19:02 Of Coal and Iron

by Manul Laphroaig, Engineer

Gather 'round, neighbors. The Christmas season is behind us, but some cold days still lie ahead, and there's still time for a hearty fireside chat and a pint. And as I raise my pint and think of fireplaces and of stockings hung by the chimneys with care, my thoughts turn to the thing that had to do with all of these and warmed the hearts and limbs of geeks of the ages past: coal.

These days, neighbors, hardly anyone gets coal in their stockings, and the coal-fed heating oven closest to you is likely in that Victorian novel on your bookshelf (unless you are in Berlin, neighbor, in which case coal might still be your winter friend). But this pint of pale ale, at least, is a reminder of the times when coal was something every geek of technology cared about.

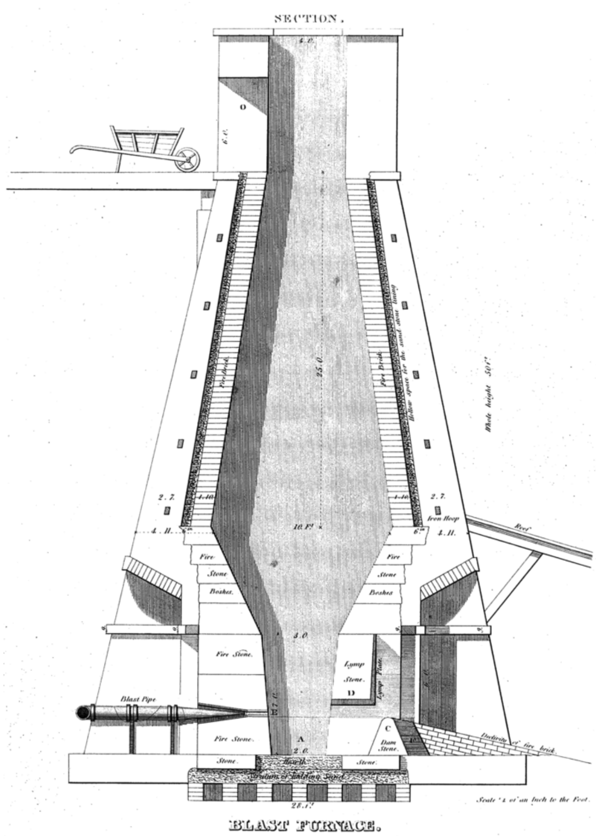

You see, neighbors, pale ale was made possible by the same thing that made the railway and the rest of the Industrial Revolution: coke, which is to coal as charcoal is to wood. Malts used to be dried with wood or peat fires, and that meant smoke and darker malts. Raw coal, although cheaper, could not be used, because hardly anyone likes their beer to smell of sulfur. Coke, on the other hand—once the process for its production got figured out, which in Europe happened in late 16th–early 17th century—was a smokeless fuel. Coke ushered in the era of lighter, "pale" malts, and by the end of the 17th century changed our idea of a neighborly pint. Which was nothing compared to how coke changed the ideas of distance and physical neighboring.

Chances are, neighbor, that you are reading this thanks to the Network of Networks, otherwise known as the Internet, and that a few of your other favorite things also need connectivity. But of course the Internet was not the first physical network of networks. It wasn't even the first network of metal that made the far things and places previously unreachable—except to the very few and at a great expense—reachable on the cheap. That network was the railway, and it would not have happened without coke—and, of course, its best friend, iron.

Just how exciting was that railway network? you might ask. Jules Verne's *Around the World in Eighty Days*, an engraving from which graces this edition's cover, was prompted by the news report that the world's public transport network of railways and steam boat routes was almost complete for circumnavigation, missing just some 140 miles in India. This was the news of the age—and the book became Verne's most popular one, prompting many real-life journeys around the globe.

In Europe the process for smelting iron[2] with coke was figured out around the beginning of the 17th century. The inventor of record, Abraham Darby (also called Abraham Darby the Elder, as his son and grandson of the same name continued

[2] It goes something like this. Iron in nature tends to be all tied up in oxides, but, given the choice, oxygen really prefers carbon. So if you heat it all up in a scene that's just right, like a blast furnace, iron gets reduced out. Just think of $2Fe_2O_3 + 3C \rightarrow 4Fe + 3CO_2$ as nature's distracted boyfriend meme—except that iron and carbon remain best friends, and the intricacies of their relationship have been the subject of countless bedside books of the geeks of the early 1900s, such as H.M. Howe's *Iron, Steel, and Other Alloys*, which you'll find in the feelies. This is true steampunk, neighbors, and truer romance of the elements is yet to be written, despite the fact that the iron obtained through smelting was called "pig iron."

to further the relationship of coal and iron), was inspired by seeing coke being used in malt ovens. Before then, smelting iron required charcoal. This was good enough for swords and similar items of expensive blacksmithing, but rather limited the amount of iron one could smelt.

