
B

Bayesian Games

- ▶ [Game Theory and Social Networks](#)

Bayesian Networks

- ▶ [Probabilistic Graphical Models](#)

B2B Marketing

- ▶ [Business-to-Business Marketing](#)

Belief

- ▶ [Trust in Social Networks](#)

Belief Propagation

- ▶ [Fraud Detection Using Social Network Analysis, a Case Study](#)

Benchmarking for Graph Clustering and Partitioning

David A. Bader¹, Henning Meyerhenke², Peter Sanders², Christian Schulz², Andrea Kappes², and Dorothea Wagner²

¹School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA, USA

²Karlsruhe Institute of Technology (KIT), Institute of Theoretical Informatics, Karlsruhe, Germany

Synonyms

[Algorithm evaluation](#); [Graph repository](#); [Test instances](#)

Glossary

Benchmarking Performance evaluation for comparison to the state of the art

Benchmark Suite Set of instances used for benchmarking

Definition

Benchmarking refers to a repeatable performance evaluation as a means to compare somebody's

work to the state of the art in the respective field. As an example, benchmarking can compare the computing performance of new and old hardware.

In the context of computing, many different benchmarks of various sorts have been used. A prominent example is the Linpack benchmark of the TOP500 list of the fastest computers in the world, which measures the performance of the hardware by solving a dense linear algebra problem. Different categories of benchmarks include sequential vs. parallel, microbenchmark vs. application, or fixed code vs. informal problem description. See, e.g., Weicker (2002) for a more detailed treatment of hardware evaluation.

When it comes to benchmarking algorithms for network analysis, typical measures of interest are solution quality and running time. The comparison process requires the establishment of widely accepted *benchmark instances* on which the algorithms have to compete. In the course of the 10th DIMACS Implementation Challenge on Graph Partitioning and Graph Clustering (Bader et al. 2012), we have assembled a suite of graphs and graph generators intended for comparing graph algorithms with each other. While our particular focus has been on assembling instances for benchmarking graph partitioning and graph clustering algorithms, we believe the suite to be useful for related fields as well. This includes the broad field of network analysis (which includes graph clustering, also known as *community detection*) and various combinatorial problems.

The purpose of DIMACS Implementation Challenges is to assess the practical performance of algorithms in a respective problem domain. These challenges are scientific competitions in areas where worst case and probabilistic analysis yield unrealistic results. Where analysis fails, experimentation can provide insights into realistic algorithm performance. By evaluating different implementations on the assembled benchmark suite, the challenges create a reproducible picture of the state of the art in the area under consideration. This helps to foster an effective technology transfer within the research areas of algorithms, data structures, and implementation techniques as well as a transfer back to the original applications.

Introduction

Graph partitioning and graph clustering (or community detection) are ubiquitous subtasks in many application areas. Generally speaking, both techniques aim at the identification of vertex subsets (*clusters*) with many internal and few external edges. In this work we concentrate our description on aspects important to the field of network analysis, in particular on community detection.

In its most general form, community detection does not require a fixed number k of clusters nor constraints on the size of the clusters. Instead, a quality function which measures both, the density inside clusters and the sparseness between them, is used. A variety of such functions has been proposed, among which the measure *modularity* has proven itself fairly reliable and largely in accordance with human intuition in the literature:

Problem 1 (Modularity Maximization) Given an undirected, weighted graph $G = (V, E, \omega)$ without parallel edges, find a partition \mathcal{C} of V which optimizes the modularity objective function:

$$Q(\mathcal{C}) := \sum_{\substack{\mathcal{C} \in \mathcal{C} \\ \{u,v\} \in E \\ u,v \in \mathcal{C}}} \omega(\{u,v\}) \frac{\left(\sum_{v \in \mathcal{C}} s(v) \right)^2}{\sum_{e \in E} \omega(e)} - \frac{\left(\sum_{v \in \mathcal{C}} s(v) \right)^2}{4 \left(\sum_{e \in E} \omega(e) \right)^2}.$$

Here we assume $e = \{u, v\} \in E$ is a multiset (i.e., a self-loop $u = v$ is allowed) and that the strength $s(v)$ of a node v is the sum of the weights of its incident edges. Recently some criticism towards modularity has emerged. Fortunato and Barthelemy (2007) demonstrate that global modularity optimization cannot distinguish between a single community and a group of smaller communities. Berry et al. (2011) provide a weighting mechanism that alleviates the resolution-limit problem, but others remain (Good et al. 2010; Lancichinetti and Fortunato 2011). That is one reason why the 10th DIMACS Implementation Challenge had a second graph clustering category apart from modularity maximization. In this second category, algorithms are compared

with respect to four different objective functions, with the goal to explicitly invite clustering algorithms that are not based on a specific objective function. These objective functions are based upon performance (van Dongen 2000), intracluster density, and intercluster conductance and expansion (Kannan et al. 2004).

