Collected Blog Posts of

AllAboutLean.com

2019

Christoph Roser

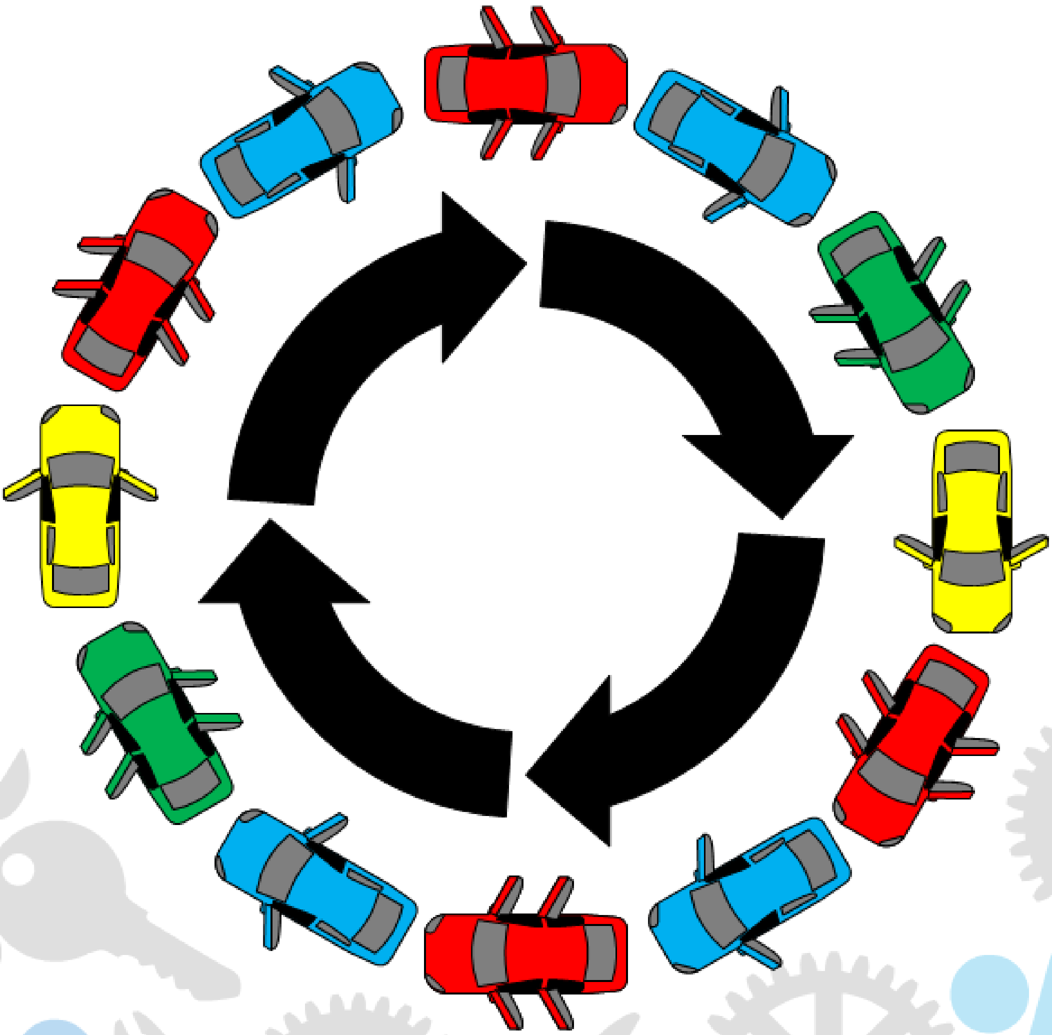

Collected Blog Posts of AllAboutLean.com 2019

Christoph Roser

AllAboutLean.com Publishing
Offenbach, Deutschland
2020

AllAboutLean.com Publishing, Christoph Roser, Hafeninsel 14, 63067 Offenbach, Deutschland, printed by Amazon Media EU S.à r.l., 5 Rue Plaetis, L-2338, Luxembourg

ISBN of the different versions

- 978-3-96382-025-0 Paperback in color
- 978-3-96382-026-7 Kindle eBook
- 978-3-96382-027-4 Free PDF for download at AllAboutLean.com

Bibliographic information published by the Deutsche Nationalbibliothek: The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available on the Internet at http://dnb.dnb.de.

Other Books by Christoph Roser

"Faster, Better, Cheaper" in the History of Manufacturing: From the Stone Age to Lean Manufacturing and Beyond, 439 pages, Productivity Press, 2016. ISBN 978-1-49875-630-3

Fertigungstechnik für Führungskräfte. 2. überarbeitete und erweiterte Auflage, 293 pages, AllAboutLean Publishing, 2019. ISBN 978-3-96382-004-5 (Manufacturing fundamentals textbook for my lectures, in German)

Collected Blog Posts of AllAboutLean.com 2013, AllAboutLean Publishing, 2020. ISBN 978-3-96382-007-6

Collected Blog Posts of AllAboutLean.com 2014, AllAboutLean Publishing, 2020. ISBN 978-3-96382-010-6

Collected Blog Posts of AllAboutLean.com 2015, AllAboutLean Publishing, 2020. ISBN 978-3-96382-013-7

Collected Blog Posts of AllAboutLean.com 2016, AllAboutLean Publishing, 2020. ISBN 978-3-96382-016-8

Collected Blog Posts of AllAboutLean.com 2017, AllAboutLean Publishing, 2020. ISBN 978-3-96382-019-9

Collected Blog Posts of AllAboutLean.com 2018, AllAboutLean Publishing, 2020. ISBN 978-3-96382-022-9

Collected Blog Posts of AllAboutLean.com 2019, AllAboutLean Publishing, 2020. ISBN 978-3-96382-025-0

Preface to the 2013–2019 Collection of Blog Posts

Having successfully written my award-winning blog, AllAboutLean.com, for over six years now, I decided to make my blog posts available as collections. There will be one book of collected blog posts per year, from 2013 to 2019. This way you can store these blog posts conveniently on your computer should my website ever go offline. This also allows you a more professional citation to an article in a book, rather than *just a blog*, if you wish to use my works for academic publications.

This work is merely a collection of blog posts in chronological sequence, and hence does not make a consistent storyline but rather fragmented reading. I am also working on books that teach lean manufacturing. These will also be based on my blog, but they will be heavily edited and reworked to make a consistent storyline. The one I am currently writing focuses on pull production, and hopefully it will be available soon.

The blog posts in this collection are converted into a book as closely as I can manage. However, there are a few changes. For one, on my blog, image credits are available by clicking on the images. This does not work in printed form, hence all images in this collection have a caption and a proper credit at the end of this book. Besides my own images, there are many images by others, often available under a free license. I would like to thank these image authors for their generosity of making these images available without cost. Detailed credits for these other authors are also at the end of this book.

Additionally, a few images had to be removed due to copyright reasons. These are, for example, images from Amazon.com. My blog also includes videos and animations. However, the print medium is generally not well suited to videos and animations, and I do not even have the rights to all videos. Hence, I replaced these videos with a link to the video, and edited the animated images. On digital versions of this book (Kindle eBook, pdf, etc.), these links also should be clickable. No such luck with the print version, unfortunately.

Since my goal is to spread the idea of lean rather than getting rich, I make my blog available for free online. Subsequently, I also make this book available as a free PDF download on my website. However, if you buy it on Amazon, they do charge for their eBooks. If you want a paper version … well … printing and shipping does cost money, so that won’t be free either.

I would like to thank everybody who has supported me with my blog, including Christy for proofreading my texts (not an easy task with my messy grammar!), Madhuri for helping me with converting my blog posts to Word documents, and of course all my readers who commented and gave me feedback. Keep on reading!

As an academic, I am measured (somewhat) on the quantity of my publications (not the quality, mind you!), and my Karlsruhe University of Applied Science tracks the publications of its professors. In other words, one of my key performance indicators (KPI) is the number of publications I author. Hence, I will submit these collected blog posts as publications. On top of that, I will submit every blog post in this book as a book section too. Hence, I will have over three hundred publications including seven books, with me as an author, in one year! It will be interesting to see the reaction of the publication KPI system on this onslaught 🙂. I just want to find out what happens if I submit over three hundred publications in one year 🙂. I don’t know if I will get an award, or if I will get yelled at, but it surely will be fun. I will keep you posted.

Table of Content

1 Hoshin Kanri – Part 1: The To-Do List

Christoph Roser, January 01, 2019, Original at https://www.allaboutlean.com/hoshin-kanri-1/

Figure 1: White Chess Pawn (Image MichaelMaggs under the CC-BY-SA 2.5 license)

This week I will look at Hoshin Kanri (方針管理, policy management). The word is often used as a sort of miracle cure for the problems in your organization. The tool itself, however, is rather mundane, although it did significantly help Toyota. This, of course, did not stop the West from over-complicating and over-hyping it. This post is the start of a small series on Hoshin Kanri.

1.1 Introduction

Figure 2: PDCA and ToDo List (Image Jim.belk in public domain)

Hoshin Kanri is a way to organize and prioritize your main objectives. You could see it as a glorified cross between a to-do list and a [PDCA (Plan, Do, Check, Act)](). It is actually well described as a prioritized to-do list and PDCA combination. Please don't get me wrong, both to-do lists and PDCAs are very, VERY useful tools, and combining them in a Hoshin Kanri makes them even stronger. Just don't expect any magic out of it.

Figure 3: It's (not) magic! (Image Hamilton Public Library in public domain)

Hoshin Kanri is sometimes oversold to make it look larger than life. Nobody would do this with a to-do list, as you all know what it is, but Hoshin Kanri has a fancy Japanese name combined with an often linteled understanding, making it a suitable word for impressing others.

Hoshin Kanri is a very useful tool, but as with pretty much all lean tools, they only work if they are used. Doing a Hoshin sheet once is not too difficult. Sticking with it, on the other hand, is much harder. Just having the sheet won't help. This is the same as the fancy home trainer you may have. If you don't use it, then it is a waste of money. Same with the bread maker that may be sitting in the bottom cabinet of your kitchen (and to all readers who eat homemade bread every day after exercising: Apologies for doubting you, I applaud your rigor).

1.2 The (Prioritized) To-Do List

Figure 4: A checklist (Image Clker-Free-Vector-Images in public domain)

The first part of Hoshin Kanri is a prioritized list of the things that you want or have to achieve. These are sometimes called *Hoshin items*. Depending on the environment you are making the Hoshin Kanri for, this may be a list for your own development (a *personal* Hoshin), or a list based on the goals you get from your superiors.

The important part is to prioritize. What are your most important to-dos that you want on the list. These should be the most relevant to the overall progress or yourself (personal Hoshin) or your company. For example, my wife made me the waste management specialist in our home … meaning I have to bring down the trash … but I would not put this on my personal Hoshin as it is not of significance to my development.

Figure 5: Too much… (Image Wolfgang Zwanzger with permission)

Be careful not to overdo it. **There should be only three to six main points** (*Hoshin items*), each of which could include a few sub-points. Less is more here. Do not get lost in too much detail.

Furthermore, **these items should be based on a process, not on a target**. It would be easy to simply say that your "defect rate has to be below xx part per millions." However, this would fall short of the potential of Hoshin Kanri. Instead you could try to "strengthen management to implement better countermeasures against repeated problems" or "develop the quality skills of your people."

Please do not neglect the topics and directions of your Hoshin items. Many other documents and web pages skip over that and focus on the more agreeable roll out across your departments (which we will do later too). However, having the right Hoshin items is crucial. Personally, I found it quite challenging to get good and process-based Hoshin items for my own personal Hoshin, and I am still not sure if they are good.

1.3 What Could Be on Your List?

Figure 6: Manager in Suit (Image Netfalls with permission)

There are a number of different topics that could be on your list. As mentioned above, these should be limited, and should be focused on the process rather than a target. Below is a list of possible headers under which you could list sub-points. The first list is geared toward a corporate Hoshin, and the second list more toward a personal one. Please do not use all! Less is more here, and focus on the important ones will help you to get them actually done.

- **Health and Safety**: If there is any issue with the health and safety of your workers, then this should be on your list, maybe even as the first topic.
- **Environment**: This is also found on Hoshin Kanris at Toyota, sometimes also quite high up. The environment could be the protection of nature, reduction of waste, or elimination of pollution, but also good cooperation with the local community. Toyota plants in Japan put in quite a bit of effort to get along with the locals.
- **People Development**: Help your people grow. Develop their professional skills, but also their interactions as a group as well as their adherence to the company philosophy.
- **Quality**: Toyota is famous for its quality, and this success did not come by itself. If quality is an issue in your company, then this could be on there too.
- **Lead Time**: Similar to quality, this could be on your Hoshin.
- **Cost**: This completes the "quality – time – cost" triangle. However, while every plant could improve quality, time, and cost, you may want to focus only on one or two areas. The tiger that tries to catch all ducks will get none. Similarly here, the less points you have, the more you can focus on them.

Figure 7: Scuba Diver (Image Joakant in public domain)

If you want to do a personal Hoshin Kanri, the bullets may look different. Below are possible suggestions, three of which are from my own Hoshin. Again, less is more!

- **Learn**: For me, part of the joy of life is to learn new things. Hence, on my personal Hoshin I put “Learn” as the first category. Do you want to learn Chinese? Put it on the list. However, merely putting it on the list won’t make you speak Chinese; the hard part comes afterward, spending lots of time studying the language. Hoshin Kanri merely helps you where to put your focus.
- **Teach**: This is also part of my own Hoshin. While it is my job as a professor to teach, I like teaching. Interestingly enough, all sub-points on my own Hoshin here were outside the normal university classroom.
- **Have Fun**: This was the third point of my own Hoshin. For me, it is important to have fun every now and then.
- **Health**: This is not (yet) part of my own Hoshin, but depending on your health it may be important enough to include it. Luckily, I am still healthy … although my wife keeps on reminding me to lose weight … maybe I should add it to my Hoshin … but next year is good enough … right?
- **Get Along**: Maybe you want to make new friends or strengthen the relations with the one you have? Get to know your neighbors? Here you go!
- **Your Pet Peeve**: Do you want to save the environment, end world hunger, cure cancer, or just get a different president/premier/chancellor elected next time? Whenever it is important enough for you to take action and make the cut for the limited space on the list, add it!

There are more possibilities, both for your company and for your personal Hoshin Kanri. Make sure to select no more than six top headers, if possible even less. Each header could have roughly one to four process based sub-topics. **I cannot emphasize this enough: focusing on the key topics is important for success!**

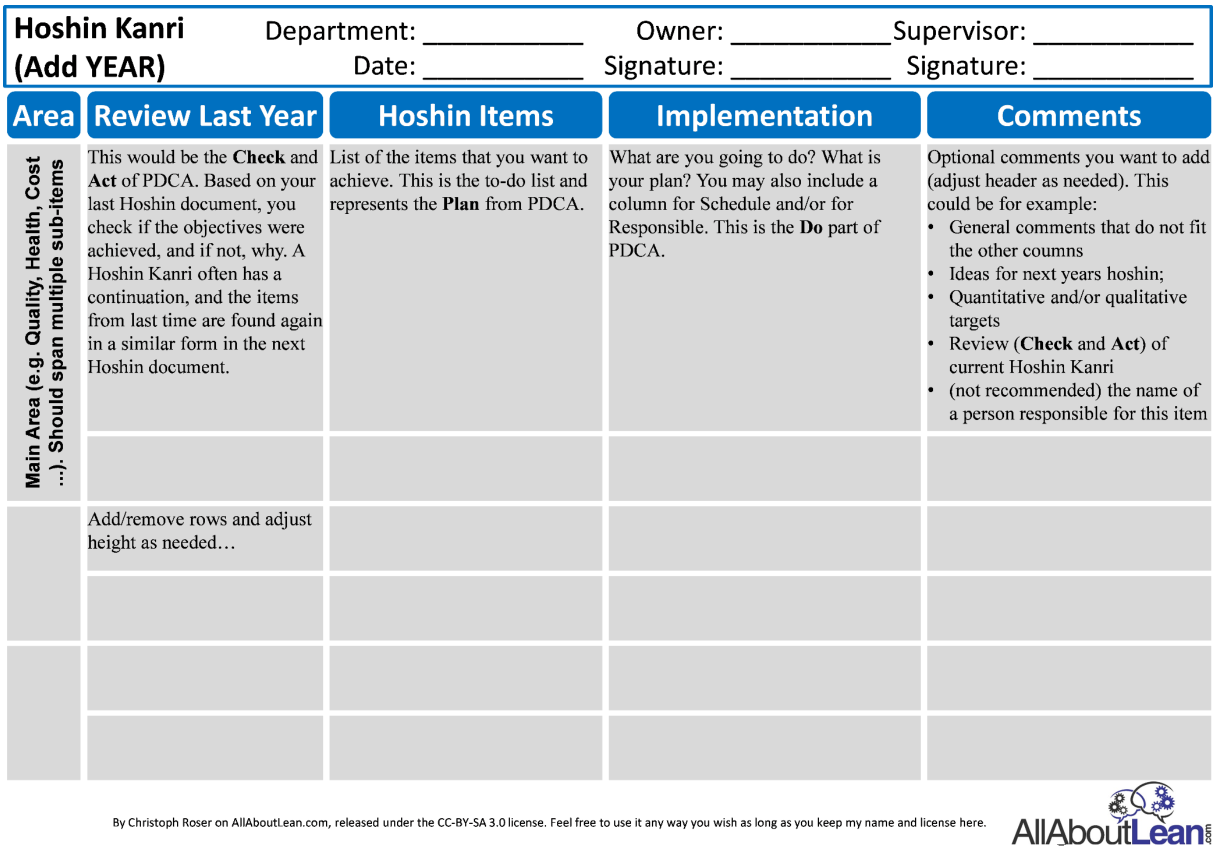

Figure 8: Hoshin Kanri Template (Image Roser)

I have created a blank PowerPoint Hoshin Kanri Template for you to use, available at https://www.allaboutlean.com/wp-content/uploads/2019/01/Hoshin-Kanri-Template.pptx. Hopefully it helps you. You can edit the PowerPoint to match the document to your needs.

In my next post I will show you how to marry this to-do list with the PDCA. Until then go out, think about what you really want to achieve in your position at the company or with your life altogether, and organize your industry!

P.S.: Many thanks to Isao Yoshino for his input!

2 Hoshin Kanri – Part 2: PDCA

Christoph Roser, January 08, 2019, Original at https://www.allaboutlean.com/hoshin-kanri-2/

Figure 9: Chess Move (Image Kyle Gese in public domain)

In my first post on Hoshin Kanri I explained the details of making the list for the Hoshin. This now has to be combined with a PDCA (Plan, Do, Check, Act). The rigor of PDCA gives value and life to what would otherwise be a simple action list. Let me show you:

2.1 PDCA

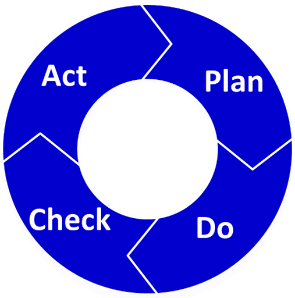

Figure 10: PDCA Circle (Image Roser)

PDCA is one of the most important tools in lean (or in any kind of process). I have written a whole series on PDCA, starting with The Key to Lean – Plan, Do, Check, Act! Just to review it briefly, the sequence is as follows:

- **Plan**: Define scope, define target, analyze the situation, understand the problem, develop one or ideally more solutions, select the best one for implementation.
- **Do**: Implement, define the new standard, train your people in the new standard, ensure the standard is followed.
- **Check**: Did the implementation actually work? Is the problem fixed? A nice presentation is no proof of a solution!
- **Act**: If it did not work (good enough), why not? What do you need to do to achieve the targets?

While the Plan and Do parts are easy and commonly done all over the world, the Check and Act are much harder. I sometimes have the feeling that management is often not interested in the actual outcome and is satisfied with a nice presentation.

If PDCA is done well, it can develop into a continuing series of PDCA loops until the problem is solved. In this case, PDCA continues with the next problem.

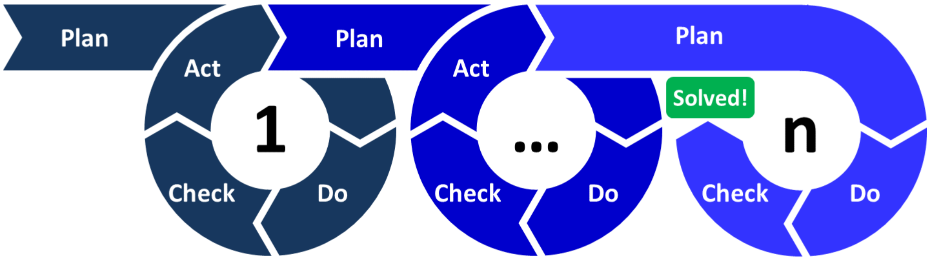

Figure 11: PDCA Circle Sequence (Image Roser)

2.2 Overlapping PDCA with the To-Do List to Get Hoshin Kanri

Figure 12: PDCA ToDo List (Image Jim.belk in public domain)

In my previous post we looked at the to-do list that will go into the Hoshin Kanri. On this list are the rows in that document. PDCA represents the columns. The initial to-do list would be the column corresponding to the Plan of PDCA.

The columns do not necessarily need to be labeled Plan, Do, Check, and Act. There may be also more than four columns (or less), as long as PDCA is represented. In fact, there is not a standard Hoshin Kanri that is used everywhere, no matter what you hear otherwise. As always in lean, the document has to fit YOUR needs, the solution has to fit YOUR problem. Just copying something from someone else may not help you much.

Below is a selection of possible columns for Hoshin Kanri. These may help as suggestions to see what you may need.

- **Review of last Hoshin**: This would be the Check and Act of PDCA. Based on your last Hoshin document, you check if the objectives were achieved, and if not, why. A Hoshin Kanri often has a continuation, and the items from last time are found again in a similar form in the next Hoshin document.
- **Hoshin Items**: List of the items that you want to achieve. This is the to-do list from the last post. Feel free to group it into overarching topics (Quality, Health, Cost …) with a small number of sub-points for each. You may also include a separate column for **targets**, although not all Hoshins do. They are not necessarily quantitative, but may be qualitative. Together with the Hoshin items, this is the Plan from PDCA.
- **Implementation Plan**: What are you going to do? What is your plan? You may also include a column for **Schedule** and/or for **Responsible**. This is the Do part of PDCA.
- **Evaluation**: Did it work? Is the problem solved? This column may also be on your next Hoshin; see the top bullet “Review of last Hoshin.”

2.3 Some Additional Items

Figure 13: Chess Set (Image Alan Light under the CC-BY-SA 3.0 license)

Like any document, you can add a number of additional topics. Not all of them are needed every time you have to decide which ones may be helpful in your case.

- **Title**: Give the document a title, i.e. "Final Assembly Hoshin 2019" or "Personal Hoshin 2015" or similar.
- **Date**: For which period (which year?) is the Hoshin?
- **Owner**: Who is in charge of the Hoshin?
- **Department**: To which department does the Hoshin belong? This should be the department of the owner.
- **Supervisor**: Who will review the Hoshin with the owner and give feedback?
- **Signature Supervisor**: A space for the supervisor to sign. This is to show that the Hoshin is completed for this term (year). Please note this does not mean that all problems are solved, however.
- **Vision**: What are the company (or your) overarching goals or guidelines? What is the corporate philosophy? Ideally, the items on the Hoshin should reflect this vision.

2.4 How Many Hoshins?

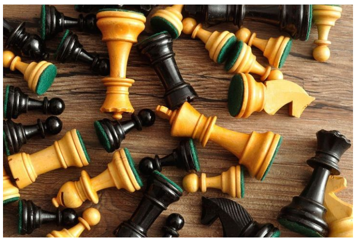

Figure 14: Jumbled Chess Pieces (Image Marietjieopp with permission)

The power of Hoshin Kanri lies in the focus on the key points. Hence, the number of Hoshin Kanri documents that you are responsive for at the same time should be kept at an absolute minimum, ideally one. You may have a second personal Hoshin besides your corporate one to also improve yourself outside the industry context. Yet, within the industry there should be a maximum of one Hoshin document per person at a time.

Management must resist the temptation to create different Hoshins for different projects or topics. Similarly, if you are unfortunate enough to have two supervisors who both can tell you what to do, try to avoid having two separate Hoshin documents. If you put everything on your Hoshin that anybody would like to have, then you end up with 40+ different key items, all of which are top priority. Does this sound familiar? If everything is top priority, then nothing is. If nothing is priority, nothing will get done. Again: Limit the number of priority topics that make it on your Hoshin.

2.5 What Period Should a Hoshin Cover?

Figure 15: Four Seasons (Image Roser adapted from Cherubino under the CC-BY-SA 4.0 license)

Most Hoshin Kanri documents that I know cover one year. This is usually a good duration, since one year allows for quite a bit of improvement activity. This duration is also long enough to see the results and review the outcome.

Again, resist the temptation to do Hoshins more frequently. You won't get twice as much done if you have two Hoshins per year; you merely increase the organizational overhead, and hence reduce the actual improvement capacity.

2.6 Who Should Have a Hoshin?

Figure 16: Crazy Office People (Image Poznyakov with permission)

A personal Hoshin for improving yourself is possible for anyone. A corporate Hoshin is mostly for managers.

In general, the person having a Hoshin Kanri should be able to influence an area under his control, making not only short-term decisions but also longer-term strategic decisions or changes.

