CS533

Modeling and Performance Evaluation of Network and Computer Systems

The Art of Data Presentation

(Chapters 10 and 11)

WORCESTER POLYTECHNIC INSTITUTE

Introduction

It's not what you say, but how you say it. – A. Putt

- An analysis whose results cannot be understood is as good as one that is never performed.
- General techniques
  - Line charts, bar charts, pie charts, histograms
- Some specific techniques
  - Gantt charts, Kiviat graphs …
- A picture is worth a thousand words
  - Plus, easier to look at, more interesting

Outline

- Types of Variables
- Guidelines
- Common Mistakes
- Pictorial Games
- Special Purpose Charts
- Decision Maker's Games
- Ratio Games

Types of Variables

- Qualitative (Categorical) variables
  - Have states or subclasses
  - Can be ordered or unordered
    - Ex: PC, minicomputer, supercomputer → ordered
    - Ex: scientific, engineering, educational → unordered
- Quantitative variables
  - Numeric levels
  - Discrete or continuous
    - Ex: number of processors, disk blocks, etc. is discrete
    - Ex: weight of a portable computer is continuous

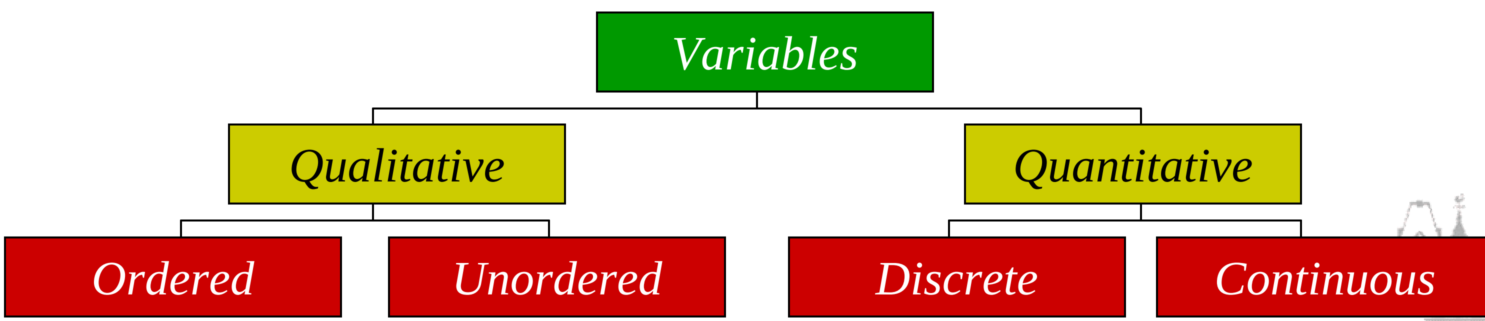

Outline

- Types of Variables
- Guidelines
- Common Mistakes
- Pictorial Games
- Special Purpose Charts
- Decision Maker's Games
- Ratio Games

Guidelines for Good Graphs (1 of 5)

- Again, "art" not "rules". Learn with experience. Recognize good/bad when see it.
- *Require minimum effort from reader*
  - Perhaps *most* important metric
  - Given two, can pick one that takes less reader effort

Ex:

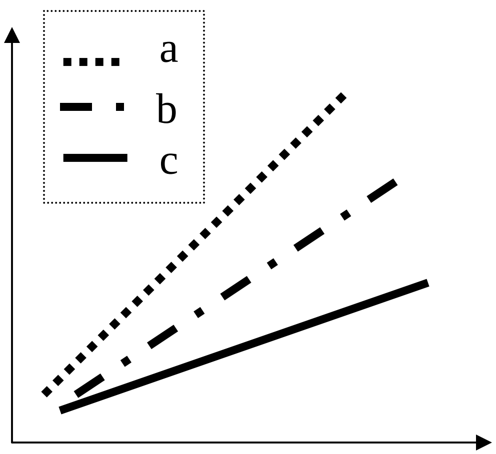

Direct Labeling

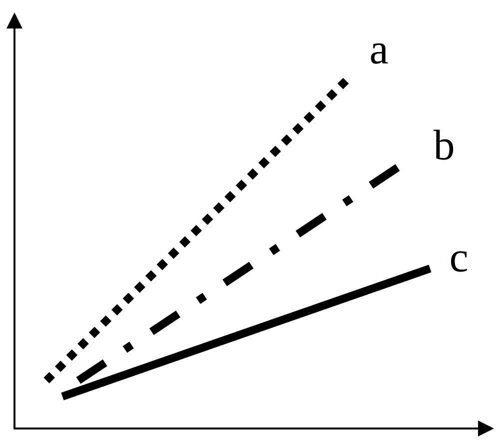

Legend Box

Guidelines for Good Graphs (2 of 5)

- *Maximize information*
  - Make self-sufficient
  - Key words in place of symbols
    - Ex: "PIII, 850 MHz" and not "System A"
    - Ex: "Daily CPU Usage" not "CPU Usage"
  - Axis labels as informative as possible
    - Ex: "Response Time in seconds" not "Response Time"
  - Can help by using captions, too
    - Ex: "Transaction response time in seconds versus offered load in transactions per second."

