TPC-W

A Benchmark for E-Commerce

Daniel A. Menascé • *George Mason University* • *menasce@cs.gmu.edu*

Scalability is one of the many challenges in designing an e-commerce site. Providers need to know how the site's performance — measured in terms of response time and throughput — varies as more and more users access it. Will performance degrade significantly beyond a given load level, or will the site deliver acceptably even as the load surges? What is the maximum number of transactions that can be processed per second? Can we upgrade the site in a straightforward way (by adding more servers, for example, or by replacing existing servers with more powerful ones) to support higher traffic volumes, or is an architectural change required?

When choosing an e-commerce site's hardware and software configuration, we need to know how a specific combination of Web servers, commerce servers, database servers, and supporting hardware will handle a desired load level. Benchmarks let us compare competing alternatives.

Researchers have extensively studied workloads on Web sites that provide information and have characterized their performance at the level of HTTP requests.[1,2] My colleagues and I have also conducted studies to understand e-commerce site workloads and to search for invariants that cut across more than one type of e-commerce site.[3] However, there is currently only one available benchmark for e-commerce sites: TPC benchmark W, designed by the Transaction Processing Performance Council (www.tpc.org). In this tutorial, I will explore TPC-W's main features, advantages, and limitations.

Characterizing Workload for E-Commerce

As illustrated in Figure 1 (next page), a benchmark specification deals with three main components: the system to be evaluated, called the *system under test* (SUT); the *workload* submitted to the SUT — specifying the type of requests as well as the intensity of the workload; and one or more *performance metrics* obtained by some type of monitoring and evaluation of the SUT.

In e-commerce, customers interact with the site through *sessions*, which are sequences of consecutive requests to execute e-business functions (`search`, `add to cart`, `pay`, and so on) during a single visit to the site. One way to capture the navigational pattern within a session is through the customer behavior model graph (CBMG),[4] which describes how users navigate through the site, which functions they use and how often, and the frequency of transitions from one e-business function to another.

The example CBMG in Figure 2 (next page) shows that customers can be in several different states — `home`, `browse`, `search`, `select`, `add`, and `pay` — and might transition between them as indicated by the connecting arcs. (The numbers on the arcs represent transition probabilities.) Although not explicitly represented in the figure, transitions to the `exit` state are indicated by arrows that leave one state and do not lead to another. For example, the probability that a user will move to the `exit` state from the `browse` state is 0.15.

We can detect several patterns of user behavior at a site. Some visits come from occasional buyers, for example, who spend time "window shopping" but very seldom buy anything. On the other hand, frequent buyers might visit a site knowing what they want, so that they can select one or more items, order, and pay for them with just a few clicks. Different customer behavior patterns generate different loads on the IT resources of a site.

TPC-W: An Overview

TPC-W mimics the activities of an e-tailer, a bookstore in this case, in a controlled Internet commerce environment. The benchmark's main features include

- Multiple online browser sessions,
- Dynamic page generation with database access

 1089-7801/02/$17.00 ©2002 IEEE

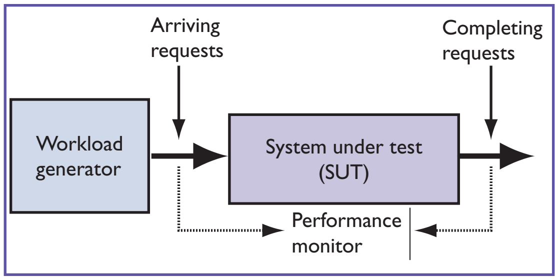

Figure 1. Benchmark components. A workload generator submits requests to the system under test (SUT), whose performance is evaluated.

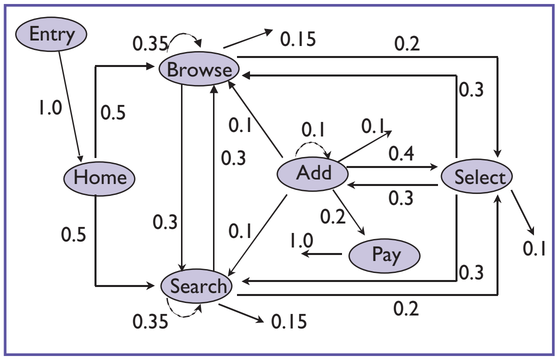

Figure 2. Example customer behavior model graph (CBMG). The numbers on each arc indicate the probability that a user will move from one state to another. Arrows that leave a state without leading to another show a move to the `exit` *state.*

and update,

- Authentication through secure sockets layer (SSL) version 3 or transport layer security (TLS),
- Payment authorization through an emulated payment gateway emulator (PGE), which is not part of the SUT,
- Databases consisting of multiple tables with a wide variety of sizes, attributes, and relationships,
- Enforcement of ACID properties (atomicity, consistency, isolation, and durability) on database transactions, and
- Online transaction execution resulting in contention on data access and update due to concurrency.