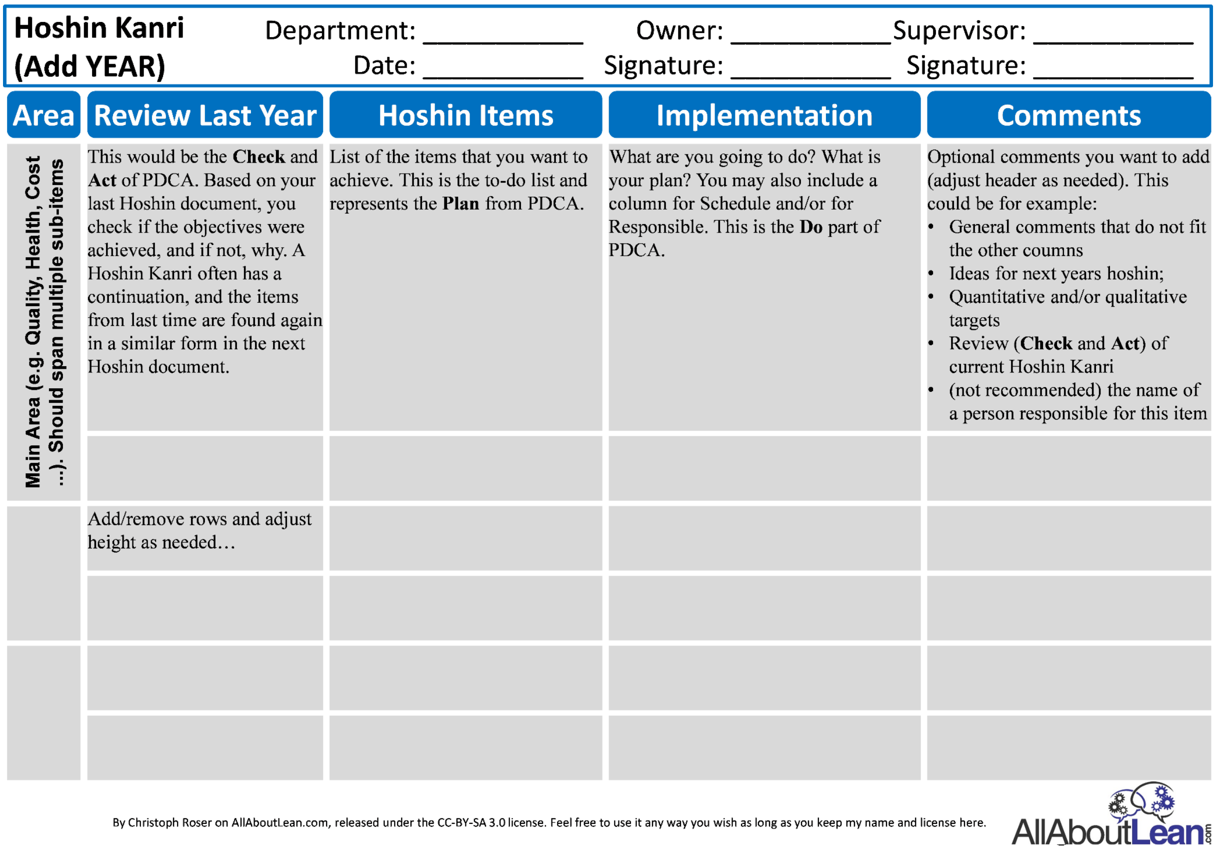

Hoshin Kanri (Add YEAR) Department: ___________ Date: ___________ Owner: ___________ Signature: ___________ Supervisor: ___________ Signature: ___________

Area	Review Last Year	Hoshin Items	Implementation	Comments
Main Area (e.g. Quality, Health, Cost …). Should span multiple sub-items	This would be the **Check** and **Act** of PDCA. Based on your last Hoshin document, you check if the objectives were achieved, and if not, why. A Hoshin Kanri often has a continuation, and the items from last time are found again in a similar form in the next Hoshin document.	List of the items that you want to achieve. This is the to-do list and represents the **Plan** from PDCA.	What are you going to do? What is your plan? You may also include a column for Schedule and/or for Responsible. This is the **Do** part of PDCA.	Optional comments you want to add (adjust header as needed). This could be for example: • General comments that do not fit the other counms • Ideas for next years hoshin; • Quantitative and/or qualitative targets • Review (**Check** and **Act**) of current Hoshin Kanri • (not recommended) the name of a person responsible for this item
	Add/remove rows and adjust height as needed…			

Figure 17: Hoshin Kanri Template (Image Roser)

Now you know how a Hoshin Kanri is structured. I also created a blank PowerPoint Hoshin Kanri Template for you to use, available at https://www.allaboutlean.com/wp-content/uploads/2019/01/Hoshin-Kanri-Template.pptx. Hopefully it helps you. You can edit the PowerPoint to match the document to your needs.

I hope this article was helpful for you. In my next post I will explain how different Hoshin Kanri documents influence each other across the hierarchy. Until then, **go out and organize your industry!**

P.S.: Many thanks to Isao Yoshino for his input!

3 Hoshin Kanri – Part 3: Hierarchical Hoshins

Christoph Roser, January 15, 2019, Original at https://www.allaboutlean.com/hoshin-kanri-3/

Figure 18: White Chess Set (Image White Spirit with permission)

Hoshin Kanri can be used individually, but its full potential is shown across the levels of corporate hierarchy. The goals of a Hoshin Kanri should be derived from the Hoshin Kanri of the next-level hierarchy above. This post is part of a larger series on Hoshin Kanri. Let's look at the hierarchy structure:

3.1 The Top-Level Hoshin Kanri

Figure 19: Private Jet and Limousine (Image Artem Alexandrovich with permission)

The top-level Hoshin Kanri starts … well … at the top. This is the document for top executives. Their items for the Hoshin Kanri are based on long-term corporate strategy, corporate culture, and corporate vision. In which overall direction do they want to move the company? This direction is used to develop the Hoshin Kanri items for the top executives.

I find two things important here. First: **We are talking long term!** The items on the Hoshin Kanri this year should be similar to the items last year and will be similar to the items next year. A company does not move like a sports car and cannot zigzag across different directions. The larger the company, the more inertia to change there is. If the top level changes the goals and the direction every year, then a totally confusing message will be received at the bottom. "Left, right, drop everything and do this. No, cancel that. Do that instead. Hurry up and wait …" Depending on where you work this may sound familiar. The longer the CEO maintains the same direction, the more likely the company will actually move in that direction. Toyota has maintained the same direction for over half a century!

Figure 20: Long-term vision of SpaceX and Elon Musk (Image NASA/Pat Rawlings in public domain)

It is **okay to have big goals, or even goals that seem impossible at the moment**. They probably will not come true this year, or next year, but maybe it will happen in ten years or in thirty. For example, Elon Musk founded SpaceX in 2002 with the goal of reducing space transportation costs and enabling the colonization of Mars. It took them six years for their first liquid-propellant rocket to reach orbit. They are still years, if not decades, away from the colonization of Mars, but it is their ultimate goal.

Second, **these goals should be process oriented**, not necessarily a numeric target. Rather than setting a goal of "fewer than xx accidents per year," foster a safety-conscious culture and enable an environment where safety is important. These process-oriented goals are often much better for defining the direction than a simple number (which may be fudged anyway).

3.2 Following the Hierarchy Down …

Figure 21: People in Office (Image Rido81 with permission)

The next level in the hierarchy derives their goals from the level above. Of course, this is not a simple copying of the goals, but an adaptation of the goals of the boss to their own area of responsibility. The superior has to avoid too-detailed instructions, but rather let the subordinate think about it himself. Avoid micromanagement, no matter how tempting it may be. Having process-oriented goals help and allow the subordinate more freedom to do what he thinks is right to achieve the goals.

With each level down, these goals become more focused on the area of responsibility. Usually, most actual implementation happens at the bottom level of the hierarchy.

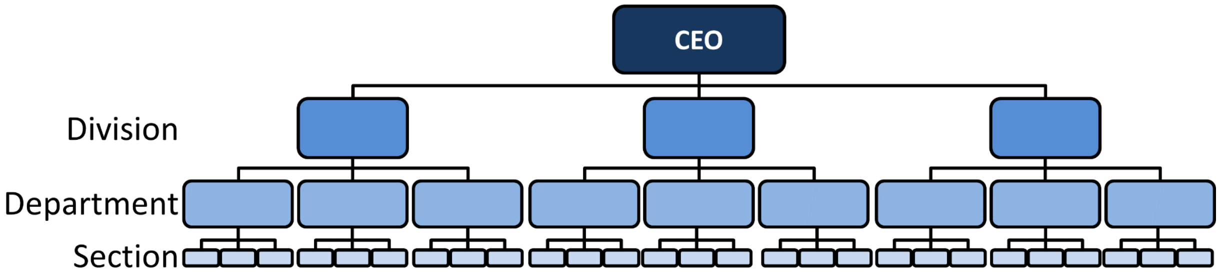

Figure 22: Hierarchy Structure (Image Roser)

3.3 Review of the Hoshin Kanri

The most often forgotten part of PDCA is the actual Check and Act. Hence, the review of the Hoshin Kanri is crucial for its success. This review not only helps the subordinate, but also gives valuable input to the superior about the current state, which he eventually needs for the evaluation of his own Hoshin Kanri. Here are some suggestions for the reviewer of the Hoshin Kanri.

The Hoshin should be reviewed at a minimum twice. The first review is when a new Hoshin is created. Here the reviewer should give input on the items on the list. Are they based on a process (good) or a target (not so good)? Do they reflect the needs and the Hoshin of the next level up (ideally the Hoshin of the reviewer)? The second review is when the period is up and the evaluation of the outcome begins (the Check and Act parts of PDCA). Toyota also uses a mid-term evaluation to check on the current progress, and managers may have additional reviews regularly to check on the progress. It is not a fire-and-forget approach.

As always, whenever possible **have the review not in a meeting room but at the location where the events are happening** (in lean-speak: Go to the Gemba!). Evaluating the outcome is much easier on site, and it will also be much easier for you to understand and to give input.

Figure 23: Understanding manager (Image Mangostar with permission)

Second … and this may be difficult for some … **be accepting of bad news!** No matter how much you dislike it, failures will happen. In order to learn from these failures and improve the next attempt at fixing it, it is important to talk honestly about these failures. If you punish people for making mistakes, they will hide their mistakes. Your people and you can both learn a lot from mistakes. Actively encourage the reporting of mistakes. It is for a reason called *trial and error* and not *trial and success*. I could go on about this topic for hours, but: Do not blame your people for mistakes!

3.4 How Not to Do It

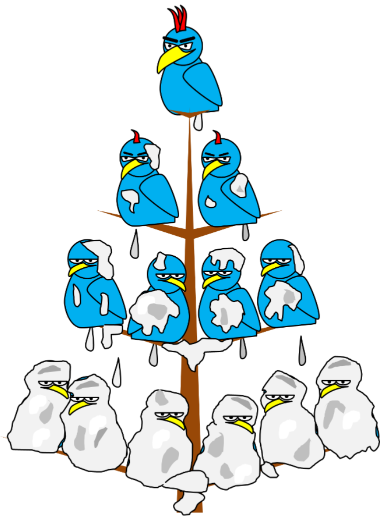

Figure 24: Bird Shit Policy Deployment (Image Roser)

Deploying goals across hierarchy can be very helpful, but done incorrectly it can also drive a company insane. Giving inconsistent and frequently changing goals will mess up the process. Punishing people for bad news will encourage and reward liars. Company morale will suffer, and people will get stuck in a gridlock, doing more about appearance without having much time for actual improvement.

I have heard of one story where a consultant asked a forklift driver out of curiosity what his corporate targets were. Without skipping a beat, he said 5% EBIT (earning before interest and taxes). The consultant was a bit confused about this, and inquired how he will achieve this. The forklift driver said he had no idea, he doesn't even know what EBIT means, but these are his targets. It turned out that these were also the targets of his supervisors, his supervisors' supervisor, and so on, because the CEO wanted to have a 5% EBIT. Obviously it does not work that way!

3.5 Summary

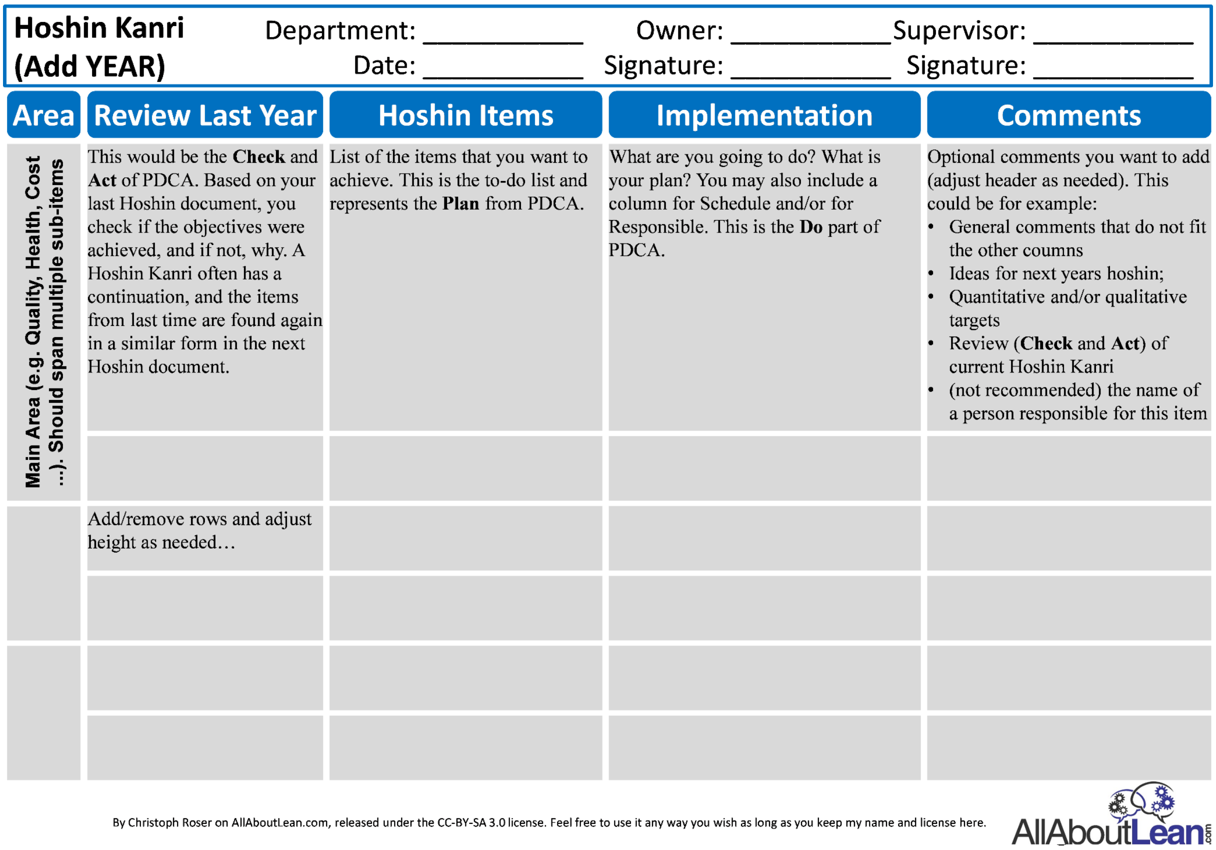

Hoshin Kanri (Add YEAR) Department: ___________ Date: ___________ Owner: ___________ Signature: ___________ Supervisor: ___________ Signature: ___________

Area	Review Last Year	Hoshin Items	Implementation	Comments
Main Area (e.g. Quality, Health, Cost …). Should span multiple sub-items	This would be the **Check** and **Act** of PDCA. Based on your last Hoshin document, you check if the objectives were achieved, and if not, why. A Hoshin Kanri often has a continuation, and the items from last time are found again in a similar form in the next Hoshin document.	List of the items that you want to achieve. This is the to-do list and represents the **Plan** from PDCA.	What are you going to do? What is your plan? You may also include a column for Schedule and/or for Responsible. This is the **Do** part of PDCA.	Optional comments you want to add (adjust header as needed). This could be for example: • General comments that do not fit the other counns • Ideas for next years hoshin; • Quantitative and/or qualitative targets • Review (**Check** and **Act**) of current Hoshin Kanri • (not recommended) the name of a person responsible for this item
	Add/remove rows and adjust height as needed…			

Figure 25: Hoshin Kanri Template (Image Roser)

I have created a blank PowerPoint Hoshin Kanri Template for you to use, available at https://www.allaboutlean.com/wp-content/uploads/2019/01/Hoshin-Kanri-Template.pptx. Hopefully it helps you. You can edit the PowerPoint to match the document to your needs.

Successful companies are successful not so much because of strategy creating, but rather strategy execution. Overall, Hoshin Kanri and its use across different levels of hierarchy can help you with deploying your strategy. However, while it is useful too, it can only help you focus on the important topics, but you still have to do them. Simply putting an item on the Hoshin Kanri will not magically fix it, but only prioritize the issue and subject it to a PDCA process. Now **go out, deploy your strategies, and organize your industry!**

P.S.: Many thanks to Isao Yoshino for his input!

4 Hoshin Kanri – Part 4: The X-Matrix?

Christoph Roser, January 22, 2019, Original at https://www.allaboutlean.com/hoshin-kanri-x-matrix/

Figure 26: Is more better? (Image Dr Jacek Filek in public domain)

When googling Hoshin Kanri, you will sooner or later come across an X-Matrix. It is a visually very impressive tool, but I am in serious doubt about its usefulness. It focuses on the creation of the Hoshin items, but to me this approach is overkill, and – even worse – may distract the user from actually following the PDCA, especially the Check and Act parts. While the article is highly critical, I hope reading it and understanding the shortcomings help you better understand how Toyota thinks.

4.1 Introduction

Setting the right goals and filtering them through the organization is important in Hoshin Kanri. In my first post I talked in detail about this as the "to-do list."

Figure 27: Sneaky Consultant X-Matrix (Image bramgino with permission)

These simple to-do lists can be modified to be quite elaborate, eventually leading to the X-Matrix. They are also sometimes called x-matrices (instead of matrix) or target-mean matrices. The origins of this matrix is a bit fuzzy, but it seems that Japanese professor Yoji Akao (1928–2016) and Bob King (owner of a quality lean-related publishing house) were involved.

In any case, looking at the matrix, I have the strong feeling that this was not developed by a practitioner in the field but more likely by an academic or consultant who has the need to impress others with fancy methods. But before I go into details of the criticism, let me briefly show how the matrix works.

4.2 The Fields in the Matrix

The X-Matrix has a number of fields. To fill it out, you start at the bottom (usually called "South") with the long-term objectives. Next comes the annual objective (left, or "West"), then the top-level priorities (top, or "North"), and finally the targets to improve (right, or "East"). At the end on the very far right (running out of compass directions here …) comes the people responsible for the different tasks. Let's have a (critical) look at the steps. In the corner intersections, you mark down if the points are related, or even how related they are.

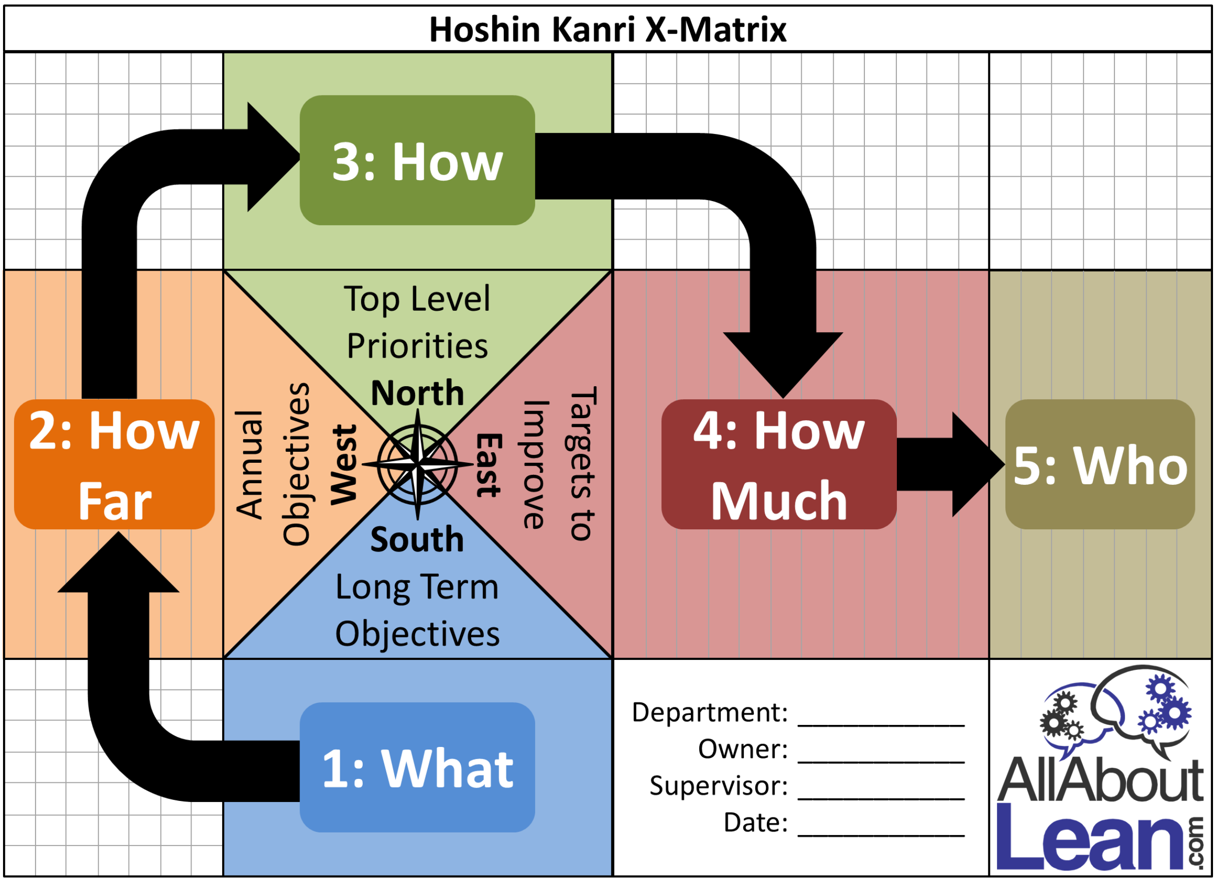

Figure 28: Hoshin X Matrix Overview (Image Roser)

- **South: Long-Term Goals**: The first step is the long-term goals. What is the overall direction you want to move your company (department)?

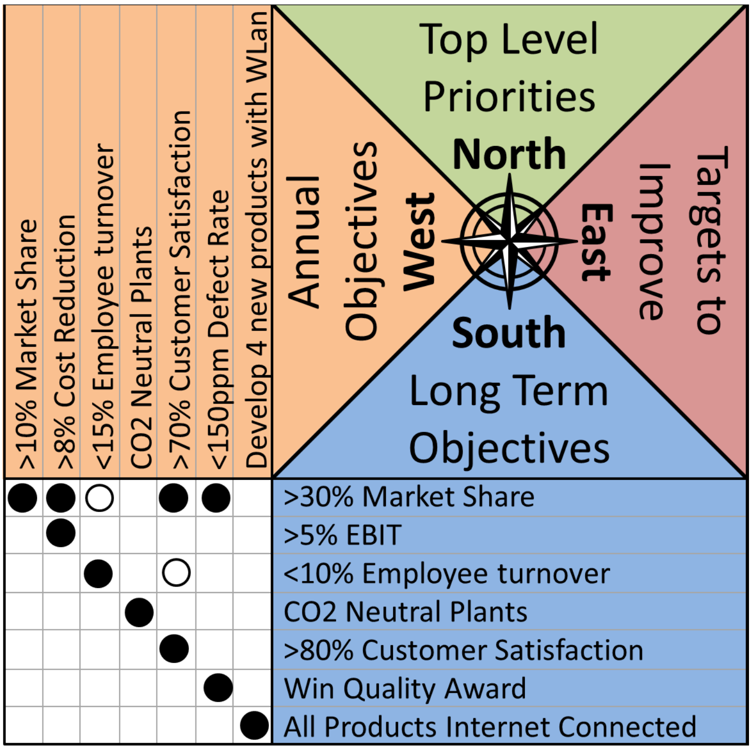

Figure 29: Hoshin X Matrix Detail of Corner (Image Roser)

- **West: Annual Objectives**: Out of the long-term objectives the annual objectives are developed. What do you want to achieve this year? In the matrix between the long-term goals and the annual objectives, you mark which long-term goal is aligned with which annual goal.
- **North: Top-Level Priorities**: Next you develop the different activities you want to do to achieve the annual results. In the matrix in the corner, you again connect the previous annual objectives with the different priorities to achieve these objectives.
- **East: Targets to Improve**: Based on the top-level priorities, you create (numeric) targets to achieve this year. Again, in the field between the top-level priorities and the targets, you mark which priority influences which target. Some examples also have a matrix that connect the targets to improve back to the long-term goals. (Which may lead to inconsistencies if an entry in a step 1 long-term goal going through steps 2 and 3 does not influence any entry step 4 targets to improve, but the matrix leading back to step 1 sees a connection. Best not to use a matrix here.)
- **Responsible**: On the very far right you add which person is responsible for what top-level priorities.

Like the “normal” Hoshin Kanri, this document is done at different levels in the hierarchy, starting with the top executive. These are named rather straightforward as top-level matrix, second-level matrix, and third-level matrix.

4.3 Criticism

Figure 30: Critical Person (Image OpenClipart-Vectors in public domain)

From my point of view, there are a lot of things flawed with the X matrix.

- **Long-term goals not long-term enough**: The long-term goals I see online usually describe long term as three to four years … which for me coming from Toyota is too short. I would much prefer to have here the truly long-term visions or the foundations of the corporate philosophy here.
- **Often redundant focus on numeric goals**: Most examples I have seen online are overly numeric. The sequence could be as follows (actual example): *1)Long-Term Objectives: Gain market share to 40% -> 2) Annual Goal: Increase market share to 25% -> 3) Top-Level Priorities: Develop growth strategies -> 4) Targets to Improve: Increase market share to 25%.* For this I do not need a fancy matrix! In this example, 2) and 4) are absolutely identical. In many examples I see online, the connection between the different fields feel forced and often redundant. In the original Hoshin Kanri, the focus is much more on having a process rather than having numeric targets.
- **Diluting responsibilities**: The Hoshin Kanri I know are always documents for one person. THIS person is responsible for the document, and has to work on implementing it. While he may work with his subordinates, the Hoshin document is his responsibility. This is different in the X-Matrix. Already as part of the design, the responsibilities are handed out right away to others. The person who made the X-Matrix is already no longer responsible. This to me is a very wrong mind-set!
- **Where's the PDCA?**: Probably the biggest gripe I have is that the X-Matrix distracts from the PDCA! In the original Hoshin Kanri at Toyota, the PDCA is clearly part of the process. The X-Matrix, on the other hand, is devoid of any hints of the PDCA. While many articles about the matrix mention PDCA, it just feels like it is not there. Some articles see the X-Matrix as the first step to get the items to fill out a proper Hoshin Kanri, but they rarely go into detail for the actual Hoshin Kanri afterwards. Even if the X-Matrix is in preparation for an actual Hoshin, many of the fields feel redundant with the Hoshin. But again, if the PDCA is nonexistent, then the approach will have serious problems implementing it.

4.4 When to Use the X-Matrix?

Personally, I would try to avoid the X-Matrix, as I think it is introducing unnecessary overhead while losing some of the power of the original Hoshin Kanri. However, there are a few instances where it may be of use:

First, if you are already using it and it works for you. If that's the case, then keep on using it. If you manage to use the X-Matrix successfully (that means with a PDCA) and you want to continue, then I am not going to stand in your way. To me, lean is whatever works, not a set of dogma-like methods.

Figure 31:Your new consultant must be really good … because he sparkles … (Image Glen Scarborough under the CC-BY-SA 2.0 license)

Second, you may be in the unfortunate position that your boss or your client wants flashy and fancy methods. In this case, the X-Matrix may be a nice thing to bedazzle your client or boss. It won't really make the works easier, but it looks sooo much cooler than the normal Hoshin Kanri to-do list and PDCA crossover. Just make sure that the implementation actually happens and the PDCA includes the C&A … but then, with this kind of boss or client, a fancy presentation with colorful slides may be substituted for actual progress. I don't like it, but then … some people do want to be lied to.