WORCESTER POLYTECHNIC INSTITUTE

Guidelines for Good Graphs (3 of 5)

- *Minimize ink*
  - Maximize information-to-ink ratio
  - Too much unnecessary ink makes chart cluttered, hard to read
    - Ex: no gridlines unless needed to help read
  - Chart that gives easier-to-read for same data is preferred

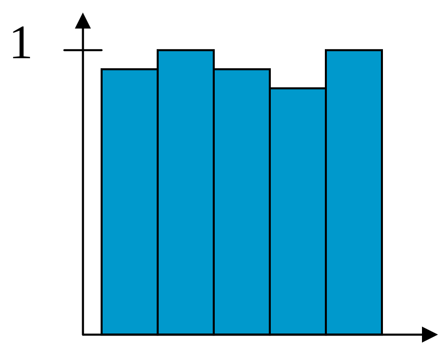

Availability

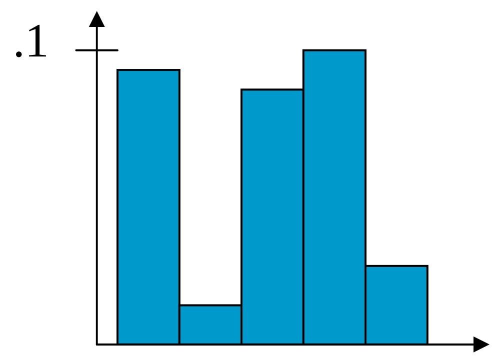

Unavailability

- Same data
- Unavail = 1 – avail
- Right better

Guidelines for Good Graphs (4 of 5)

- *Use commonly accepted practices*
  - Present what people expect
  - Ex: origin at (0,0)
  - Ex: independent (cause) on x-axis, dependent (effect) on y-axis
  - Ex: x-axis scale is linear
  - Ex: increase left to right, bottom to top
  - Ex: scale divisions equal
- Departures are permitted, but require extra effort from reader so use sparingly

Guidelines for Good Graphs (5 of 5)

- *Avoid ambiguity*
  - Show coordinate axes
  - Show origin
  - Identify individual curves and bars
  - Do not plot multiple variables on same chart

Guidelines for Good Graphs (Summary)

- Checklist in Jain, Box 10.1, p. 143
- The more “yes” answers, the better
  - But, again, may consciously decide not to follow these guidelines if better without them
- In practice, takes several trials before arriving at “best” graph
- Want to present the message the most: accurately, simply, concisely, logically

Outline

- Types of Variables
- Guidelines
- Common Mistakes
- Pictorial Games
- Special Purpose Charts
- Decision Maker's Games
- Ratio Games

Common Mistakes (1 of 6)

- *Presenting too many alternatives on one chart*
- Guidelines
  - More than 5 to 7 messages is too many
    - (Maybe related to the limit of human short-term memory?)
  - Line chart with 6 curves or less
  - Column chart with 10 bars
  - Pie chart with 8 components
  - Each cell in histogram should have 5+ values

WORCESTER POLYTECHNIC INSTITUTE

Common Mistakes (2 of 6)

- *Presenting many y-variables on a single chart*
  - Better to make separate graphs
  - Plotting many y-variables saves space, but better to requires reader to figure out relationship
    - Space constraints for journal/conf!

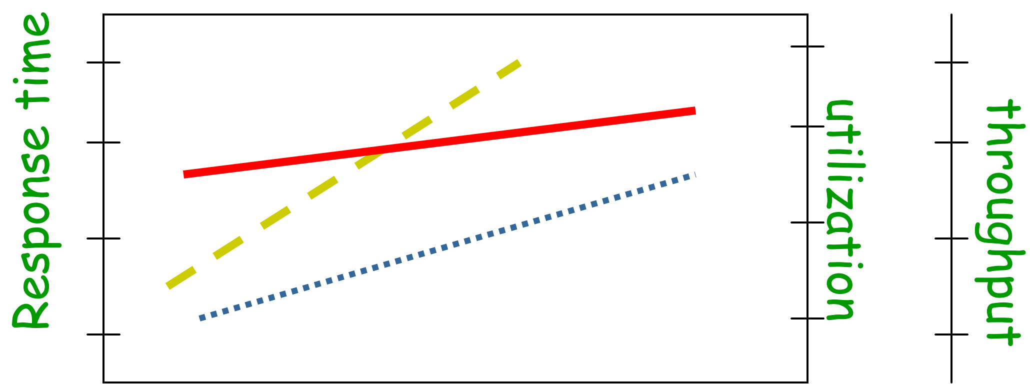

Common Mistakes (3 of 6)

- *Using symbols in place of text*
- More difficult to read symbols than text
- Reader must flip through report to see symbol mapping to text
  - Even if "save" writers time, really "wastes" it since reader is likely to skip!

