3. TPC-C -- The Standard Benchmark for Online transaction Processsing (OLTP)

Francois Raab, Company x.

The Transaction Processing performance Council introduced the TPC Benchmark™ C in August 1992. At the time the TPC had two other OLTP benchmarks, TPC-A and TPC-B. The TPC continued to support and publish results on TPC-A, its first OLTP benchmark until December 1995. TPC-A simulates all the major functions of a simple OLTP system and was, until its retirement by the TPC, accepted by the industry as the leading tool for comparing systems. Since then, TPC-C has replaced it in that role and gained even greater recognition.

You can read a good description of TPC-C in sections 1.1and 1.2 of the specification attached here. Rather than repeat that discussion, I will focus on how TPC-C benchmarks are done, and how TPC-C compares to the earlier TPC-A benchmark.

3.1 TPC-C's Basic OLTP Components

TPC-C was designed to carry over many of the characteristics of TPC-A, the TPC's standard version of DebitCredit. Therefore, TPC-C includes all the components of a basic OLTP benchmark. For the benchmark to be applicable to systems of varying computing powers, TPC-C implementations must scale both the number of terminals and the size of the database proportionally to the computing power of the measured system. To test whether the measured system is fully production-ready, including efficient recovery capabilities, the database must provide what are defined as the ACID properties: atomicity, consistency, isolation, and durability.

To facilitate independent verification of the benchmark's results, the test sponsor must release, in a fall disclosure report, all information necessary to reproduce the reported performance. This performance, which is a measure of the throughput of the system, must be reported, along with the total cost of the system.

The system's total cost is close to the true cost of the vendor-supplied portion of an OLTP system. It includes the cost of all hardware and software components; maintenance costs over five years; and sufficient storage to hold the data generated over a period of 180 eight-hour days of operation at the reported throughput. It only prices one hour-day of recovery log data, since logs can be archived off-line. Table 3.1 summarizes the TPC-C storage requirements.

Table 3.1. Priced storage capacity for 180 days (in 1,000 bytes).

Table Name	Initial Population (KB)	8-hour Increment (KB)	Total Priced (KB) (for 180 days)
WAREHOUSE	0.089	0	0.089
DISTRICT	0.950	0	0.950
CUSTOMER	19,650	0	19,650
HISTORY	1,380	253	45,540
ORDER	720	132.5	23,850
NEW-ORDER	72	0	72
ORDER-LINE	16,200	1,324.9	238,482
STOCK	30,6C0	0	30,600
ITEM	8,200	0	8,200

The basic characteristics shared TPC-C characteristics can be summarized as follows:

- Multi-user benchmark that requires a remote terminal emulator (RTE) to emulate a population of terminals,
- Database size that scales with the throughput of the system under test (SUT),
- Requires support of the ACID properties of transactions,

-
- Requires publication of a full -disclosure report, and
-
- Requires reporting of both performance (transaction throughput) and price/performance (system price divided by throughput) metrics.

TPC-C involves a mix of five concurrent transactions of different types and complexity that are executed either on-line or queued for deferred execution.

Having multiple transactions of different types compete for system resources stresses many components of the measured system and is one of the most significant extensions that TPC-C added to the basic TPC-A OLTP benchmark model. A second significant extension was to increase the complexity of the database's structure. The TPC-C database is composed of nine types of tables with a wide range of row sizes and cardinalities. Therefore, there is greater diversity in the data manipulated by the five transactions and there is greater database contention. The input data to the TPC-C transactions include some of the basic characteristics of real-life data input. For example, popular items are ordered more frequently than other items.

In modeling more realistic environments, TPC-C reduced a number of artificial limitations commonly found in other benchmarks. For example, to promote the use of fully functional terminals or workstations and screen management software, TPC-C required all terminal inputs and displays to be usable by human operators. Presentation services were specifed, and all screens must be formatted using labeled input and output fields. The user interface must provide all the common screen-manipulation features, including moving forward or backward through the input fields. Any physical database design technique that can be used to improve the performance of a real-life application, such as transparent partitioning or replication of data, is allowed in TPC-C.

