Reliability of Data Storage

RAID 10

The simplest way to ensure data reliability is to "mirror" your hard drive - that is, have 2 paired disks where each is a copy of the other. The only way a pair can fail is if both of the drives fail. This is very fault tolerant as a pair fails only if both drives fail at the same time. If you have more data than a single pair can manage you simply link pairs in sequence, with your array working only all the pairs are working. This kind of setup is called RAID 10 or RAID 1+0.

For example suppose you had 3 full drives of data to worry about . To support them in RAID 10 you would actually need 6 disks consisting of 3 linked pairs. A linked pair fails only when both of its drives fail, and the entire array works only when all 3 linked pairs work:

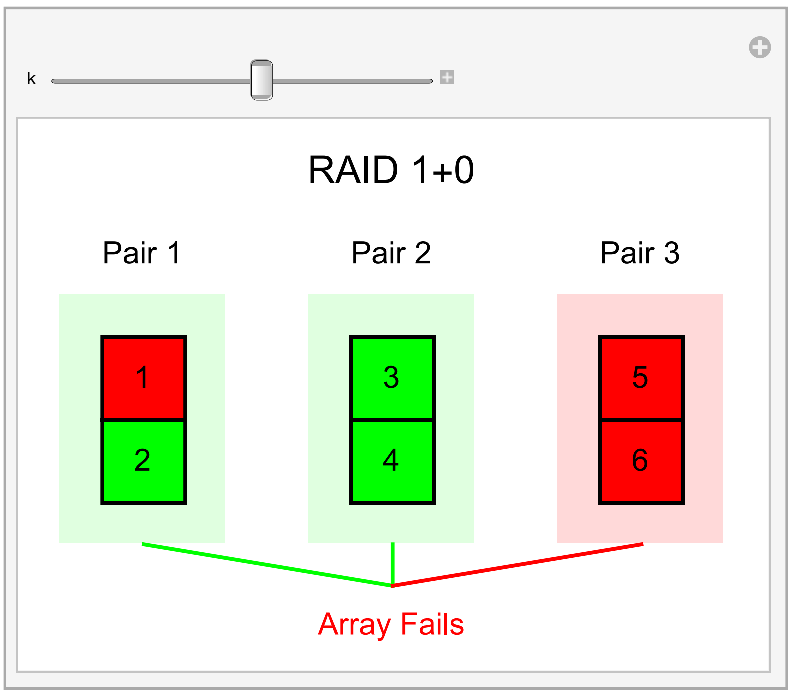

RAID 5

RAID 5 works in a different way to secure data. Essentially rather than pairing drives it requires only a single extra drive - the single extra drive is the sum (parity) of the other 3. So if you had 3 drives plus a pairty drive, if the three drives stored the bits {1,2,7} the parity drive would store a 1+2+7=10. If any one of the drives fail its values can be reconstructed from the parity drive - if the middle drive fails and the bit is {1, X ,7,10} you can see that X must be 2.

This is not as secure as RAID 10 as if any two drives fail the array will fail, but it uses fewer drives as well:

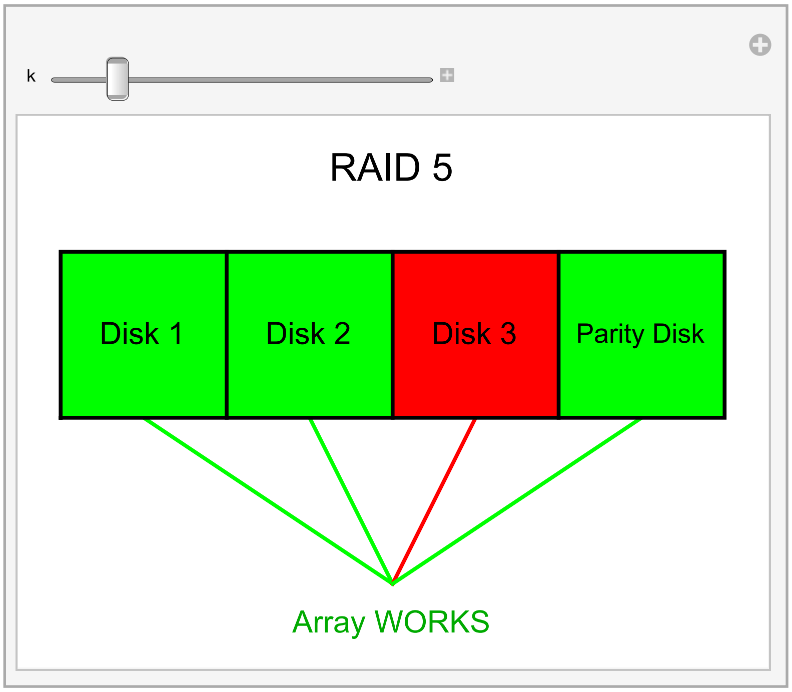

Which is better: RAID 10 or RAID 5?

If you have the same money to spend to build a RAID 10 or RAID 5 array you can get more reliable drives if you go with RAID 5 as you can spend more on each disk. Let's say that for a RAID 10 array you can get 6 cheap disks with a mean lifespan of 5 years; for the same price you can build a RAID 5 array using 4 more expensive disks with a mean lifespan of 7 years:

```
cheapdisk = ExponentialDistribution[1 / 5];
gooddisk = ExponentialDistribution[1 / 7];
```

We can define a RAID 10 array by describing which of its drives must work for the array to work:

```
raid10 = ReliabilityDistribution[ (a ⋁ b) ⋀ (c ⋁ d) ⋀ (e ⋁ f), {{a, cheapdisk},
     {b, cheapdisk}, {c, cheapdisk}, {d, cheapdisk}, {e, cheapdisk}, {f, cheapdisk}}];
```

Now we can ask what is the likeihood that a RAID 10 array will fail in under 3 years:

```
NProbability[ x ≤ 3.0, x ≈ raid10]
```

```
0.494826
```

This is still just better than using 3 drives together with no backups at all:

```
NProbability[ x ≤ 3.0, x ≈ ReliabilityDistribution[
   a ⋀ b ⋀ c, {{a, cheapdisk}, {b, cheapdisk}, {c, cheapdisk}}] ]
```

```
0.834701
```

Likewise we can define a RAID 5 array by describing the cases in which it will work:

```
raid5 = ReliabilityDistribution[ (a ∧ b ∧ c) ∨ (a ∧ c ∧ d) ∨ (b ∧ c ∧ d) ∨ (a ∧ b ∧ d),
   {{a, gooddisk}, {b, gooddisk}, {c, gooddisk}, { d, gooddisk}}];
```

We can see the average time to failure for both types of arrays:

```
{Mean[raid10], Mean[raid5]}
```

$\left\{\frac{7}{2}, \frac{49}{12}\right\}$

```
N[%]
```

```
{3.5, 4.08333}
```

We can also look at the medoan time to failure:

```
{Median[raid10], Median[raid5]} // N
```

```
{3.02753, 3.41122}
```

In both cases the means are noticeably larger than the medians, which is common in distributions with long tails.

Most importantly we can ask what is the chance that a RAID 5 array will last longer than a RAID 10 array:

```
Probability[ x > y, {x ≈ raid5, y ≈ raid10}]
```

$\frac{93\,980\,824}{171\,337\,155}$

```
N[%]
```

```
0.548514
```

So our RAID 5 arrays are on average a bit better than the RAID 10 arrays. We can take a look at what fraction of the arrays survive over time:

```
Plot[ Evaluate[ {SurvivalFunction[ raid10, x], SurvivalFunction[raid5, x]}],
 {x, 0, 10}, PlotLegends → {"RAID 10", "RAID 5"},
 AxesLabel → {"time", None}, PlotLabel → "Fractions of arrays surviving"]
```