But, if you are not yet using it and your boss is at least somewhat reasonable, then my advice is to avoid the X-Matrix and rather put the effort into the classical Hoshin Kanri. In my next post I will tell you a bit about the history of Hoshin Kanri and how the Kanri Noryoku program saved Toyota. Until then, stay away from overcomplicated tools, do the normal Hoshin Kanri, and **go out and organize your industry!**

P.S.: Many thanks to Isao Yoshino for his input!

5 Hoshin Kanri and the Kanri Noryoku Program: Rejuvenating Toyota

Christoph Roser, January 29, 2019, Original at https://www.allaboutlean.com/kanri-noryoku-program/

Figure 32: Hoshin Kanri Noryoku Program (Image Roser)

In my previous posts I explained how Hoshin Kanri works. This post looks at how Toyota embeds Hoshin Kanri as part of their overall management structure. Toyota started this in 1979 when director Masao Nemoto started the Kanri Noryoku Program (管理能力プログラム), usually shortened to KanPro.

5.1 Linguistics

Figure 33: True North (Image Hike The Monicas under the CC-BY-SA 4.0 license)

The term *Hoshin Kanri* comes from Japanese. It is written as 方針管理. The first half is the word 方針 (*hoshin*) for "policy; plan; course; principle; objective; but also as magnetic needle." The second half is 管理 (*kanri*) for "control; management." Hence the translation is *policy management*, although in lean it is more often translated as *policy deployment*, but even the more exotic *compass management* has sometimes been used. This is in contrast to Nichijou Kanri (日常管理,everyday management) for the daily management.

In the Kanri Noryoku Program, Kanri (管理) is the same as in Hoshin Kanri for control and management. *Noryoku* (能力) stands for ability or faculty (or, less common, a temple employee). プログラム is the Japanese katakana writing for "program." Hence, altogether the Kanri Noryoku Program is a management capability improvement program, and that was exactly what it was.

5.2 Events Leading to the Kanri Noryoku Program

5.2.1 Management by Objectives

Historically, Hoshin Kanri started out as what we now call management by objectives (MBO). Management by objectives starts with specific objectives at the top level of the hierarchy, which then are used to develop the objectives for the next hierarchy level down, and so on. The idea is that everybody's objectives support the objectives above and give the company an unified direction.

While management by objectives was first used by management guru Peter Drucker in his 1954 book The Practice of Management, the idea is much older and dating to at least 1926.

5.2.2 Total Quality Control and Hoshin Kanri

The concept of management by objectives eventually made it to Japan. Toyota picked up the idea around 1963–1964, shortly after they started with Total Quality Control (TQM). However, Toyota found that management by objectives fell far short of their expectations. Hence they modified it by combining management by objectives with PDCA (Plan, Do, Check, Act). This was called Hoshin Kanri. Subsequently Toyota rolled out a program called "cross functional management." This aimed to integrate the different departments (product planning, design, preparation, procurement, production, and sales) across the different functions (quality, cost, technical development, production, sales, HR) using a matrix. Hoshin Kanri was a part of this *cross functional management* initiative.

	Product Planning	Product Design	Product Preparation	Procurement	Production	Sales
Quality	++	++	++	++	++	++
Cost	++	+	++	++	+	+
Technical Development	+	++	+	0	0	+
Production	0	+	++	0	++	+
Sales	++	+	0	0	+	++
HR	+	+	+	+	++	+

Cross-Functional Management

Departmental management

Figure 34: Toyota policy deployment (Image Roser, after Isao Yoshino)

Please note that this is the origin of the Hoshin Kanri at Toyota. Other sources claim other origins, for example that Bridgestone Tire named the planning techniques of Deming Prize winning companies Hoshin Kanri in 1965; another source claims Hewlett-Packard Japan developed it in 1976. Japanese Professor Yoji Akao(1928-2016) also claims to have invented Hoshin Kanri. I have no idea which of these claims are true.

5.3 The Kanri Noryoku Program: Rebirth of Hoshin Kanri

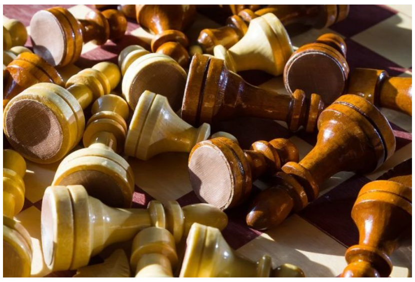

Figure 35: Fallen Chess Pieces (Image banderchenno with permission)

Hoshin Kanri continued to come along this cross functional management. However, Toyota started to slack off in the aftermath of the 1973 oil crisis, and their focus on total quality management declined. Ten years after they won the Deming Award for Quality, their focus was weakening. Management got stuck in the daily squabble and lost sight of the big picture. If this sounds familiar, it is. To me it seems the leadership of a lot of companies would benefit from more focus on the big picture.

As a response, in 1979 the director in charge of quality control, Masao Nemoto (根本正夫), introduced a program called Kanri Noryoku Program (管理能力プログラム), usually shortened to KanPro.

Due to this Kanri Noryoku Program, Masao Nemoto is considered to be one of the great people at Toyota, and may even have contributed as much to its success as Taiichi Ohno. He was involved early on in Total Quality Control at Toyota. Eventually he retired as a senior managing director in 1982.

Figure 36: Business Hierarchy (Image Gajus with permission)

Altogether, 2,000 key leaders at Toyota had to develop their own Hoshin Kanri on A3 paper. They had to think what their most important goals were, develop an action plan for each goal, implement it, and then check back if it worked or not. Here (again) you find the PDCA (Plan, Do, Check, Act) cycle that is so important to Toyota's success. The managers were trained in

- How to use PDCA (not only the theory but actually on using it correctly)
- How to identify key issues and problems
- How to develop targets and goals
- How to make plans for improvement
- How to implement
- How to analyze results
- How to review the Hoshin Kanris of their own people
- How to use an A3 sheet, including how to use the A3 for a presentation for management

Figure 37: Circle of chess pawns (Image ronstik with permission)

Every six moths, their Hoshin Kanris were reviewed with their superiors. The focus of this review was on the cause of the problems (not like in the West where the focus is often finding someone to blame). In particular the reporting of mistakes was encouraged as a valuable learning opportunity. Another focus for managers was on developing their own subordinates.

The Kanri Noryoku Program refreshed the vigor for Hoshin Kanri, and also strengthened the use of A3 sheets for project management. All presentations in this program had to be done using

A3 sheets, which helped to focus on the key points. Isao Yoshino compares this to doing 5S on a presentation (sources below). Furthermore (unlike PowerPoint) made all points visible all the time.

5.4 The Legacy of the Kanri Noryoku Program

Figure 38: Black chess pieces on board (Image STILLFX with permission)

While the Kanri Noryoku Program itself lasted only for two years, it had a lasting impact on Toyota. It refreshed Toyota and helped managers improve their focus. The method also made its way down through the ranks, and employees started using it after observing it with their superiors. It continued to help the development of subordinates. Managers at Toyota were not afraid of presenting failures and also were not punished but praised for being honest about mistakes (although unfortunately with the current CEO this seems to be changing). The A3 also took firm hold in Toyota, and has been used regularly ever since.

5.5 Sources

There are few articles on the Kanri Noryoku Program. The following is a selection of articles that helped me write this post.

- Toyota Leadership Lessons: Part 7 – Insights into how "respect for people" & "continuous improvement" became the pillars of the Toyota Way by Katie Anderson on her excellent blog. Katie also does Japan Study Tours together with Hoshin Kanri guru Isao Yoshino. If you are interested, check out the offer on her blog.
- Isao Yoshino reflects on the role of management at Toyota, an interview of Isao Yoshino on Planet Lean. I also learned from Mr. Yoshino directly during a workshop at the ELEC conference 2018 in Portugal.
- How the A3 Came to Be Toyota's Go-To Management Process for Knowledge Work by John Shook and Isao Yoshino.

6 Should You Use Physical or Digital Kanban Cards?

Christoph Roser, February 05, 2019, Original at https://www.allaboutlean.com/digital-kanban/

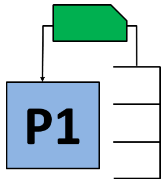

Figure 39: One Process Kanban Loop (Image Roser)

Kanban and similar pull systems like CONWIP or POLCA are basic parts of lean production. They limit the maximum number of parts by attaching a sort of token (i.e., the kanban card) to the part, and return to the beginning when the part leaves the system.

These kanban can be physical cards or digital representations. In this post I look into when you should use a physical kanban and when you should use a digital kanban.

6.1 Introduction

The kanban card or CONWIP card goes around the kanban or CONWIP loop repeatedly. Information on a (digital or physical) fresh card arriving signals the start of production. The information then travels with the part through the production process and the corresponding inventories. When the part leaves the loop, the information is returned to the beginning and the process starts anew.

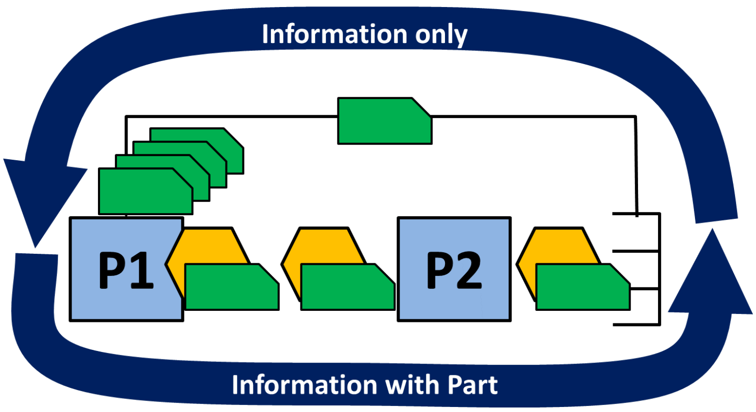

Figure 40: Kanban Loop (Image Roser)

6.2 The Speed of the Card

For the outward leg (downstream), the information is attached to the part and moves with the part. Hence, on the downstream path the information cannot be faster than the part. Regardless if it is physical information or digital information, the speed is the same as the corresponding part. Therefore, we cannot influence the speed through our choice of data or physical information.

On the return leg, however, the information is on its own. Therefore, on the way back the (physical or digital) information should go as quickly as possible. The faster the movement of the information, the shorter the replenishment time and the faster the turnaround time. Once the information is back in the production queue, the information has to wait and the speed is no longer relevant.

Digital information can move much faster than physical information. Hence, the digital information will always be faster – once the information is in the system. The picture may be more mixed, however, if we take the surrounding actions into consideration. Regardless if the system is physical or digital, the information starts with the removal of the part.

A physical card has to be transported back to the origin by someone. Usually the cards are collected at a post box, and transported back to the first machine at regular intervals.

For digital systems this would require the additional work of scanning (barcode, RFID chips, or similar) or typing. This may happen while the part is taken out of the inventory, or shortly afterwards for multiple removed parts as a batch. Once the data is back in cyberspace, the speed can be (almost) lightning fast. On the other end it may appear on a monitor, or as a print out, which again would have to be transported.

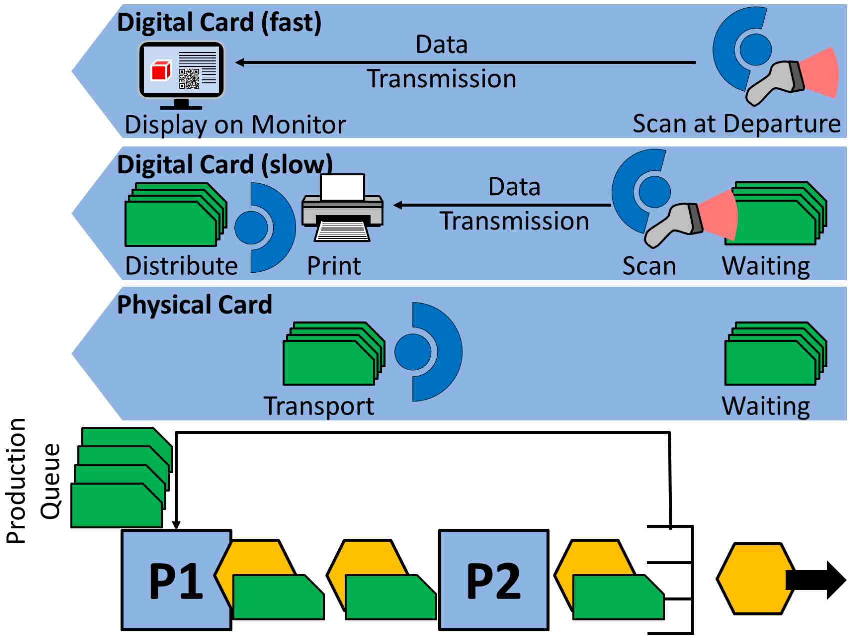

Figure 41: Digital vs Physical Kanban (Image Roser)

Depending on how the collection and distribution of this digital information is managed, the digital information may be slower than physical information. In any case, on **short distances** there won't be much difference. A physical card may take one or two hours, and a digital information may also take (depending on the set-up) thirty minutes to two hours. Not too much difference for a replenishment time usually measured in days.

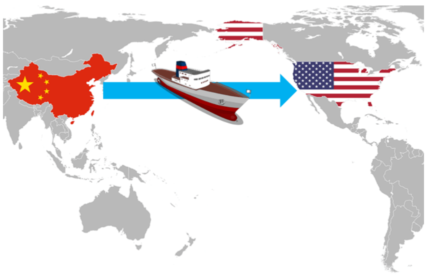
Figure 42: Shipping from China to the USA (Image Roser)

This becomes very different, however, on **longer distances**. Here the electronic transmission of data significantly outpaces that of physical data. Assuming you are receiving your goods in America by ship from a supplier in China, would you send the reorder information by postal mail? By ship this would add weeks to the replenishment time. By air it still would add days. A digital transmission, on the other hand, would be almost instant, and even with handling take only minutes or hours.

Hence **there is a strong case to use digital transmissions for long distances**. It does not have to be as far apart as America and China; even if your supplier sits in the next village over, a digital transmission may be beneficial. But before you make the decision based on the speed of the information flow alone, please note that there are more things to take into consideration.

6.3 Understanding the Situation

Another important factor is the ability to understand the current situation. Which jobs or cards are where? Here it is important to make a distinction between the manager's view from the office and the operator's view from the shop floor.

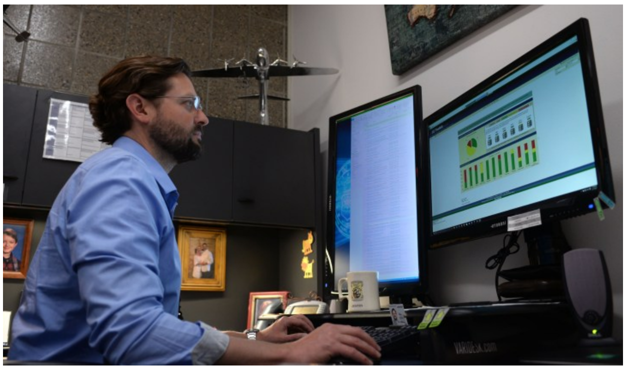

Figure 43: A manager accessing data (Image Thomas Karol in public domain)

The manager likes to click on his computer and see the current situation in his ERP system. This works anywhere in the world as long as he has a connection to the database. This is often even the preferred way for a manager to access data, and many shop-floor-related managers spend way too little time on the shop floor.

Figure 44: A worker accessing data (Image style-photographs with permission)

The worker, on the other hand, wants to see the situation on the shop floor, without requiring him to log into a system that he may or may not be familiar with, or for which he may not even have access rights. Hence he prefers the physical version of information.

Managers sometimes have a tendency to see themselves as more important than the worker, and think if it is good for them, it will be good for everybody. Unfortunately, this is not necessarily true. Especially with shop-floor-related data, the people on the shop floor have a much more frequent and much more urgent need to understand the situation than a manager who may not even look at the data but insists that it must be accessible. This ties in very closely with visual management.

In sum, a physical representation of data is often much more beneficial for the shop floor where the workers frequently need to know the upcoming orders to provide material and plan the manning of the machines.

On the other hand, if there is no shop floor but only logistics, then this is somewhat less relevant. The truck driver will not open his truck to check what is loaded but rather refer to the loading papers. Hence if you are not on a shop floor, then there is much less benefit of physical data.

6.4 Continuous Improvement

Figure 45: Kaizen (Image Rawpixel.com with permission)

Lean lives and breathes continuous improvement. As for digital versus physical information, the ability to do continuous improvement is much, MUCH easier in physical systems. Changing the handling of information (think kanban cards) is much easier if it is a paper card than if it is a digital system. Verifying your workload and inventory levels to adjust the manning or the number of cards is also much easier in physical systems.

Of course, this can also be done in digital systems. However, here you need a programmer or specialist for the digital system. And those are always in much higher demand than they are in supply. You have to get the programmer, convey him what you want him to do, and then hope that he understood what you actually wanted. Even if he got you what you wanted, you may not able to use trial and error to try things out. Usually, continuous improvement grinds to a standstill if computer systems are involved.

6.5 How About … Both?

You may be thinking, If both physical and digital systems have advantages, could I do both systems and get the advantages of both? Don’t!

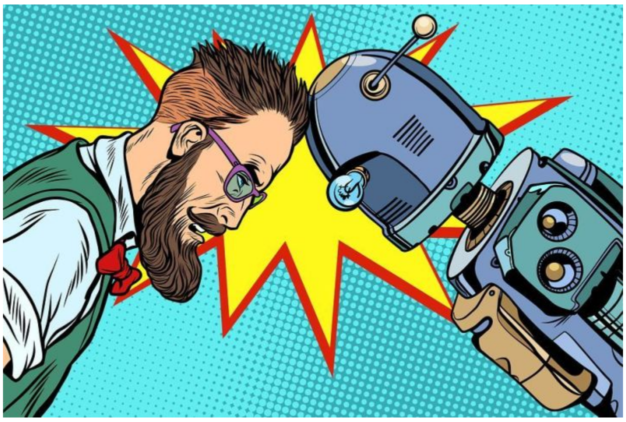

Figure 46: Robot Fight (Image studiostoks with permission)

Seriously, don’t do it! Obviously you will have twice the work by creating two systems. But the much bigger problem is that these two systems will have differences. The digital systems says A and the physical system says B. What do you do? The operator on the shop floor has to get a single clear signal, not two conflicting messages. While everything works in theory, the practical situation will be much more chaotic. Do your people a favor and do not have cards on top of a digital system.

It is fine to print out paper versions of the digital system, however. It is also possible but less common to scan cards to update the physical system. But only one can be the master of the data. Either the cards are just a dumb copy of the digital world, or the digital world is just a dumb counter of the physical cards. If there are differences between the systems, one of them has to give way to the other … and the physical information is more likely to be the correct one.

6.6 Summary

So, what should you do? My recommendation is to use digital information for longer distances (across different plants). This will give you a faster replenishment time, and since it is probably not a production line but a logistics process, the visual management is not quite as critical. If yours is information flow within the plant, physical information may be much easier to understand and improve. Now **go out, get the information flowing, and organize your industry!**

7 How to Make CONWIP Loops

Christoph Roser, February 12, 2019, Original at https://www.allaboutlean.com/conwip-loops/

Figure 47: CONWIP-go-round with different jobs: like a roller coaster with different people (Image MaxPixel in public domain)

CONWIP (Constant Work in Progress) is an easy way to establish pull production for custom-made products. Traditionally there is only one large loop for the product. However, there may also be situations where it is sensible to split a longer CONWIP loop into smaller segments. Let's have a look at the details.

7.1 Introduction

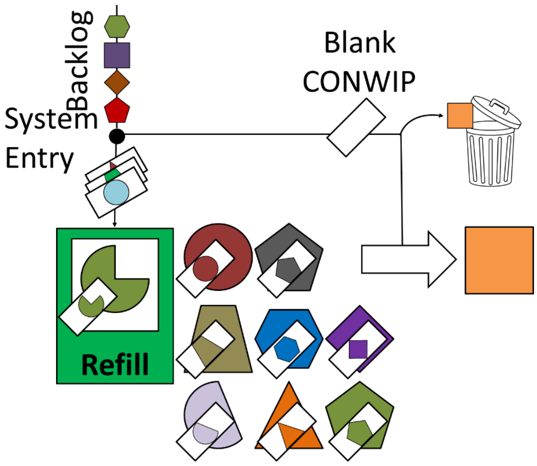

Figure 48: CONWIP system (Image Roser)

The most famous pull system is kanban. However, since the part information is constantly attached to the kanban card, it is not suited for small quantities, exotics, custom orders, or in general any product that you do not want to have in stock (Make to Stock, or MTS) but rather produce only when you need one (Make to Order, or MTO).

CONWIP is used to create pull production for Make to Order parts. A card goes around similar to a kanban card, but the product information is removed after completion, and the next product to be produced is taken from a backlog list of open orders. I wrote a whole series on CONWIP starting with Basics of CONWIP Systems. This post looks in detail at the placement of the CONWIP loops. I discussed a bit of it already in my post Frequently Asked Questions on CONWIP Systems – Part 2, but I felt like I need to go into more detail.

7.2 One Single All-Encompassing Loop

One option that is always possible with CONWIP is to make one all-encompassing loop. This is easiest for straight flow shops.

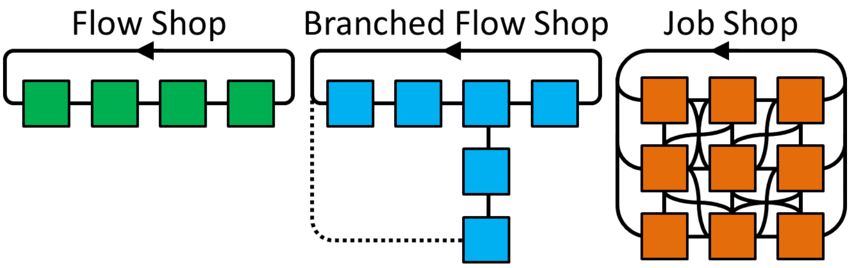

Figure 49: CONWIP Loops Branched Flow Job Shop (Image Roser)

If there are branches (e.g., if the product needs sub-components), then the CONWIP card would have to initiate the production or purchase of these sub-components too. The available CONWIP card gets assigned a job from the backlog, which automatically releases additional work orders for the sub-components, but these are all associated with the single CONWIP card for the end product. No new CONWIP cards are introduced for the sub-components. The start of the work on the sub-component may be timed to start at the same time, later, or even earlier than the work on the main component to make sure the component is likely to be ready when it is needed.

If the system is a job shop, the routing of the product is different for different products, and the final route may not even be known in advance, then it is also possible to have a single all-encompassing CONWIP loop. The product just moves through the system with the CONWIP card.

This is actually the approach as intended by the inventors of CONWIP, Hopp and Spearman. However, while it is the easiest to set up, there may be situations where you could benefit from smaller loops. It is also sometimes suggested to use a separate CONWIP loop for every possible routing. However, for job shops this would quickly turn impractical. For flow shops on the other hand it would rarely be needed.

7.3 Loops for Different Segments

It is also possible to make loops for different segments or groups of machines. The challenge is now managing the transition from one loop to the next loop. The output of the preceding loop is now the backlog for the next group.

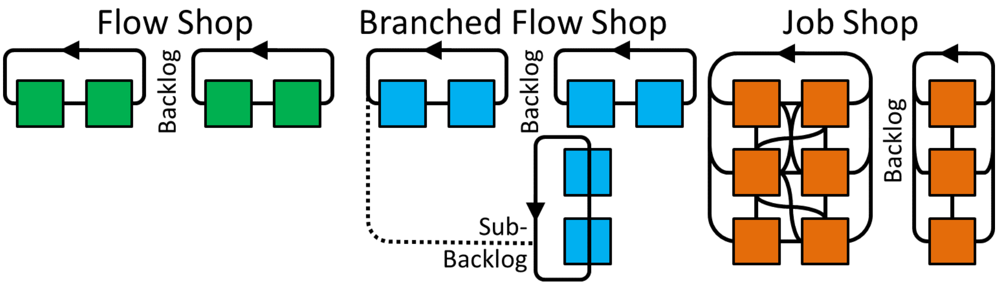

Figure 50: Split CONWIP loops Branched Flow Job Shop (Image Roser)

This is extra work. Hence, you should not do this just because it is possible, but only if there is a valid reason for it. This is based on my previous post Frequently Asked Questions on CONWIP Systems – Part 2. A couple of the situations are exceedingly rare for Make to Order CONWIP cards.

- **Supermarket in front of the customer**: This is a common place for the CONWIP loop to end. The loops the customer uses for your products may not be relevant for you.
- **Supermarket for large distance between processes:** This may help you to reduce fluctuations by creating separate buffers in separate CONWIP loops.

- **Supermarket for change of responsibility:** Definitely! The ability of humans to blame others may be easier to handle if they are on separate CONWIP loops. So, unless you enjoy the blame-game, a separate loop may be helpful, although it is not a cure-all solution.
- **Supermarket in the case of high demands on flexibility and reaction time:** Shorter loops make you more agile and allow for re-prioritization at different backlogs. Be aware that this is more work though!
- **(Rare)Supermarket for process-specific lot-size differences:** A rare situation for Make to Stock, and exceedingly rare and unlikely for Make to Order. You would need an order that you split into different batches with separate CONWIP cards, and subsequently have a machine with a different batch size. If this applies to your product, consider revising the batching and CONWIP rules rather than splitting a loop into two.
- **(Rare) Supermarket when creating different variants:** While this is common for Make to Stock, it is exceedingly rare for Make to Order products, since the definition of a Make to Order product is that it is assigned for a specific customer.

7.4 Loops for Individual Machines

Finally, you can give every single machine its own CONWIP loop. This is of course the most complex solution. However, it gives you the finest level of individual workload of the machines and allows frequent re-prioritization. This is actually now becoming very similar to the POLCA (Paired-Cell Overlapping Loops of Cards with Authorization). See my series of posts starting with "What Is POLCA?" for more details.

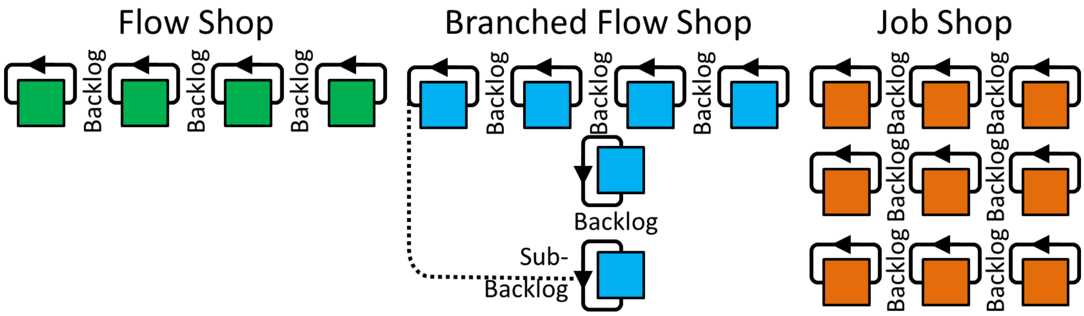

Figure 51: Single Process CONWIP loops Branched Flow Job Shop (Image Roser)

7.5 Which One Is Right for You?

That is a tough question. The answer is of course, "It depends." Try to find a compromise between the additional effort of creating and maintaining separate loops, and the benefit these loops give you. Do not underestimate the effort to maintain loops! It is easy to make a few lines on paper, but it takes constant work and vigilance to keep a loop in good condition. Here especially the re-prioritization can be challenging (although there are many companies that do it successfully using POLCA). For job shops you have the additional problem of managing the routing. You would have to do that anyway, but it adds another layer of complexity to the problem.

