Lesson Plan

Titus Barik

Software Engineering as a Human Activity
March 16, 2011

Topic: Substantiating Programmer Variability

Lesson Objectives

In a previous lecture, we examined the paper, "Exploratory Experimental Studies Comparing Online and Offline Programming Performance "by Sackman et. al. This paper is often cited in support of the fact that there are "orders of magnitude" differences between the best and worst programmers.

However, this paper contains several internal validity issues, yet, subsequent papers that identify these problems are often ignored.

In this lesson, we will examine the following papers:

  Substantiating Programmer Variability by Bill Curtis
  Programmer Variability by Thomas E. Dickey
  Programmer Performance and the Effects of the Workplace
    by Tom DeMarco and Tim Lister

At the end of the lesson, the students should be able to:

critically evaluate the validity issues in the Sackman paper given this additional evidence.

explain how long-tailed distributions (such as Pareto) can have an effect on statistical analysis

provide evidence of externalities that affect programmer variability

show how replication studies can be used to validate or invalidate prior research

Outline

Place lecture outline on white board for students:

  - brain storming activity
  - relationship between papers
  - bingo game on required reading
  - interpretation of Sackman data set
  - counterpoints by Curtis and Dickey
  - explanation of Pareto distribution
  - alternative explanations by DeMarco
  - group activity: analysis of data by DeMarco

Introduction [EFFECTIVE QUESTIONING]

To begin this lecture, start with a 30 second - 1 minute brain storming session. Here, half the group should try to identify positive metrics that are associated with "good" programmers. The other half should try to identify negative metrics that are associated with "bad" programmers.

They will do this brain storming activity individually. Then, each student is asked what their more important criteria is for either category. Typically, students will identify factors such as efficiency, good comments, documentation, and so on.

Yet, the Sackman paper studies only one aspect of variability: debugging performance.

Relationship Between Papers

With so many papers, it can be somewhat confusing. Draw a diagram of the relationship between all of the papers on the white board.

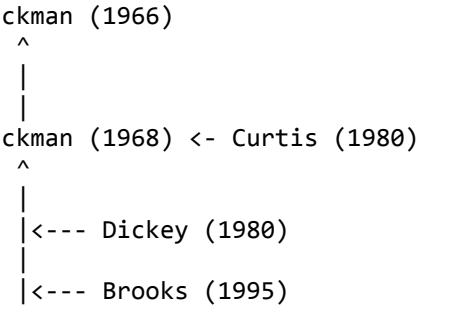

When most individuals cite the 28:1 variability in programmer performance, they cite the 1968 summary paper as evidence. However, the full details are actually found in the 1966 Sackman paper.

Curtis presents an argument that programmer variability exists, but cannot replicate the 28:1 ratio. He does a replication study where the same language is used.

Dickey provides a critical analysis of the 1966 paper, and demonstrates that several validity issues exist with the Sackman paper.

DeMarco provides an alternate explanation in that the work place is more important than the actual employee.

Bingo Game [ACTIVE LEARNING]

Students are now split into groups of 3-4 members. Each member is provided

with a 3x3 bingo board, which they can write on. They are given a Bingo Quiz (see attachment) with "fill in the blank" and "true false questions".

The teams are given 14 questions, but only 9 spaces are available on the bingo board. The teams will pick any 9 questions, answer them, and place the question number and answer in one of the 9 boxes.

This task should take approximately 5-10 minutes. At the end of the time period, a game of bingo is played.

The instructor randomly picks a question. If the team has that question on their board, they mark the board position. This process continues until a team is able to mark 3 in a row. The team should call "bingo".

The instructor will then check to see if the team answers their selected 3 in a row correctly. If so, they win the game. The instructor may elect to eliminate the team if they are incorrect and continue the game until another team receives Bingo.

Interpretation of Sackman Data

Having refreshed the students' memory on the paper, we can now begin to interpret the data set. Here, many issues with the paper can be identified:

- control group issues (different tasks, time sharing / batch)
- language issues (assembly versus FORTRAN)
- selecting outlier data points (28:1) and generalizing
- and so on.

For this, it is best to use the Dickey paper, and then go through the items one by one and find the original data set in Sackman.

Curtis: replication study, corrects language issue, but then finds 8:1 variability. One programmer was unable to compete the study.

Pareto Distribution

In this section, demonstrate that human activity is sometimes a Pareto distribution or other long-tailed distribution. Use the PowerPoint graph to demonstrate this.

Note that such distributions can have a finite mean but infinite variance.

Give examples of Pareto distributions in action:
- wealth
- bug fixing (fixing 20% of top bugs reduces complaints by 80%)

Workplace [ACTIVE LEARNING]

In this section, summarize the DeMarco paper. This paper describes the effects of programmer performance as a function of workplace, and demonstrates that

the workplace has a large effect on programmer performance.

After explaining Figure 3 (spread of performance), split up the group and have each group explain the remaining results. These results are:

- Fig 4. Correlation Between Teammates
- Fig 5. Spread of Defect Counts
- Fig 6. Work Time in Minutes

Each group then spends 5 minutes devising a way to explain their figure.

Each group is then given 1-2 minutes to present their figure.

Conclusions

At the end of the lecture, go through the outline and summarize the main points.

- how Sackman came to be cited so frequently
- the internal validity issues in the paper
- counterpoints to the Sackman study
- long-tailed distributions with respect to human activity

[END]