Not only trees take a while to grow, and Britain's timber was already in scarce supply by 1700s, but charcoal doesn't pile up so well with iron ore. So coke both saved the trees and allowed for much larger blast ovens, resulting in much cheaper iron, in much larger quantities. It was initially not as good as hand-hammered *wrought iron*, but it was good enough, and there was enough of it to be poured into casts, at a fraction of the cost. So much, in fact, that one could make buildings, bridges, and railroads out of it.

In some 50 years cast iron made its way from pots and pans to what we now call *critical infrastructure*. It went from the first coke-powered blast furnaces set up by Abraham Darby in 1709 to the icons of the Industrial Revolution such as the Crystal Palace of the London's Great Exhibition of 1851 and the great cast iron bridges such as the 2.75-mile long Tay Bridge of 1879 across the Firth of Forth.

The time cast iron took to get adopted for major infrastructure projects was not accidental, as chemical impurities of coke were still larger and less controllable than those of charcoal, and defects such as those caused by gas bubbles were inherent in the casting process. Also, cast iron is hard and compresses well, but is brittle, because it still contains a fairly large amount of carbon and slag, in a heterogeneous alloy structure, which is one of the many subtle and fascinating phases of the relationship between iron and carbon. So cast iron was not without its downsides.

But the choice between infrastructure you can afford right now and the one you can't is pretty easy, and so is the employer's choice between labor that can be had on the cheap and the expert labor that's scarce. The march of the cheap technology cannot be stopped—think of Javascript and IoT.

Who said IoT? Neighbor, what is that bottle over on that shelf right next to the divine nectar of Islay? Indeed, it is the Glenrothes scotch, and so suitable for the story I am going to tell, for the first of its kind, they say, was distilled on the same day it happened. Give me a generous pour, neighbor, and take another, for the story is not a happy one.

This is the story of a great feat of infrastructure, the engineer knighted for it, and not surviving it by even a year. This is the story of the Tay Bridge.

Beautiful Railway Bridge of the Silvery Tay!
With your numerous arches and pillars
in so grand array,
And your central girders,
which seem to the eye
To be almost towering to the sky.
The greatest wonder of the day,
And a great beautification to the River Tay,
Most beautiful to be seen,
Near by Dundee and the Magdalen Green.
– William McGonagall, 1879

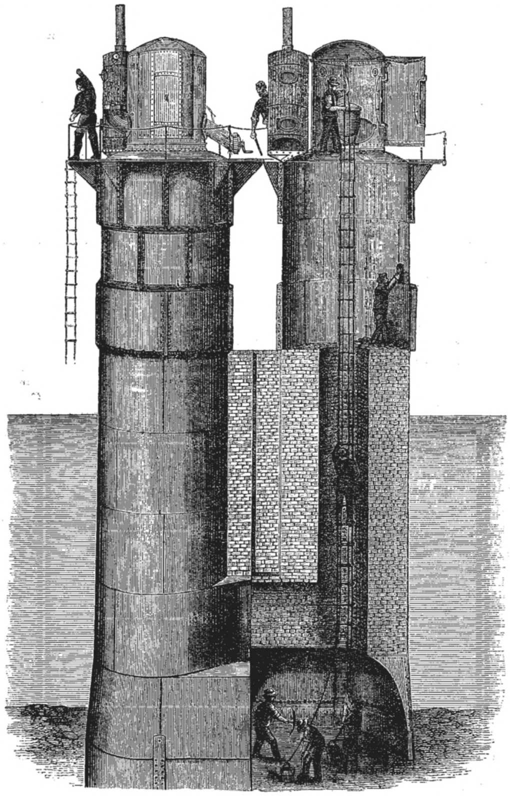

The Tay Bridge was designed by Sir Thomas Bouch, the inventor of the railway ferry and the lattice girders of railway bridges, the design you can still see on the Manhattan Bridge, San Francisco–Oakland Bay Bridge, and many smaller bridges. The famous Eiffel Tower uses the same lattice principle.

The Tay Bridge exemplified the engineering approach that brought Sir Thomas to fame and knighthood: that it was the duty of the engineer to accomplish his work without extravagance and waste, making it solid and substantial, but only just as solid and as substantial as required by the circumstances. Through Sir Thomas' designs, many clients in need of railway connectivity were able to actually afford it. In his projects he used the cheapest technologies, like cast iron columns for bridges, and used advice on the wind loads from experts such as the Astronomer Royal—whom we'd now call data scientists or perhaps climate scientists—to get the safety allowances just right for the specific tasks rather than the excessive one-size-fits-all. This approach brought him fame, and, eventually, knighthood, a week after Queen Victoria on June 20, 1879, crossed the celebrated Tay Bridge, an engineering marvel of the day and an economical one at that.

