
Public Computer Workstation Use: Visualizing Occupancy

Jeremy Buhler
University of British Columbia, Canada

Abstract

Occupancy rates for public computer workstations at the University of British Columbia Library vary considerably depending on the desirability of the work space, convenience of the location, time of year, and type of computer. This paper describes a method of calculating and comparing occupancy to inform decisions about the number and placement of workstations. An online dashboard created with Tableau Public makes it easier to compare use patterns and establishes a foundation for target-based resource management (see <https://googl/SsztxI>). The method can be applied to other resources and services such as laptop and equipment lending programs.

Introduction

The University of British Columbia (UBC) Library maintains 590 public computer workstations distributed across 11 branches. Despite this large investment, in recent years there has been limited systematic review of the library's public computer infrastructure. We know the computers are frequently used, but do we know that we have the right number? Are they in the locations where they can provide the greatest benefit to library users?

In UBC Library's current financial environment these are pressing questions. The budget for public workstations is insufficient to upgrade all computer hardware as often as desired and we know from surveys that aging computers frustrate many UBC Library users. Could UBC Library still meet demand if it maintained fewer computers and invested more in regular upgrades? If not, is demand so demonstrably high that it would be foolish *not* to increase the workstation budget?

This paper describes a way of using data about public computer workstation use to calculate occupancy and help allocate resources where they are likely to provide the greatest return on investment. Emphasis is on the practical application of a simple dataset using tools that are free (Tableau Public) and widely available (Microsoft Excel).

Preparing the data

Library activities are often monitored with rudimentary counts plotted over time: the number of reference questions, gate count, loans, etc. Likewise, computer workstation use is sometimes measured by the number of login sessions or the number of unique users during a time period. One limitation of this method is that it hides fluctuations within the time period that may affect user experience. Total counts flatten peak periods, for example, and do not reveal how often a user might be turned away because all computers are in use. A more nuanced and informative measure of computer workstation activity is occupancy, the number of workstations in concurrent use at any given time.

At UBC Library each login and logoff event on public workstations is timestamped and recorded in a local database along with the computer ID and location. Structuring this data as shown in Figure 1 facilitates visualization and analysis with Tableau software, which in turn can convey a story about workstation activity to a wider audience. Note that:

- *login* and *logoff* events are recorded on separate rows;
- *login* events are identified with a positive number (1), and *logoff* with a negative (-1);
- rows are sorted in chronological order by *timestamp*.

Figure 1: Structure of source data about computer workstation use

Login/ off	Computer ID	Location	Timestamp
1	32	Level 3 Stacks	04/01/2016 8:13
-1	32	Level 3 Stacks	04/01/2016 8:26
1	12	Level 3 Stacks	04/01/2016 8:49
1	25	Level 3 Stacks	04/01/2016 8:53
-1	25	Level 3 Stacks	04/01/2016 9:00
1	14	Level 3 Stacks	04/01/2016 9:02
1	32	Level 3 Stacks	04/01/2016 9:07
1	25	Level 3 Stacks	04/01/2016 9:10
1	7	Level 3 Stacks	04/01/2016 9:13

Once the source data is structured in this way it is relatively easy to calculate additional values, including occupancy, using spreadsheet software such as Excel. UBC Library's analysis of workstation activity requires the three new columns shown in Figure 2:

- *Occupancy level.* The number of workstations in concurrent use at any given time.
- *Minutes at this level.* The amount of time elapsed before the occupancy level changed.
- *Session minutes.* The duration of each user session (calculated on login rows only).

Figure 2. Source data with additional fields calculated in Excel

Login/ off	Computer ID	Location	Timestamp	Occupancy level	Minutes at this level	Session minutes
1	32	Level 3 Stacks	04/01/2016 8:13	1	13	13
-1	32	Level 3 Stacks	04/01/2016 8:26	0	23	
1	12	Level 3 Stacks	04/01/2016 8:49	1	4	49
1	25	Level 3 Stacks	04/01/2016 8:53	2	7	7
-1	25	Level 3 Stacks	04/01/2016 9:00	1	2	
1	14	Level 3 Stacks	04/01/2016 9:02	2	5	31
1	32	Level 3 Stacks	04/01/2016 9:07	3	3	30
1	25	Level 3 Stacks	04/01/2016 9:10	4	3	12
1	7	Level 3 Stacks	04/01/2016 9:13	5	0	91

Occupancy level is calculated as the running total of the *login/off* column. Each login event increases the occupancy level by one, while each logoff decreases it by one. The value shows the number of concurrent sessions at any given time during the reporting period. *Minutes at this level* is the difference between successive rows' timestamps (assuming rows are sorted in chronological order). Calculating *minutes at this level* makes it possible to determine how much of the time a particular location is above 75% occupancy, for example. *Session minutes* records the duration of each session, that is, the time difference between successive login and logoff events on the same workstation. *Session minutes* can be used to compare activity among workstations: does average session length vary depending on workstation position or type? Which workstations are most often in use?