7.6 On the Hand-Over

Regarding the hand-over of the material, there are two options. You could consider completed parts outside of the previous loop, and the CONWIP card is sent back as soon as the part is completed at the last process within the loop.

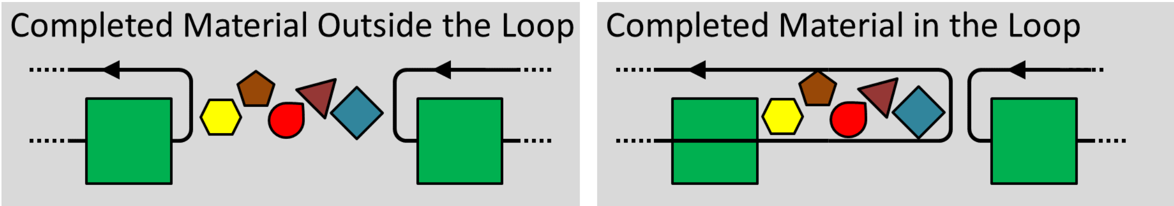

Figure 52: CONWIP Material Hand Over (Image Roser)

My preference and recommendation, however, would be to **keep the CONWIP card attached with the parts in the loop until the subsequent process or customer actually takes the part** (or in the worst case, the part is scrapped or disassembled if the customer's order was canceled). Otherwise you risk overproduction, and your system is no longer a pull production, losing all the benefits of pull production.

So, I hope this was interesting to you, even though it went into a more-detailed aspect of pull production. In any case, **go out, make sure your material is flowing in pull production, and organize your industry!**

8 Delivery Sequences: FIFO, LIFO, and Others

Christoph Roser, February 19, 2019, Original at https://www.allaboutlean.com/fifo-lifo-etc/

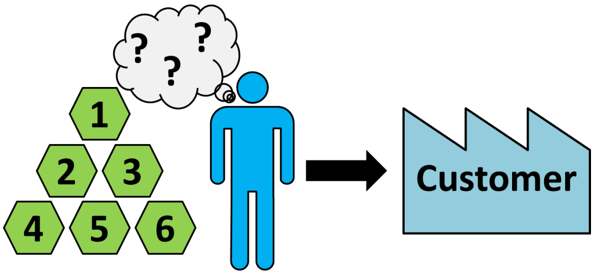

Figure 53: Pick One of Identical Parts (Image Roser)

Sometimes, when you need a part or product from your inventory, you may have more than one item in stock. Which one do you pick? In this blog post I want to present a few strategies for choosing which item to take. The most famous one is FIFO, but there are more options available. In my next post I will present similar strategies if you need to start production before the item becomes available.

8.1 Introduction

As said above, this post looks into strategies for deciding which part out of a larger group of parts to pick. The demand could be a customer order or a machine that needs parts to continue working. In any case, you need a part or product, and you have this item readily available. If you have only one item (or only exactly as many as you need), then the rule is simple: Take the one item since you have no other choice. However, if you have more inventory than you need, you have to choose which one to take.

Here we need to distinguish two fundamental different situations:

1) You have (nearly) **unlimited capacity** to bring material to the place where it is needed. This could be, for example, a delivery process where your ability to ship stuff from your warehouse is much larger than the demand of the customer or the process that consumes the material. In this case you can send anything the customer needs as long as you have it in stock, and the question is only which products to send if you have more products of that type than needed. In other words, if the customer needs three type A's and you have ten, which three out of ten A's will you send?

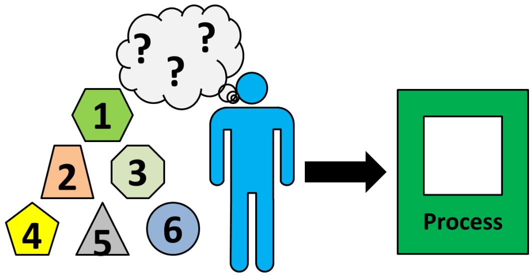

Figure 54: Pick One of different Parts (Image Roser)

2) You have **limited capacity** to bring material to the place where it is needed. These are often production processes. You cannot produce all at the same time, but have to make a sequence for production. Here you must look not only at the one product type that the customer at the end wants, but must also consider all products that have to go through the process. In other words, if the customers need three type A, five Type B, and one Type C product, which one do you start with? This will be discussed in the next post.

8.2 FIFO – First In, First Out

The most famous and best known method is FIFO (first in, first out). The first item that entered the inventory is the first item that will be removed. The item that is waiting in inventory the longest will always be taken first.

The big benefit of FIFO is that the parts maintain the sequence in which they arrived. This has multiple advantages. If there is a design change, the not-changed parts are sold first. If you need to update all parts, it is easier to determine which parts to update. Similarly, if you notice a systematic error, it is easier to track which ones to fix, and it may even help you understand the cause of the error. This sequence also avoids parts getting much older than others.

FIFO works with a random-access inventory, where you can get any part at any time. However, the advantage of FIFO is that it can also work with a sequential-access inventory, where you can get a part only if the part in front of it is removed. This is even another advantage since through the design of the inventory you can force your people to always take the next item in a queue. Even if it is not forced, you can make it easier to simply take the next one instead of jumping the line.

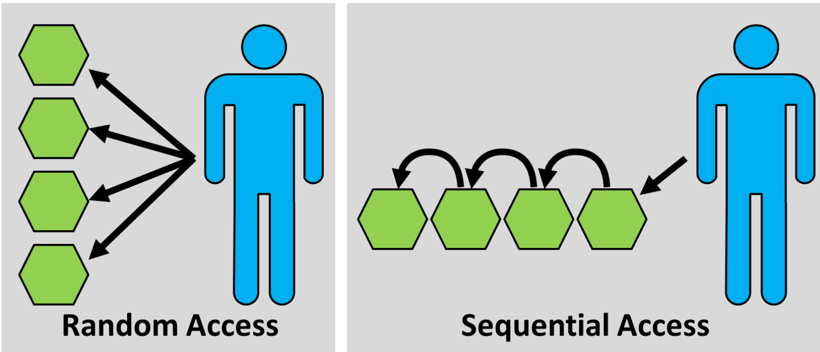

Figure 55: Random and Sequential Access (Image Roser)

You often find this in the supermarket. The spice rack is always restocked from the rear, so the "oldest" spice is always at the front. This avoids spices getting excessively old and losing their flavor.

Figure 56: Spice Rack in Supermarket as a FIFO example (Image Roser)

8.3 LIFO – Last In, First Out

Figure 57: Sand Stockpile as a LIFO example (Image Peter Craven under the CC-BY 2.0 license)

LIFO (last in, first out) is exactly the opposite of FIFO. You always remove the item that has spent the shortest time in the inventory. The big disadvantage is that the oldest items will get older and older and eventually expire before being used. Hence, LIFO is a rather bad strategy. **Do not use LIFO if you can avoid it!**

It is used only if the type of storage forces LIFO on you (i.e., due to the structure of the inventory, you cannot access the oldest item without removing the others first).

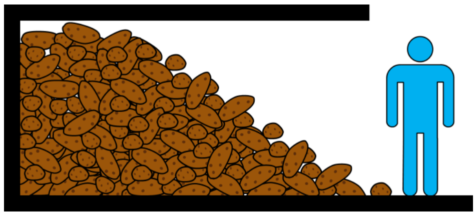

Figure 58: Potato Cellar as a LIFO example (Image Roser)

A common example is piles of bulk materials like coal or sand. The material at the very bottom of the pile is the oldest. However, to get to it you would have to remove all the other material. Other examples are a stack of material (e.g., CDs on a spindle, boxes on top of each other) or storage locations that are accessible only from one side (e.g., a potato cellar or other bulk food storage on some farms).

Even if the system forces LIFO on you, it is common to have multiple LIFO storages and periodically empty them completely to avoid excessive aging of the goods. Hence, the potato cellar has to be completely empty once per year, and you use up one of your sand piles completely before moving to the next.

8.4 FEFO – First Expiry, First Out

With FEFO (First expiry, first out) you would need to track an expiration date of your goods. The strategy is always to remove the material with the earliest expiration date. This strategy requires random-access storage – or lots of moving goods around to get to the item you want. For most durable goods, this is not necessary. However, if you deal in perishable goods, this may make sense, and may be an even better alternative that FIFO.

Figure 59: Is it fresh? (Image TerriC in public domain)

For example, if you make yogurt, you may make a batch from very fresh milk and then a second batch from just barely fresh-enough milk. Even though the second batch is newer, you may give it an earlier expiration date. With FEFO this second batch would be sold before the first batch since it expires earlier. Hence overall you will be able to sell fresher products to the customer and extend the shelf life at the customer. This comes at the expense of the additional effort of setting and tracking an expiration date.

You may even use slightly different related expiration dates. Besides the **Expiration Date** when the product is likely to go bad, there may be a **Best Before Date** before which quality is next to certain. Some companies also use an **End of Life Date,** after which the product may be dangerous, or a **Removal Date** by which the material has to be removed from stock for whatever reason. They all work like FEFO; just substitute the type of date you want.

8.5 FEMAL – First Expiry Minimum Available Lifetime

Another variant is FEMAL for First Expiry Minimum Available Lifetime. You deliver the part that has the first expiration date (like FEFO), but with the additional condition that the expiration date must be at least as big as the expiration date desired by the customer.

For this there is even a slightly modified variant if your expiration date depends on the use at the customer. An example in reality would be a polymer resin that lasts 6-12 months if stored at -18°C; but lasts only 4-30 days if stored at 21°C. If you take the resin out of storage and bring it into the factory, it ages faster. Returning it to storage afterwards results in an overall shorter expiration date. In this case there is FESAL (First Expiry Shortest Ambient Lifetime). Here you update the expiration dates based on the previous history of the product. Many thanks to Juan Carlos Viela for the update in his (Spanish language) article on FIFO or FEFO or, FESAL.

8.6 HIFO – Highest In, First Out

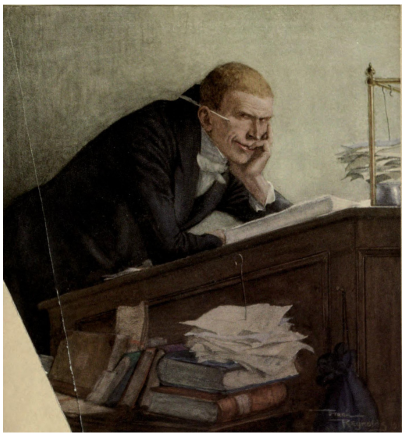

Figure 60: Twisting the numbers … (Image Frank Reynolds in public domain)

Another strategy is HIFO (highest in, first out). If you purchased, produced, or otherwise obtained identical products at different costs, you use the most expensive (highest price) first. **The advantage: NONE!** Luckily, this method is very rare.

Seriously, this has no advantage! Instead it has the disadvantages of missing out on all the benefits of FIFO or FEFO.

It is still sometimes done for bookkeeping reasons. Depending on how you do your accounting, it may look better on paper to first get rid of the most expensive items from a group of identical items. Your inventory value may be (on paper) lower if you use HIFO. In reality, it has mostly disadvantages, but on paper it may look better. Depending on the country you are in, it may reduce your taxable income.

8.7 LOFO – Lowest In, First Out

LOFO (Lowest in, first out … but LIFO was already taken) is exactly the opposite of HIFO. The goods that you purchased cheapest are sold first. This is extremely rarely used if you want to increase the book value of your inventory or your taxable income. Overall, while HIFO is rare, LOFO is mostly an exercise in academic rigor to include all possibilities, no matter how stupid they are.

8.8 First One Found (Not Really a Strategy)

Figure 61: Where is it ? (Image WavebreakMediaMicro with permission)

The last strategy is not really a strategy, but a common approach in industry. If a part is needed, the worker simply takes the first one he comes across. If the material is organized, this may be the closest one. If the material is less organized, this may require some searching until he finds an item somewhere.

This "strategy" happens usually if there is no real strategy in place. It misses out on all the benefits of FIFO or FEFO, and is a sign of sloppy management.

Overall, if you don't know what to use, use FIFO. Alternatively, for perishable goods you may use FEFO. Stay away from the rest unless they are forced upon you. Now, **go out, get your material in order, and organize your industry!**

9 Production Sequences: FCFS, EDD, and Others

Christoph Roser, February 26, 2019, Original at https://www.allaboutlean.com/fcfs-edd-etc/

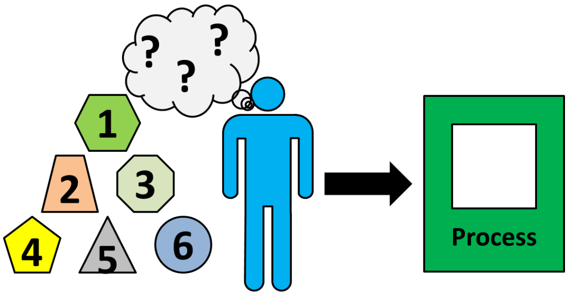

Figure 62: Pick One of different Parts (Image Roser)

In my last post I looked at delivery sequences like FIFO, LIFO, etc. This second post looks at simple production sequences where you do have to deal with limited production capacity. If you cannot make everything at once, you need a sequence in which you process the parts.

9.1 Introduction

As mentioned in the previous post, we now have to deal with limited capacity. We cannot process all items at once. This could be most production systems, where you typically start with one item, and as the item proceeds through the system you start to process more items. Another less-common example would be a truck, ship, or plane that cannot fit everything you want to deliver. Some items would have to wait for the next transport. Overall, you have **limited capacity**. In comparison, the capacity of a warehouse to ship items usually exceeds the demand of the customers, and hence the warehouse has a much higher capacity than the need of the customer.

9.2 FIFO (First In, First Out) and FCFS (First Come, First Served)

The first strategy is the easiest: FIFO (first in, first out) and FCFS (first come, first served), which are two names for the same thing. The orders are processed in the sequence they arrive. The people at the supermarket get to the checkout in the sequence they have arrived at the queue.

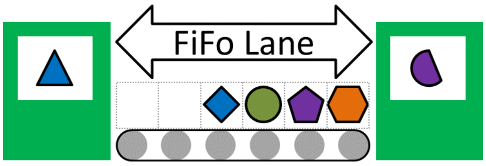

Figure 63: FIFO lane with different parts (Image Roser)

This is an excellent approach to sequence material between processes in a flow line. I have written a lot of posts on this topic exactly since it is so common and useful. See Theory and Practice on FiFo Lanes for a start and have a look through my list of posts.

Maintaining the sequence makes it much easier to find and fix problems and changes. It is also possible to pretty much force FIFO though the use of a FIFO-type inventory system like a rolling lane. If it is easiest for the worker to pick up the first part, he will do so.

FIFO will be the backbone of any flow production. Only at selected locations like at the beginning and in supermarkets will the sequence be defined and changed (see my posts Ten Rules When to Use a FIFO, When a Supermarket). Yet even at these locations, FIFO can be used. At a supermarket the kanban cards can go back to the process in the sequence that the parts are taken out (hence FIFO). At the beginning of your value stream, you could process the orders in the sequence they arrive (hence also FIFO).

However, especially for make-to-order products and at time of potential stock-outs, a pure FIFO does not always make sense. At least at some location you may want to change the sequence. Here are a few simple alternatives.

9.3 EDD – Earliest Due Date

Out of the pool of available open jobs or parts to process, the sequence is based on the due dates. All parts or jobs in the pool for that machine are evaluated, their due dates are compared, and the job or part with the earliest due date is processed first, followed by the second-earliest due date and so on.

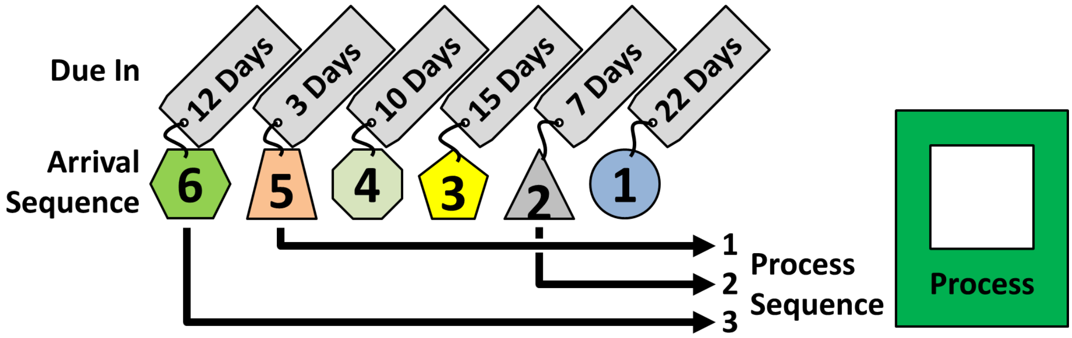

Figure 64: Earliest Due Date Example (Image Roser)

This is helpful if the sequence of the parts and the sequence of the due dates differ significantly, as in the example above. Job #5 is due in three days, which is the earliest due date. Hence job #5 should take precedence over all other jobs to ensure that it is ready when the customer wants it. Next would be job #2, which is due in seven days, and so on.

For this to work, you would need a due date for all of your jobs (obviously) and a way to figure out what job has the earliest due date. If you have your open orders in an Excel file, you simply sort by due date. However, if the worker has to look at the paperwork of all parts on the shop floor, a lot of time will be wasted in figuring out what to do next.

This method is a valid approach as long as all of your parts have a similar lead time (i.e., take a similar time to be processed).

9.4 SPT – Shortest Processing Time

Another strategy is SPT (shortest processing time). You add all the (estimated) processing times for each job to see how many minutes (or hours) of work each part needs until completion. Let me show you an example before I tell you why **this method is very flawed**.

In the image below, job #6 has the shortest remaining processing time of thirty minutes, and hence would go first. This is followed by job #4 with the second-shortest processing time, and so on.

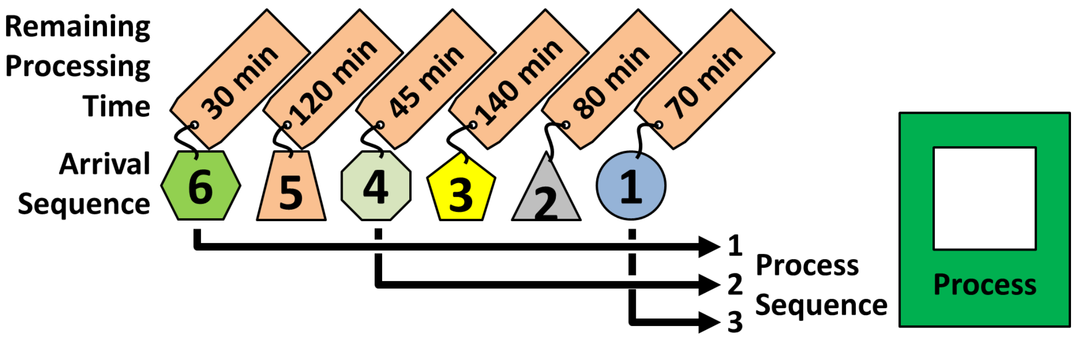

Figure 65: Shortest Processing Time (Image Roser)

The advantage of this method is that it gets the material out quickly. Like when you are doing your chores, it is a satisfying feeling to complete a few short chores. However, this is a one-time effect, since once the parts with the shortest processing times are out, you are stuck with the long processing times. The overall lead time will not really budge, since whatever time you save for the quick parts will be added to the lead time of the complex parts.

A complex part with a lot of processing time may potentially be stuck in the system forever since there are always new jobs arriving with a shorter processing time. This is the big drawback. It makes the production feel good about churning out parts, but neglects the needs of the customers for more complex parts. I strongly recommend you not use this.

Even worse would be a longest processing time (LPT) method, which luckily I have not found in literature. You would start everything but complete nothing! Hence again **stay away from production sequences that are solely based on work content!**

9.5 ODD – Operational Due Date

Figure 66: Savior … or master of disaster? (Image Boffy b under the CC-BY-SA 3.0 license)

The operational due date uses more-advanced scheduling methods to figure out the latest a job has to start at a process in order to make the deadline. Often this is based on ERP systems. This system calculates backwards from the due date of the completed product, including estimates of processing times, waiting times, and safety time buffers to estimate when a product has to be processed where.

The sequence would then be created individually for every process or process group based on the estimation of when the part has to pass through this process group to make the deadline. The part with the earliest operational due date for the next process goes first, and so on.

Such a method can have an overview of the entire production system, and magically determine a good sequence. However, it can also magically mess it up big time. Personally, I am wary of such systems. They can work, but they also can make a god-awful mess.

Another flaw of this method is that jobs may be started long before necessary just because capacity is available, building up inventory and tying up capital. Overall it can work, but it also may not. In any case, it usually requires an ERP system or a really good paper-based system.

9.6 MODD – Modified Operation Due Date

The MODD (modified operation due date) is an acronym used for a number of different sequencing techniques based on the operational due date. They modify the ODD method using additional conditions, weights, or other considerations in order to improve the scheduling quality.

One MODD approach uses ODD as the main sequence, except the job becomes late and the processing time is larger than the remaining time until the due date. In this case the shortest processing time (SPT) of the late jobs takes precedence over the ODD rule.

9.7 AI – Artificial Intelligence

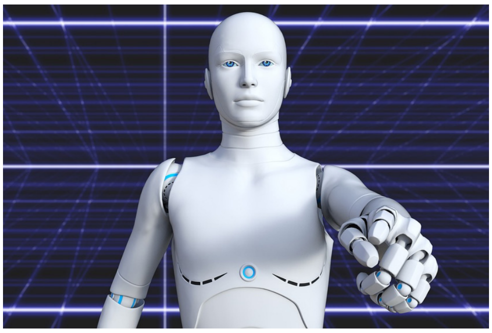

Figure 67: Humanoid Robot (Image MaxPixel in public domain)

Scheduling is a complex task, and in the past human experience often saved the day and avoided a delay. While not yet mainstream, artificial intelligence systems are starting to be used, learning the complexities of the particular production systems and giving suggestions on when to do what. This is still under development, but maybe it is common in ten years. We will see.

9.8 Loudest Yeller First (Not Really a Strategy)

Figure 68: Also a scheduling strategy … (Image Minerva Studio with permission)

Yet another "strategy" is to simply listen to whoever is making the most noise. The customer that calls and complains gets prioritized. The boss calling and doing a round of yelling will lead to a re-sequencing of the schedule. This "strategy" is of course not so hot for the overall system performance, but depending on your work environment, it may be necessary to save your career.

9.9 More Complex Sequencing Strategies

The above are merely a selection of simple sequencing strategies. Much-more complex strategies exist to optimize certain aspects of the production. These may, for example, optimize changeover times (see my series on Changeover Sequencing), leveling (see my series on leveling), using prioritization (see my series on How to Prioritize Your Work Orders), Just in Sequence (see my JIT blog posts series), and workload balancing (see my series … *wait … I have not yet written that …*) Apologies, it is on my to-do list! In the meantime, **go out, manage your production schedule, and organize your industry!**

10 Production Control with COBACABANA

Christoph Roser, March 05, 2019, Original at https://www.allaboutlean.com/cobacabana/

Figure 69: Copacabana Beach (Image bisonlux under the CC-BY 2.0 License)

The Copacabana is a very nice beach in Rio de Janeiro. Spelled slightly differently, COBACABANA is a production control approach. Here, COBACABANA stands for *Control of Balance by Card Based Navigation* (sometimes also abbreviated to *COBA*). It is an approach to manage a job-shop workload of custom orders using paper cards. A lot of paper cards, in fact, which also makes the method a bit complex, and I am doubtful if this method is practical. Let me show you how it works.

10.1 Introduction

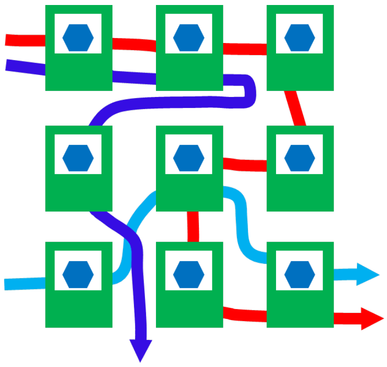

Figure 70: Job shop (Image Roser)

COBACABANA was developed by Martin Land and improved by Matthias Thürer (Side note: I do like his book on Card-Based Control Systems quite a bit).

COBACABANA aims to help with the management of job shops, especially the task of keeping your processes busy without overloading one or idling the other. This is a daunting task that many have tried but few succeeded at. Furthermore, it does so using only paper cards. As a result, the method is rather complex.

COBACABANA breaks the flow of a customer order into two main parts. When the customer order arrives (or is accepted), it first goes into a **pool of open jobs** waiting for production capacity. Once the order is released for production, the job **travels through the shop floor**, which is the second part.

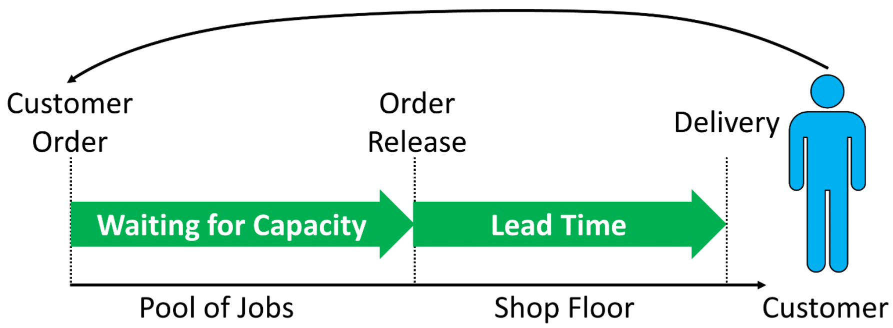

Figure 71: Two Phases COBACABANA (Image Roser)

A push system would simply release these open jobs into the factory. A pull system, however, controls and limits the amount of work on the shop floor. COBACABANA is a pull system, as it releases work into the shop floor only if capacity is available.

10.2 The Release from the Pool of Open Orders

A customer orders an item. This item would need to be processed at different machines within the system. The time needed for this order is estimated for every machine. Assume an order of parts would need six hours of milling, eight hours of hardening, and four hours of grinding. This order first goes into the **pool of open orders** before being released for production.