In contrast to graph clustering, the term graph partitioning usually implies that the number of partitions is fixed and the task is to partition the vertex set into blocks of (almost) equal size. Its main application is not network analysis but the preprocessing of data for parallel computing. The objective functions used for the partitioning sub-challenges are the number of edges between the blocks and the maximum communication volume (Çatalyürek and Aykanat 1996) of the partition.

Participants of the challenge were invited to submit solutions to the different challenge categories on graph partitioning and graph clustering. This way different algorithms and implementations were tested against the benchmark instances. Thereby future researchers are enabled to identify techniques that are most effective for a respective partitioning or clustering problem – by using our benchmark set and by comparing their results to the challenge results and to those in the literature published afterwards.

In this entry we describe the benchmark suite and its assembly process. Moreover, we sketch some of the results obtained by the challenge participants using the benchmark graphs.

Key Points

Collecting the instances for the benchmark suite was performed with two main aspects in mind, diversity of source applications and diversity of instance sizes. Moreover, some graphs have been frequently used in previous work, whereas others are new or fairly recent.

Some instances are based on real-world inputs, while others have been created using a generator. The generated graphs also vary in how closely they resemble real-world counterparts.

All instances have been long term archived with public access (Bader et al. 2012).

The solutions generated by the challenge participants using the benchmark suite constitute a valuable picture of the state of the art in graph partitioning and graph clustering. To better suit algorithms that do not explicitly optimize a traditional objective function (and to circumvent known flaws in these traditional objective functions), additional criteria to assess the quality of the submitted clusterings were evaluated.

Moreover, a nondiscriminatory way to assign scores to solvers that takes both running time and solution quality into account was used.

Historical Background

Previous DIMACS Implementation Challenges addressed a large variety of algorithmic problems, several of them involving graphs and networks. Graph repositories similar to our benchmark suite exist as well. However, they often lack the size, broadness of source applications, and connection to a quality-driven competition.

An example repository widely used in combinatorial scientific computing is Chris Walshaw's graph partitioning archive (Soper et al. 2004). It stores 34 graphs and the best-known graph partitions computed for these graphs. This archive has substantially simplified the improvement of graph partitioning algorithms over the last decade. Today, however, the instances contained therein have to be deemed rather small and also somewhat limited in terms of application areas. For example, there are no social networks contained in this archive.

The University of Florida Sparse Matrix Collection (Davis 2008), maintained by Tim Davis, is broader in terms of application areas and matrix sizes. Although social networks have been included recently as well, most matrices stem from technical applications.

Graph collections focusing on scale-free graphs such as social networks do exist, e.g., Arenas (2009) and Newman (2012). One of the more prominent examples is the Stanford

Large Network Dataset Collection SNAP (). These collections are often somewhat limited in the number of stored graphs or their variability. Also, and maybe more importantly, in most cases, the collections lack a significant comparison of how a larger number of different algorithms perform on the data.

Proposed Solution and Methodology

With the 10th DIMACS Implementation Challenge and its graph collection, we addressed both of these issues. Our collection contains more than 100 graphs of various origins and assembled in different categories. (In addition we link to the Walshaw archive and the University of Florida Sparse Matrix Collection.) We took care that our collection contains instances best suited for partitioning in technical applications as well as instances particularly intended for clustering and related network analysis tasks. Additionally the challenge results provide guidance as to how different algorithms perform on different classes of graphs.

The driving considerations during the assembly of the graph collection were to include a sufficiently large variety of application sources (thereby instance structures) and graph sizes. In that line of reasoning, we identified three higher-level classes from which to select: random graphs, generated graphs with close resemblance to data from real-world applications, and actual real-world data. Our intent is to offer a good diversity in order to provide a meaningful benchmark for network analysis and graph partitioning algorithms.

With generators at hand, an experimenter can scale to (nearly) arbitrary graph sizes, retaining the general structure of the graphs while increasing their sizes. This is, for example, important when performing weak-scaling studies for the experimental analysis of parallel algorithms. It is worth mentioning that, since the instance sizes are only limited by architectural constraints, generators provide a means to “grow” instance sizes with future architectural improvements. This way the current state of a collection does not age

as quickly as without generators. More details on the graphs generated to resemble real-world inputs follow below.

Generated random graphs offer an additional benefit. They are usually easier to analyze with theoretical methods than other graph types. The Erdős-Rényi (ER) model, for example, has experienced significant consideration in theoretical works. Many important properties of these graphs were proved in this long course of research; see, e.g., Bollobás (1985). Due to the lack of resemblance of typical ER graphs to real-world inputs, an active line of research is developing more realistic models such as R-MAT (Chakrabarti et al. 2004) and BTER (Seshadhri et al. 2012). The random graphs we included and their generators are described later in this section in more detail.

Finally, real-world graphs add the necessary confidence that an algorithm’s performance in terms of running time and quality on the collection resembles its performance on the represented real-world applications. The real-world graphs we included are also described below.

In the remainder of this section, we first explain the preprocessing performed to unify the instances in our collection. After that, the individual categories of the collection are described.