The Tay Bridge used an ingenious and cost-effective structural scheme, which combined cast iron columns with wrought-iron cross-bracing. It combined the strengths of the two kinds of materials: the cheapness and hardness of cast iron, and the tensile strength of the more expensive wrought iron. Unlike cast iron, wrought iron could bend without breaking, as the slag in its microstructure was shaped by hammering and rolling (i.e., *working* it, hence *wrought* in its name) into fibers.[3] The wrought-iron braces and tiebars stabilized the open-lattice piers by linking the cast iron columns. The structure had to be light enough to carry the weight of the lattice girders and itself, given the limited support the tricky river bed could offer. The maximum windload observed across the Firth of Forth was taken into account, too, rather than adding an arbitrary allowance.

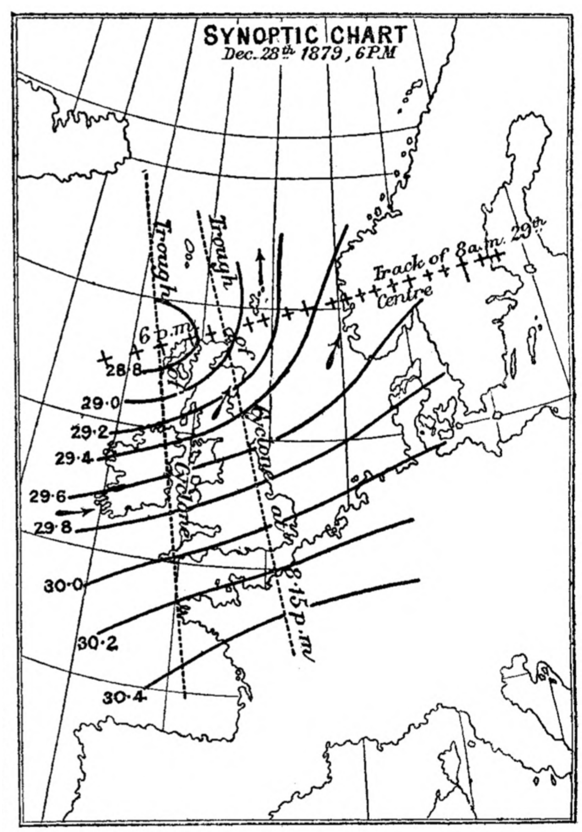

[3]These days, wrought iron is a thing of the past, because mild steel gives the same structural properties without the slag, due to its iron-carbon structure layering of iron allotropes. But at the time steel production still could not compete with wrought iron.

Then, on Sunday the 28th of December 1879, the Tay Bridge collapsed to high winds as a train was passing through it, killing all aboard.

Beautiful railway bridge of the silv'ry Tay
Alas! I am very sorry to say
That ninety lives have been taken away
On the last sabbath day of 1879
Which will be remember'd
for a very long time.
– William McGonagall, 1880

What brought the bridge down? Was it poor design or flaws in the workmanship? An inquiry board set up to investigate the deadly collapse brought to light many things, such as the ingenious practices of the foundry workers to disguise the casting flaws they considered minor by filling them in with a paste of beeswax, iron filings, etc., that appeared to be metal when burnished. Another practice that turned out to be common among moulders was to cast the holes for bolts when casting the columns, rather than drilling them afterwards. This made the holes conical rather than cylindrical, putting more load from the bolt on the narrow edge end, crushing the bolt's thread, allowing extra play for the bolted tiebars, and weakening the overall lattice structure as a result. As the windload calculations were traced to the authoritative books of the day and redone, questions were raised whether the wind speeds in the respective formulas were meant to be instantaneous maximal values at a point or average values calculated over time or over the length of a bridge's span, which were smaller.

Sir Bouch was known for designs that optimized costs. The makers of the bridge's columns added their own optimizations to the casting processes: casting bolt holes while the column was cast was much cheaper than boring them afterwards. Bolts, in turn, were cheaper than pins. During the inquiry it transpired that Sir Bouch did not know that the bolt holes were cast as a common practice, while the casters did not think the difference important. In turn, the casters had concerns about the attachment of tying braces, *"knowing how treacherous a thing cast iron is"*, but assumed the engineers knew and compensated for the weaknesses with redundancy.

The bridge as built was the sum of many independent optimizations, from the overall design to lower its weight to the labor of casting its iron columns. All of these optimizations were made in good faith, from the chief engineer down to the foundry foreman and the bridge maintenance inspector, each acting within their normal layers of competence and trusting the judgment of experts in other layers. With so many people involved, layers of engineering abstraction once again became boundaries of competence.