To keep an overview of the workload, each process has a set of **acceptance cards** representing a certain workload that is already in the pool for this order. This is shown below. In this example, each card represents two hours of work. There are already open orders in the pool equivalent to ten hours of milling, eighteen hours of hardening, and eight hours of grinding. For easier visualization, each process has its own color.

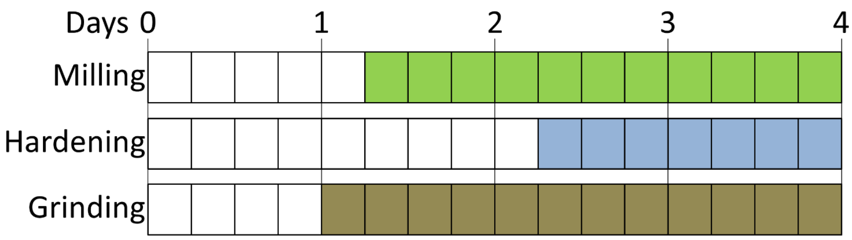

Figure 72: Acceptance Stack Original (Image Roser)

With the new order, we need another six hours of milling, eight hours of hardening, and four hours of grinding. Hence we remove three cards (six hours) from milling, four cards (eight hours) from hardening, and two cards (four hours) from grinding. These nine cards are attached to the open order in the pool.

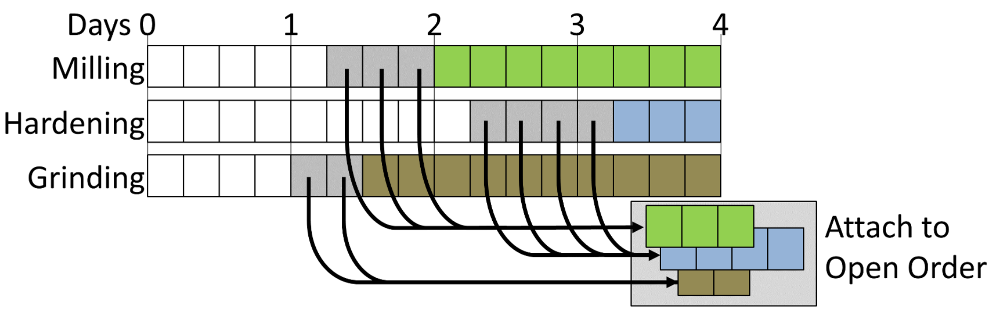

Figure 73: Acceptance Stack Removing (Image Roser)

Afterwards, the acceptance cards stack would look like the image below. Hence, the person managing the stack always has a good overview of the work that is already committed to the customer but not yet released.

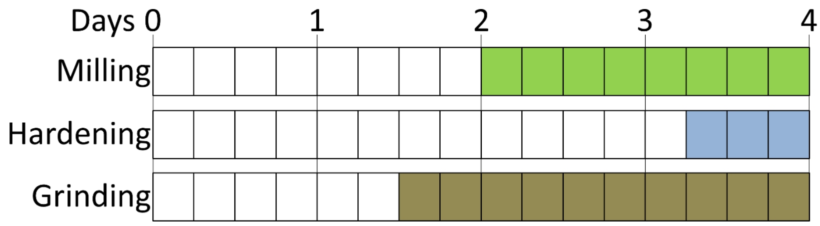

Figure 74: Acceptance Stack After (Image Roser)

10.3 The Shop Floor Release Process

A similar method is used for the release of the orders. The backbone of every pull system is a limit on the work in the system. Most pull systems merely count the number of jobs, but COBACABANA actually measures the workload. The approach is similar to the pool of open orders above, except that there is a fixed number of cards. These cards are called **Release Cards**. The total number of cards represents the maximum amount of work you want to put into the system at the same time.

An example is shown below. Each card represents a certain amount of work. For simplicity's sake, I also assumed each card is the equivalent of two hours' worth of work (although the original author suggested 1% of the desired maximum workload – in which case you would end up with one hundred cards per station).

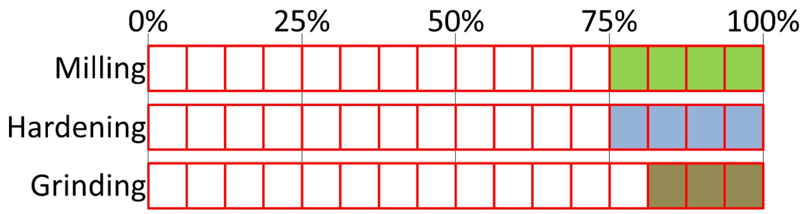

Figure 75: Release Stack Original (Image Roser)

When releasing a job from the pool to the shop floor, the acceptance cards go back to the pool. The release cards are removed from the release stack and attached to the open order. You must have enough release cards to release the order; otherwise the order cannot be released and has to wait in the acceptance pool. The jobs are released according to their planned release date, with the hope that the most urgent job will be the next one in line for the shop floor.

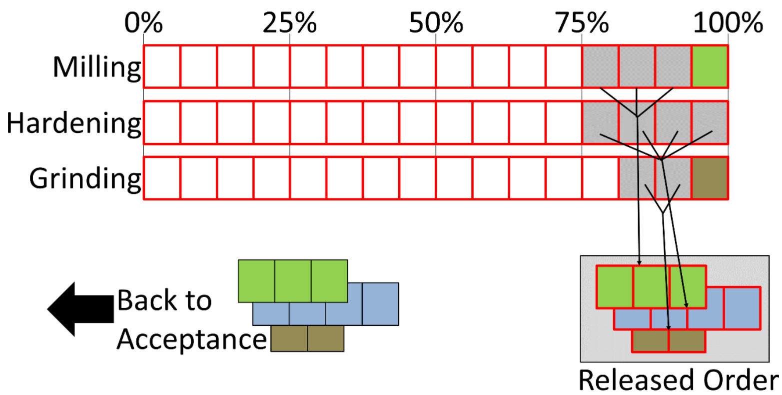

Figure 76: Release Stack Removing (Image Roser)

In our example, we just barely had enough cards for hardening to release the order to the shop floor. The release stack afterwards is shown in the image below. Hence, the order release always has a good overview of the current workload of each station (although again at the cost of quite a lot of cards).

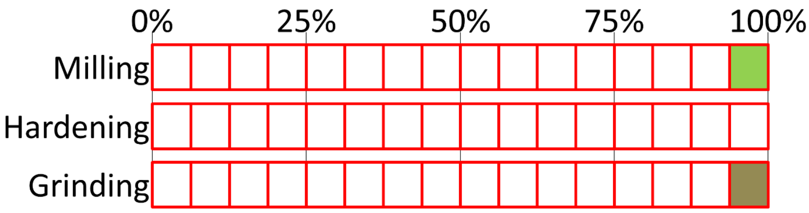

Figure 77: Release Stack After (Image Roser)

Once the job is completed, the release cards go back to the release stack to be available for the next jobs to be released.

10.4 Some Tweaks

Since the method was originally presented by Land in 2009, the method has been adjusted a bit. Quite a bit actually. While it does have the same name, consider it COBACABANA 2.0 (or even 3.0).

Literature describes a use to **estimate a due date** based on the data from the release pool under the assumption that the lead time on the shop floor is relatively stable (although I am a bit doubtful of that).

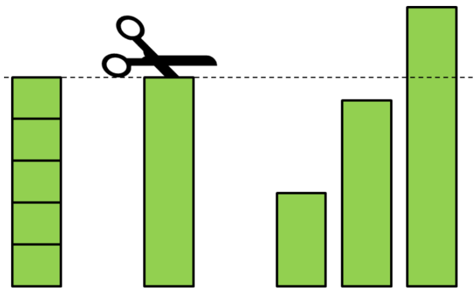
Figure 78: COBACABANA Card Sizes (Image Roser)

To avoid the inflationary large number of cards for different jobs, Thürer recommends instead using only a **single custom card per job and workstation whose size (i.e., length) represents the workload for this station.** This significantly cuts down on the number of cards, but requires the cutting of custom-sized cards for every job. Since this update to the method, they have figured out that cutting cards is cumbersome, and now recommend **cards in different standard sizes**. Here it seems three sizes (small, medium, large) is good enough in terms of accuracy. You simply pick the card whose size matches your workload best.

Additional cards are added. A single **pool card** per open order in the pool is created. The acceptance cards from above are given to the salesperson. The height of the stack of acceptance cards at the salesperson gives him an estimate on how long the delay of orders in the pool will be. The pool card stays with the open order.

On the shop-floor side, an **operation card** was added to the release cards. The release card is cut to the correct size and stays with the planner. The height of the stack of release cards represents the workload. New jobs are released only if there is enough space in the stack for them. The operation card travels with the job to the processes. After completion, the operation card returns to the planner, who then also releases the release card back into the stack.

10.5 Review

Figure 79: Manager with Playing Cards (Image logoboom with permission)

Overall, the approach looks a bit theoretical. I have the feeling that there is not yet any beneficial real-world application, and that not all kinks and issues of the method have yet been fully debugged. There is definitely ongoing work (e.g., workload representation changing from *many single cards of same size* to *custom cut card size* to *different standard card sizes…*)

The number of cards (or the cutting of custom card sizes) and its complexity is an issue for me, albeit in my opinion there simply are no good AND easy solutions for job shop control.

I do like the focus on a purely paper-based system without any ERP system that most people don't really understand. In that aspect, COBACABANA is quite unique. It is the only purely paper-based method that manages the workload of a job shop.

In any case, there seems to be an ongoing stream of publications on the topic, mostly by Thürer, and maybe in the future there will be more updates and changes to make the system easier to use. In the meantime go out, use kanban, CONWIP, or POLCA, and organize your industry!

10.6 Selected Sources

- Thürer, Matthias. Card-Based Control Systems for a Lean Work Design: The Fundamentals of Kanban, ConWIP, POLCA, and COBACABANA. Productivity Press, 2017.
- Land, Martin. "Cobacabana (Control of Balance by Card-Based Navigation): A Card-Based System for Job Shop Control." International Journal of Production Economics 117 (2009): 97–103.
- Thürer, Matthias, Mark Stevenson, and Charles W. Protzman. "COBACABANA (Control of Balance by Card Based Navigation): An Alternative to Kanban in the Pure Flow Shop?" International Journal of Production Economics 166 (August 1, 2015):143–51. https://linkinghub.elsevier.com/retrieve/pii/S0925527315001620.

11 The Toyoda Model G Loom (with Videos)

Christoph Roser, March 12, 2019, Original at https://www.allaboutlean.com/toyoda-model-g/

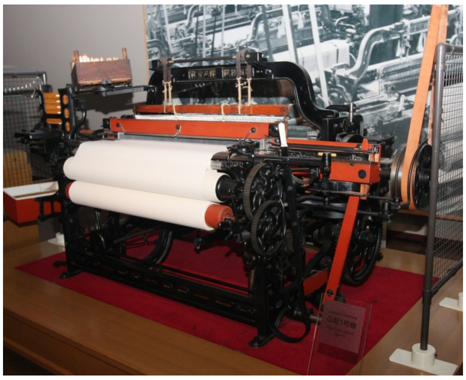

Figure 80: Toyoda Model G automatic Loom (Image Morio under the CC-BY-SA 3.0 license)

Toyota Motor originated from the Toyoda loom factory, where Sakichi Toyoda invented looms. Probably the most famous one is the Toyoda Model G Automatic Loom. This loom touches on many points that are part of the Toyota Production System and lean manufacturing. During my last visit to Japan in September, I made some videos detailing many of the features of the Toyoda Automatic Loom from 1924. Be advised: Lots of images and videos ahead!

11.1 Introduction to the Model G

Figure 81: Toyoda Model G Automatic Loom (Image Roser)

The Toyoda Model G loom, introduced to the world in 1924 by Sakichi Toyoda, was probably the most advanced loom of its time. Its quality and productivity were unparalleled. Using many mechanical gadgets, one unskilled worker was able to supervise thirty to fifty looms simultaneously. Hence it already incorporated the idea of karakuri kaizen and jidoka.

Licensing the loom to the Platt brothers (United Kingdom) generated significant cash flow for Toyoda. Legend has it that these funds were used to establish the Toyota Motor Corporation in 1937, although in reality these funds were spent on a large bonus for the employees, and Toyota Motor was established using conventional financing (and they changed the name from Toyo**d**a to Toyo**t**a).

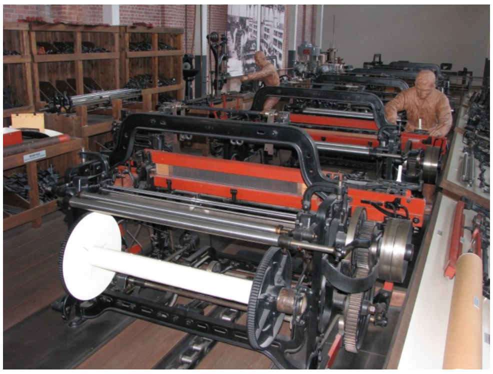

Figure 82: Model of the assembly line (Image Roser)

While the Platt brothers initially praised the loom, they had difficulties manufacturing it due to their low-quality production. Toyoda, on the other hand, produced it very successfully and built its first assembly line for the Model G loom in 1927.

The following images are from the highly recommended Toyota Commemorative Museum of Industry and Technology in Nagoya, Japan.

11.2 Preparation of the Shuttle

The shuttle is the device that moves back and forth transporting a spool of yarn. In preparation, the yarn has to be threaded through a small hole in the shuttle. Previously this was done by putting the yarn near the hole and then sucking it through with your mouth – breathing in all the dust and dirt around the hole.

Toyoda used a neat trick by making a small, open channel in the other direction. The yarn is simply pulled into this channel. A length of yarn continues to be pulled until it breaks. When breaking, the yarn's elastic tension rapidly snaps it in the other direction, shooting it though the hole where it was supposed to go in the first place. Watch the video below for a demonstration.

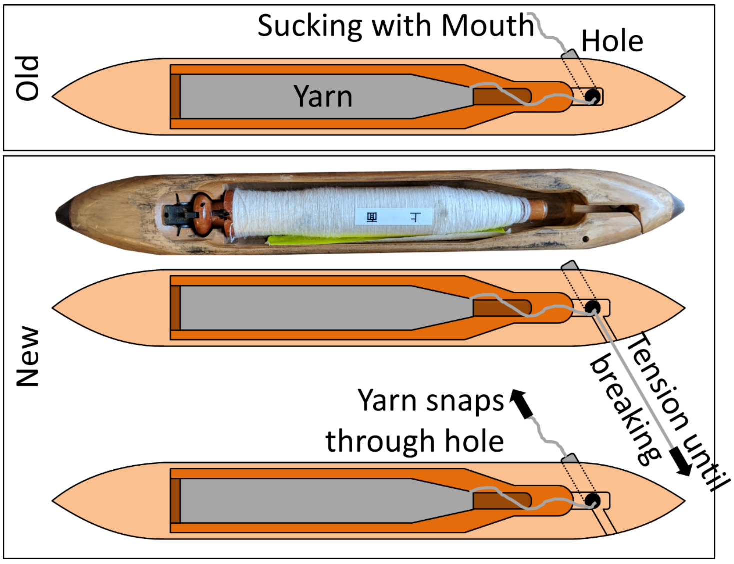

Figure 83: Toyoda Model G Threading Shuttle Diagram (Image Roser)

The Video by AllAboutLean.com is available on YouTube as "Toyoda Model G Shuttle Preparation" at https://youtu.be/1LKWuY1R3uI

11.3 Automatic Shuttle Change

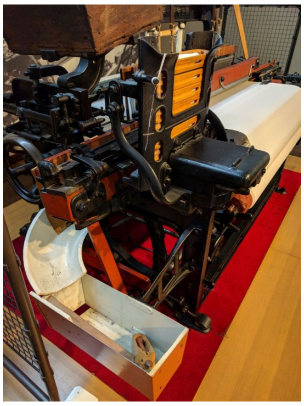

Figure 84: Toyoda Model G Automatic Loom Detail Automatic Shuttle Change (Image Roser)

Another neat feature is an automatic shuttle change. Previously, attendants always had to monitor the shuttle to put in a new one when the old one was out of yarn. Forgetting this wasted machine time and degraded the quality.

The Model G loom has a feature that detects an empty shuttle, and exchanges it with a new one.

The spool of yarn inside the shuttle has a slot. A pin in the machine presses against a piece of wood, attempting to push it into this slot. As long as there is yarn on the spool, the yarn prevents the wood from entering the slot. If the yarn is used up, the wood can enter the slot. A mechanism detects this, and ejects the old shuttle while pushing in the new shuttle. Watch the video for the different mechanisms in action. Please note that this is a mixture of videos of a full-sized Model G loom, different museum visualizations, and even some older looms preceding the Model G while the technology was developed.

The Video by AllAboutLean.com is available on YouTube as "Toyoda Model G Auto Shuttle Change" at https://youtu.be/AQbqmwiwIYw

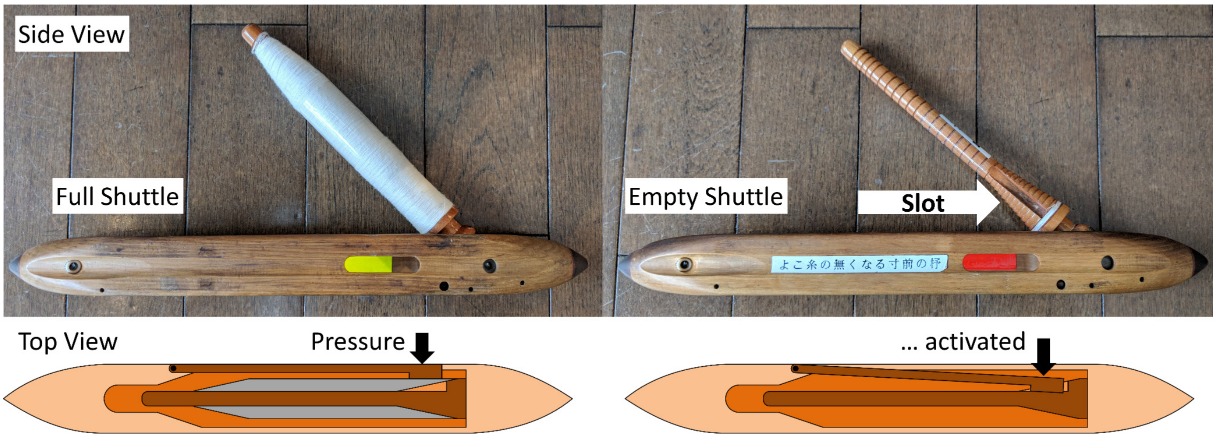

Figure 85: Model G Automatic Shuttle Change Diagram (Image Roser)

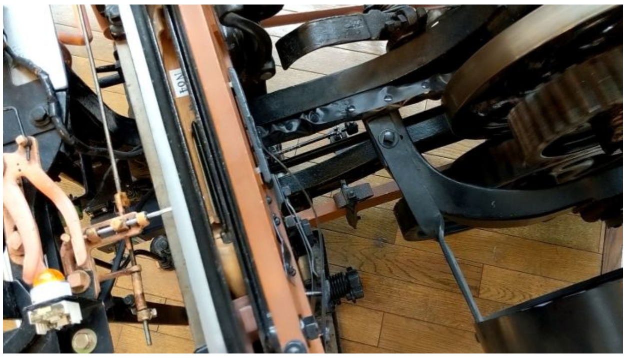

Figure 86: Pin in the lower left sensing status of shuttle in machine (Image Roser)

11.4 The Warp-Break Auto-Stop Mechanism

One major problem with looms is breakage of yarn. The two types of yarn are called warp and weft. Both of them can break, creating significant quality and productivity problems. Previously, many workers were kept busy merely looking out for broken warps.

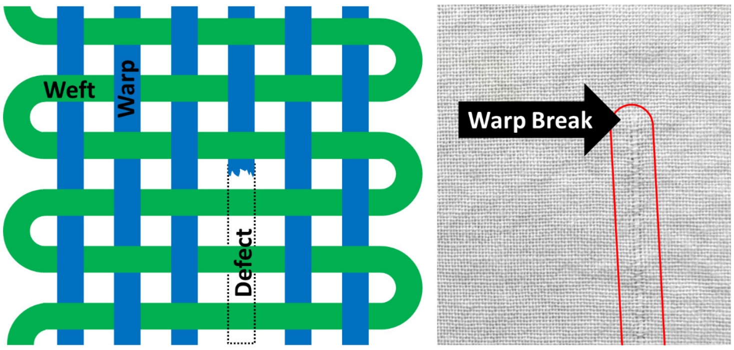

Figure 87: Warp Break Model and Cloth (Image Roser)

Figure 88: Toyoda Model G Automatic Loom Detail of Warp Break Stop (Image Roser)

The Model G loom solved this nicely. A small piece of sheet metal is hanging from every warp. The tension of the warp keeps this sheet metal up. If the warp breaks, however, the tension goes away and the sheet metal drops down.

This dropped sheet metal then blocks a bar moving back and forth underneath the sheet metals. The stopped bar subsequently activates another mechanism that eventually turns the machine off. Watch the video to see the mechanism in action. Please note that this is a mixture of videos of a full-sized Model G loom, different museum visualizations, and even some older looms preceding the Model G while the technology was developed.

The Video by AllAboutLean.com is available on YouTube as "Toyoda Model G Warp Break Auto Stop" at https://youtu.be/PdGcfHucmKc

11.5 The Weft-Break Auto-Stop Mechanism

Figure 89: Toyoda Model G Weft Fork Mechanism (Image Toyota with permission)

Just as a break in the warp disrupts quality and productivity, so does a break in the weft. Here too, Toyoda solved this problem with a device that stops the machine automatically if the weft breaks. This is done using a weft fork.

The shuttle with the yarn runs within another part back and forth. This part is called the race. The race moves back and forth. In the picture, the race would be the reddish wood in the lower half.

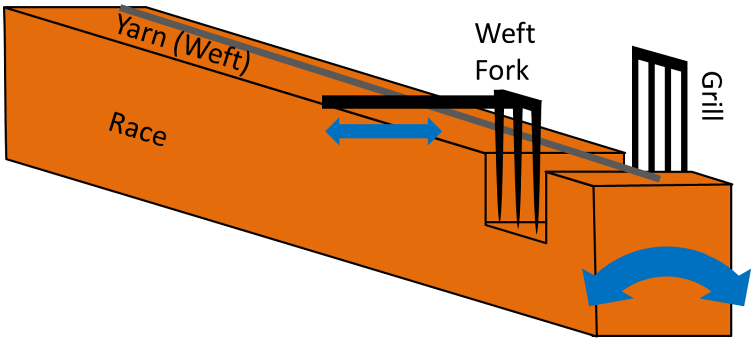

Figure 90: Toyoda Model G Weft Fork Schematics (Image Roser)

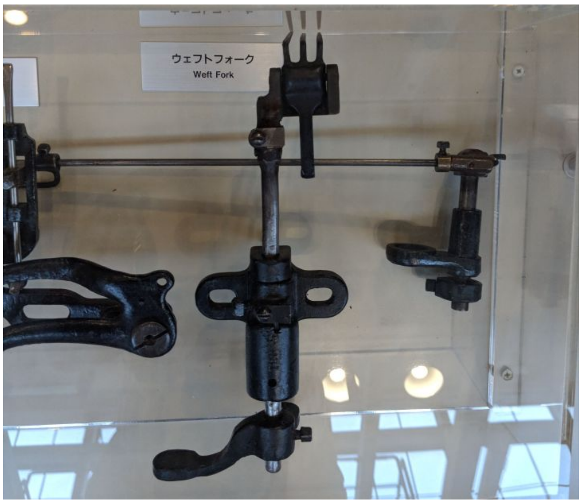

Figure 91: Parts for the weft fork (Image Roser)

The weft's three-pronged fork squeezes the yarn (the weft) against a metal grill. In the image, you see the weft fork in the center and the grill to the right of the center. The fork moves through the gap of the race.

If there is yarn (the weft), the fork cannot go through the grill due to the tension of the weft. However, if the weft is broken, the fork can pass through the grill. A subsequent mechanism detects this movement and turns off the machine. Unfortunately, there was no display model for the weft fork, hence I have no video for this mechanism.

11.6 Visualization

One famous feature of the Toyota Production System is the andon. In modern times this is often a digital display board telling you how much was produced, and how much you still need to achieve the target. Simpler forms of andons are stacked lights in different colors, indicating the status of the machine. Green is usually good, orange is a warning, red is a bigger problem, and so on.

The Toyoda Model G has a very early pre-digital, even pre-electric version of the andon. Not all looms but at least one per cluster of looms had a movable indicator that could be either off, white, or red as shown below. The indicators were not flat metal, but had a perpendicular metal too, so that the status could also easily be seen from the side. You could even signal both white and red, but this would be confusing signals from different sides.

Figure 92: Pre-electric andon off, white, or red (Image Roser)

Overall, the Toyoda Model G loom included many of the aspects that the Toyota Production System is famous for. I hope you liked this deep dive into production history, and I hope you also liked the videos. Many thanks to the team from Lightworks for giving away a free and powerful video editing software. Also many thanks to the Toyota Commemorative Museum of Industry and Technology for their awesome exhibits! Now **go out, learn from history, and organize your industry!**

Important Addendum: I have gotten 5+ comments so far from somewhere around India: They want to buy the loom, and are asking me for a catalogue, the price of the loom and how to purchase it. **This is a museum piece. It has not been produced for 80 years, and I do not sell these looms!** Please stop asking.

12 150 Years after the Birth of Albert Kahn

Christoph Roser, March 21, 2019, Original at https://www.allaboutlean.com/150-years-albert-kahn/

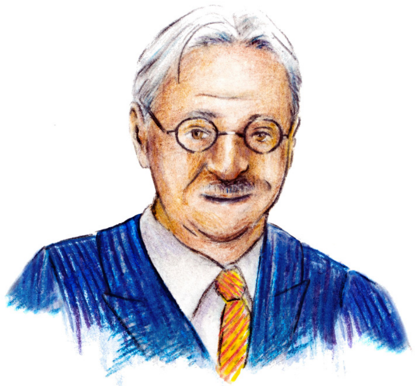

Figure 93: Albert Kahn Portrait (Image Roser)

Albert Kahn (1869–1942) is an often unknown but very influential figure in the history of manufacturing. An architect by trade, he revolutionized industrial architecture, and is often nicknamed the "Architect of Detroit." Most modern factories have a design that goes back to his innovations. Since he was born exactly 150 years ago on March 21, 1869, it is a good time to look at his achievements.

12.1 Early Beginnings

Albert was born in Rhaunen in Prussia (now Germany) on March 21, 1869. His father Joseph was a rabbi, and the family moved to Detroit in 1880 when he was eleven. He wanted to become an artist, but it turned out that he was color blind.

As a teenager he worked for a local architect as an office boy, initially without pay (modern architect students may still be familiar with this work without payments, unfortunately). However, one of the architects started to mentor him, and in 1891 Kahn won a one-year scholarship to travel Europe, during which he visited with another architect, Henry Bacon, who later designed the Lincoln Memorial.