Preprocessing

Graph partitioning and, with some exceptions, also graph clustering are usually applied to undirected graphs. A common preprocessing step is therefore to *symmetrize* the graph prior to partitioning, i.e., to make the graph undirected by including an undirected edge between two vertices a and b if and only if there exists an edge from a to b or from b to a .

If both directions are present in the original graph, there are several possibilities to assign a weight to the resulting undirected edge. We chose the following approach: if the input graph is unweighted, an edge between two vertices is considered as the information that they are related in some sense, independent of the strength of this connection. Hence, if an edge in both directions exists, it is translated into an unweighted, undirected edge. On the other hand, if the input graph is weighted, we add up the edge weights of both

directions, as the connection between the vertices is typically stronger if both directions exist. Analogously, the weight of parallel edges is summed up only in case of weighted networks; otherwise, parallel edges are removed. Self-loops are always removed (with the exception of some synthetic Kronecker graphs, for which versions with self-loops and parallel edges and versions without exist). Only a handful of the real-world networks included in the benchmark contain (few) parallel edges and self-loops; therefore, this decision does not alter the structure of the networks considerably.

Furthermore, the graph format used in the benchmark is a slight extension of the format that some well-established partitioners such as Metis (Karypis and Kumar 1999) use. This format supports only integer edge weights, but some of the real-world graphs use fractional weights in very different orders of magnitude. It would have been possible to define an extension of this format to allow for fractional weight. However, this might have prevented some solvers to enter the challenge. Multiplying the edge weights by a suitable power of 10 to get integer weights would have been another approach. Yet, as the edge weights are of very different ranges, each graph would have needed its own normalization, and without rounding, the resulting edge weights could be too large to fit in standard integer types. Although we are aware that this causes a loss of information, we felt that the neatest solution was to make the respective graphs unweighted. One of the benchmark graphs (`cond-mat-2005`) originally contains edges with a weight `inf`. As their meaning is not clear and none of the objective function is well defined in case of infinite weights, we discarded these edges as well as all edges with an edge weight of 0.

Random Graphs

The *Erdős-Rényi random graph* generator in the collection creates graphs according to the well-known $G(n, p)$ model (which is very similar to the original model proposed by Erdős and Rényi but was actually devised by Gilbert (1959)). The included graphs have been generated with

$p = \frac{1.5 \ln n}{n}$, where the value of p is chosen with the intent to obtain connected graphs with high probability. This class of graphs is well studied in theory. It is also known that typical Erdős-Rényi graphs do not resemble real-world graphs. The class was included nevertheless for its theoretical importance and due to the easy generation of large graphs with high average degree.

The graphs in the category *Kronecker* are generated using the *Graph500 benchmark* (Bader et al. 2010). This benchmark’s purpose is to measure the performance of computer systems when processing graph-structured workloads. More specifically, our instances are derived from an R-MAT generator which is part of the benchmark. R-MAT graphs (Chakrabarti et al. 2004) are generated by sampling from a perturbed Kronecker product. They are scale-free and reflect many properties of real social networks. All files have been generated with the R-MAT parameters $A=0.57$, $B=0.19$, $C=0.19$, and $D=0.05$ and edge factor 48, i.e., the number of edges equals $48n$, where n is the number of vertices. The original Kronecker files contain self-loops and multiple edges. These properties are also present in real-world datasets. However, as some tools cannot handle these “artifacts,” we present “cleansed” versions of the datasets (yielding simple graphs) as well.

Delaunay and random geometric graphs are taken from KaPPa (Karlsruhe Parallel Partitioner). Here, *rggX* is a *random geometric graph* with 2^X nodes. Each node represents a random point in the unit square and edges, connecting nodes whose Euclidean distance is below $0.55\sqrt{\ln n/n}$. This threshold is chosen in order to ensure that the graph is almost connected. The graph *DelaunayX* is the Delaunay triangulation of 2^X random points in the unit square.

Generated Graphs with Real-World Structure

Each graph in the star mixture section of the benchmark represents a starlike structure of different graphs S_0, \dots, S_t . Here the graphs S_1, \dots, S_t are weakly connected to the center

S_0 by random edges. The total number of random edges added between each S_i and S_0 is less than 3% out of the total number of edges in S_i . The graphs are mixtures of the following structures: social networks, finite-element graphs, VLSI chips, peer-to-peer networks, and matrices from optimization solvers. These graphs were submitted by Safro et al. and are included into the benchmark because they are potentially hard graphs for graph partitioning.

Two classic random models of social networks are *preferential attachment* (Barabási and Albert 1999) and *small world* (Watts and Strogatz 1998). In the context of graph clustering, *planted partition* or $G(n, p_{in}, p_{out})$ graphs are frequently used to validate algorithms (Lancichinetti and Fortunato 2009). These networks do not exhibit common properties of real-world social networks like a power-law degree distribution. However, their use is typically motivated by the knowledge of a ground-truth clustering that is used in the generation process and can be used to compare algorithms independent of specific objective functions. We included one graph of each category in the benchmark set as we deemed it interesting to see to what extent algorithmic behavior on these graphs coincides with the behavior on real-world data.

