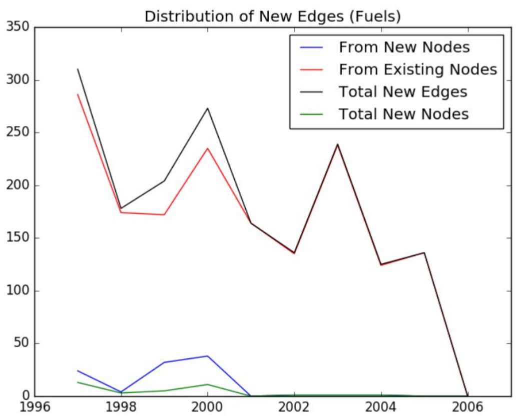

Fig. 4. The distribution of new edges by year for the fuel network

model, in which the dominating nodes are regional trade powerhouses. In this model, the new edges appear likely to form between a low degree node and specific high degree nodes (based on the location of the low degree node). These results appear to suggest that the global trade network is growing denser and more connected. Indeed, 93.65% of the new edges between 1996-2005 closed a triad.

C. Centrality Results

Our modified version of betweenness centrality (i.e., edge betweenness centrality that also takes into account the degrees of the nodes attached to each edge) corroborates both results from [2] and the general economic narrative of globalization. For instance, Table II displays the highest ranked edges in 1997 for the fuels trade, where we see that trade was largely scattered to "Areas, nes [not elsewhere specified]". By 2005, we saw in Table III, the top few edges by our betweenness centrality metric included various countries in Oceania and Southeast Asia, including Australia, Thailand, Malaysia, Cambodia, and China.

Additionally, we saw that top ranked countries by our metric had significantly higher betweenness values in

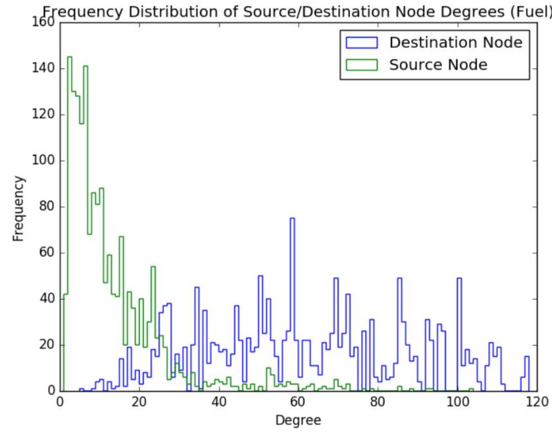

Fig. 5. The distribution of degrees for source/destination nodes for all added edges between 1997 and 2006

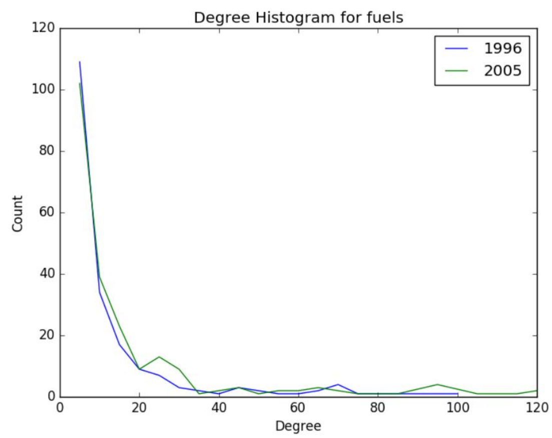

Fig. 6. Histogram of degree count for fuels

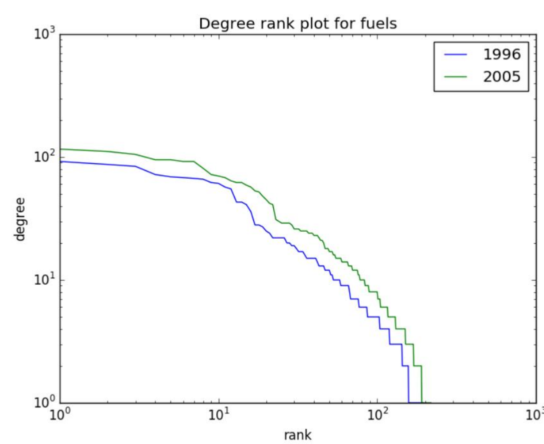

Fig. 7. Degree rank for fuels over time

TABLE II
WEIGHTED BETWEENNESS CENTRALITY (1997)

Source	Destination	Weighted Btw. Cent.
Ethiopia	Finland	0.1445
Ethiopia	Areas, nes	0.1359
Areas, nes	Mongolia	0.1349
Gabon	Areas, nes	0.1230
Areas, nes	Kazakhstan	0.092
Saint Vincent	Areas, nes	0.090
Finland	Colombia	0.086
Switzerland	Trinidad and Tobago	0.081
Trinidad and Tobago	North Korea	0.077