Figure 94: Kahn (front left on the table) in his firm in 1896 (Image unknown author in public domain)

Although he never formally graduated from an architecture school, he opened his own architecture firm in 1895 at the age of twenty-six. While first working with two partners George W. Nettleton, and Alexander B. Trowbridge, he became the sole owner in 1905. He renamed his firm Albert Kahn Associates, which is still active. (I asked the firm for permission to use some more recent photos of Kahn, but got no reply)

12.2 Reinforced Concrete Replaces Brick and Wood

Figure 95: Kahn System Logo (Image Trussed Concrete Steel Company for editorial use)

While reinforced concrete has been in use since 1850, the technique was flawed, and beams occasionally broke. Albert's younger brother Julius Kahn (1874–1942), who also worked at his architecture firm, developed a new approach to reinforced concrete that produced much more reliable concrete trusses. Julius eventually started his own firm, the Trussed Concrete Steel Company, in 1903.

12.3 Let There Be Light …

While Julius improved the technical details, Albert used the new opportunities to create a new approach to industrial architecture. Previously, industrial buildings were shaped by the needs of steam and water power. A central power plant supplied the entire building with power through rotating shafts. Longer shafts meant more friction and also more vibration, and hence the building was built as close to or even around the power source, often using multiple floors to reduce the length of these wooden shafts.

Figure 96: First Cromford Mill Photoshopped (Image chevin in public domain)

This can be seen here in this (photoshopped) image of one of the earliest cotton mills, the Cromford Mill by Richard Arkwright. The mill was five floors high, with a waterwheel in the center. The current mill is only three floors, since the upper floors burned down. Luckily, Photoshop can fix this for you (clumsily). Notice how all the floor space is clustered around the water wheel?

Also notice how they tried to add as many windows as possible, but due to the brick-and-mortar construction could only use 20%–25% of the wall space for windows? This was at a time when electric light didn't exist, and gas lights were quite dangerous in a cotton factory. Hence they tried to maximize the natural light, albeit limited by the brick construction.

Figure 97: Packard Plant 1903–1910 (Image unknown author in public domain)

Using modern reinforced concrete, Albert Kahn was able to break through these constraints. The first building where he used this technique was the Packard Motor Company No 10. Plant, designed in 1903 and completed in 1910. Notice how much larger the windows are?

Nowadays we are used to large windows. But back then having such large windows was revolutionary. Now 70%–80% of the walls could be windows. An additional bonus was that the construction was highly fire resistant, compared to the conventional structures using wood.

Figure 98: Ford Highland Park Crystal Palace 1908–1910 (Image Andrew Jameson under the CC-BY-SA 3.0 license)

This also impressed others, and Henry Ford hired Albert Kahn to build his Highland Park plant between 1908 and 1910. Due to these large windows, this building was nicknamed "Crystal Palace."

Please note that while this was revolutionary in the USA, I do know of at least one contemporary building in Germany that also has extremely large windows. While Kahn was building the Packard plant, the famous Steiff toy factory also built a new factory hall between 1904 and 1908.

Figure 99: Steiff factory building 1904–1908 (Image Zacharias L. under the CC-BY-SA 3.0 license)

They used a curtain wall construction where a glass facade hangs in front of the supporting pillars. This allowed even more windows. Since the workforce consisted mostly of unmarried women, this building was quickly named “Aquarium of Virgins” by the locals.

In any case, the German architect is not recorded in history, but Albert Kahn is well known for his achievement. Soon all over the world people used reinforced steel to build factories with large windows.

12.4 Let There Be Space

Figure 100: Section of the River Rouge Plant with sawtooth roof (Image Detroit Publishing Co. in public domain)

Due to the success of the Highland Park plant, Henry Ford hired Albert Kahn again for the largest factory in the world: the Ford River Rouge Complex, built between 1917 and 1928. Here Albert Kahn shed another historic baggage of industrial architecture and created single-floor steel-support buildings. Now it was possible to have a wide floor space and wide space between supports, and the windows for light and ventilating were simply built into the roof, giving the factory its now iconic sawtooth shape. This shape is to us immediately recognizable as a factory building.

Figure 101: Inside River Rouge (Image Alfred T. Palmer in public domain)

Inside there is usually no wall, and only a few pillars to support the roof. The space between the walls is much larger than in previous buildings, since you no longer had to be close to a wall with a window for air and light. Hence you have much more flexibility in placing your machines, paths, storage, and any other layout related issues. Larger products like aircraft could not even be produced without such a wide span between supports.

The material transport is also easier, since there is only a single floor. Multi-floor factories almost always suffer from a bottleneck in the elevators. Visual management is also easier, since you can see the entire shop floor.

12.5 Let's Keep On Building …

Figure 102: William L. Clements Library (Image Michael Barera under the CC-BY-SA 4.0 license)

Overall, Kahn built over one thousand industrial buildings for Henry Ford, and hundreds more for other companies all over the world. On top of that, there were hundreds more non-industry buildings. He also built numerous military installations during World War II. His firm employed up to six hundred people, which is gigantic for an architecture firm. He built 19% of all industrial buildings in 1937, and in 1941 he received the eighth-highest salary in the US: $486,936. Not bad for a poor immigrant!

He said that he wanted most to be remembered for his work on the William L. Clements Library at the University of Michigan in Ann Arbor. He died aged seventy-three on December 8, 1942 in Detroit.

If you go to your shop floor, it is quite possible that the basic design was influenced by Albert Kahn. In any case, I hope this little excursion in history was interesting to you. Now, go out, use the space Kahn gave you, improve your layout, and organize your industry!

13 The Power of Six: Relation between Time and Money in Manufacturing

Christoph Roser, March 26, 2019, Original at https://www.allaboutlean.com/power-of-six-1/

Figure 103: Time is Money (Image Alexas_Fotos in public domain)

Time is money. You know that. But with respect to product cost and lead time, there is a rule of thumb that estimates this relation. Let me present to you the "**Power of Six,**" discovered by Rajan Suri. This gives you a rough estimate of how the lead time of your products influences the cost and vice versa. This first post looks at the original work, and my next post applies this rule also to segments of the value stream.

13.1 Introduction

Figure 104: Running Rabbit (Image Malene Thyssen under the CC-BY-SA 3.0 license)

One very big aspect of manufacturing is time. How fast can you get your products to the customer (i.e., what is the delay between the order of the customer and the delivery to the customer)?

This metric has a huge impact on the success of a company. Usually, customers like to get their stuff fast, and a faster delivery often increases customer satisfaction.

There are even a number of different overlapping terms for this time. The lead time is usually from the start of production till completion. The replenishment time is the time to replace a consumed item in your inventory. The turnaround time or response time is the time between a customer order and the delivery of this item.

However, what is less known is that a faster response time also usually decreases cost. If you have a faster response time, then you probably have less inventory, less fluctuation, less handling, and many other expenses. Hence, on average, being faster reduces your production cost. The Power of Six gives an estimate of this relation.

Traditional cost accounting has difficulties with this. They cannot estimate the effect of reduced fluctuations, and see only parts of the effect of reduced inventory. Accountants take the easy way out by considering anything they do not know as zero. Hence, for them, reducing fluctuations, etc. holds no value – which is clearly wrong. For more on this see my post The Problems of Cost Accounting with Lean.

13.2 Time and Money

Figure 105: Time vs Money on Scale (Image 3Dman_eu in public domain)

There are two elements that the Power of Six brings into relation. The first one is the turnaround time, the second one is product cost.

The Power of Six originated with Quick Response Manufacturing (QRM). One key metric is the **Manufacturing Critical Path Time** (MCT). Since the MCT is used for quite a few different things, it has a very detailed definition. The definition is:

Manufacturing Critical Path Time MCT: The typical amount of calendar time from when a customer submits an order, through the critical path, until the first end item of that order is delivered to the customer.

For our purposes, however, we can simplify this to **the average time between receiving an order and delivering the (first part of the) order to the customer**. This is based on real time, not working time, and includes for example all off-shifts, weekends, plant holidays, etc. You could for example take the average days between receiving the order and delivery. Do not take only rush orders, as this will be a biased sample and will be far from your average duration. It is important that this includes the entire value chain. If you fulfill a customer order merely by grabbing the make-to-stock item from the shelf, this rule of thumb won't work.

The **cost is simply the total product cost** including overhead, materials, and everything else.

13.3 The Power of Six

Figure 106: Power of Six Image (Image Petr Kratochvil in public domain)

Assume you want to reduce your cost by reducing your turnaround time (or MCT). Let's assume you have the following variables:

- C_0 is the current product cost
- C_1 is the desired new product cost
- T_0 is the current turnaround time
- T_1 is the new turnaround time needed

The relation according to the power of six is as follows:

$$\frac{T_1}{T_0} = \left(\frac{C_1}{C_0}\right)^6$$

and that is it. If you have a desired cost reduction, then you can simply calculate the needed percentage reduction of your turnaround time. For example, if you want to reduce your cost by 5% (i.e., reduce it to 95% of your previous cost), your formula would be

$$\frac{T_1}{T_0} = \left(\frac{C_1}{C_0}\right)^6 = \left(\frac{95\%}{100\%}\right)^6 = (95\%)^6 = 74\%$$

Hence you can estimate that you would have to reduce your turnaround time by around 26% (i.e., to 74% of the original value) to achieve a 5% cost saving, assuming you only influence the turnaround time and no other cost saving levers.

It would also work if you start with the reduction of the turnaround time. Assuming you re-shore supply from China back to the USA, your turnaround time would reduce by three months shipping time from eight months to five months. You can turn the formula around as shown below by not taking the ratio of the costs to the power of six, but the ratio of the times to the power of 1/6th.

$$\frac{C_1}{C_0} = \left(\frac{T_1}{T_0}\right)^{\frac{1}{6}} = \left(\frac{5}{8}\right)^{\frac{1}{6}} = (63\%)^{\frac{1}{6}} = 92\%$$

Hence, if you reduce your turnaround time by 38% to 63%, you would reduce your cost by around 8% to 92% of the previous cost. The overall relation is shown in the graph below (albeit I presume the prediction becomes less accurate as your improvements approach 100%).

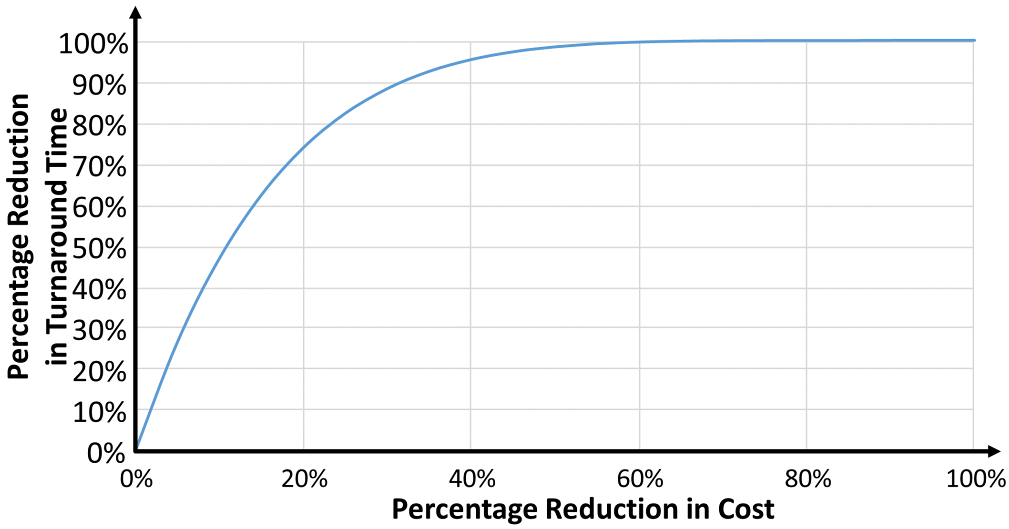

Figure 107: Power of Six Relation (Image Roser)

By the way, it also works the other way round. If you outsource some production to China, you extend your supply chain by three months. Let's do the calculation, assuming you go from a three-month turnaround to a six-month turnaround time you increase your turnaround time to 200%:

$$\frac{C_1}{C_0} = \left(\frac{T_1}{T_0}\right)^{\frac{1}{6}} = \left(\frac{6}{3}\right)^{\frac{1}{6}} = (200\%)^{\frac{1}{6}} = 112\%$$

Hence by doubling your turnaround time, the cost would increase to 112% of the cost, an increase of 12%. Therefore, **production in China would have to be at least 10.7% cheaper just to break even.** (Different number because the way percentages work. If you increase 100€ by 12% you get 112€. Decreasing 112€ by 12% would give you 98.56€. Decreasing it by 10.7% would give you 100€ again.) The graph below is an extension of the graph above continuing the relationship if you make the turnaround time worse.

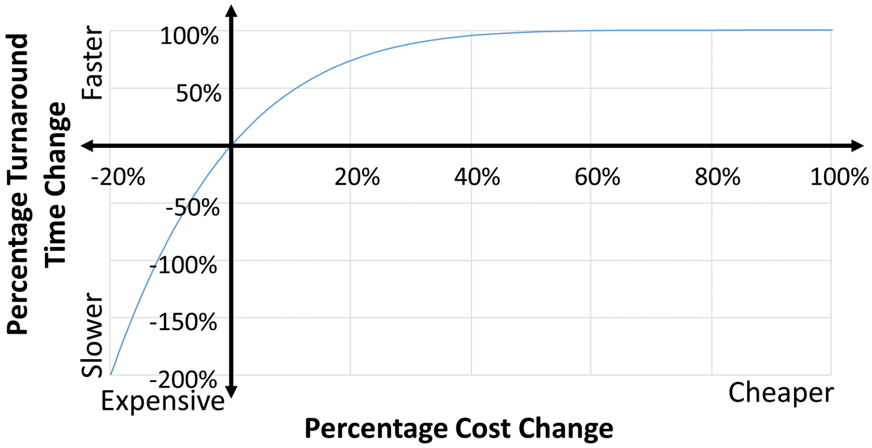

Figure 108: Power of Six Relation in both directions (Image Roser)

13.4 Accuracy of the Power of Six

This rule is a rough rule of thumb. To show you the accuracy, I have plotted the raw data below from the original paper from Tubino and Suri 2000 (full source below). This is to give you a feeling of the underlying change, but is by no means an accurate estimate of the performance after the improvements.

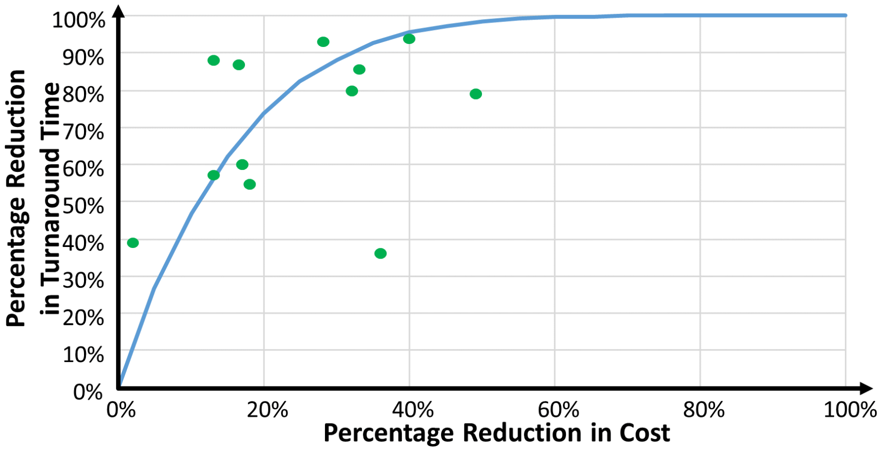

Figure 109: Accuracy of the Power of Six (Image Roser)

This original Power of Six rule applies to the entire value stream. However, by extrapolation it should also work for segments of the value stream. I will talk about this in more detail in my next post. Until then, **go out, reduce your lead time, and organize your industry!**

P.S.: Many thanks to Rajan Suri for his input and help, and of course for finding this relation in the first place!

13.5 Sources

- Suri, Rajan. It's About Time: The Competitive Advantage of Quick Response Manufacturing. 1 edition. New York: Productivity Press, 2010. Pages 165–167.
- Suri, Rajan. MCT Quick Reference Guide. Suri Consulting and Seminars, LLC, 2014. Page 11.
- Tubino, Francisco, and Rajan Suri. "What Kind of 'Numbers' Can a Company Expect After Implementing Quick Response Manufacturing? – Empirical Data from Several Projects on Lead Time Reduction." In Quick Response Manufacturing 2000 Conference Proceedings, 2000.

14 The Power of Six: Time and Money for Parts of Your Value Stream

Christoph Roser, April 02, 2019, Original at https://www.allaboutlean.com/power-of-six-2/

Figure 110: Time and Money (Image MaxPixel in public domain)

In my last post I presented you Rajan Suri's Power of Six – the relation between turnaround time and product cost. In this post I extend his work to apply it to not the entire value stream but to segments of the value stream. Enjoy!

14.1 Power of Six for Segments of Value Chain

Figure 111: Power of Six Image (Image Petr Kratochvil in public domain)

The Power of Six is used to calculate improvements across the entire value chain. As a recap, here is the Power of Six relation again from the previous post:

- C_0 is the current product cost
- C_1 is the desired new product cost
- T_0 is the current turnaround time
- T_1 is the new turnaround time needed

$$\frac{T_1}{T_0} = \left(\frac{C_1}{C_0}\right)^6$$

To consider the entire product cost, you would need to consider the entire value chain. The original source is a bit fuzzy on that, as they merely looked at the value chain underneath of their control, for products where this part of the value chain was the largest part of the value chain. Yet, considering the overall accuracy of the data this approach works.

But how do you approach if you only want to improve a small segment of the value chain? For example, if you have a series of kanban loops and want to improve only one loop. Take for example the simple system below.

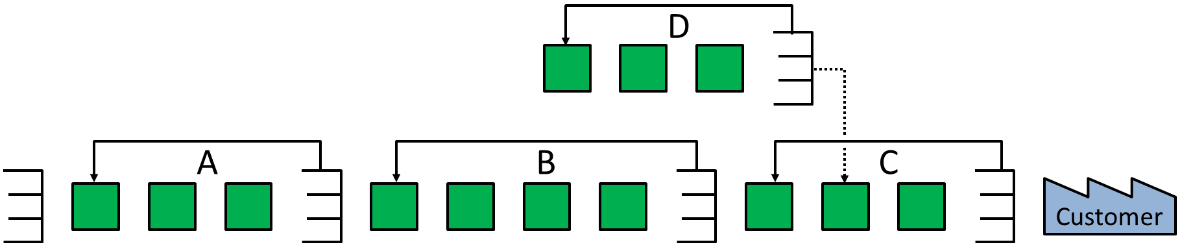

Figure 112: A connection of pull loops (Image Roser)

14.2 A Couple of Reasons Why This Would Not Work

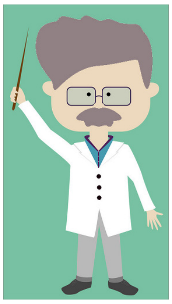
Figure 113: Scientist (Image MaxPixel in public domain)

Okay, I rarely do this, but let me get a bit academic and tell you all the things that should not work … before telling you later why, with respect to the accuracy of the Power of Six, it is good enough. So if you want, you can skip this section and jump to the next one for how it would work.

There are a couple of pitfalls here. First, for the kanban calculation, you do have the replenishment time. However, if you use the replenishment time, you ignore the time the material is waiting in the supermarket. This would have to be included.

Second, assume you want to improve the kanban loop D above. However, loop D is not part of the critical path. Improving loop D will not improve the overall turnaround time, which would be loop A, B, and C. The original Power of Six would consider here only improvements along loops A, B, and C. Yet, improving loop D will also be beneficial; it is just that the Power of Six rule can no longer calculate it.

You could also believe that the power of six would work for subsegments, where you relate the turnaround time for that segment $T_{0,S}$ and $T_{1,S}$ and its impact on the value add within this segment $C_{0,S}$ and $C_{1,S}$, where

$$\frac{T_{1,S}}{T_{0,S}} = \left(\frac{C_{1,S}}{C_{0,S}}\right)^6$$

Mathematically speaking this relation is incorrect. First let me explain why this is so, before farther down telling you why it is good enough. Assume the total time T_0 is the sum of the time of the segment $T_{0,S}$ you are improving and the time of the remainder outside of the segment $T_{0,R}$. Similarly the total cost C_0 is the cost of the value add of the subsegment $C_{0,S}$ under analysis plus the remaining costs $C_{0,R}$. Similar applies to T_1 and C_1.

$$T_0 = T_{0,S} + T_{0,R}$$

$$T_1 = T_{1,S} + T_{1,R}$$

$$C = C_{0,S} + C_{0,R}$$

$$C_1 = C_{1,S} + C_{1,R}$$

Hence the formula would be

$$\frac{T_{1,S} + T_{1,R}}{T_{0,S} + T_{0,R}} = \left(\frac{C_{1,S} + C_{1,R}}{C_{0,S} + C_{0,R}}\right)^6$$

which would solve to something messy like

$$\frac{T_1}{T_0} = \frac{C_{1,S}^6 + 6C_{1,S}^5 C_{1,R} + 15C_{1,S}^4 C_{1,R}^2 + 20C_{1,S}^3 C_{1,R}^3 + 15C_{1,S}^2 C_{1,R}^4 + 6C_{1,S} C_{1,R}^5 + C_{1,R}^6}{C_{0,S}^6 + 6C_{0,S}^5 C_{0,R} + 15C_{0,S}^4 C_{0,R}^2 + 20C_{0,S}^3 C_{0,R}^3 + 15C_{0,S}^2 C_{0,R}^4 + 6C_{0,S} C_{0,R}^5 + C_{0,R}^6}$$

and definitely not into something nice like the formula we had a little bit earlier

$$\frac{T_{1,S}}{T_{0,S}} = \left(\frac{C_{1,S}}{C_{0,S}}\right)^6$$

Hence, purely mathematically speaking, this approach would not work. But see farther down below.

14.3 The Correct Approach Using the Power of Six

If you improve only a segment of your value stream, and if this segment is on the critical path, you would have to estimate the share of your improvements with respect to the entire turnaround time to get an estimate of the improvement of the entire cost.

Let's take the example below. You improved segment B in your value stream, reducing the turnaround time for this segment from six days to four days, improving it by two days. Since the entire turnaround time across the value stream is twenty days, your overall improvement is still only two days, and the new turnaround time comes down to eighteen days.

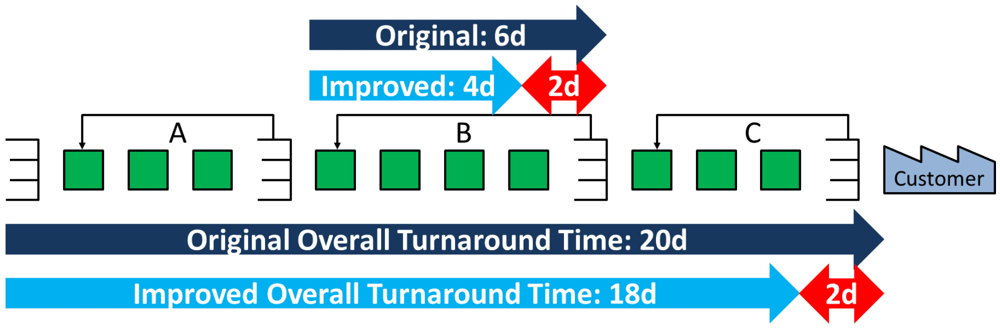

Figure 114: Power of Six for sections of the value stream (Image Roser)

Hence your overall reduction of the turnaround time was by 10% to 90% of the original value, and your cost should go down by around 1.74% to 98.26% of the original value as shown below.

$$\frac{C_1}{C_0} = \left(\frac{T_1}{T_0}\right)^{\frac{1}{6}} = \left(\frac{18}{20}\right)^{\frac{1}{6}} = (90\%)^{\frac{1}{6}} = 98\%$$

This approach would work, although it requires you to understand the turnaround time for the entire value stream, which may sometimes be tricky. It also does not work for subsegments that are not along the critical path (as for example segment D in the example farther up).

14.4 A "Good Enough" and Practical Estimation

Figure 115: Good Enough (Image Free Clip Art and Roser under the CC-BY-SA 4.0 license)

A little bit earlier I introduced you this equation, and told you that it is not mathematically proper. Here it is again and also its reverse from:

$$\frac{T_{1,S}}{T_{0,S}} = \left(\frac{C_{1,S}}{C_{0,S}}\right)^6$$

$$\frac{C_{1,S}}{C_{0,S}} = \left(\frac{T_{1,S}}{T_{0,S}}\right)^{\frac{1}{6}}$$

However, I am not an mathematician but an engineer, and 1+1≈2 is often good enough for me. And, luckily, this equation is good enough for pretty much all cases. Let's compare the accuracy of this equation for a segment with the "correct approach" based on the improvement of the total turnaround time above.

Let's take the example from above again and assume your total turnaround time is twenty days. You optimize one six-day segment of your entire value stream to reduce your turnaround time by two days to four days (i.e., the total improvement would reduce turnaround from twenty to eighteen days). Using the Power of Six properly, this would estimate a cost improvement of 1.74% to a new cost of 98.26% (and I am aware that this number of digits is excessive for the accuracy of the method).

If we look only at the section itself, we have a reduction by two days from six to four days. This would reduce the cost in this segment by 6.53% to 93.47%.

$$\frac{C_{1,S}}{C_{0,S}} = \left(\frac{T_{1,S}}{T_{0,S}}\right)^{\frac{1}{6}} = \left(\frac{4}{6}\right)^{\frac{1}{6}} = (66.66\%)^{\frac{1}{6}} = 93.47\%$$

If we assume a linear relationship of the cost and hence assume that this segment that has six of the twenty days turnaround time has also 6/20th of the cost, then we can estimate the overall cost improvement as follows:

$$1 - \frac{C_{1,S}}{C_{0,S}} = \left(1 - \left(\frac{T_{1,S}}{T_{0,S}}\right)^{\frac{1}{6}}\right) \cdot \frac{T_{1,S}}{T_1} = (1 - 93.47\%)^{\frac{1}{6}} \cdot \frac{6}{20} = 6.53\% \cdot 30\% = 1.96\%$$

So overall, the "correct" approach would give us an improvement of 1.74%, and the simplified approach would give us 1.96%. For me this is good enough, especially considering the accuracy of the method.

Just to make sure that this is not a fluke, I tested different improvements with different shares of the overall value stream. Unless you manage to eliminate nearly 100% of the time in your segment, the overall estimate is to me close enough. I tested an overall improvement of 10%, 30%, 60%, and 90% for different shares of the segment of the overall turnaround time. This is shown below, where I compare the segment based calculation with the correct approach (the red lines).