Although we do not store dynamic graphs, three so-called frames (static instances within the same dynamic sequence) from three dynamic mesh sequences each are included in our collection. These sequences resemble two-dimensional adaptive numerical simulations. The generator is explained in some detail by Marquardt and Schamberger.

Computational task graphs model temporal dependencies between tasks to be solved, here for applications working on data streams. The generated graphs can be used for performance analysis of algorithms and the development of improved hardware parameters. These graphs have been submitted by Ajwani et al.

Real-World Graphs

The benchmark includes a large number of real-world networks stemming from many different applications. Since scientific computing is a major application area using graph partitioning,

we included graphs that have been used in *numerical simulations* of various kinds.

The partitioning of *road networks* is an important technique when it comes to preprocessing for shortest-path algorithms (Bauer et al. 2010). The graphs that can be found in this section are road networks from whole continents, e.g., Europe, as well as from whole countries, e.g., Germany. These graphs were submitted by Kobitzsch and are based on data from the OpenStreetMap project.

Parallel direct methods for solving linear systems yield another important application of graph partitioning. We therefore included a subset of graph representations of matrices from the University of Florida Sparse Matrix Collection (Davis 2008).

In the context of graph clustering, the analysis of social networks is one of the most important applications. The part of the benchmark suite especially addressed to clustering algorithms reflects this by including a variety of real-world social networks. Most of these are taken from the webpages of Newman (2012) and Arenas (2009) and have been previously used to compare and evaluate clustering in the context of modularity maximization.

A special subcategory of social networks are *coauthorship networks*. In a scientific context, coauthorship networks link scientists that have coauthored at least one publication. The DIMACS benchmark includes coauthorship graphs from the field of astrophysics, condensed matter and high-energy theory, network science, and computer and information science. Closely related to these are *copaper* and *citation* networks. Copaper graphs are compiled analogously to coauthorship graphs by linking papers if they share at least one author. In contrast to that, citation networks link papers with another if one cites the other. The benchmark set contains graphs of both kinds based on publications in computer and information science.

Graph clustering has also been successfully applied to *web graphs*, where edges link webpages based on hyperlinks. A subset of the web graphs we included is gathered by the Laboratory for Web Algorithms in Milano by domain-wise

crawls performed between 2000 and 2007. In the context of the challenge, these networks are particularly interesting due to their size; in fact, the graph combining 12 monthly snapshots of the .uk domain comprises over 3 billion edges, which makes it the largest network in the whole benchmark set.

Apart from these, the part of the benchmark set explicitly addressed to clustering contains a variety of (mainly) small networks from various application areas such as biology and political science. All of these are well known in the modularity-based clustering community. For details on particular networks and references, we refer to the challenge webpage (Bader et al. 2012).

The graphs in the *redistricting* category represent US states. They are used for solving the redistricting problem, i.e., determining new electoral boundaries, for example, due to population changes. Each node represents a block from the 2010 census. Two nodes share an edge if their blocks are adjacent.

Illustrative Example

As running times would have been prohibitive for the whole set of benchmark instances, participants of the competition were only required to submit clusterings for a subset of instances, the final challenge testbed. This subset was announced 2 weeks before the deadline. To illustrate the performance of different algorithms on graphs from different categories, Table 1 shows the best modularity values achieved by the submitted solvers on the final challenge testbed. From the 15 solvers in this category, 2 clearly lead the field. *CGGci_RG* (Ovelgönne and Geyer-Schulz 2012) iteratively combines several high-quality clusterings to find a solution with higher quality. In contrast to that, *VNS* (Aloise et al. 2012) uses the metaheuristic *variable neighborhood search*, a variant of local search. With few exceptions, *VNS* achieves the best results on networks with up to approximately 100,000 vertices but is outperformed by *CGGci_RG* in larger networks. An interesting observation is that

ParMod (Çatalyürek et al. 2012), a technique based on recursive bipartitions, attains the best modularity values on two graphs. Neither the size nor the density of these graphs is exceptional, but unlike the majority of graphs used for this competition, they exhibit a mesh-like structure.

In addition to quality, running time is also an important aspect when choosing an algorithm for a certain application. This is why the DIMACS Challenge included a second subchallenge for each objective function, where both quality and speed contributed to the final scores. More specifically, the scoring is based on the *Pareto Count* of a submitted algorithm on an instance, i.e., the number of competing algorithms that are both faster and achieve a higher quality. In this category, a fast agglomerative solver named *RG* (Ovelgönne and Geyer-Schulz 2012) obtained the best scores. While the differences in running time might not seem very important in the context of small instances, they were in fact huge on larger instances. Considering, for example, the raw running times on the web graph *uk-2002*, *RG* needs approximately 13 min to compute a clustering, which is more than 600 times faster than the running time of *CGGci_RG*, while the difference in modularity is less than 0.001. This running time can be further improved by using parallel algorithms. For example, one of the submissions is able to cluster this instance in only 30 s by using a GPU (Fagginger Auer and Bisseling 2012), with a modularity that is still larger than 0.97.