The major characteristics that TPC-C added beyond TPC-A can be summarized as follows:

- Multiple types of transactions of varying complexity
- On-line and deferred execution of transactions
- More complex database structure, resulting in
  - Greater diversity in the data that are manipulated
  - Higher levels of contention for data access and update
- Input data that include basic real-life characteristics, such as:
  - Non-uniform patterns of data access to simulate data hot spots
  - Data access by primary as well as secondary keys
- More realistic requirements, such as:
  - Terminal input/output with full-screen formatting
  - Required support for basic features of users' interface
  - Required application transparency for all database partitioning
- Transaction rollbacks

3.2 Comparing TPC-A and TPC-C Processing Requirements

It is useful to compare the processing requirements for TPC-C with those of TPC-A. There is a long history of TPC-A results going back to 1990. A simplistic approach assumes that an average TPC-C transaction is ten times more complex than the average TPC-A transaction. That ratio was the goal of the original TPC-C designers. If they archieved that goal, a 1 tpsA system would be a 60 tpmA system and a 6 tpmC system. But, giventhat only the New-Order transactions are counted (i.e., only 44 percent of all the TPC-C transactions), a 1-tpsA system is by this simple logic a 2.6 tpmC system (2.6 = 0.44 x 6). These simplistic calculations are wildly wrong.

During the design of the benchmark, the benchmark designers assumed that an average TPC-C transaction would be ten times more complex than the TPC-A transaction. In other words, the path length ratio, X, was assumed to be around ten. Several vendors published results on the same model computer for both TPC-A and TPC-C. With that data, actual path length ratios can be computed. Given that the number of transactions executed for TPC-C is equal to 2.27 times the reported tpmC, and given that the number of transactions executed per minute for TPC-A is equal to 60 times the reported tpsA, the path length ratio can be computed as $X = (tpsA \times 60)/(tpmC \times 2.27)$. TPC-C results submitted by vendors six month after the release of the new benchmark have shown the path length ratios, X, to be between 4.87 and 9.53 with a mean of 6.9, as summarized in Table 3.2.

Table 3.2 Submitted TPC-C results as of February 93 and the computed ratio between a tpcA transaction and an average tpcC transaction. It suggests that the average TPC-C transaction is about seven times more complex than the TPC-A transaction.

System Name	tpmC	tpsA	Tps C (=tmpC x 2.3 / 60)	Path Length Ratio X
HP 9000 Series 800 Model H40	406.65	74.93	15.4	4.87
IBM AS/400 Model 9404 E10	33.81	9.92	1.3	7.76
IBM AS/400 Model 9406 E35	54.14	14.00	2.1	6.83
IBM AS/400 Model 9406 E70	268.76	54.90	10.2	5.40
IBM RS/6000 Model 570	356.45	128.50	13.5	9.53
			mean	7

The wide range of path length ratio between TPC-A and TPC-C confirms that the two benchmarks exercise different aspects of the systems under test. A given system might be very efficient at processing the short TPC-A transaction but comparatively less efficient at scheduling and processing the five different transactions of TPC-C. Conversely, systems that were never considered to be top performers with TPC-A might be well equipped to process TPC-C transactions.

Hardware and software have improved enormously since these 1993 measurements. While the path length of a TPC-C transaction in 1998 is only a fraction of that of 1993, it is likely that those improvements would have also been reflected in TPC-A if that benchmark had continued to be used. Consequently, and although no public data is available to substantiate this claim, this comparison of processing requirements between TPC-C and TPC-A is still approximately correct.

3.3 A Measure of Business Throughput

The throughput of TPC-C is a direct result of the level of activity at the terminals. Each warehouse has ten terminals, and all five transactions are available at each terminal. The RTE maintains the required mix of transactions over the performance measurement interval. This mix represents the complete processing of an order as it is entered, paid for, checked, and delivered. More specifically, the required mix is defined to produce an equal number of New-Order and Payment transactions and to produce one Delivery transaction, one Order-Status transaction, and one Stock-Level transaction for every ten New-Order transactions.