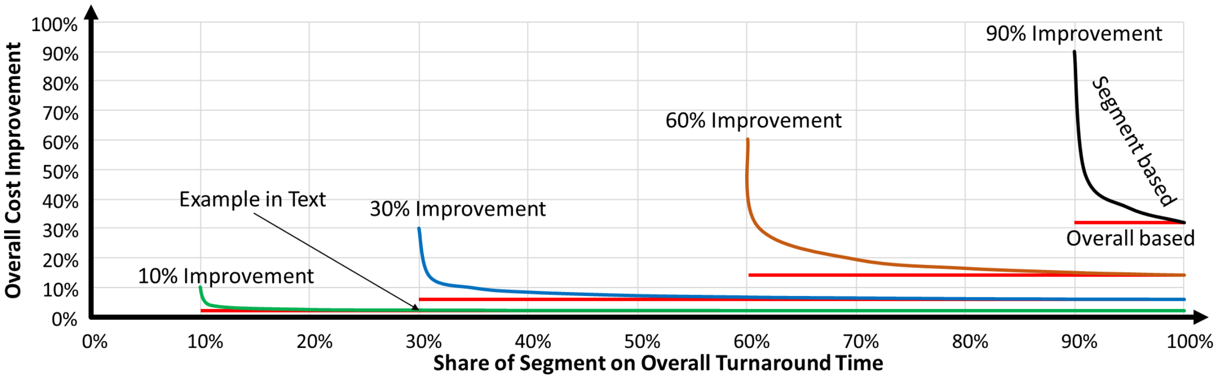

Figure 116: Power of Six Segment Approach Accuracy (Image Roser)

Naturally, if the overall improvement is 60%, the segment cannot be less than 60% of the overall. Even then the segment would have to reduce its time to zero to achieve an overall improvement of 60%. Hence the graphs below always end on the left side.

It can clearly be seen that unless you eliminate a segment completely, the segment-based calculation and the overall approach have rather similar results, again considering the accuracy of the method. Hence this "good enough and practical" application of the power of six is probably good enough for your cases. An additional benefit is that this segment based approach allows the calculation of the improvements of segments not on the critical path. In this case you merely calculate the improvement for the value add based on this part of the segment.

14.5 Conclusion for Segments

Hence you can use the Power of Six rule also for segments. The equation would stay the same, you merely apply it to part of the value stream.

- $C_{0,S}$ is the current value add in your **value stream segment**
- $C_{1,S}$ is the desired new value add in **your value stream segment**
- $T_{0,S}$ is the current turnaround time for your **value stream segment**
- $T_{1,S}$ is the new turnaround time needed for your **value stream segment**

The relation according to the power of six is as follows:

$$\frac{T_{1,S}}{T_{0,S}} = \left(\frac{C_{1,S}}{C_{0,S}}\right)^6$$

or in its inverse form

$$\frac{C_{1,S}}{C_{0,S}} = \left(\frac{T_{1,S}}{T_{0,S}}\right)^{\frac{1}{6}}$$

14.6 How to Get There?

You can start with a financial target and estimate the needed reduction in turnaround time, or you can start with the reduction in turnaround time and estimate the financial benefits. In any case, you would need to reduce turnaround time eventually. This, of course, is not that easy (otherwise you would have done it already). The full extent of your options would exceed the scope of this post, and I have written about this in many other blog posts. See A Eulogy for Little's Law, How Product Variants Influence Your Inventory, How to Reduce Your Inventory, and many more. But somehow you need to get your turnaround time down.

I hope this post was not too mathematical for you. Now, go out, reduce your turnaround time, estimate the benefits using the Power of Six, and organize your industry!

P.S.: Many thanks to Rajan Suri for his input and help, and of course for finding this relation in the first place!

14.7 Sources

- Suri, Rajan. It's About Time: The Competitive Advantage of Quick Response Manufacturing. 1 edition. New York: Productivity Press, 2010. Pages 165–167.
- Suri, Rajan. MCT Quick Reference Guide. Suri Consulting and Seminars, LLC, 2014. Page 11.
- Tubino, Francisco, and Rajan Suri. "What Kind of 'Numbers' Can a Company Expect After Implementing Quick Response Manufacturing? – Empirical Data from Several Projects on Lead Time Reduction." In Quick Response Manufacturing 2000 Conference Proceedings, 2000.

15 Maintaining Strong FIFO in Parallel FIFO Lanes

Christoph Roser, April 09, 2019, Original at
https://www.allaboutlean.com/strong-parallel-fifo/

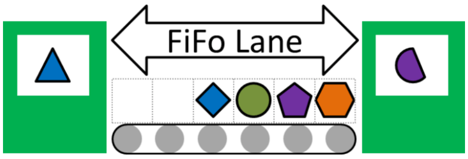

Figure 117: FIFO lane with different parts (Image Roser)

FIFO lanes are a common and easy way to manage material flow. However, sometimes there is not enough space to make a long FIFO lane. In this case the FIFO lane is often split into multiple sub-segments. This post looks at how to maintain strict (or strong) FIFO in such parallel FIFO lanes.

15.1 Introduction

FIFO stand for First In First Out, and is often used to maintain a sequence in manufacturing. A FIFO is also a buffer inventory, and hence decouples fluctuations (see The Three Fundamental Ways to Decouple Fluctuations). The bigger the inventory, the better the decoupling, although smaller inventories have other advantages.

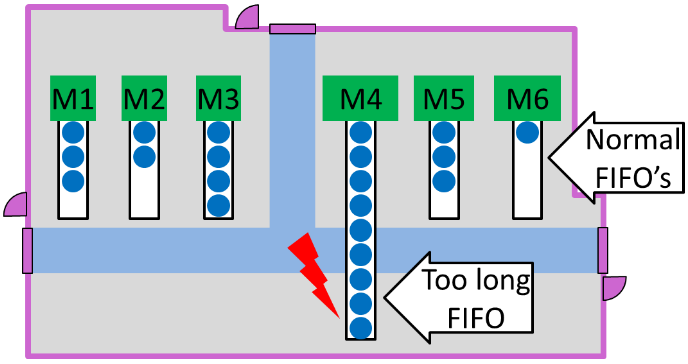

Figure 118: Too Long FIFO Layout (Image Roser)

While there are ways to calculate the length of a FIFO (see my post The FIFO Calculator), these approaches are usually not practical, and in reality somebody just estimates the length of a FIFO. In some cases, the number of slots in a FIFO may be large, and may make the FIFO too long for practical use.

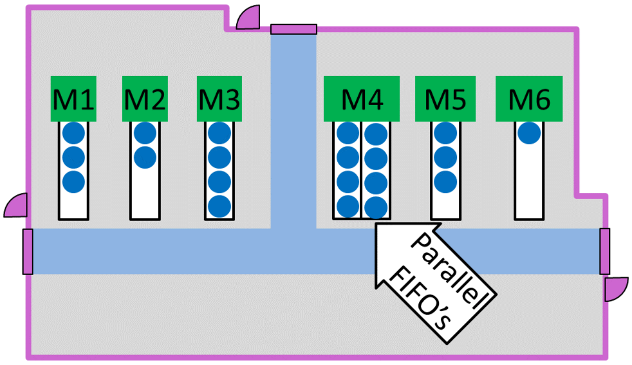

Figure 119: Parallel FIFO Layout (Image Roser)

In this case it may make sense to split the FIFO into two (or more) segments just to fit the whole thing in your floor space without blocking roads or generally being in the way. The challenge is to make sure that you still maintain the FIFO sequence.

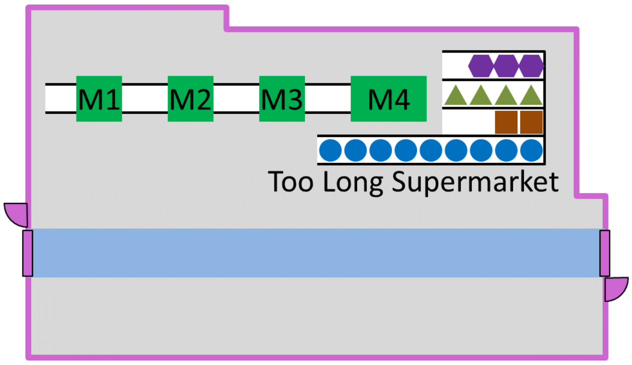

Figure 120: Too Long Supermarket Layout (Image Roser)

Similarly, a supermarket is just a bunch of FIFO lanes separated by products. For a supermarket you can also calculate the number of kanban. Here too, you may end up with a rather large number of kanban, requiring at least one of your supermarket lanes to be excessively long as shown in the image here.

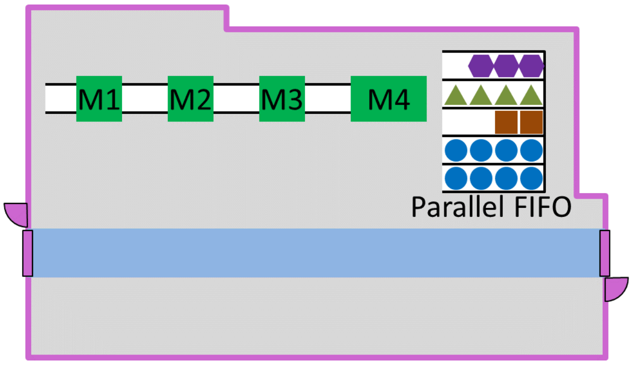

Figure 121: Parallel FIFO Supermarket Layout (Image Roser)

Here too, you may split a long single FIFO lane within a supermarket into two or more parallel lines. Again, the challenge is how to maintain the FIFO sequence. In the following I will show you approaches to maintain a strict FIFO. The options for handling parallel FIFO lanes are the same, regardless if it is "only" a FIFO, or a FIFO within a supermarket.

15.2 Data-Heavy Labeling

One option is to simply attach a sheet of paper to each item in the parallel FIFO lanes, indicating the sequence. For practical reasons this may also be a date and time of production or date and time of the item being added into the supermarket. An item can be added to any free lane in the supermarket. For the removal, however, the oldest item in the front row always has to be removed first. This approach ensures that the correct FIFO sequence is maintained.

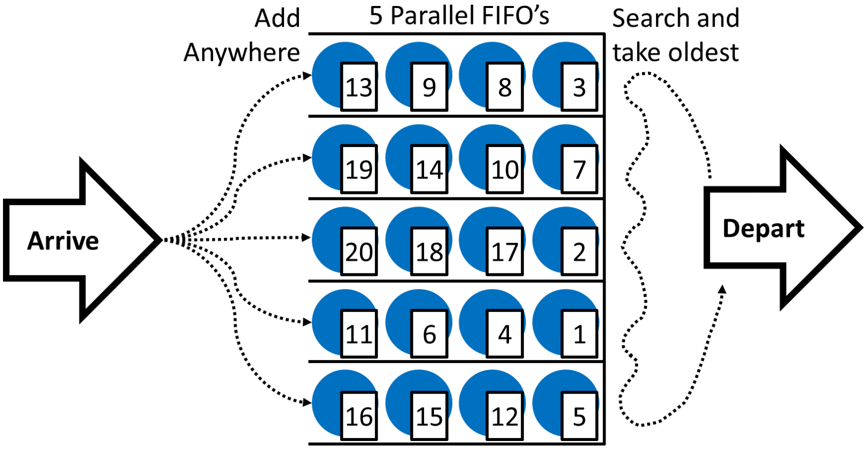

Figure 122: Labelled Parallel FIFO (Image Roser)

However, from the practical point of view, this method is usually inconvenient. First of all, someone has to label the items before adding them into the supermarket. However, if you are lucky this may be already part of the documentation of the item.

The real work is usually for the person removing the items. This person has to look at all items in the front row, and has to search the oldest one. Not only is the work time consuming, but it is also pretty boring, especially if you have to do it dozens of times during a day.

This approach is best used if you have a random access storage anyway (like a pallet shelf) that is managed by your ERP program. The ERP program knows the oldest part, and simply directs a forklift driver to the correct location to pick up the oldest part.

15.3 Fill One Row/Empty One Row

Another option is for the adding process to add into one row till it is full, and for the removing process to remove one row till it is empty. (But read the entire section before using it, there is one MAJOR caveat!)

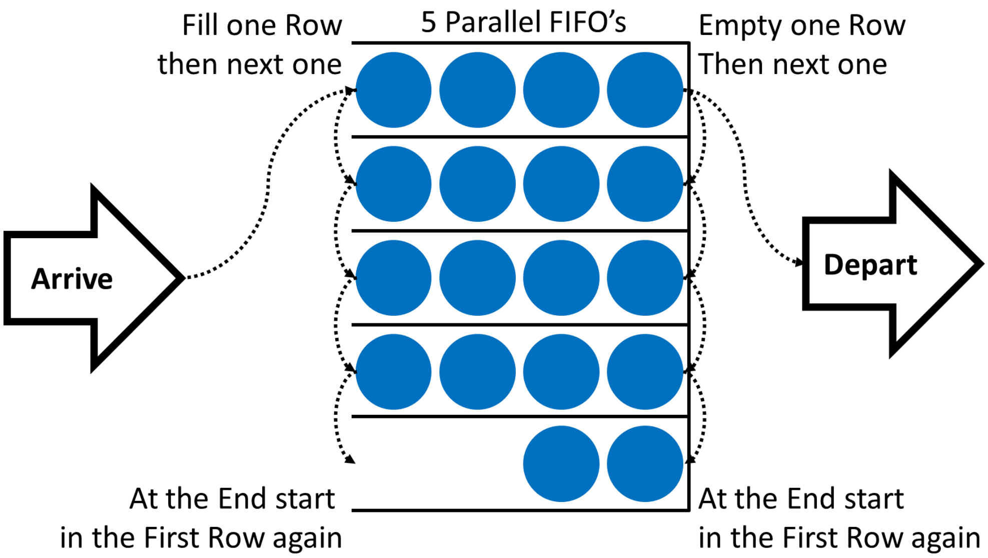

Figure 123: Incorrect Fill one Row Empty One Row Parallel FIFO (Image Roser)

15.3.1 First Two Rules

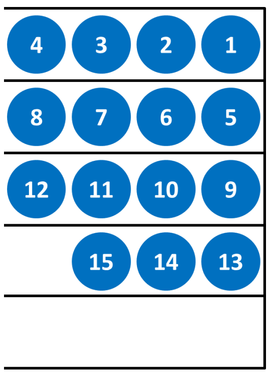

Figure 124: Animation of Parallel FIFO with Adding and Removing. The original image can be found at https://www.allaboutlean.com/strong-parallel-fifo/ (Image Roser)

The animation shows how sequential FIFOs are filled and removed again. There are two rules (and an important third one comes later):

- Add items in the row where the last item was added. If full, move to the next row. If the entire FIFO system is full, stop adding until space is available.
- Remove items in the row where the last item was removed. If empty, move to the next row. If the entire FIFO system is empty, stop removing until parts are available.

Granted, this animation has a long filling sequence, followed by a long removing sequence, and is only intended to demonstrate the process. In reality this will be more mixed – you add some, remove some, add some more, and so on.

15.3.2 The Problem with Overtaking

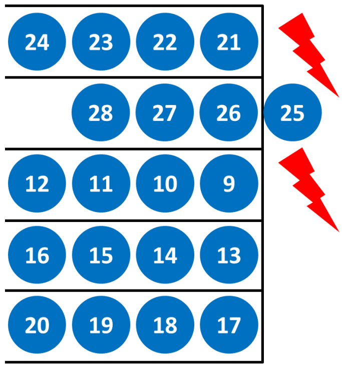

Figure 125: Animation of a Parallel FIFO Mismatch. The original image can be found at https://www.allaboutlean.com/strong-parallel-fifo/ (Image Roser)

This, however, makes the entire process tricky. If the filling process overtakes the removing process, the sequence goes out of whack! See for example the second animation. Parts 1–7 are removed, and then parts 21–25 are added. The next removal of part 7 in the second row is followed by adding part 25 also in the second row. A few cycles later, part 8 is removed, followed according to the above logic by part 25 instead of part 9.

Now you are out of sequence! It is easy to imagine that a sequential adding and removing always uses the second lane. Parts 9–20 will be forgotten and the sequence will be totally off.

15.3.3 The Important 3rd Rule

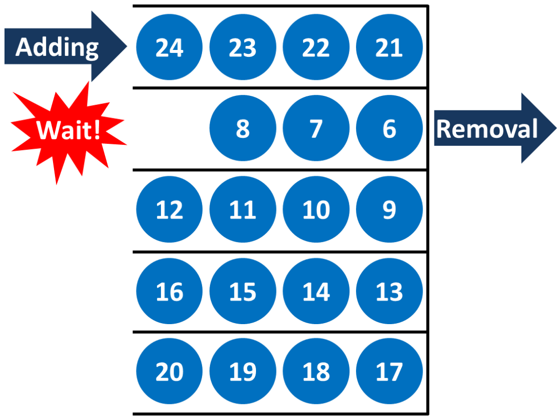

Figure 126: Ânimated Parallel FIFO Avoiding Mismatch. The original image can be found at https://www.allaboutlean.com/strong-parallel-fifo/ (Image Roser)

To avoid this problem of being out of sequence, we need an important additional rule:

The adding process can NEVER add to the lane currently assigned for removal, nor overtake that lane!

The adding lane can never add an item to the removal lane, overtaking the removal lane! Otherwise the sequence gets misaligned.

15.3.4 No Problem for Removal Catching Up with Adding

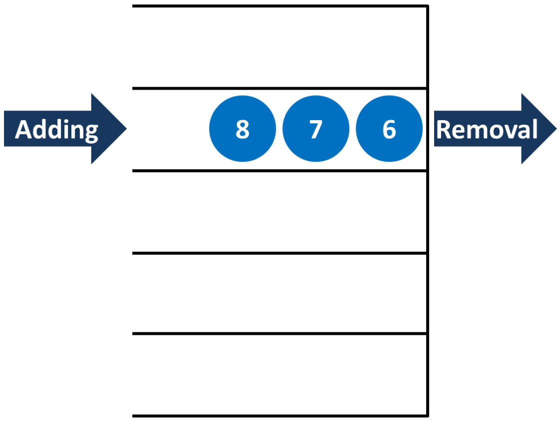

Figure 127: Animated Parallel FIFO Remove from Adding Lane. The original image can be found at https://www.allaboutlean.com/strong-parallel-fifo/ (Image Roser)

Please note that there is no similar restriction on the removing lane. The removing lane can without any problems remove items from the adding lane as shown in the animation. Hence, only the process adding to the parallel FIFO lanes needs to take care of not moving into the lane where items are currently removed.

15.3.5 Practical Set Up

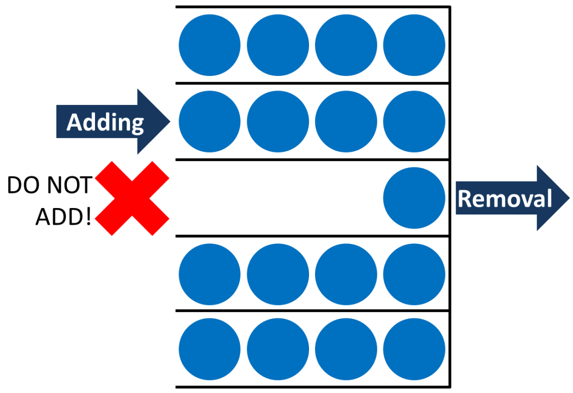

Figure 128: Parallel FIFO Do Not Add Lane (Image Roser)

The challenge is to ensure the adding process does not overflow into the lane where items are removed. This is indeed a challenge. Material handling like forklift drivers will be happy to place the material in a free spot. As shown in the image here, there is indeed a free spot, but they are not allowed to use this space.

If you have ever worked with forklift drivers, you will know that this is hard to compute for these drivers. More likely than not, they will break the rule ("just this once"), and your sequencing breaks. Unfortunately, it needs only one rule breaking to create a massive disruption in the sequence.

Hence, from the practical point of view, you need first a visual signal of which lane is forbidden. The people adding material need to know which lane is off limits. Here you have many different options. Of course, you could work with some simple signals. I have frequently seen plastic flower pots placed on top of larger material as signals.

There could also be a red lamp that is activated by a switch on the removal side. While this informs the person adding items, it does not prevent them from adding items. Hence you need quite a bit of training to make them understand why this is necessary. Better but more cumbersome to set up would be a device that actively prevents them from adding items in a blocked lane. This could be done karakuri kaizen style with two gates that are connected by rollers. If the removing process opens up one gate, the corresponding other side automatically rolls down. The removing process just should not forget to close the gate again after the lane is empty, hence freeing up the lane for the adding process.

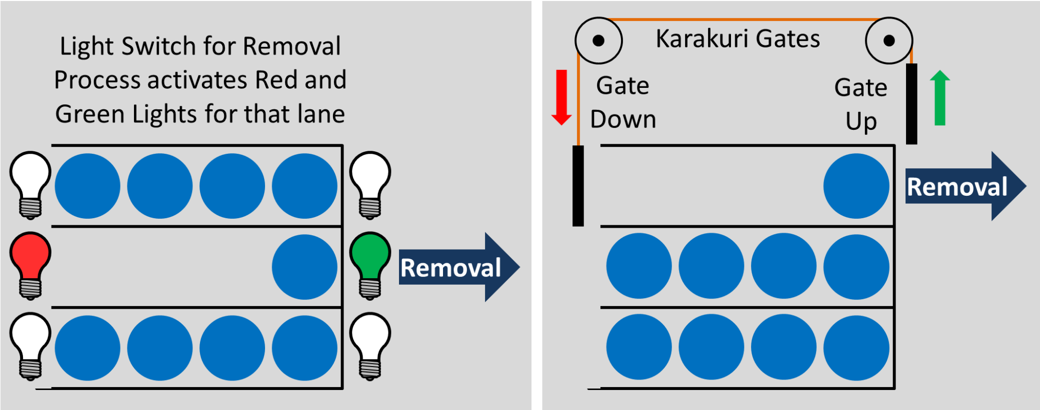

Figure 129: Parallel FIFO Practical Solution (Image Roser)

Again, there are many different options; the above are only suggestions. In all cases you have the problem that some of your FIFO space cannot be used to store material. Hence you do not use the full space available (otherwise you break FIFO sequence). To be precise, the average quantity you can store in this parallel FIFOs is half of a lane less than the full space available, since one lane is always blocked from adding material.

You could of course avoid this with the data-heavy labeling and searching for the oldest material, but again this would make the work more cumbersome. Another disadvantage of this approach is that it is not very robust, and mistakes can mess up the sequence very badly.

15.4 Adding and Taking Cyclic

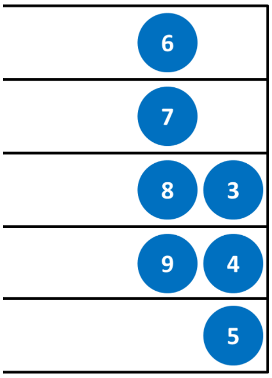

Figure 130: Animated Parallel FIFO with cyclic Addimng and Taking. The original image can be found at https://www.allaboutlean.com/strong-parallel-fifo/ (Image Roser)

Another option that maintains a strict FIFO is to add the next part always in the next adjacent lane to the previous adding. Similarly, the next part is always removed in the next adjacent lane to the previous removal. The animation here shows this sequence for this first 10 parts in the parallel FIFO example.

The disadvantage of this method is that you need to switch lanes at every single addition or removal, hence it is quite a bit of work to maintain the standard. Compared to the "Fill/Empty Row" above, you can now use the entire buffer space. Even better, if there is a mistake and an item is added in or removed from the wrong lane the sequence is only slightly disturbed. This method is more robust against mistakes than the "Fill/Empty Row".

In my next post I will present you alternatives that use the full space with less effort, in which they do not maintain a strict FIFO sequence but only a weak sequence. Until then, **go out, get your material flowing, and organize your industry!**

16 Maintaining Weak FIFO in Parallel FIFO Lanes

Christoph Roser, April 16, 2019, Original at https://www.allaboutlean.com/weak-parallel-fifo/

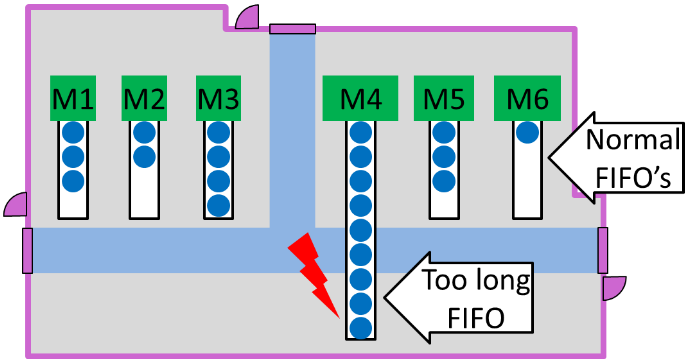

Figure 131: Too Long FIFO Layout (Image Roser)

Sometimes you would like to put more material in a single FIFO lane than the space you have available. In this case you would have to use a combination of two or more parallel FIFO lanes. In my last post I described how to maintain a strict FIFO sequence in parallel lanes. This post looks at an easier but less accurate method.

16.1 Introduction

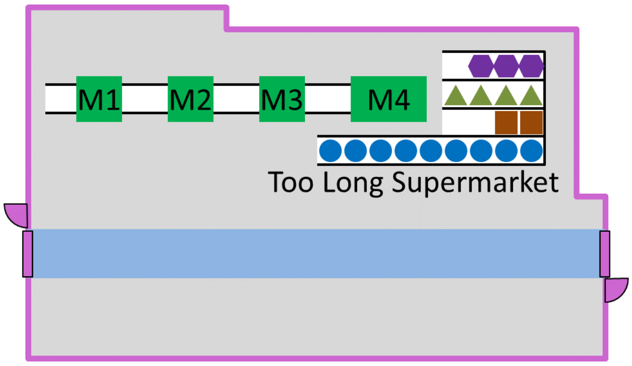

Figure 132: Too Long Supermarket Layout (Image Roser)

Please check my last post on why you may need parallel FIFO lanes in the first place instead of a long single lane between processes or in a supermarket. In my last post I presented a method that can maintain a strict FIFO sequence, although it may be a bit cumbersome to implement. Such a strict FIFO is in academia sometimes also called a **strong FIFO**.

In this post I would like to talk about a **weak FIFO** sequence. In such a weak sequence, the FIFO sequence is not perfect, but … meh … good enough. This post is based on the master's thesis of one of my students, Kaan Kalkanci. For the full source see below.

16.2 When Is It “Good Enough”?

Figure 133: Whatever … as long as it works … (Image PathDoc with permission)

A strict FIFO can help you with tracking and fixing errors, implementing traceability, and avoiding excessive aging of items while newer ones are already consumed. And, do not underestimate this, such a strict FIFO can help you a lot. With a single lane FIFO, it is very easy to implement. With two or more parallel FIFOs, however, this becomes trickier.

However, not all shop floors are on the pinnacle of manufacturing, often not even close, and tracking of errors and traceability are on a level that does not (yet) benefit much from a strict FIFO. In this case a weak FIFO with a few smaller errors could be good enough if it requires less effort. I myself have done this in one of my plants too, simply because the surrounding system does not allow you to benefit much from a strict FIFO.