TABLE III
WEIGHTED BETWEENNESS CENTRALITY (2005)

Source	Destination	Weighted Btw. Cent.
Gabon	Australia	0.251
Malaysia	Gabon	0.231
Australia	Colombia	0.227
Australia	Mongolia	0.158
Finland	Thailand	0.155
Colombia	Saint Kitts and Nevis	0.150
Thailand	Cambodia	0.139
Mongolia	Finland	0.135
Cambodia	Germany	0.125
China	Guatemala	0.118

2005 than they did in 1997, which seems to support the intuition that the world is becoming more globalized.

However, it's worth questioning what measures of centrality actually convey in this context. Intuitively, a node in our graph with many edges to a range of countries will score highly on any centrality measure used. However, it's unclear whether such a node is always "important" to the graph. In our graph, a central node with many edges corresponds to a country with many export/import relationships. If a country is importing a good from numerous other countries, it's probable that no single import is incredibly significant. In this context, a high centrality score is a measure of "import diversification" and not importance. A country with many incoming edges is importing a good from a diverse set of different countries. Intuitively, such a country may be more resistant to any shocks in the system.

D. Community Detection Evaluation

We now discuss the results of the various community detection algorithms we implemented.

1) Weighted Betweenness Centrality: Weighting edge betweenness centrality with the weight (trade value) of the edge yielded more common sense results. We first ran Girvan-Newman (with vanilla betweenness centrality) to generate 10 communities on each year of the fuel network data from 1996 to 2005. Table V contains the results of the average year-over-year purity scores and the purity scores between 1996 and 2005. One of the interesting findings of this particular algorithm was the most "stable" countries (those with the highest year-over-year purity scores) were islands in the Pacific/Oceania (see Table IV). This result seems to suggest that trade around these island nations hasn't fundamentally changed between 1996 and 2005. This seems to suggest that perhaps globalization hasn't had as much of an effect as predicted in this specific region.

2) Largest Trade and Degree-Weighted Trade: We also tried varying our centrality input to Girvan-Newman, which produced incoherent results. Modifying Girvan-Newman to iteratively remove the largest edges by trade (or trade weighted by the degrees of the respective nodes) resulted in a number of one-country communities, which for our purposes did not yield any qualitative insights into our research questions. This also explains why their scores in Table V are incredibly high. At every iteration, the majority of countries were being assigned to the same community. Thus, this community would be appear to be consistent, even though it conveys virtually no information.

3) Weighted Edge Betweenness Centrality by Degree: Combining the above few approaches, we wanted to not only capture the centrality of an edge, but also to encapsulate some notion of the importance of the edge's nodes. Thus, we ran Girvan-Newman with a degree-weighted version of weighted edge betweenness centrality. Our results here corroborated many common sense intuitions of trade communities. The algorithm detected the following such communities in 2005.

- Middle-East and Eastern European community (including Saudi Arabia, Israel, Turkey, Estonia, Belgium, etc).
- Middle-East and Asian community (including China, Pakistan, Vietnam, Ukraine, USA, Uzbekistan).
- African community (Egypt, Congo, Botswana, Morocco, etc).
- Island communities (French Polynesia, Samoa, Australia, American Samoa)
- Iran, Iraq, North Korea, UAE, Syria, Yemen, Czech
- American (South and North) community (Costa Rica, Colombia, Barbados, Canada, Mexico, Venezuela)
- Asian and Oceania community (China, Singapore, Indonesia, Japan, South Korea, New Zealand, Papua New Guinea)
- Central American community (Bahamas, Belize, Cuba, Guatemala, Haiti, Honduras, Nicaragua, Trinidad and Tobago)

As expected, we found that geography played a large role in communities, with at least one OPEC member in most of these geographic regions. What was interesting was that while Girvan-Newman reported a number of