Consequently, the question which algorithm is the “best” cannot always be answered globally. Instead, the answer often depends on application-specific parameters like the size and structure of certain instances, as well as the available hardware and a custom trade-off between quality and running time. Comparing the results of different algorithms on various benchmark instances can assist the choice of an appropriate algorithm.

Future Directions

With the graph archive of the 10th DIMACS Implementation Challenge on Graph Partitioning

Benchmarking for Graph Clustering and Partitioning, Table 1 Best modularity scores achieved by challenge participants on the challenge testbed

Graph	Modularity	Solver
as-22july06	0.678267	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
astro-ph	0.744621	VNS (Aloise et al. 2012)
audikw1	0.917983	VNS (Aloise et al. 2012)
belgium.osm	0.994940	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
cage15	0.903173	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
caidaRouterLevel	0.872042	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
celegans_metabolic	0.453248	VNS (Aloise et al. 2012)
citationCiteseer	0.823930	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
coAuthorsCiteseer	0.905297	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
cond-mat-2005	0.746254	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
coPapersDBLP	0.866794	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
email	0.582829	VNS (Aloise et al. 2012)
er-fact1.5-scale25	0.077934	comm-el (Riedy et al. 2012)
eu-2005	0.941554	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
G_n_pin_pout	0.500098	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
in-2004	0.980622	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
kron_g500-s-logn16	0.065056	VNS (Aloise et al. 2012)
kron_g500-s-logn20	0.050350	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
ldoor	0.969370	ParMod (Çatalyürek et al. 2012)
luxembourg.osm	0.989621	VNS (Aloise et al. 2012)
memplus	0.700473	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
PGPgiantcompo	0.886564	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
polblogs	0.427105	VNS (Aloise et al. 2012)
power	0.940851	VNS (Aloise et al. 2012)
preferentialAttachment	0.315994	VNS (Aloise et al. 2012)
rgg_n_2_17_s0	0.978324	VNS (Aloise et al. 2012)
smallworld	0.793042	VNS (Aloise et al. 2012)
uk-2002	0.990301	CGGCi_RG (Ovelgönne and Geyer-Schulz 2012)
uk-2007-05	0.480210	comm-el-xmt2 (Riedy et al. 2012)
333SP	0.989095	ParMod (Çatalyürek et al. 2012)

and Graph Clustering, we have introduced a comprehensive collection of graphs that can be used for the assessment of graph partitioning and network analysis algorithms.

With the archive we hope to simplify the development of improved solution techniques in these areas by allowing algorithm engineers to compare the performance of their implementations to the state of the art.

A deliberate limitation is to not consider dynamic graphs, directed graphs, nor hypergraphs. As instances of this type were not considered in the challenge, they were not included in the

collection either. We do consider all these omitted graph types useful though, and they represent interesting applications.

Of particular interest for network analysis are directed and dynamic graphs. There is no lack of data (albeit not all of them are publicly accessible). As an example, the dynamic interaction of social network users over time constitutes a dynamic graph that is of particular interest to social media enterprises and online marketers. When compiling dynamic instances into a collection, it should be considered that dynamic graphs are more difficult to assemble or generate in a

consistent way – issues such as a suitable interval length and space-saving storage formats arise.

We encourage interested colleagues to start a new collection using a similar methodology, this time focusing on the types of graphs we omitted. Such an effort would certainly be beneficial to the network analysis community.

Acknowledgment

The authors would like to thank all contributors to the 10th DIMACS Implementation Challenge graph collection. Tim Davis provided valuable guidelines for preprocessing the data. Financial support by the sponsors DIMACS; the Command, Control, and Interoperability Center for Advanced Data Analysis (CCICADA); Pacific Northwest National Laboratory; Sandia National Laboratories; Intel Corporation; and Deutsche Forschungsgemeinschaft (DFG) is gratefully acknowledged.

Cross-References

- ▶ [Clustering Algorithms](#)
- ▶ [Community Detection, Current and Future Research Trends](#)
- ▶ [Extracting and Inferring Communities via Link Analysis](#)
- ▶ [Simulated Datasets](#)
- ▶ [Social Network Datasets](#)
- ▶ [Sources of Network Data](#)