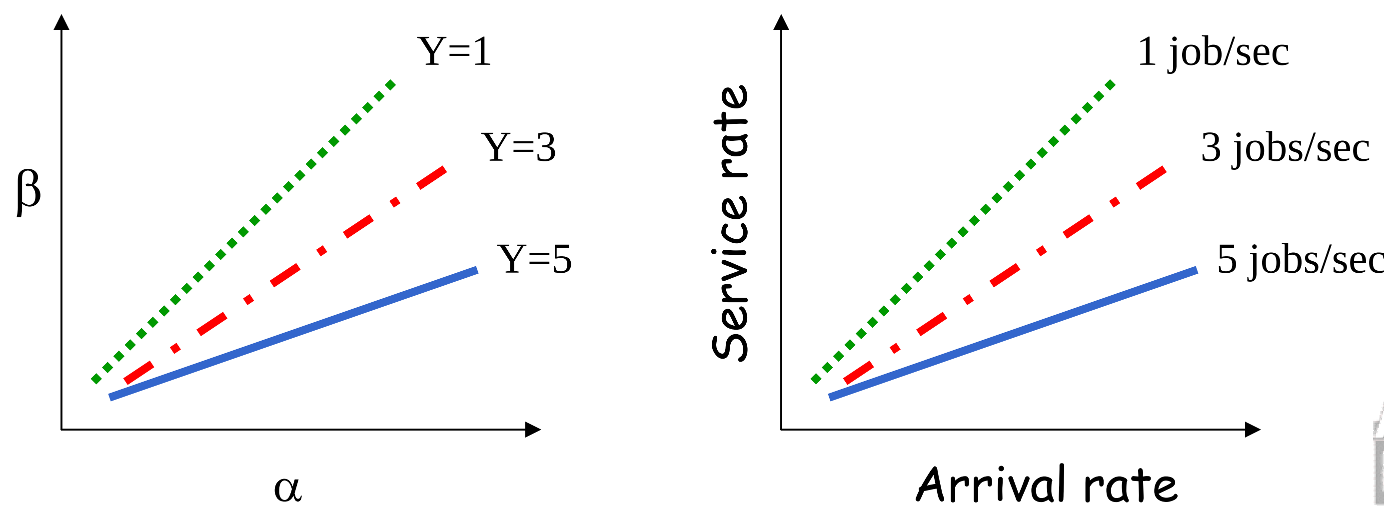

Common Mistakes (4 of 6)

- *Placing extraneous information on the chart*
  - Goal is to convey particular message, so extra information is distracting
  - Ex: using gridlines only when exact values are expected to be read
  - Ex: "per-system" data when average data is only part of message required

Common Mistakes (5 of 6)

- *Selecting scale ranges improperly*
  - Most are prepared by automatic programs (excel, gnuplot) with built-in rules
    - Give good first-guess
  - But
    - May include outlying data points, shrinking body
    - May have endpoints hard to read since on axis
    - May place too many (or too few) tics
  - In practice, almost always over-ride scale values

Common Mistakes (6 of 6)

- *Using a Line Chart instead of Column Chart*
  - Lines joining successive points signify that they can be approximately interpolated
  - If don't have meaning, should not use line chart

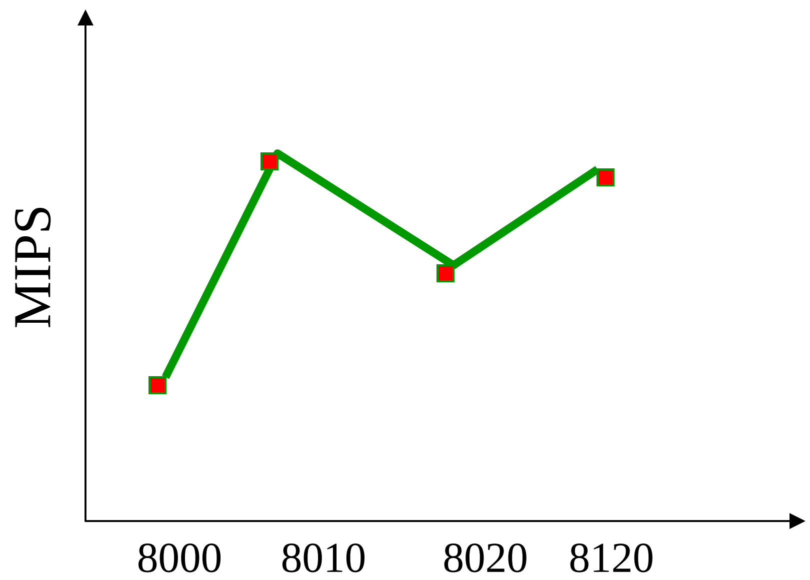

- No linear relationship between processor types!
- Instead, use column chart

Outline

- Types of Variables
- Guidelines
- Common Mistakes
- Pictorial Games
- Special Purpose Charts
- Decision Maker's Games
- Ratio Games

Pictorial Games

- Can deceive as easily as can convey meaning
- Note, not always a question of bad practice but should be aware of techniques when reading performance evaluation

Non-Zero Origins to Emphasize (1 of 2)

- Normally, both axes meet at origin
- By moving and scaling, can magnify (or reduce!) difference