TPC-W handles scalability by establishing a relationship, called the *scale factor*, between the number of concurrent sessions and the size of the store, measured in terms of the number of items in the inventory.

System Under Test

As specified by TPC-W, an SUT comprises all Web servers, commerce servers, database servers, load balancers, internal networks, and network interfaces required to implement the emulated e-commerce application. The SUT can also use caching products, as long as they are commercially available and maintain data consistency.

The TPC-W database consists of a minimum of eight tables.

- *Customer*. Customer personal and session data.
- *Address*. Customer address data.
- *Country*. Country name and exchange rate.
- *Order*. Order information, including total amount and shipping information.
- *Order_line*. One order line data per order.
- *CC_Xacts*. Credit card transaction data.
- *Item*. Description of each item in the inventory.
- *Author*. Author data.

The entity relationship (ER) diagram in Figure 3 illustrates the TPC-W database's logical design. The database must support look up, insert, and update functionality as well as the ability to commit and roll back transactions.

Workload Generation

TPC-W generates requests to the SUT using a group of emulated browsers (EB). These EBs act like users sending and receiving HTML content via browsers using HTTP and TCP/IP over a network connection. As described later, the SUT's size and scaling factor determine the number of EBs used for a given test. (The number remains constant throughout an experiment.)

Each EB submits a series of requests within a user session. The workload generated by the EBs is specified by the navigational patterns within a session, represented by a CBMG, and the workload intensity, which is specified by the number of EBs and the *think time* – the interval between when an EB receives the last byte of a result page and when it submits the first byte of its next request.

Table 1 (page 76) shows the 14 unique pages that TPC-W specifies in the full CBMG. These pages are divided into two categories: *browse* and *order*. The browse pages typically consist of nonsecure requests with few transaction-processing requirements, whereas order pages have greater processing requirements in terms of database access or secure transactions. Figure 4 shows a slightly simplified version of the CBMG for TPC-W, omitting the `admin request` and `admin confirm` states used by site administrators. Transitions to the `exit` state (not shown in the figure) can