References

- Aloise D, Caporossi G, Perron S, Hansen P, Liberti L, Ruiz M (2012) Modularity maximization in networks by variable neighborhood search. In: 10th DIMACS implementation challenge workshop. Georgia Institute of Technology, Atlanta
- Arenas A (2009) Network data sets. <http://deim.urv.cat/~aarenas/data/welcome.htm>. Online; accessed 28 Sept 2012
- Bader DA, Berry J, Kahan S, Murphy R, Jason Riedy E, Willcock J (2010) Graph 500 benchmark 1 (“search”), version 1.1. Technical report, Graph 500
- Bader D, Meyerhenke H, Sanders P, Wagner D (2012) 10th DIMACS implementation challenge. <http://www.cc.gatech.edu/dimacs10/>. Online; accessed 28 Sept 2012
- Barabási A-L, Albert R (1999) Emergence of scaling in random networks. *Science* 286:509–512
- Bauer R, Dellling D, Sanders P, Schieferdecker D, Schultes D, Wagner D (2010) Combining hierarchical and goal-directed speed-up techniques for Dijkstra’s algorithm. *ACM J Exp Algorithmics* 15: 2.3:2.1–2.3:2.31
- Berry JW, Hendrickson B, LaViolette RA, Phillips CA (2011) Tolerating the community detection resolution limit with edge weighting. *Phys Rev E* 83:056119
- Bollobás B (1985) *Random graphs*. Academic, London
- Çatalyürek ÜV, Aykanat C (1996) Decomposing irregularly sparse matrices for parallel matrix-vector multiplication. In: Ferreira A, Rolim J, Saad Y, Yang T (eds) *Parallel algorithms for irregularly structured problems*. Volume 1117 of lecture notes in computer science. Springer, Berlin/Heidelberg, pp 75–86. doi:10.1007/BFb0030098
- Çatalyürek ÜV, Kaya K, Langguth J, Uçar B (2012) A divisive clustering technique for maximizing the modularity. In: 10th DIMACS implementation challenge workshop. Georgia Institute of Technology, Atlanta
- Chakrabarti D, Zhan Y, Faloutsos C (2004) R-MAT: a recursive model for graph mining. In: *Proceedings of the 4th SIAM international conference on data mining (SDM)*, Orlando, April 2004. SIAM
- Davis T (2008) The University of Florida Sparse Matrix Collection. <http://www.cise.ufl.edu/research/sparse/matrices>. Online; accessed 28 Sept 2012
- Fagginger Auer BO, Bisseling RH (2012) Graph coarsening and clustering on the GPU. In: 10th DIMACS implementation challenge workshop. Georgia Institute of Technology, Atlanta
- Fortunato S, Barthelemy M (2007) Resolution limit in community detection. *Proc Natl Acad Sci* 104:36–41
- Gilbert H (1959) Random graphs. *Ann Math Stat* 30(4):1141–1144
- Good BH, de Montjoye Y-A, Clauset A (2010) Performance of modularity maximization in practical contexts. *Phys Rev E* 81:046106
- Kannan R, Vempala S, Vetta A (2004) On clusterings: good, bad, spectral. *J ACM* 51(3):497–515
- Karypis G, Kumar V (1999) A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J Sci Comput* 20:359–392
- Lancichinetti A, Fortunato S (2009) Community detection algorithms: a comparative analysis. *Phys Rev E* 80(5):056117
- Lancichinetti A, Fortunato S (2011) Limits of modularity maximization in community detection. *Phys Rev E* 84:066122
- Lescovec J. Stanford Network Analysis Package (SNAP). <http://snap.stanford.edu/index.html>. Online; accessed 28 Sept 2012
- Newman M () Network data. <http://www-personal.umich.edu/~mejn/netdata/>. Online; accessed 28 Sept 2012

Ovelgönne M, Geyer-Schulz A (2012) An ensemble learning strategy for graph clustering. In: 10th DIMACS implementation challenge workshop. Georgia Institute of Technology, Atlanta

Riedy EJ, Meyerhenke H, Ediger D, Bader DA (2012) Parallel community detection for massive graphs. In: 10th DIMACS implementation challenge workshop. Georgia Institute of Technology, Atlanta

Seshadhri C, Kolda TG, Pinar A (2012) Community structure and scale-free collections of Erdős-Rényi graphs. *Phys Rev E* 85(5):066122

Soper AJ, Walshaw C, Cross M (2004) A combined evolutionary search and multilevel optimisation approach to graph-partitioning. *J Glob Optim* 29(2): 225–241

van Dongen SM (2000) Graph clustering by flow simulation. PhD thesis, University of Utrecht

Watts DJ, Strogatz SH (1998) Collective dynamics of “Small-World” networks. *Nature* 393: 440–442

Weicker R (2002) Benchmarking. In: Calzarossa M, Tucci S (eds) Performance evaluation of complex systems: techniques and tools. Volume 2459 of lecture notes in computer science. Springer, Berlin/Heidelberg, pp 231–242

Biased Graph

- ▶ [Signed Graphs](#)

Bibliometrics

- ▶ [Pajek](#)

Biclustering

- ▶ [Network Anomaly Detection Using Co-clustering](#)

Biosurveillance

- ▶ [Disease Surveillance, Case Study](#)

Blackhole Pattern

- ▶ [Mining Blackhole and Volcano Patterns for Fraud Detection](#)

Blind Source Separation

- ▶ [Independent Component Analysis](#)

Board Interlocks

- ▶ [Interlocking Directorate Networks](#)

Boundary Problem

- ▶ [Sources of Network Data](#)

Boundary Spanning

- ▶ [Managerial Networking](#)

BPEL4People

- ▶ [Social Interaction Analysis for Team Collaboration](#)

Brokerage

- ▶ [Structural Holes](#)

Bump Hunting Technique

- ▶ [Semi-discrete Decomposition](#)

Business Marketing

- ▶ [Business-to-Business Marketing](#)