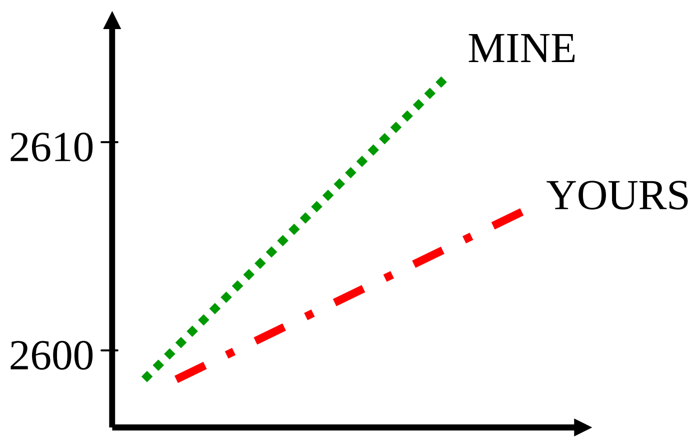

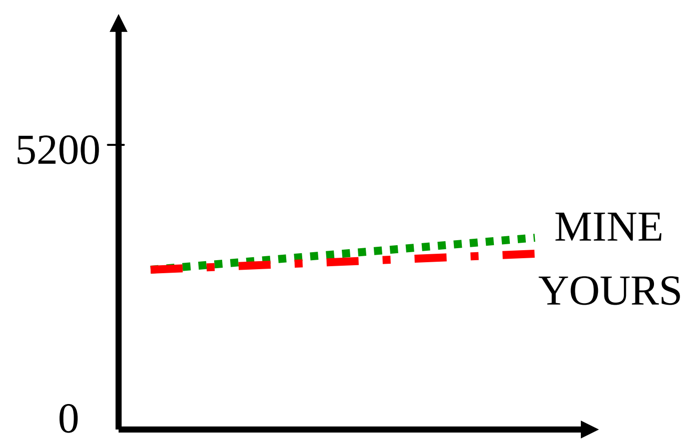

Which graph is better?

Non-Zero Origins to Emphasize (2 of 2)

- Choose scale so that vertical height of highest point is at least ¾ of the horizontal offset of right-most point
  - Three-quarters rule
- (And represent origin as 0,0)

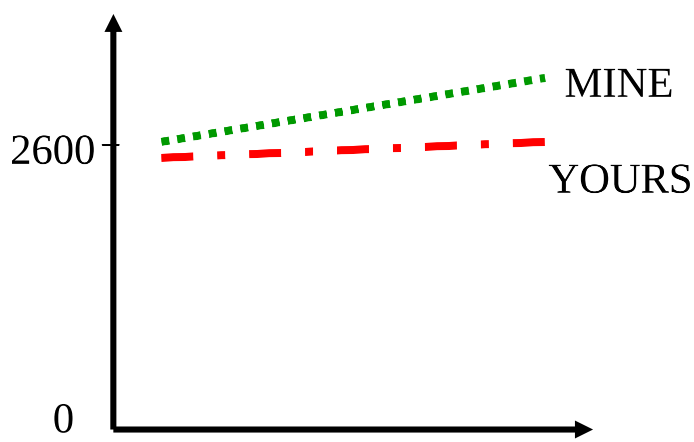

Using Double-Whammy Graph

- Two curves can have twice as much impact
  - But if two metrics are related, knowing one predicts other … so use one!

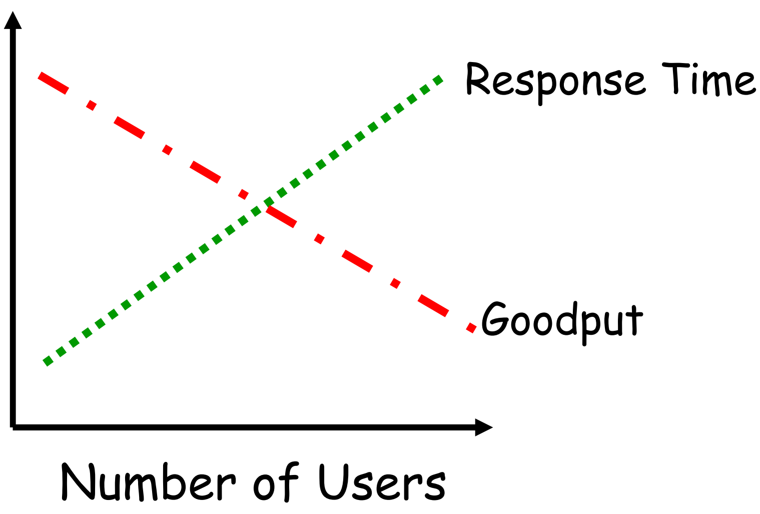

Plotting Quantities without Confidence Intervals

- When random quantification, representing mean (or median) alone (or single data point!) not enough

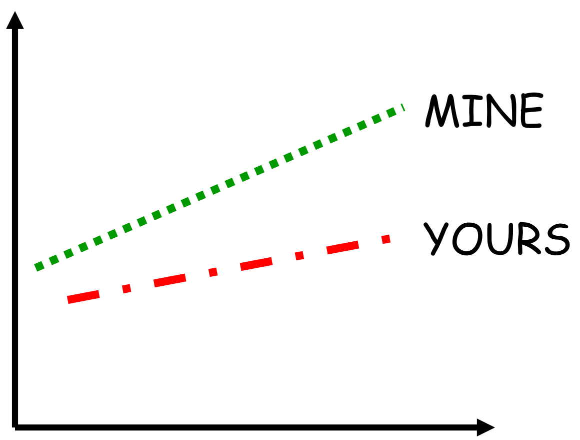

(Worse)