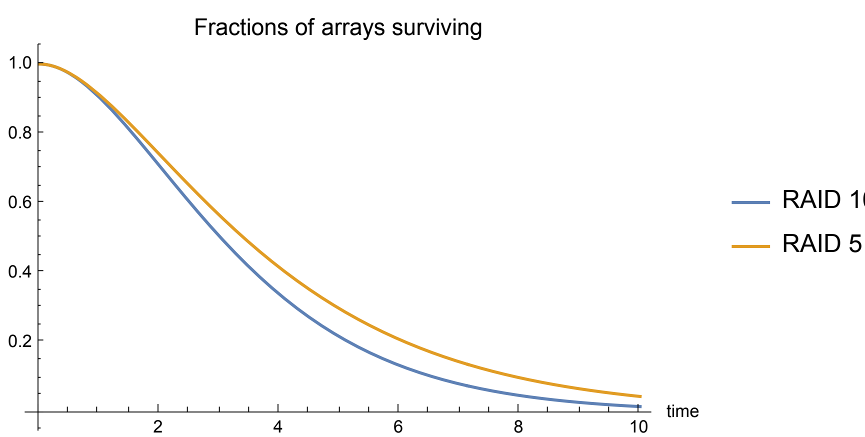

You might ask the question of how the extra lifespan of the “good disks” factors into this analysis. One way to do this is to take advtange of *Mathematica*’s Manipulate command to reconstruct the previous graph for different values of the “good disk” lifespan, say as this lifespan starts at 5 years (same as the “cheap disks”) up to 10 years:

```
Manipulate[cd = ExponentialDistribution[1 / 5];
 gd = ExponentialDistribution[1 / gdm];
 raid10a = ReliabilityDistribution[ (a ⋁ b) ⋀ (c ⋁ d) ⋀ (e ⋁ f),
   {{a, cd}, {b, cd}, {c, cd}, {d, cd}, {e, cd}, {f, cd}}];
 raid5a = ReliabilityDistribution[ (a ⋀ b ⋀ c) ⋁ (a ⋀ c ⋀ d) ⋁ (b ⋀ c ⋀ d) ⋁ (a ⋀ b ⋀ d),
   {{a, gd}, {b, gd}, {c, gd}, { d, gd}}];
 Plot[ Evaluate[ {SurvivalFunction[ raid10a, x], SurvivalFunction[raid5a, x]}],
  {x, 0, 20}, PlotLegends → {"RAID 10", "RAID 5"},
  AxesLabel → {"time", None}, PlotLabel → "Fractions of arrays surviving"],
 {{gdm, 5, "Mean failure time for good disks"}, 5, 10, .1}, SaveDefinitions → True]
```

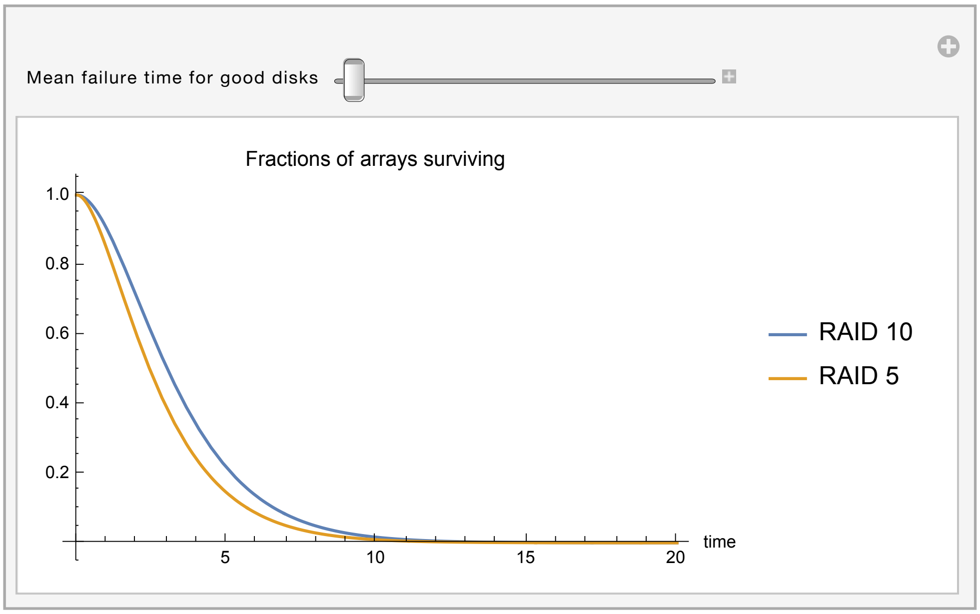

By moving the slider you should be able to see that when the disks used in the RAID 5 array have a mean failure time of around 6-6.1 years the two graphs are almost identical, and any values larger than about 6.5 years clearly puts the RAID 5 array ahead of the RAID 10 array in terms of survival rates.

Author: Dr. Christopher Moretti
Revision Date: August 3rd, 2015.