Business Process Creation

- ▶ [Web Service Composition](#)

Business Process Intelligence

- ▶ [Desire Lines in Big Data](#)

Business-to-Business Marketing

Ross Brennan
Hertfordshire Business School,
University of Hertfordshire, Hatfield, UK

Synonyms

[B2B marketing](#); [Business marketing](#); [Industrial marketing](#); [Interorganizational marketing](#); [Organizational marketing](#)

Glossary

AMA American Marketing Association

ARA Analysis Consideration of the actor bonds, resource ties, and activity links that connect the parties in an industrial network

B2B Acronym for “business-to-business”

B2C Acronym for “business-to-consumer”

Decision-Making Unit The team of managers who are involved, both directly and indirectly, in making organizational purchasing decisions

IMP Group International group of researchers who collaborate on business-to-business relationships research (see www.impgroup.org)

Marketing Mix (Approach) An approach to marketing based on the ideas that the buyer and seller operate independently and that the marketer is the active participant while the buyer responds passively to the seller’s marketing mix. The marketing mix is the

combination of stimuli (such as product design, advertising, social media, pricing, customer service, and distribution) employed by the marketer to convince the buyer to buy.

Definition

Business-to-business marketing concerns seller-initiated processes directed at creating mutually beneficial exchange transactions between organizations.

Historical Background

Marketing, in general, and B2B marketing, specifically, are relatively recently developed fields of scholarship. Recent years have seen marketing scholarship seek to emancipate itself from its underpinning disciplines, notably neoclassical economics, psychology, and social psychology. Much recent scholarship revolves around two key directions in the development of marketing thought: firstly, the attempt to develop original marketing models that do not simply borrow ideas from other social science disciplines and, secondly, to differentiate between scholarship intended to advance the field of study and scholarship undertaken to develop managerial technologies and techniques. The field of B2B marketing, as explained here, has played a key role in the development of both of these directions.

Foundations

The Changing Definition of Marketing

In 2007 the American Marketing Association defined marketing as:

The activity, set of institutions, and processes for creating, communicating, delivering, and exchanging offerings that have value for customers, clients, partners, and society at large. (<http://www.marketingpower.com/Community/ARC/Pages/Additional/Definition/default.aspx>)

This new definition replaced the 2004 AMA definition, which was:

Marketing is an organizational function and a set of processes for creating, communicating, and delivering value to customers and for managing customer relationships in ways that benefit the organization and its stakeholders.

The transition from the 2004 definition to the 2007 definition represented the continuation of a shift in marketing thought, away from definitions dominated by managerial concerns and the management of organizational processes (often understood as the marketing mix), towards a definition that acknowledges that the marketing system is embedded within, and a part of, wider social systems. Intellectual developments within the field of B2B marketing in particular had, over the preceding decades, played their part in bringing about this transition.

B2B Marketing and B2C Marketing

Until the 1960s B2B marketing was regarded as a minor offshoot of B2C marketing, and both were dominated by managerially orientated research. Marketing was seen as the sole concern of marketing and sales managers, and customers were seen as the passive – if sometimes recalcitrant – recipients of marketing messages. The job of marketing (management) was to find out what customers wanted, design profitable products to meet those needs, and then persuade customers to part with their money to buy the products (note that in the language of marketing, the category “product” subsumes both goods and services). The emphasis was on the identification of the key buying criteria used by purchasing organizations, the identification of the members of the decision-making unit (or buying center), and the development of strategies to influence those members through persuasive marketing communications, involving advertising and direct selling.

The Emergence of Modern Theories of B2B Marketing

Since the 1970s this rather naïve model of marketing has been challenged in numerous ways and in virtually all of the contexts in which marketing

takes place. It is in the B2B context that the model has been challenged most extensively and most effectively. The challenges were both conceptual and empirical (Ford 2001).

Conceptually, it is clear that customer organizations engage in complex buying processes that are simultaneous with and that interact with the marketing and selling processes of supplier organizations. It makes little sense to analyze the B2B marketing and selling process independently of the customer organization’s buying process; these processes are interlinked (Ford and Håkansson 2006). Furthermore, B2B buying and selling frequently takes place in the context of small-number exchange, meaning, simply, that in many B2B markets there are relatively few qualified suppliers and relatively few substantial buyers. In contrast to the B2C context, where the customer is usually an individual person and the supplier is often a large organization, in the B2B context the customer is an organization and may be economically very powerful (consider the buying power of Ford, Boeing, or Samsung). Consequently, the notion of a passive buyer responding to the marketing messages of an active seller is inappropriate for a wide range of B2B markets. Sometimes the buying organization will be in charge of the process (specifying and designing the product that it requires, managing the process of communication with prospective suppliers, and then subcontracting supply). In other cases the process will be managed through an equal partnership between buyer and supplier. There are also some B2B markets in which the seller-dominant model of marketing makes sense. The key point is that the nature of the B2B marketing process is contingent on a wide range of factors, such as the number of buyers and sellers, the concentration of buying and selling power, the dynamism of the economic environment, the complexity of the product, and the rate of change of technology. A simplistic conception of marketing as the process by which the seller identifies and then meets customer needs is entirely inadequate.