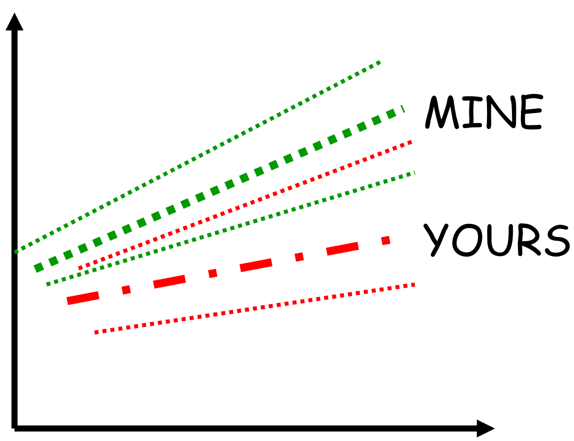

(Better)

Pictograms Scaled by Height

- If scaling pictograms, do by area not height since eye drawn to area
  - Ex: twice as good → doubling height quadruples area

(Worse) (Better)

Using Inappropriate Cell Size in Histogram

- Getting cell size "right" always takes more than one attempt
  - If too large, all points in same cell
  - If too small, lacks smoothness

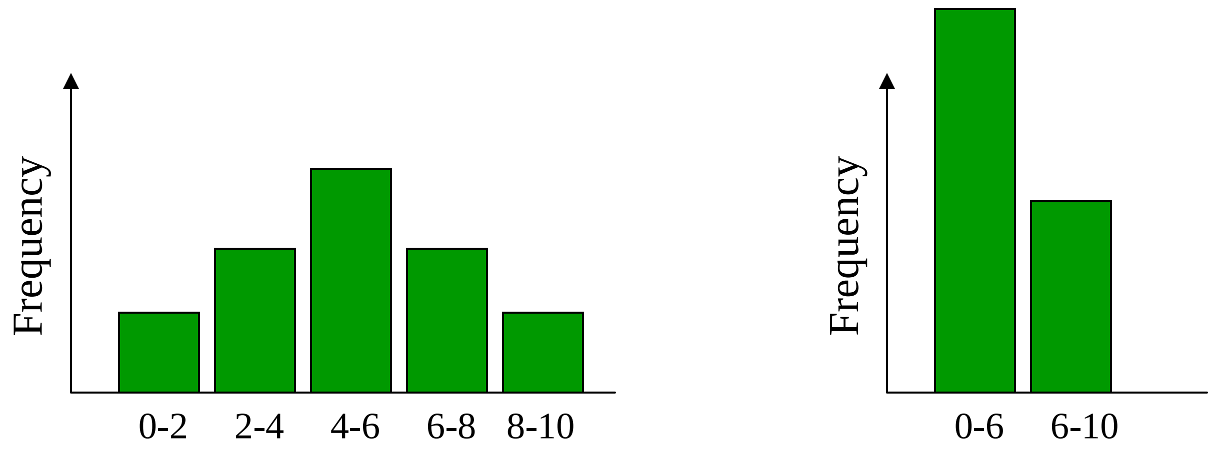

Same data. Left is "normal" and right is "exponential"

Using Broken Scales in Column Charts

- By breaking scale in middle, can exaggerate differences
  - May be trivial, but then looks significant
  - Similar to "zero origin" problem

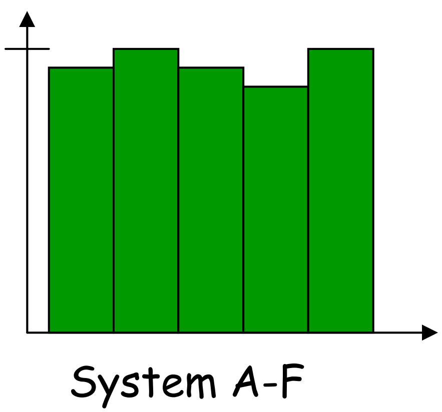

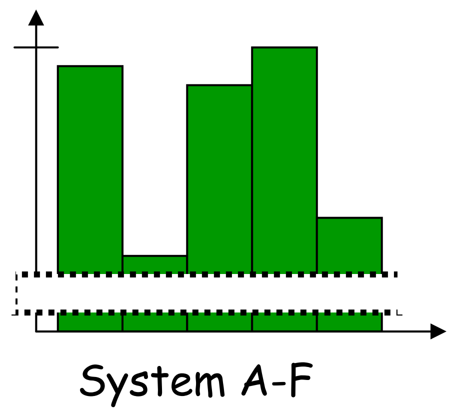

Outline

- Types of Variables
- Guidelines
- Common Mistakes
- Pictorial Games
- Special Purpose Charts
- Decision Maker's Games
- Ratio Games

Scatter Plot (1 of 2)

- Useful in statistical analysis
- Also excellent for huge quantities of data
  - Can show patterns otherwise invisible
  - (Another example next)