The tpmC metric reported by TPC-C is the number of New-Order transactions executed per minute. Given the required mix and the wide range of complexity and types among the transactions, this metric more closely simulates a complete business activity, not just one or two transactions or computer operations. For this reason, the tpmC metric is considered to be *a measure of business throughput.*

3.3.1 The OLTP Yardstick

Users of a benchmark's information and results, whether they be members of the press, market researchers, or commercial users, want to be assured that the results they see are valid measures of performance. To meet that demand, the TPC designs its benchmarks to simulate and test systems with all the necessary production-oriented features, including backup and recovery features. In addition, the TPC requires complete documentation of the benchmark's run (the full-disclosure report). These reports are available to any user and are subjected to the TPC's own internal review process. All these requirements help to ensure that users of TPC's benchmark results will receive valid, objective measures of performance.

TPC-C follows the TPC's benchmarking philosophy and methodology in all the above respects, but it also includes multiple transactions and complex operational requirements. TPC-C's performance measurement metric, tpmC, does not simply measure a few basic computer or database operations; rather, it measures how many complete business transactions can be processed per minute. For the past five years this benchmark has proven its ability to give users a more extensive, more complex yardstick for measuring the performance of OLTP systems.

Looking back at the TPC-C results published since 1994, a strong trend emerges. As the graph is Figure 3.1 illustrates, the performance of measured systems has dramatically improved while their cost has dramatically decreased.

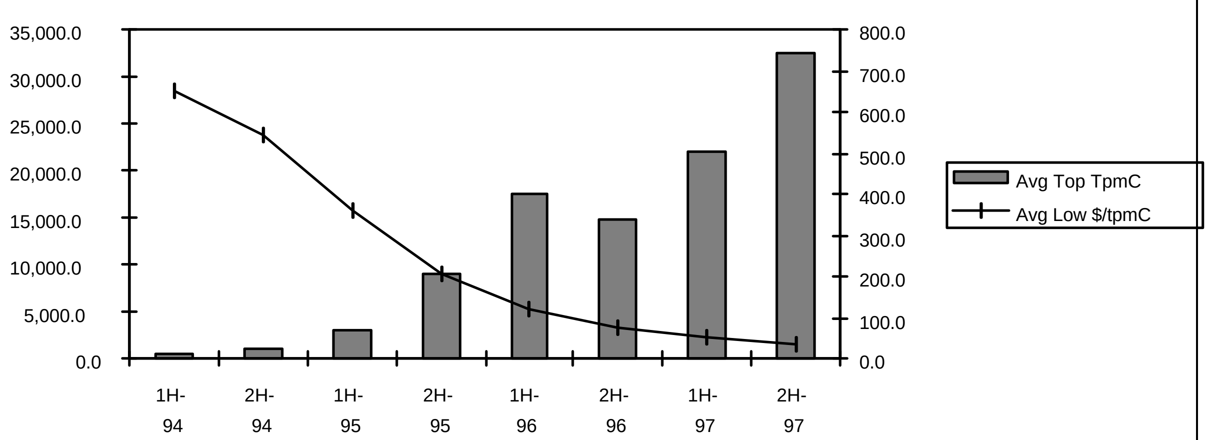

Figure 3.1 - Performance and Price/Performance evolution in TPC-C

Each bar on the graph, scaled on the left axis, represents the average between the five top performance results published during a six-month period. Averaging the five best results shows the performance trend. It lessens the impact of the very best result in its period. The curve of price performance, scaled on the right axis, is also plotted by averaging the five lowest price/performance results within a given six-month period.

The intense competition among vendors, has accelerated the trend illustrated in Figure 3.1. The net result in only four years is a sixty fold increase in performance for a twenty fold decrease in price.