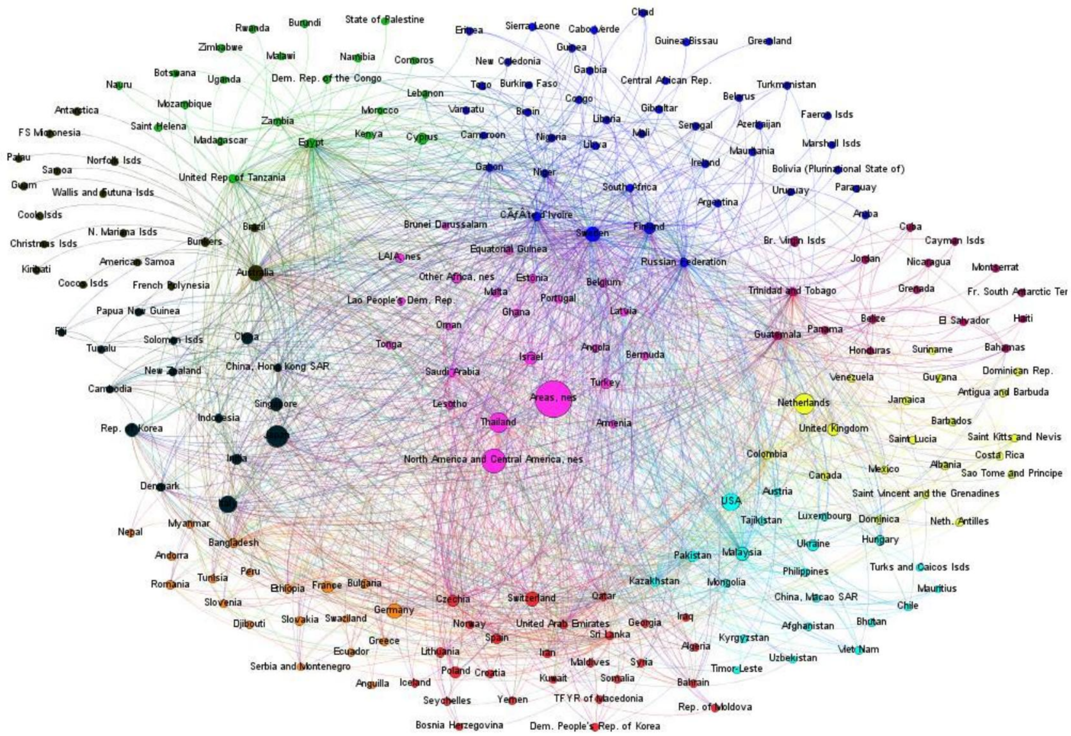

Fig. 8. Global fuels communities, 2005

communities that accorded with our intuitions of global trade, this was not as much the case in 1997. The only community that remained largely constant between the two time periods was the Island communities of Australia, French Polynesia, Samoa, etc. Additionally, our result also corroborated [2]'s result in which the US-China-Germany-UK community in 1997 fell with the emergence of a China-Oceania community by 2005.

TABLE IV
HIGHEST SCORING YEAR-OVER-YEAR PURITY AVERAGES FOR BETWEENNESS CENTRALITY

Antartica
Wallis and Futuna Isds
N. Mariana Isds
American Samoa
Guam
Palau
Nauru
New Caledonia
Australia
Norfolk Isds

VI. CONCLUSIONS AND FUTURE WORK

Our results largely confirmed a number of intuitions in mainstream economic discourse. We found that the world has become globalized over time (with the degrees of nodes and centralities of edges increasing with the number of new nodes and new edges increasing over time, suggesting that nodes are becoming more interconnected). We also corroborated some of the results of previous work, noting that China's formal entrance

TABLE V
METRICS FOR VARIOUS COMMUNITY DETECTION ALGORITHMS

Method	YOYP Average	1996-2005 Avg Purity
Weighted Btw. Cent.	0.18	0.12
Weighted Btw. Cent. by Degree	0.19	0.16
Largest Edge	0.87	0.79
Largest Edge Weighted by Degrees	0.88	0.77

to the WTO resulted in communities that were more geographically related (most notably the rise of the China-Oceania community in place of the China-US-UK-Germany community).

We also developed a metric of community volatility, which we called *purity*. From this metric we were able to quantitatively measure how much communities shifted over time, again confirming a number of economic intuitions. Future work would include a more rigorous model of trade network perturbances (in the form of new nodes and edges) with the understanding that new communities that form are often geographic in nature. Additionally, we found that while most communities changed over the 10 year time horizon that we explored, the Island communities (including Barbados, Polynesia, etc) tended to remain constant. As such, understanding the nature of those economies and how resilient they have been to global changes (price shocks, new nodes, noew edges, etc) may be an illuminating case study.

REFERENCES

[1] United Nations Comtrade Dataset http://comtrade.un.org/
[2] The Rise of China in the International Trade Network, Zhu et al, 2014 http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0105496
[3] International Trade Cascades http://snap.stanford.edu/class/cs224w-2012/projects/cs224w-007-final.pdf
[4] Predicting Crisis in Global Trade Network http://snap.stanford.edu/class/cs224w-2015/projects_2015/Predicting_Crisis_in_the_Global_Trade_Network.pdf
[5] http://www.foreign-trade.com/reference/hscode.htm
[6] https://networkx.github.io/
[7] http://snap.stanford.edu/snappy/index.html