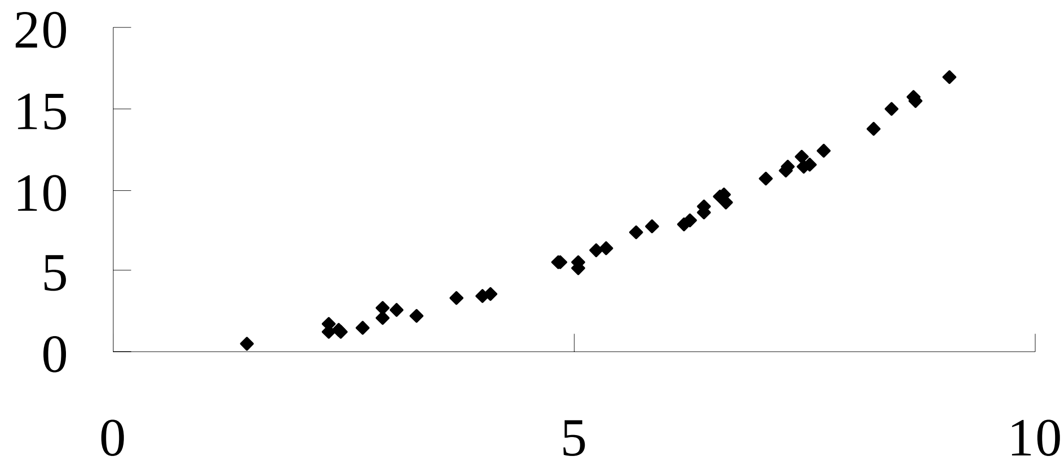

(Geoff Kuenning, 1998)

Scatter Plot (2 of 2)

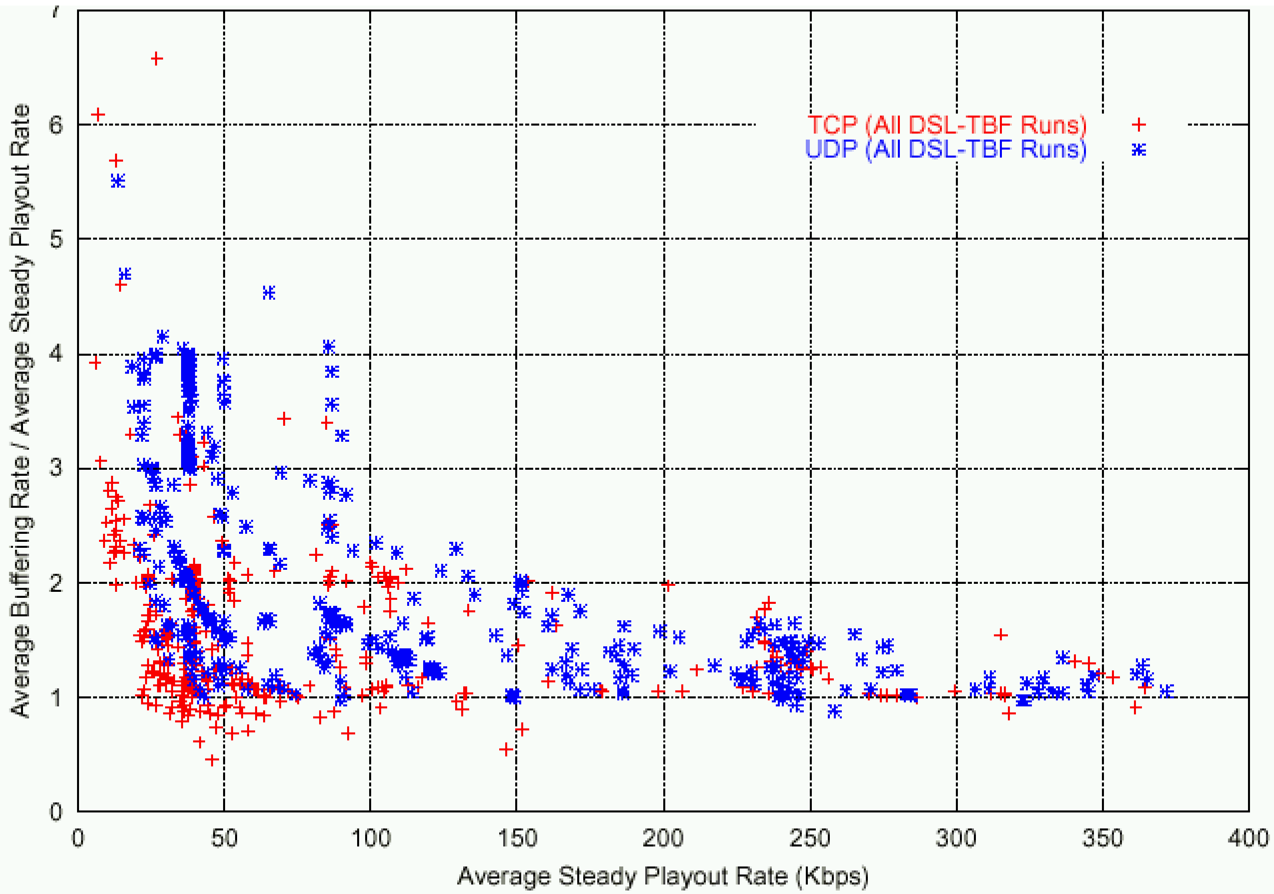

Fig. 10. Ratio of Average Buffering Rate to Average Steady Playout Rate versus Average Steady Playout Rate (All Runs)

