University of Illinois at Urbana-Champaign
Dept. of Electrical and Computer Engineering

ECE 120: Introduction to Computing

Pareto Optimization*

* *

What's Best When We Have Several Metrics?

As engineers, you will rarely have the luxury of a single metric.

How does one choose between metrics?

Imagine the following …
- You are working as an intern
- designing hardware to execute DNNs (deep neural networks, which may be useful in a variety of tasks).

* *

Example: Compare Designs Based On Area and Delay

Building on your ECE120 knowledge, you have two metrics.
- **Area**, which you have normalized from **1 to 100**.
- **Delay**, which you have also normalized from **1 to 100**.

In both metrics, **smaller is better**.

For a design **X**, **A(X) is the area**, and **D(X) is the delay**.

* *

Which Metric is More Important?

Now imagine that you have two designs, **X** and **Y**.

How do you choose between them?

Which is more important, area or delay?

* *

The Answer Depends on How the Design is Used

The answer **depends on the context** in which your design is used

- datacenter
- laptop
- mobile phone
- car or other vehicle
- space probe
- children's toy

* *

One Option is to Combine Metrics Numerically

How do you make a choice?

One option: **linearize**. Pick some weights

- actually, one weight **W** is enough
- **W** is the **relative importance** of delay compared to area

Then

- for each design **X**
- calculate **M(X) = A(X) + W D(X)**

Choose the design with the smallest M(X).

* *

Relative Importance is Not Easy to Choose in Practice

But how do you pick **W**?

What if you need designs for ALL of those applications?

As an engineer, you may not be in a position to know the right weights!

So **what can you do if you don't know the relative importance of the metrics**?

For two designs, probably just report both to your manager.

* *

Remember Dilbert? Oxygen is Good.

What if you have created 10,000 designs?

- Not by hand, but by using parameters.
- For example, does your design provide hardware for 8-bit, 16-bit, 32-bit, or 64-bit addition?

Do you report all 10,000 to your manager?

Probably not if you want a job offer.

But **what can you do**?

* *

A Design that is Worse in All Metrics is Pareto-Dominated

Pick two designs **X** and **Y**.

What if **A(X) < A(Y)** AND **D(X) < D(Y)**?

Remember that smaller is better for both area and delay.

In such a case, **do you need to report Y**?

No! **X** is better in both metrics.

We say that **design Y is Pareto-dominated by design X**.

If there were **N** metrics, **Y must be worse than some X in ALL metrics to be Pareto-dominated**.

* *

Eliminate Designs that are Pareto-Dominated by Others

You can use the idea of Pareto domination to eliminate designs.

Any design that is Pareto-dominated by any other design **can be discarded**.

Only designs that may be better in some context remain.

The remaining designs form a **Pareto curve** (a Pareto surface for more than two metrics).

* *

A Graph Illustrates Pareto-Dominance (Small is Good)

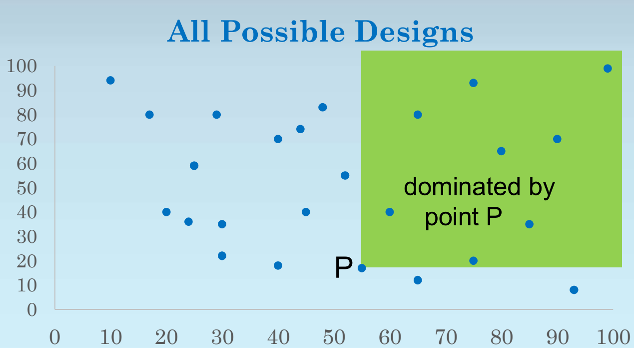

* *

A Pareto Curve (after Discarding Dominated Points)

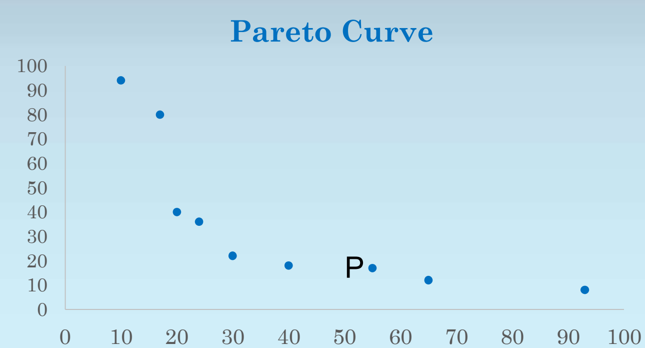

* *

Use of Pareto Curves/Surfaces is Common

In many hardware design environments, engineers run **design-space exploration** tasks (on computers, of course!):

- Given a set of parameters for a design
- Generate hardware for each possible combination of parameters
- Then use Pareto dominance to trim the results down
- And show the engineer the Pareto surface of area, delay, and power consumption.

* *

Want to Learn More about Optimization?

Take ECE490 some day.

Combines theory and practice:

- optimization algorithms,
- Implementations,
- use of libraries to solve problems.