Comparing occupancy levels: time and location

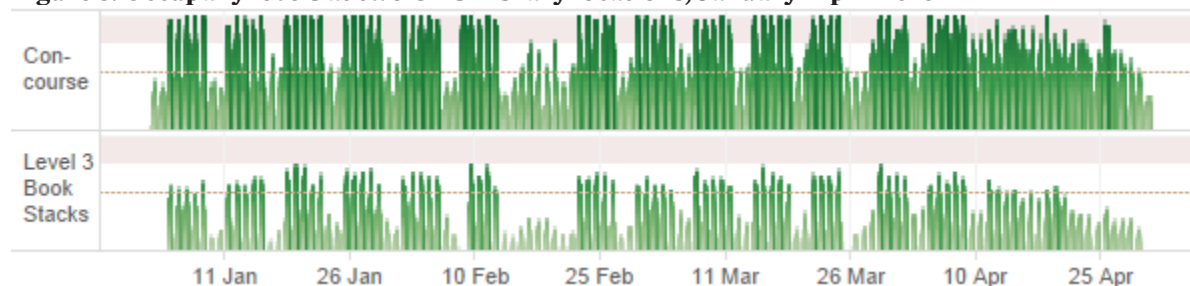
This extended dataset can tell a more compelling story about workstation use. Beyond counting the number of sessions or the number of users, it supports statements like "occupancy rarely exceeds 75% at location X" or "location Z reaches 100% occupancy at least once per weekday during the academic term." From a resource management perspective these statements are more useful but still do not answer the basic "so what?" question. Are these occupancy levels good? Are they bad? Do we need to change anything and, if so, how?

To answer these questions the data must be understood in relation to a comparator or benchmark. The comparison could be against a previous time period, the occupancy at another

location, or a pre-established target. At UBC Library a useful starting point for comparison is location. The UBC Library dataset includes workstation data

for more than 25 locations, and visualizing the data in Tableau software (Figure 3) makes it easier to see, quantify, and communicate differences among them.

Figure 3. Occupancy levels at two UBC Library locations, January–April 2016



The height of the green bars in Figure 3 represents occupancy levels at two locations in UBC Library's Irving K. Barber Learning Centre. The Concourse is a high traffic area on the entrance level with 24 workstations. The Level 3 Book Stacks has 29 workstations and is a quieter, less visited area separated from the Concourse by a flight of stairs. In each row, the dotted line represents 50% occupancy and the bottom of the pink band indicates 75%. Where the green line reaches the top of the pink band, all computers at that location were occupied. It is easy to see at a glance that workstations in the Concourse were often 100% occupied during this period. It is also striking that occupancy never exceeded 75% in the Level 3 Book Stacks. Even at the busiest times only 22 of 29 computers were in use.

Target-based resource management

Comparing activity between locations is an easy first step toward rationalizing the distribution of resources. The data in Figure 3 reveals that in 2016 UBC Library could have moved seven workstations from Level 3 to the Concourse, reducing wait times at one location with almost no impact on the other. But comparison among locations will not always reveal such obvious opportunities. A more robust approach is to establish targets for workstation occupancy. Here is an example of what a workstation occupancy target might look like:

During the busiest four consecutive weeks of the year, weekday occupancy for the busiest six hours of the day should be at least 50% three quarters of the time, and 100% no more than 5% of the time.

This sample target describes an acceptable occupancy range that represents a balance between

user desires (that a workstation always be available) and financial responsibility (that use be high enough to justify the investment). The balance will likely be different for each institution. A library with a limited budget for workstations, for example, might set the occupancy bar higher and accept that during busy periods students will wait for workstations more often. Regardless of the actual target, this approach helps answer questions about allocation as well as about the total number of workstations required.

A target can also account for seasonal variations in activity. Like many universities there is less activity on the UBC campus during the summer. To ensure resources are available when demand is high, the sample target is based on the busiest four consecutive weeks of the year. Further, within those four weeks only the busiest six-hour block is used in calculations. Early mornings and weekends, periods of relative inactivity, are excluded so they will not unintentionally skew results and obscure the demand during peak periods.