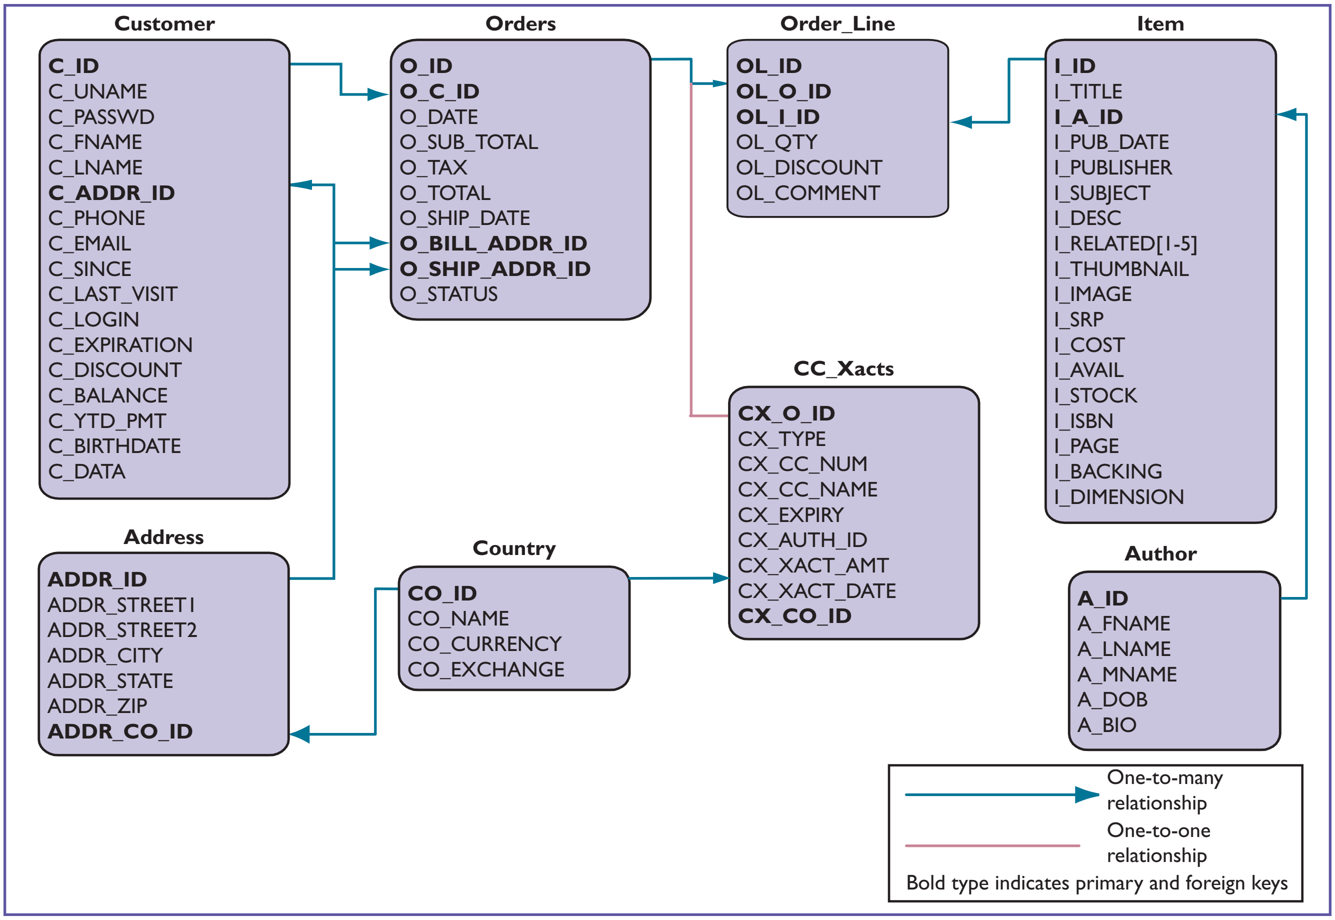

Figure 3. Entity relationship diagram for the TPC-W database.

occur from any state except `entry`.

Figure 4 does not show transition probabilities in the CBMG because TPC-W defines three different mixes of Web interaction types based on the ratio of browse-to-buy activities: primarily shopping, browsing, and product ordering. Table 1 shows the percentage of browse and order interactions for each type of interaction mix. TPC-W enforces the indicated percentages by specifying the transition probabilities for the CBMG of each Web interaction mix type. The values in Table 1 can be obtained from these probabilities using the methods I have described elsewhere.[4]

In a user session, the EB sends the SUT multiple requests, each separated by a unique think time derived from an exponential distribution with a mean between seven and eight seconds, truncated at 10 times the mean value. TPC-W defines a user session's duration as the time elapsed between the EB's first executed transaction and the current time. The *user session minimum duration* (USMD) is based on an exponential distribution with a mean of 15 minutes, truncated at 60 minutes. At the end of a session, the EB closes any open SSL sessions and TCP connections and starts a new

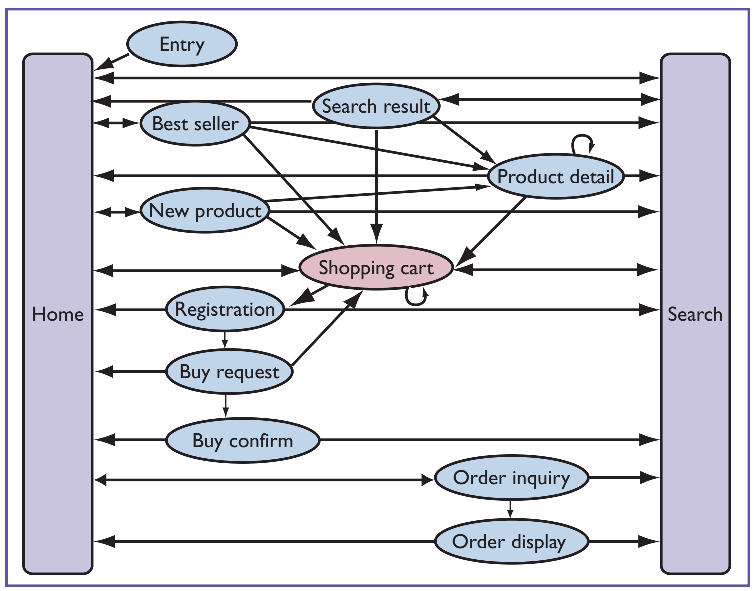

Figure 4. CBMG for TPC-W. The graph shows possible transitions between states during a TPC-W user session. No transition probabilities are indicated.