Empirically, many studies have demonstrated that, frequently, buying and selling organizations engage in what can be described as “buyer-seller

relationships.” Such relationships involve, but are by no means uniquely defined by, repeated exchanges of products for money. Repeated exchanges are necessary but insufficient conditions for the existence of a B2B buyer-seller relationship. In addition to repeated exchanges, a buyer-seller relationship may be characterized by a range of factors such as extensive sharing of commercial and technical information, extensive social interaction between personnel from the buying and selling organizations, tolerance for occasional underperformance by the relationship partner, and emergence of mutual trust which involves both cognitive and affective components. Many empirical studies have sought to measure key relational variables, such as trust, commitment, power, relationship strength, adaptation, and information exchange, within B2B buyer-seller relationships.

The IMP Group: The Interaction and ARA Models

According to the ARA (actors, resources, activities) model of the IMP Group, B2B buyer-seller relationships lead to the development of bonds between the parties (actors) to the relationship, ties between their resources, and links between their activities (Håkansson and Snehota 1995). The IMP Interaction Model was an influential, relationship-based approach to B2B marketing developed in the 1980s. The Interaction Model sought to identify the important characteristics of dyadic (i.e. one to one) business relationships. In addition to the relationship characteristics mentioned already, the concept of the relationship atmosphere was an important contribution from the IMP approach. Although never defined entirely satisfactorily, the concept of the relationship atmosphere seeks to capture the essence of “how the relationship feels” at a given point in time; short-term underperformance by one party to the relationship may damage the relationship atmosphere but, in the case of a robust relationship, is unlikely to lead to relationship termination. Various attempts have been made to develop, either empirically or conceptually, life cycle models of B2B

buyer-seller relationships. Such models, defining relationships in terms of stages through which they progress, remain controversial. While there is clearly some sense in which a B2B relationship is initiated and then, if successful, develops and matures (and possibly is eventually terminated), there is little or no evidence that this process occurs in a predictable series of readily defined “stages” with clear boundaries.

The Industrial Network Perspective

The industrial network school of thought emerged from the IMP Group approach to B2B marketing (Axelsson and Easton 1992; Araujo and Easton 1996). The change of emphasis is from a dyadic perspective (factors affecting the relationship between two organizations in a relationship) to a network perspective (factors affecting the relative position of all of the actors within a defined network space). Although the dyadic perspective remains influential, and case studies of specific buyer-seller relationships remain a popular method of empirical investigation in this field, the industrial networks perspective is considered a more comprehensive approach to the analysis of B2B marketing. B2B buyer-seller relationships are necessarily embedded in a network of relationships, and behavior at the level of the individual dyad will often only be comprehensible when viewed through a network lens. For example, a marketing organization may choose not to develop further an apparently successful relationship with a key customer for reasons that are inexplicable at the dyadic level. Network-level effects may provide a rationale for this behavior (for example, if the customer is a major rival to another, more influential customer, so that the further development of one relationship might endanger the other).

Future Directions

The question of how the further development of information technology will affect B2B marketing remains open. This has been an enduring question, investigated for each new

wave of IT. Currently, much debate surrounds the role of social media (such as Facebook, Twitter, and LinkedIn) in B2B marketing. With each new wave of IT, similar questions arise, notably whether the pattern of relationships and networks in B2B markets will be fundamentally disrupted by new forms of interaction, and whether new media will be complements to or substitutes for personal interaction in B2B relationships and networks.

In addition, lively debate continues around the question of strategy within industrial networks. On one hand, it has been argued that meaningful strategic action within complex networks is virtually impossible, since no actor can hope to predict the outcome of its actions because they are mediated by the actions of countless other actors. On the other hand, business managers in B2B organizations clearly believe that they are engaged in strategic decision-making that makes a difference to business performance.

Cross-References

► [Inter-organizational Networks](#)

► [R&D Networks](#)
► [Supply Chain Networks](#)

References

- Araujo L, Easton G (1996) Networks in socioeconomic systems: a critical review. In: Iacobucci D (ed) *Networks in marketing*. Sage, Thousand Oaks, pp 63–107
- Axelsson B, Easton G (1992) *Industrial networks: a new view of reality*. Routledge, London
- Ford D (ed) (2001) *Understanding business marketing and purchasing*. Routledge, London
- Ford D, Håkansson H (2006) IMP – some things achieved: much more to do. *Eur J Mark* 40(3/4):248–258
- Håkansson H, Snehota I (1995) *Developing relationships in business networks*. International Thompson Business Press, London

Recommended Reading

- Brennan R, Canning L, McDowell R (2011) *Business-to-business marketing*. Sage, London
- Ford D, Gadde L-E, Håkansson H, Snehota I (2011) *Managing business relationships*. Wiley, Chichester
- Håkansson H, Snehota I (1990) No business is an Island: the network concept of business strategy. *Scand J Manag* 14(3):177–200