3.3.2 The Evolution of TPC-C

One of the difficulties often met by benchmarks is that the tight framework they use to insure their integrity eventually becomes a limit to technological innovation. During the last five years TPC-C has proven to be remarkably adept at allowing the use of new and creative technologies within its existing framework. This ability to accept evolution was best illustrated by the introduction of a web-based user interface. As the Internet was growing in popularity solutions were developed to replace user terminals with web browsers. Vendors were looking for an avenue to showcase the price/performance advantage of that new approach. After careful examination, it became clear that the TPC-C requirements, as they stood, allowed such a radical shift in the end-user connectivity paradigm. Since then, the great majority of published TPC-C results have used a web based user interface.

As with any performance benchmark, TPC-C bares the wrinkles of its long life. It now shows signs of aging. In the early years of the benchmark, vendors improved the software to better handle TPC-C features. Now that most of the “low hanging fruit” has been harvested. But, the pace of performance improvements has not slowed. Performance engineers have shifted to using new hardware (clusters and giant SMPs) and it have also optimized their systems for the benchmark rather than for general applicability to real-life operations.

This shift in focus is one of the basic laws that govern benchmarks. It resulted in the obsolescence of TPC-A and it is the motivation behind the TPC’s on-going effort to revise its benchmarks. In the last five years many small adjustments were made, too many to list here. More substantive changes have also been made since TPC-C Revision 1.0. They are outlined here.

- Revision 1.1, June 1993
    - Clarification and strengthening of isolation requirements. This change put the emphasis on better defining the required levels of isolation and, in particular, the ability to support repeatable reads.
    - Clarification of checkpoint requirements. The maximum interval of sixty minutes between any two checkpoints was clarified.
    - Inclusion of customer spareable components in the pricing. This new pricing model reflected what an increasing number of vendors were offering.
- Revision 2.0, October 1993
    - Addition of General Implementation Guidelines (targeting ‘benchmark specials’). Although apparently obvious, a written rule to disallow ‘benchmark specials’ was not included in the TPC-C specification until this major revision.
    - Addition of new disclosure requirements and numerical quantity summary. The most relevant numerical data were now to be included as part of the executive summary of the benchmark results, rather than scattered around the thick Full Disclosure Report.
- Revision 3.0, February 1995
    - Addition of new transaction monitor requirements for routing transactions. Limited transaction monitor functions are exercised by the benchmark. To avoid the unrealistic implementation of these functions in the operating system, this new requirements was added to require a full featured transaction monitor if such functions were used.
    - Small changes in schema to accommodate future requirements. In anticipation of the benchmark evolving toward the processing image data, new ‘hooks’ were added in the schema. As it turned out, these ‘hooks’ remain unused today.
    - Elimination of the distinctions between WAN and LAN. As the line between WAN an LAN technologies was becoming blurred, this distinction was removed.
    - Exclusion of the emulated terminals from the pricing. Alll vendors were pricing the lowest cost terminal available on the market (the Liberty, Freedom 1). Still, those terminals accounted for a substantial portion of the reported price metric. So, the terminals were excluded from the pricing. This change refocused the price on more crucial components of the configuration.

In late 1997, the TPC initiated a project to examine and actualize all major aspects of the TPC-C specification. This project, which will result in the publication of Revision 4.0 of the benchmark, aims at making substantial changes in the following areas:

- Database layout
    - Orders are made of 45 items (three times more than now).
    - Each customer starts with 3 orders (so the database is much larger).
- New transaction profiles
    - New Customer_Inquiry transaction
    - All New_Order, Payment and Order_Status are preceded by a Customer _Inquiry
    - Payment transactions refence the last previous payment
- New referential integrity requirements
    - Automatic enforcment of constraints
    - Covers all primary-foreign key relationships
- Tighten response time constraints and cycle times
- Tighten recovery requirements encouraging the use of RAID for all data

The above changes are designed to rebalance the read-write ratio toward more reads. They will also reduce the number of terminals and disk spindles necessary to achieve a given level of performance ten fold. The will increase the tpmC path length by a factor of ten. With the publication of Revision 4.0 of TPC-C, scheduled for the

Q3 of 1998, the benchmark should challenge vendors and promote further improvements in real-life technology for a few more years.

The views expressed in this chapter are those of the author and may differ from the official views of the TPC.