Table 1. Web interaction mixes in the CBMG. (percent)

Web interaction	Browsing mix (WIPSb)	Shopping mix (WIPS)	Ordering mix (WIPSo)
Browse	**95.00**	**80.00**	**50.00**
Home	29.00	16.00	9.12
New products	11.00	5.00	0.46
Best sellers	11.00	5.00	0.46
Product detail	21.00	17.00	12.35
Search request	12.00	20.00	14.53
Search result	11.00	17.00	13.08
Order	**5.00**	**20.00**	**50.00**
Shopping cart	2.00	11.60	13.53
Registration	0.82	3.00	12.86
Buy request	0.75	2.60	12.73
Buy confirm	0.69	1.20	10.18
Order inquiry	0.30	0.75	0.25
Order display	0.25	0.66	0.22
Adm. request	0.10	0.10	0.12
Adm. confirm	0.09	0.09	0.11

Table 2. Top TPC-W results for catalogs of 100,000 items

Rank	Company	WIPS	$/WIPS
1	A	6,045	76.67 US$
2	B	10,439	106.73 US$
3	C	7,554	136.80 US$
4	D	6,272	195.59 US$

(as of 4 April 2002)

user session. Each EB must generate a new USMD for each new session.

Metrics

TPC-W measures two main metrics: Web interactions per second (WIPS) during a shopping mix and cost-per-WIPS, which is the ratio between the SUT's total price (including hardware, software, and three-year hardware and software maintenance) and the WIPS value. When reporting a value for the WIPS metric, a vendor must specify the cardinality (number of rows) in the `Item` table. A WIPS value measured for a database with 100,000 items, for example, might be reported as 7,000@100,000. The secondary metrics, WIPSb and WIPSo, measure the number of Web interactions per second during browsing and ordering mixes.

Table 2 shows some actual metrics (with company names omitted) for databases with 100,000 items, as posted at the TPC site on 4 April 2002. The top-ranking system (system A) processed 6,045 WIPS at US$76.67/WIPS, which translates to a total cost of US$464,000. System D, on the other hand, costs about 2.65 times as much and performs only 3.8 percent better.

To pass the benchmark, an SUT must meet the required response times — five seconds for Buy Confirm pages and three seconds for Search Requests, for example — at least 90 percent of the time on each of the 14 pages.

Using the Response Time law[5] to establish a relationship between the average response time R and the number of EBs and the average think time $\overline{Z}$, we get

$$R = \frac{\text{No.EBs}}{\text{WIPS}} - \overline{Z}\,. \tag{1}$$

Using $\overline{Z} =$ 7 seconds, we would get an average response time of 1.27 seconds for 50,000 concurrent users on system A (1.27 = 50,000 / 6,045 – 7). System D would provide an average response time of 0.97 s for the same number of concurrent users and average think time.

TPC-W places upper and lower bounds on the value of the WIPS metric that can be reported. Let us examine these bounds by rewriting equation 1 as

$$\text{WIPS} = \frac{\text{No. EBs}}{R + \overline{Z}}\,. \tag{2}$$

We can use equation 2 to determine the maximum possible theoretical WIPS value by setting the response time equal to zero and $\overline{Z}$ equal to 7 (the minimum allowable average think time). We thus find that

$$\text{WIPS} < \frac{\text{No. EBs}}{7}\,. \tag{3}$$

On the other hand, if the response time is too high, the WIPS value will be very small, according to equation 2. To prevent overscaling the SUT, the reported WIPS must be at least 50 percent of the maximum theoretical value from equation 3. The WIPS value must thus satisfy the following constraint:

$$(\text{No. EBs} / 14) < \text{WIPS} < (\text{No. EBs} / 7). \tag{4}$$

Scalability

TPC-W is designed with scalability in mind. The idea is to maintain the ratio between the transactional load placed on the site, as indicated by the number of concurrent EBs, and the initial database size, as represented by the cardinality of the various database tables. Figure 5 indicates how we can obtain each table's cardinality as a multiple of the cardinality of another table or, in the case of the `Customer` table, as a multiple of the number of EBs. The cardinality of the `Country` table is fixed at 92 rows. The cardinality of the `Item` table, TPC-W's scaling factor, can be one of the following values: one thousand, ten thousand, one hundred thousand, one million, or ten million rows.