The combined effect of these good faith optimizations was wilder and more deadly than anyone could predict. Although the inquiry board members disagreed on whether the bridge as designed would have stood if its workmanship were perfect or close, it was abundantly clear that continuing the business of cast iron structures as usual was too risky. Several major bridges and viaducts were abandoned and redesigned or condemned and eventually replaced. Cast iron designs gave way to more expensive wrought iron (think Eiffel Tower), and then the steel industry caught up and made wrought iron obsolete.

The stone pier stumps of the original Tay Bridge, though, are still visible next to the new bridge.

BEAUTIFUL new railway bridge of the
Silvery Tay,
With your strong brick piers and buttresses
in so grand array,
And your thirteen central girders,
which seem to my eye
Strong enough all windy storms to defy.
–William McGonagall

And so ends this story of coal, iron, and critical infrastructure, neighbors. But all of this had happened before, and it will all happen again.

Although our networks are not of iron and carbon, we too have had miraculous breakthroughs that, like coke, allowed us to scale them far beyond the limits any sane economist would've thought possible. Our networks and artifacts too are subject to the same real world forces that favor engineering them on the cheap, and our choices of materials by brittleness and the skill needed to work them are eerily similar.

Our boundaries of competence are as strong as ever, and our drive to optimize on both sides of an abstraction boundary is just as disastrous. Nor have we any lack of "evidence-based" expert advice that looks so authoritative in a book or in powerpoint, but may not even use relevant metrics.

Indeed, our hardware has more kinds of Spectres than a Victorian ghost novel.

It is hard to fault the CPU engineers who, in pursuit of affordable performance, introduced the cache. The cache is and will likely remain one of the breakthrough computing inventions that delivered miraculous improvements on a budget, suddenly making the impossibly huge computations actually economical. The cache allowed programmers to be effective without honing the finer skills of understanding and hand-optimizing the memory footprint of their algorithms. Just as with cast iron, much larger edifices could suddenly be constructed without rare and extraordinary skill, their occasional defects ignored or polished over.

Then came speculative execution. Quite hard to get right and quite impossible to fully understand, it became another miracle, creating another layer of abstraction that just worked and was assumed perfect by all the designs above it. Graduate-level architecture textbooks extolled its virtues without quite explaining how it could be tractably implemented or meaningfully explored in an actual CPU on one's desk.

Just as with the Tay Bridge, independent good-faith optimizations piled up until no one could exactly understand the effects of their composition and predict their results. Instead, we replaced understanding with cost metrics and supposedly authoritative benchmarks, trusting them to capture everything that matters, just as poor Sir Bouch did, and forged on, optimizing the hell out of everything we could.

Every profession has its temptations that are subtle and hard to resist, and that pave the road to hell not just with good intentions but with high-grade ingenuity in pursuit of these intentions. Optimizing to benchmarks as if these benchmarks represented reality is ours. It calls to our competitive spirit and entices us with the beauty of the well-defined contest. It helps us show off miracles of clever winning solutions.

Miracles create a taste for more miracles. Optimizations create an appetite for more optimizations across the board. Since the combined effects of optimizations become hard to understand, metrics and benchmarks proliferate, become the proxy of reality, and eventually get mistaken for the whole of reality. This works for a while, with a feverish build-up of critical dependencies and their proliferation. Then reality strikes back and reminds us that composition is a really, really hard problem, and that measuring a system in any number of ways is no substitute for understanding how it works across the layers, from top to bottom.

Who needed exact understanding of CPU optimizations when the benchmarks all agreed that miraculous improvements have been achieved? Who would argue with the carefully curated sets of computations-that-mattered, and which millions of dollars in pure engineering effort have been spent to tune CPUs to? Certainly not the former students who spent their advanced architecture courses calculating weighted averages of instruction mixes to assert that one ISA was superior to another.

It is said that generals always prepare to fight the previous war. Just in case we are tempted to feel superior to these proverbial generals, let us remember that several generations of CS and CE students have been made to reenact the benchmark battles of the RISC vs CISC war in lieu of an actual education in their contemporary CPU microarchitectures.

Just as poor Sir Bouch, we allowed the metrics that have been useful to a point to get entrenched in our thinking and our processes. We forgot that, unlike math and mechanisms, metrics have no life of their own and will borrow it from other things. Bouch's countryman, the economist Charles Goodhart, formulated a mild version of this observation as "When a measure becomes a target, it ceases to be a good measure." But as we see, neighbors, the truth deserves much harsher words: *metrics are vampires*. When allowed, they will drink the profession's lifeblood, and, if the hapless engineers are too unlucky, will take lives as well.

We've had our fair warnings. So far our Tay Bridge moments have been largely bloodless. They will keep coming, though, because metrics, benchmarks, and layers of abstraction tend to extract their cost as soon as we mistake them for reality or chase them too doggedly.

Remember the bridge over the silvery Tay, neighbors, watch your allowances, trust the experts and the metrics only so far as the wind can blow them, and be sure you understand the workmanship and the optimization shortcuts of at least two layers down. Amen.