Measuring performance relative to targets: Tableau dashboard

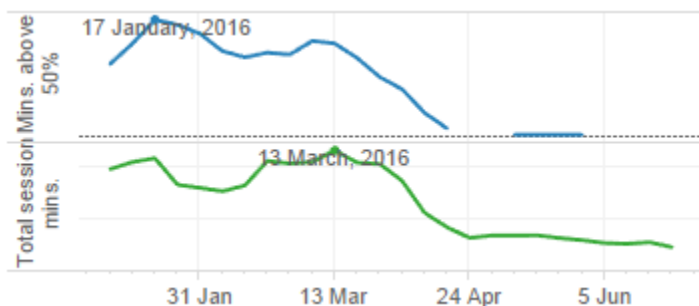
The dataset in Figure 2 contains enough information to measure whether a location meets the target, but not in a format that is easy to read. To make the data more accessible UBC Library used Tableau software to create a dashboard that graphs activity patterns and calculates occupancy for each location. The dashboard viewer can explore the data by selecting among locations, limiting calculations to the desired time period, and estimating the impact of reducing the number of workstations. For reference, a version of the UBC Library dashboard with data from three

locations is available to view and download in Tableau Public (see <https://goo.gl/SsztxI>).

UBC Library’s workstation dashboard is designed to measure performance relative to the sample target statement above. For each location, a bar graph similar to the one shown in Figure 3 provides an overview of activity. Another chart indicates how often occupancy was greater than 25%, 50%, and

75%, and how often it was 100%. Optional filters can adjust the occupancy calculations to include only the busiest four-week period and six-hour weekday block, in alignment with the occupancy target. Since use patterns vary among branches, there are also line graphs to help identify the busiest periods for each location, not just for the library as a whole (Figure 4). Once the date filters have been set, a bar graph indicates whether activity was within the target range.

Figure 4. Weekday activity levels: four week moving average.
Labelled dates indicate the start of the busiest four-week period.



For locations where activity was below the target, the Tableau dashboard also provides a planning tool to help determine the optimal number of workstations. The viewer can estimate the impact of removing workstations by changing the number with an on-screen slider; this, in turn, shows a before-and-after picture comparing actual occupancy levels to what they *would have been* with the smaller number.

Figure 5 shows a before-and-after comparison for UBC’s Level 3 Books Stacks. This location currently has 29 workstations, a number that exceeds demand as was already shown in Figure 3. By adjusting the slider the viewer can find the number of workstations that would have met the target in 2016: *at least 50% full three-quarters of the time and 100% full no more than 5% of the time.*

Figure 5. Level 3 Book Stacks occupancy comparison, peak periods, Jan–Apr 2016

With 29 workstations (current)		With 20 workstations (adjusted)	
■ 75-99% full	0.3% of the time	■ Above capacity	1.5% of the time
■ 50-74% full	40.3% of the time	■ 100% full	1.9% of the time
■ 25-49% full	48.4% of the time	■ 75-99% full	37.2% of the time
■ Below 25% full	10.9% of the time	■ 50-74% full	40.5% of the time
		■ 25-49% full	16.3% of the time
		■ Below 25% full	2.6% of the time

Reducing from 29 to 20 workstations represents a 31% drop at this location. Even so, at 2016 activity levels, the remaining workstations would only have been fully occupied 3.4% of the time during peak periods (the busiest six-hour block on weekdays during the four busiest weeks). Likewise, with 20

workstations the location would have been at least 50% full about 81% of time—compared to 41% of the time at the current level. Presenting this information in a clear, interactive, and accessible form allows a wide audience to explore different scenarios and may facilitate conversations about change.

Despite its practical value, UBC Library's current workstation dashboard remains a rudimentary tool in at least two respects. First, it does not calculate the impact of adding workstations: it can identify *where* additional workstations may be needed but not *how many* are required. Second, the dashboard does not account for changes in workstation activity levels: before-and-after estimates are based on past data and will be skewed in environments where activity levels are trending up or down. These limitations could almost certainly be addressed by someone with the requisite mathematical and software skills. In the meantime even this limited tool stands to improve UBC Library's understanding of activity levels and the management of its resources.

The combination of data and dashboard supports target-based resource management at UBC Library in several practical ways. It is currently being used in a review of the library's computer workstations, but the approach also applies to other resources where session or loan start- and end-times are available. At

UBC Library the method also informed decisions about a laptop lending program, and plans are underway to apply it to equipment lending programs.

Conclusion

Evidence-based management of UBC Library's public computer workstations will help improve the library's return on investment in this area. Using accessible tools, the relatively simple dataset described in this paper can help libraries understand their computer use in terms of occupancy rather than the number of sessions, providing a more nuanced picture of activity. Combining visual reporting and monitoring tools like Tableau with local occupancy targets can help libraries make better decisions about workstation numbers and placement, taking some of the guesswork out of planning and making it easier to measure performance over time.

—Copyright 2017 Jeremy Buhler