According to Figure 5, for example, if 100 EBs are used:

- the `Customer` table should have 288,000 rows,
- the `Order` table should have 259,200 rows,
- the `Address` table should have 576,000 rows,
- the `Order_Line` table should have 777,600 rows, and
- the `CC_Xacts` table should have 259,200 rows.

If the database has 100,000 items, then there should be 25,000 authors. So, as the online store supports a higher number of concurrent users, the size of the database that supports the store operation has to scale up accordingly.

Limitations

As with any benchmark, TPC-W does not accurately represent some specific applications. An auction site's workload characteristics could be much different from an online bookstore's, for example, as the auction site might exhibit surges in requests as closing times near for specific items.

TPC-W is oriented toward business-to-consumer e-commerce; it does not model business-to-business transactions well. TPC-W also ignores the presence of robots in the workload of actual e-business sites and is designed for workloads composed of requests generated by human beings who expect HTML pages as a result. As Web services become more prominent,[6] Web service providers will also need benchmarks designed for users that are computer programs rather than humans.

TPC-W results are obtained in a controlled Internet environment, in which the EBs are typically connected to the SUT via very high speed switches (at 1 Gbit per second, for example). When users access Web sites through WANs and low-speed

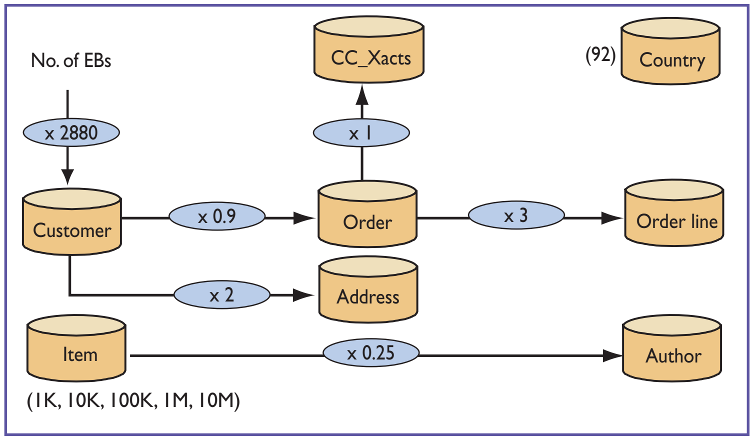

Figure 5. Database scalability. We can compute the cardinality of the various database tables as a function of the number of EBs and items in the catalog.

connections, resources at the Web site (such as TCP connections and threads) remain tied up for much longer periods, and site throughput decreases.

Despite its limitations, however, TPC-W is very well designed. It tests many of the important elements of most e-commerce applications, and e-commerce solution providers should seriously consider using it.

References

1. V. Almeida et al., "Characterizing Reference Locality in the WWW," *Proc. IEEE Conf. Parallel and Distributed Information Systems*, IEEE CS Press, Los Alamitos, Calif., 1996, pp. 92-103.
2. M. Arlitt and C. Williamson, "Web Server Workload Characterization: The Search for Invariants," *Proc. 1996 ACM Sigmetrics Conf. Measurement and Modeling of Computer Systems*, ACM Press, New York, 1996, pp. 126-137.
3. D.A. Menascé et al., "In Search of Invariants for E-Business Workloads," *Proc. Second ACM Conf. Electronic Commerce*, ACM Press, New York, 2000.
4. D.A. Menascé and V.A.F. Almeida, *Scaling for E-Business: Technologies, Models, Performance, and Capacity Planning*, Prentice Hall, Upper Saddle River, N.J., 2000.
5. P.J. Denning and J.P. Buzen, "The Operational Analysis of Queuing Network Models," *ACM Computing Surveys*, vol. 10, no. 3, Sept. 1978, pp. 225-261.
6. D.A. Menascé and V.A.F. Almeida, *Capacity Planning for Web Services: Metrics, Models, and Methods*, Prentice Hall, Upper Saddle River, N.J., 2002.

Daniel Menascé is a professor of computer science, the co-director of the E-Center for E-Business, and the director of the MS in E-Commerce program at George Mason University. He received a PhD in computer science from UCLA and has published five books, including *Capacity Planning for Web Services: Metrics, Models, and Methods* and *Scaling for E-Business: Technologies, Models, Performance, and Capacity Planning* (Prentice Hall, Upper Saddle River, N.J.). He is a Fellow of the ACM and a recipient of the A.A. Michelson Award from the Computer Measurement Group.