Box and Whisker's Plot

- Shows (range, median, quartiles) all in one:

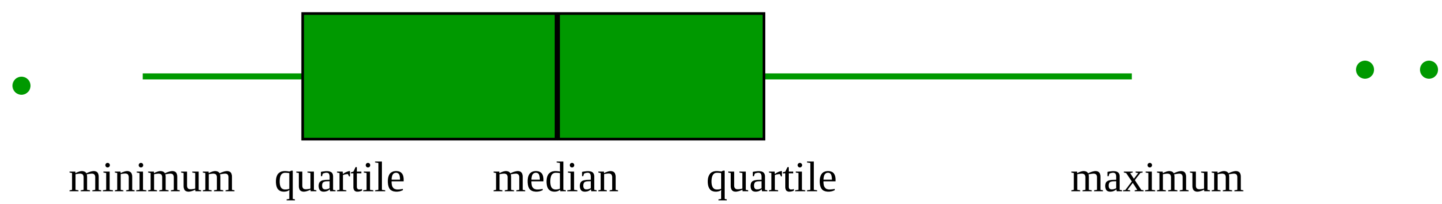

- Variations:

(Geoff Kuenning, 1998)

WPI
WORCESTER POLYTECHNIC INSTITUTE

Stem and Leaf Display

- "Histogram-lite" for analysis w/out software
- Scores: 34, 81, 75, 51, 82, 96, 55, 66, 95, 87, 82, 88, 99, 50, 85, 72

```
9| 6 5 9
8| 1 2 7 2 8 5
7| 5 2
6| 6
5| 1 5 0
4|
3| 4
```

Gantt Charts (1 of 2)

- Resource too high is bottleneck
- Resource too low could be underutilization
- Want mix of jobs with significant overlap
  - Show with Gantt Chart
- In general, represents Boolean condition … *on* or *off*. Length of lines represent busy.

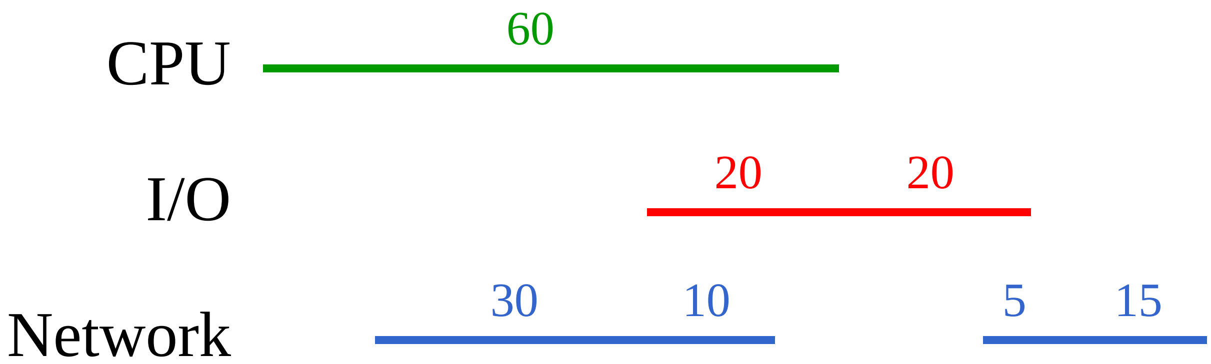

(Example 10.1 Page 151 Next)

Gantt Charts (2 of 2) - Example

A	B	C	D	Time
0	0	0	0	5
0	0	0	1	5
0	0	1	0	0
0	0	1	1	5
0	1	0	0	10
0	1	0	1	5
0	1	1	0	10
0	1	1	1	5

A	B	C	D	Time
1	0	0	0	10
1	0	0	1	5
1	0	1	0	0
1	0	1	1	5
1	1	0	0	10
1	1	0	1	10
1	1	1	0	5
1	1	1	1	10

- Pattern is A and not-A first
- Rest are not-R and R

(Jain, Example 10.1
Page 151)

WPI
WORCESTER POLYTECHNIC INSTITUTE

Kiviat Graphs (1 of 2)

- Also called "star charts" or "radar plots"
- ½ are HB, ½ are LB
- Note, don't have to have all at 100% can be "10% busy", say
- Useful for looking at balance between HB and LB metrics ("Star" is best)

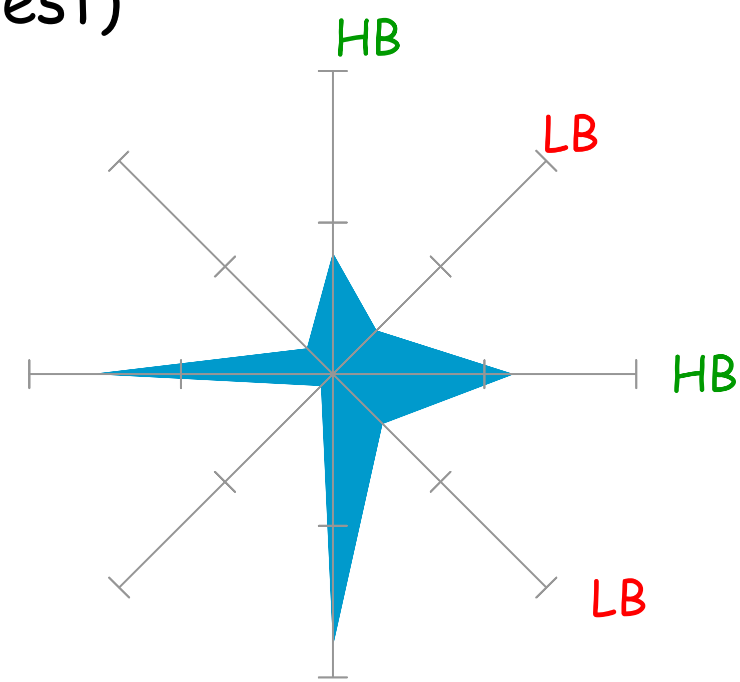

(Geoff Kuenning, 1998)

Kiviat Graphs (2 of 2)

- Commonly occurring shapes can be useful to characterize system
  - "CPU keelboat" (CPU bound) (fig 10.19)
    - (A shallow, covered riverboat for freight)
  - "I/O wedge" (I/O bound) (fig 10.20)
  - "I/O arrow" (CPU + I/O) (fig 10.21)
- Most for data processing, but can be applied to other systems. Ex: network

HB Metrics	LB Metrics
App throughput	App response time
Link utilization	Link overhead
Router utilization	Router overhead
% packets arrive	# duplicates
% implicit acks	% packets with error

Outline

- Types of Variables
- Guidelines
- Common Mistakes
- Pictorial Games
- Special Purpose Charts
- Decision Maker's Games
- Ratio Games

Decision Maker's Games

- Even if perf analysis is correctly done, may not convince decision makers (boss, conference referees, thesis advisor…)
  - Box 10.2, p. 162 has list of reasons
- Most common:
  1) "*More analysis*." This is <u>always</u> true. Does not mean analysis done is not valuable.
  2) "*Alternate workload*". Since based on past, can always be questioned as good future workload
- Lead to endless discussion ("rat holes"). Can "head off" criticism by stating this.

Outline

- Types of Variable
- Guidelines
- Common Mistakes
- Pictorial Games
- Special Purpose Charts
- Decision Maker's Games
- Ratio Games

WORCESTER POLYTECHNIC INSTITUTE

Ratio Games (Ch 11)

If you can't convince them, confuse them. – Truman's Law

- A common way to play games with competitors
- Two ratios with different bases cannot be compared or averaged
  - Doing so is called "ratio game"
- Knowledge of "ratio games" will help protect ourselves, avoid doing

Games with Base System

- Beware!
  - Normalize each system's performance for each workload by system A and average ratios
  - Normalize each system's performance for each workload by system B and average ratios

System	Work-load 1	Work-load 2	Average
A	20	10	15
B	10	20	15

System	Work-load 1	Work-load 2	Average
A	2	0.5	1.25
B	1	1	1

Games with Ratio Metrics

- Choose a metric that is ratio of two other metrics. Power = thrput/respTime

Network	Thrput	RespTime	Power
A	10	2	5
B	4	1	4

- Suggests that A is better.
- But maybe it should be:

$$power = thrput/respTime^2$$

→ $Power_A = 2.5$, $Power_B = 4$

Games with Relative Performance

- Metric may be specified but can still get ratio game if two are on different machines
- MFLOPS, System X-Y, accelerators A-B

Alternative	Without	With	Ratio
A on X	2	4	2.00
B on Y	3	5	1.66

(Base systems are different)

Games with Percentages (1 of 2)

- Percentages are really ratios, but disguised
  - So can play games

Test	A Runs	A Passes	A %	B Runs	B Passes	B %
1	300	60	20	32	8	25
2	50	2	4	500	40	8
Total	350	62	18	532	48	9

- A is *worse* under both tests

  → but it looks *better* in Total!

Games with Percentages (2 of 2)

- Percentages
  - Have bigger psychological impact
    - 1000% sounds bigger than 10-fold
  - Are great when both original and final performance are lousy
    - Ex: payment was $40 per week, is now $80
- When used, base should be *initial*, not *final* value
  - Ex: Price was $400, now $100
    - Drop of 400%! But that makes no sense

Strategies for Winning Ratio Game (1 of 2)

- (Again, don't do these, just be aware of them so no-one does them to you)
- If one system is better by all measures, a ratio game won't (usually) work
  - Although, remember percent-passes example!
  - And selecting the base also lets you change the magnitude of the difference
- If each system wins on some measures, ratio games might be possible
  - May have to try all bases

Strategies for Winning Ratio Game (2 of 2)

System	Work-load 1	Work-load 2	Base B	Base A
A	20	10	1.25	1
B	10	20	1	1.25

- For LB metrics, use your system as the base
  - Ex: response time
- For HB metrics, use the other system as a base
  - Ex: throughput
- If possible, adjust lengths of benchmarks
  - Run longer when your system performs best
  - Run short when your system is worst
  - This gives greater weight to your strengths

Extra Credit for Next Class

- Bring in one either notoriously bad or exceptionally good example of data presentation
  - The bad ones may be more fun
- From proceedings, technical documentation, newspaper …
- Make copies before class or send to me and I'll make copies
- We'll discuss why good/bad