On the other hand, you do not want your sequence to go totally out of control. Imagine getting an old product version out of your parallel FIFO system that you stopped making six months ago. While you may not mind getting a few stragglers in the hours or days after you changed the design, you do not want them long after the design change. Hence, what you may need is a weak FIFO, where the sequence is not perfect, but … meh … good enough.

Of course you have the option to ignore FIFO completely. If your shop floor is total chaos and major problems are crashing over you multiple times per day, a FIFO sequence for the (probably rare) parallel FIFO lanes may not be at the front of your mind. Work your problems, establish FIFO where it is easier (and a **single** FIFO lane is about as easy as it gets), and try to get things under control (good luck!). Eventually, however, you may come to the issue of FIFO quality if you have parallel lanes.

16.3 What Did We Look At?

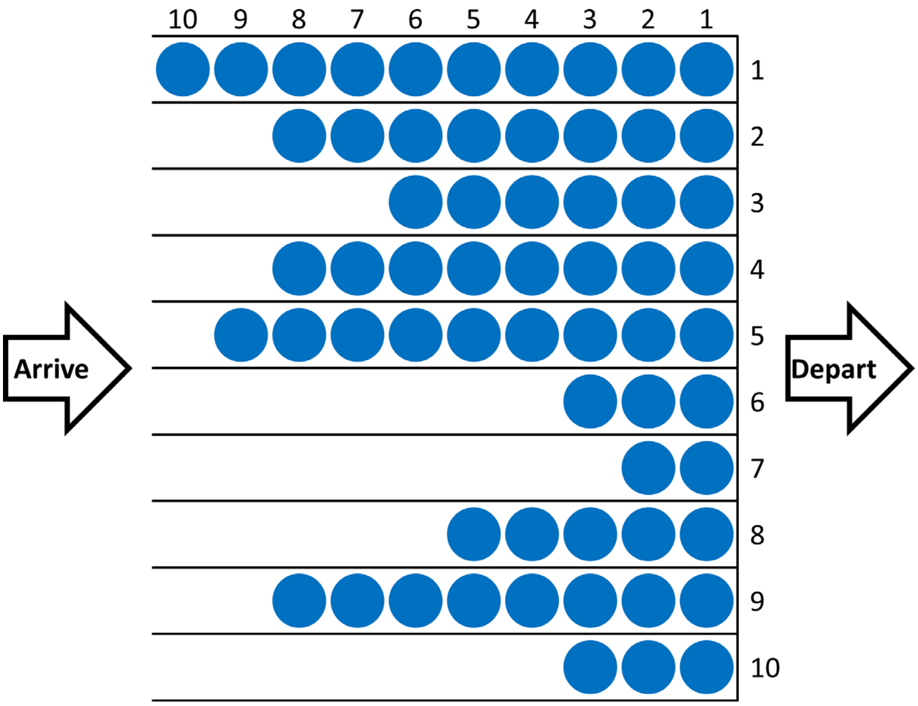

Figure 134: Parallel FIFO Example by Kalkanci (Image Roser)

My master's student looked at the quality of the sequence (i.e., the error in the sequence). If a part was added as the 37th part into the FIFO, but left it as the 39th part, then you have an error of 2. We measured the mean squared error between the position in the arriving sequence and the departing sequence.

We simulated a simple system with a single exponentially distributed random arrival and a single exponentially distributed random departure, and for a low, medium, and high utilization of ten parallel FIFO lanes with capacity ten each (details in the thesis if you are interested). We compared five different strategies to add and remove parts for a total of twenty-five strategy combinations.

- The first strategy was to add or remove items randomly. Let's call this **Add Random** for adding and **Take Random** for removing.
- The second strategy was to add and remove items from the FIFO lane with the largest inventory. Let's call this **Add Max** for adding and **Take Max** for removing.
- The third strategy was to add and remove items from the FIFO lane with the smallest inventory. Let's call this **Add Min** for adding and **Take Min** for removing.
- The fourth strategy was to add (or remove) items always from the lane where the last item was added (or removed), until it was full (or empty), in which case the next lane was used. Let's call this **Add Repeat** for adding and **Take Repeat** for removing.
- The fifth strategy was to always use the next lane from the previous one for adding or removing items respectively. When arriving at the last lane, the system starts over from the first lane. Let's call this **Add Cyclic** for adding and **Take Cyclic** for removing. Please note that the "Add Cyclic" – "Take Cyclic" combination gives a strict FIFO as described in my previous post.

To spare you all the mathematical and technical details, the result was:

It does not really matter in what sequence you add the parts to the parallel FIFO, as long as you always take the next part from the next adjacent lane of the previous one (Take Cyclic from above)

The graph below shows the mean squared errors of different combinations of adding and removing methods. Removing by “take cyclic” was by far the best, regardless of the method in which you added the parts to the system. The average “out of sequence” error was around 5. While not perfect, it avoids stragglers that are significantly out of sequence.

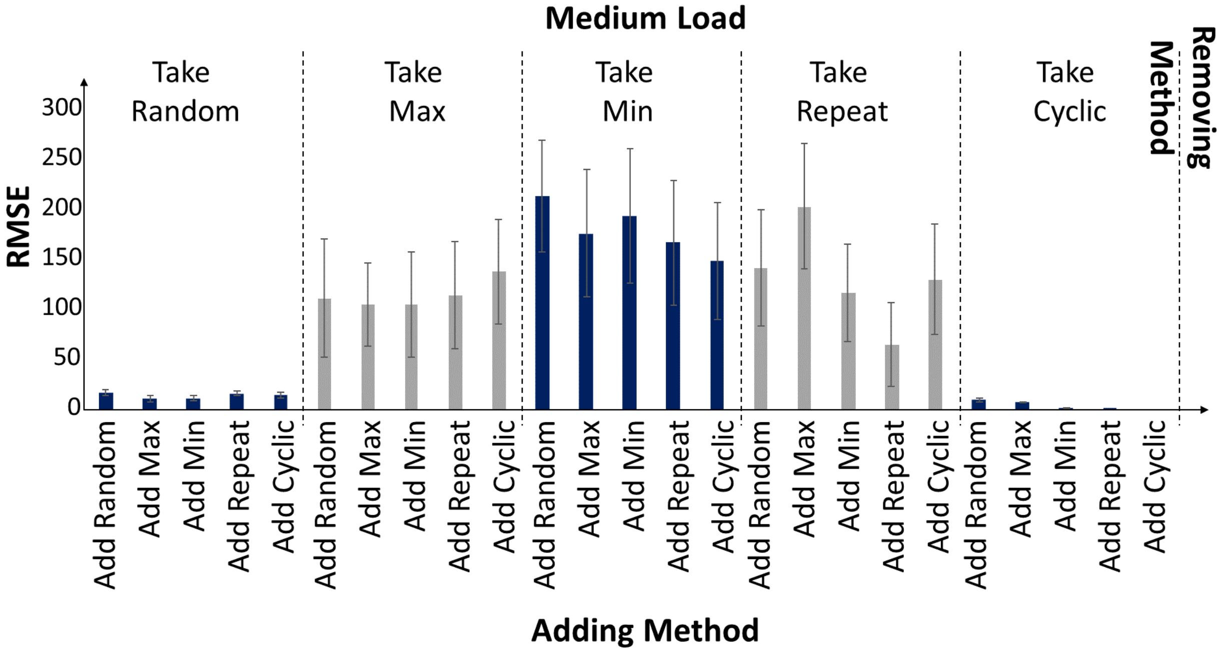

Figure 135: Parallel FIFO results from Kalkanci (Image Roser)

If you want to make it even a bit better, adding the parts also in a cyclic fashion (“add cyclic”) was the very best combination that maintained a strict FIFO sequence as shown in my last post. However, in this case you would need a system to ensure the cyclic behavior not only for the removal but also the adding of the part, hence double the work. Below is the average root mean squared error of the different simulation combinations under different loads.

		Removing Method					
		Random	**Max**	**Min**	**Repeat**	**Cyclic**	**Ø**
Adding Method	**Random**	15.9	106.4	196.0	125.7	9.6	**90.7**
	Max	12.7	91.4	185.8	168.6	6.7	**93.1**
	Min	12.3	99.5	198.5	139.3	1.5	**90.2**
	Repeat	14.6	168.9	160.9	122.0	8.8	**95.0**
	Cyclic	15.6	129.2	208.6	99.1	0	**90.5**
	Ø	**14.2**	**119.1**	**190.0**	**130.9**	**5.3**	**91.9**

Another method that worked pretty well was to remove the parts in a random order. For every removal, you pick a random lane. This had an average sequencing error of around 14, which may or may not be acceptable. At a first glance the benefit of taking parts from a random line may be that you do not need any system to pick the lane; the employee simply chooses any lane. However, humans are pretty bad at generating randomness, and there is a huge difference between “pick any lane” and “pick a random lane.” In case of doubt the employee may simply go for the lane closest to him or easiest to reach. Hence, I would not really recommend this method.

16.4 Summary

If you do not need a perfect (strong) FIFO sequence from your parallel FIFO lanes, but a weak FIFO is good enough, **simply have the parts removed always from the next lane to the previous part removal**. Hence the removal cycles lane by lane through the entire set of parallel lanes over and over again. While this will not give you a perfect FIFO sequence, it will be a reasonably good sequence, and good enough for many cases, especially if it is easier to implement than a perfect FIFO sequence.

Okay, I hope this problem was relevant to at least some of you. I definitely faced this issue on the shop floor – and the wrath of the purists that wanted a strict FIFO no matter what. In any case, **go out, get your parts in an approximate sequence, and organize your industry!**

16.5 Source

Kalkanci, Kaan: "*Entwicklung und Simulation von Ein- und Auslagerungsstrategien zur Sicherstellung des FIFO-Prinzips bei dezentralen parallelen Materiallagerungen*," Master Thesis, Karlsruhe University of Applied Sciences, Karlsruhe, Germany, 2019.

17 Relation between Quantity and Cost in Manufacturing

Christoph Roser, April 23, 2019, Original at https://www.allaboutlean.com/quantity-and-cost/

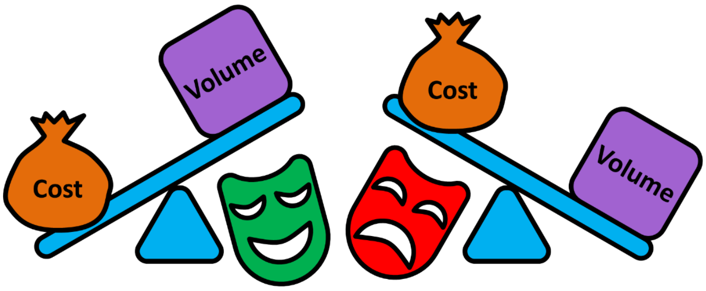

Figure 136: Cost Volume See Saw (Image Roser)

As you surely know, it is more efficient to produce larger quantities. This is the economy of scale. In a recent post I talked about the Power of Six, a rule of thumb for the relation between lead time and cost. In this post I will show you a rule of thumb for the relation between quantity and cost. Credit for this rule goes to Juan Carlos Viela.

17.1 Economy Of Scale

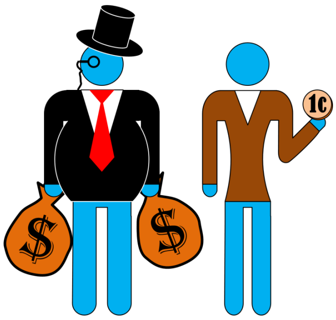

Figure 137: Rich and Poor (Image Roser)

The economies of scale are a well-known trend between the quantity produced and the cost per item, or more generally the cost benefits of larger enterprises. This initially applies to the entire enterprise, where larger companies can often produce more efficiently than smaller companies. However, it also applies to individual products, where producing larger quantities is likely to reduce the cost per item.

Some of these benefits can be calculated (or more accurately, **estimated**) by cost accounting. Others are there but are hard to grasp quantitatively.

17.1.1 Some Things Accounting Can Figure Out

Figure 138: Angry senior man pointing his finger at somebody (Image Minerva Studio with permission)

There are a few effects of economy of scale that accounting can figure out and quantify. Most of them are fixed costs that are spread across more products sold. Examples of these are:

- **Tool or Machine Utilization**: The more parts you make with a tool or machine, the smaller the share of the cost of each part for the tools or machines.
- **Development Cost**: The cost to develop the product is also shared among more parts, reducing the cost per part if you produce more.
- **Material Cost**: Economies of scale also apply to your suppliers, and they are likely to offer you discounts for larger purchases.
- **Management Overhead**: The more products you make, the smaller the share of overhead cost for each part … usually. However, be aware that management costs also go up with more parts, even though they increase slower than the number of parts, as you get more overhead personnel and they may be paid more (since they manage a bigger product volume). In the worst case, you can even have a *diseconomy of scale*, where the overhead increases faster than the product quantity.
- **Marketing**: Marketing always connects only to few real customers compared to the number of people who see the advertisement. The more customers you have, the more effective your marketing.

17.1.2 Some Things Accounting Cannot Figure Out

Figure 139: Angry businessman (Image Minerva Studio with permission)

Some other aspects are definitely there, but are difficult to grasp by cost accounting. Hence cost accounting usually assumes them to be zero – which they are not! See my post The Problems of Cost Accounting with Lean for more. Examples here are:

- **Worker Experience**: The more parts of a certain type a worker makes, the faster he can make them. While this is true for product groups, it is also true for individual parts.
- **Smaller Buffer Inventories**: With larger quantities you may need a bit more buffer inventory, but the buffer inventory usually increases slower than the product quantity. Overall the cost per part will go down. See also my post How Product Variants Influence Your Inventory.
- **Improved Flow**: Producing more parts is likely to improve both the material flow and the information flow. This has a whole lot of synergy effects that are hard to quantify like faster detection of quality problems.
- **Leveling**: The more parts you make, the smaller the fluctuations in your value chain. Customer orders will be more regular, production will be smoother, and leveling will improve. Hard to quantify, but a significant effect.

17.1.3 It Is Bidirectional!

Figure 140: Two Way Traffic (Image MarianSigler {bla}, Mediatus in public domain)

Please be aware that this economy of scale works in both directions. If you produce more, the cost per item is likely to go down. If you produce less, the cost per item is likely to go up.

These curves are not always linear. If you need a second machine, you may have a jump in fixed costs, and until the second machine is also well utilized your cost may go up a bit with more production.

Also particularly bad is that this curve is not the same in both directions. Especially if you have to reduce your quantity, you will find out that your management overhead, number of employees, and other factors will not go down quite as quickly as you would hope. This often causes real problems for companies in a downturn.

17.2 Estimating the Effect of Volume on Cost

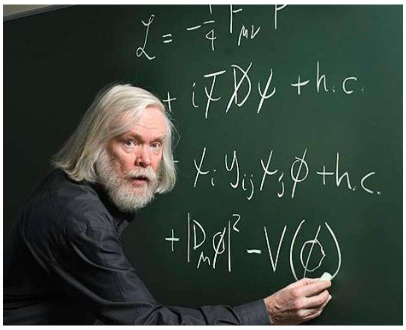

Figure 141: Scientist on Blackboard (Image Open Knowledge Foundation under the CC-BY 2.0 license)

Juan Carlos Viela discovered a simple formula that can estimate the change in the cost of a product if the quantity changes. The formula is quite simple:

$$\Delta C = -\frac{V}{k}$$

- ΔC is the percentage change in cost.
- ΔV is the percentage change in volume.
- k is a constant. See below for more information on this value.

For example, if you increase your volume by 10% and assume a constant k of 4, your cost per item would go down by 2.5%.

$$\Delta C = -\frac{V}{k} = -\frac{10\%}{4} = -2.5\%$$

Similarly if you decrease your production volume by 10%, your cost may go up by 2.5%.

17.2.1 Validity

Please be aware that this is only an estimation, and the accuracy of this estimation may vary depending on your situation. Furthermore, this formula is used only for small changes in volume of up to ±20%. Hence if your production volume increases or decreases by more than 20%, the formula is no longer good.

17.2.2 What is my Value k?

The big question of course is: What is your value of k? By experience for many cases, a value of k could be around 3 to 5. Hence your cost increases by 1/3rd to 1/5th of the volume decrease; or decreases by 1/3rd to 1/5th of the volume increase. This relation is shown in the graph below for values of k of 3, 4, and 5, showing the expected change in the cost for a change in volume.

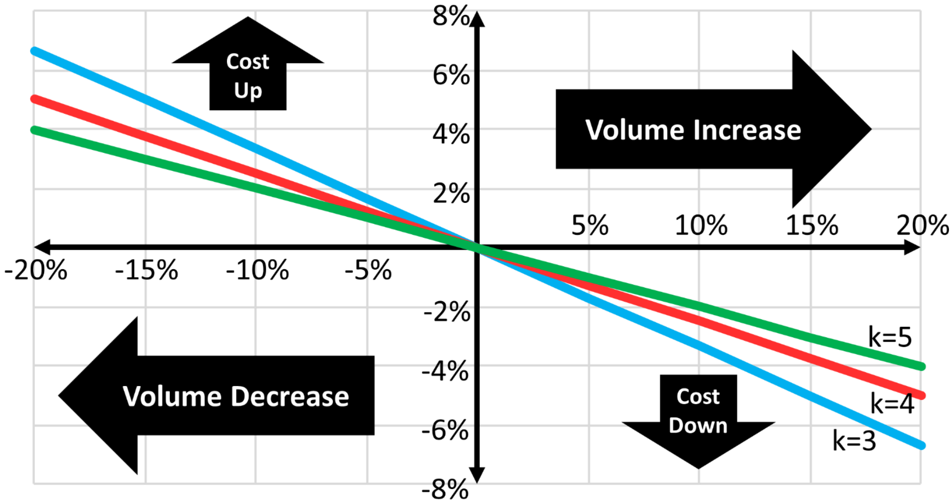

Figure 142: Cost Volume Relation Graph (Image Roser)

Again, please note that this is an estimate that fits a lot of situations for quantity changes between -20% and +20%. It may not necessarily fit yours, but if you have no idea otherwise it may help as a first indicator. In any case, I hope this helps you to estimate the impact of smaller volume changes on the cost. Now, **go out, get the cost down, or the volume up, or both, and organize your industry!**

P.S.: Many thanks again to Juan Carlos Viela for sharing this relation and giving me permission to write about it.

18 Diseconomies of Scale

Christoph Roser, April 30, 2019, Original at
https://www.allaboutlean.com/diseconomies-of-scale/

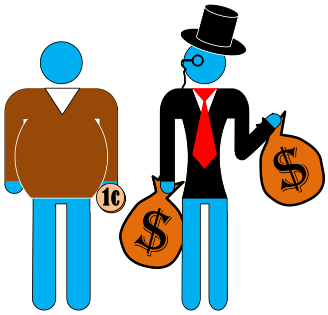

Figure 143: Bigger means poorer? (Image Roser)

The economy of scale is well known. The larger a company gets, the more efficient it becomes. However, this trend does not go on indefinitely, and eventually turns into a diseconomy of scale. In this post I will look at the diseconomies of scale, and also the very related Parkinson's Law.

18.1 Economies of Scale

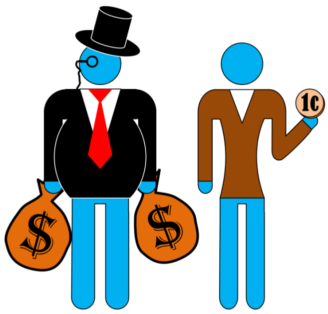

Figure 144: Bigger means richer? (Image Roser)

Economies of scale are well known. I wrote a bit about them in my last post. There are lots of different causes, from machine utilization, buying in bulk, more efficient use of overhead, reduced development cost per part, and so on.

However, if we would follow this thought logically, then a larger company is always more profitable than a small one, overtaking or bankrupting the small one. In the end there would be only one large "World Corp." company, since it can make things much better, faster, and cheaper than any smaller company.

Obviously, this is not true, thanks to diseconomies of scale.

18.2 Diseconomies of Scale

Diseconomies of scale are effects where the costs go up with the size of the company, and profitability goes down. There are multiple reasons for this. Most of them hare hard to quantify, and hence often ignored by conventional cost accounting.

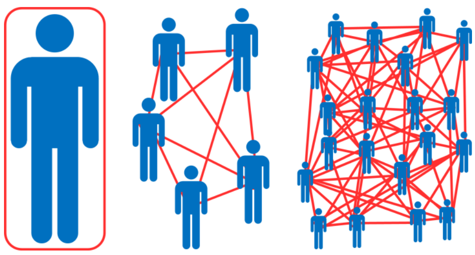

Figure 145: Network Connections … (Image Roser)

First of all, there is **communication**. If you are a one-man business, then you know everything any employee of your company (i.e., you) knows about the business. If you hire more employees, there is an additional need to communicate with each other to keep up to date. For small groups this is not a problem. However, the effort becomes larger for larger groups and eventually prohibitive. There will be subgroups that have little or no interaction with other subgroups. As an effect, the left hand does not know what the right hand is doing. If you work for a larger company, you are surely familiar with many examples of this. Lots of departments trying to do what is best for them, which may not be what is best for the company. Sales selling tons of products that lose money, because … hey … it's a sale!

These larger structures also lead to **larger hierarchies**. Communication (either top down or bottom up) takes longer, and much information may be lost on the way. Top executives who are clueless about what really is happening on the shop floor are common. People at the bottom also often feel disconnected from the top, and sometimes even left alone to do whatever they please. In any case, since the information has to flow much further and through many more hands (possibly with signatures, too!), the information flow is much slower. As a result, the entire company is much slower.

This also leads to a **duplication of efforts**, where multiple people are doing the same thing. Depending on the situation it may be not only double work, but an effort afterward to integrate or merge these two different approaches. Wikipedia has a nice example of General Motors developing two CAD systems independently for two different branches of the company. Later they had to merge it into one at a significant additional cost.

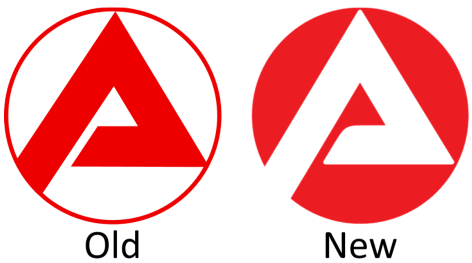

Figure 146: Updating the logo … for 10 million € (Image Bundesanstalt für Arbeit for editorial use)

Another thing that happens is **organizational rot**. At one point a section or department was created, which may or may not be necessary. Even if it was necessary, the department usually exists much beyond the need for the department. Minor topics may be given to new departments, much beyond the actual need of the company. I know a large company that has a separate department in charge of the corporate color (i.e., what color the logo of the company is; if you ever want a job that has no stress at all, this one is it!). Managing the company color means most often not doing anything. Other companies seem to have departments with apparently the only task to keep themselves and others busy. As an example, the German unemployment office "*Bundesagentur für Arbeit*" (around 100,000 employees) updated their logo in 2005. You can see both the old one and the new one here. Total price tag including all associated expenses was around 10 million Euro. I hope you do like the new one better …

Figure 147: The top one is cheaper AND better! (Image IFCAR and OSX in public domain)

Inter-company competition (sometimes called cannibalization) also hampers profitability. If two branches of the same company try to undercut themselves by price, the company as a whole will lose money. Common examples are multi-brand automotive companies, where one brand competes with another one. At GM, cheaper Buicks compete with other brands. At Volkswagen in 2010, the much-cheaper Skoda Superb Combi won a ranking in a famous magazine against the much-pricier Volkswagen Passat. As a reward, the Skoda manager in charge had to leave …

Another one would be the **Ringelmann Effect**. It is scientifically proven that in larger groups, people don't work as hard as they do on their own, because their contribution makes less of a difference. There are many more effects, but overall the economy of scale eventually starts to reverse into a diseconomy of scale. When this happens is not always clear. A well-managed company may be able to benefit from economies of scale for quite some time while they grow, whereas an ill-managed company may feel the effect at an much earlier stage.

18.3 Parkinson's Law

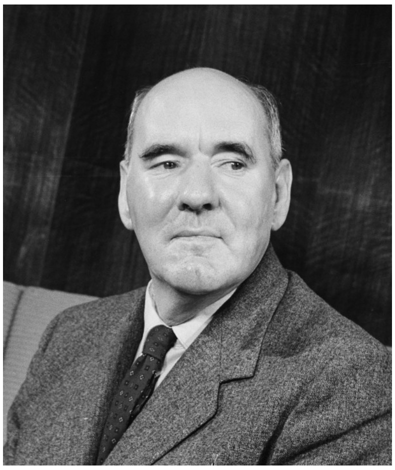

Figure 148: Cyril Northcote Parkinson (Image Wim van Rossem, Anefo under the CC-BY-SA 3.0 Netherlands license)

The whole idea is also related to Parkinson's law, named after Cyril Northcote Parkinson. This is a somewhat **loose, satirical, and humorous** observation that has to be taken with a grain of salt. The law states that

- the work will fill the available time;
- employees prefer subordinates to rivals; and
- employees together create even more work.

His conclusion was that the manpower of organizations grows 5–6% per year regardless of external influences. While by no means a scientific relation, many small observations seem to fit this guideline.

18.4 Preventing Diseconomies of Scale

Preventing diseconomies of scale is not that easy. Here are a few suggestions that may (or may not) help. Please bear in mind that this list won't solve all your problems, nor is it complete.

First of all, you could **keep the company small**, or at least stop it from growing too big. While this goes against the convention of industry and is often unpopular with shareholders, some companies choose this path.

Figure 149: Keep Calm and Create Urgency (Image Roser)

Keep a **sense of urgency**. Many good managers try to reflect the risks of business somewhat onto the employees, keeping them aware that there are many risks which could undermine the success of the company. If the manager would tell his people that they are making tons of money and are doing very well, the people may slow down a bit (since everything is fine) or may even ask for some of the "tons of money." If the employees are aware that there are always threats to the business, it may keep them a bit more on their toes. On the other hand, do not get them into panic mode either, as this would also be disruptive. Find a middle ground between overly relaxed and panicking.

Figure 150: Magnifying glass (Image Roser)

Regularly **evaluate the necessity** of different programs, sections, projects or other parts of the company. Try not to have holy cows that need to be protected beyond their use. Sometimes a department has to shrink or be shut down, a project could be stopped, production of a product could end, a plant could close, and so on. Be aware, of course, that it is much easier to open/start/enlarge something than to close/end/reduce something. Also, if at all possible, avoid firing your people, since this will definitely hurt morale in the company, and the people will spend a lot of their time worrying about their jobs instead of working. Also, if you reduce the workforce too much, there won't be any capacity left for improvements, which again will hurt you in the long run.

Reduce meetings in quantity, duration, or number of participants. If you are like most companies I know, then you have too many meetings. Similar goes for email. Sending an email to everyone, even if it is only CC, should be limited for only a few emails. Make sure you do not "reply to all." Overall, streamline communication.

So, these are a few suggestions with which you could work against the diseconomy of scale. What? … You expected them to be easy? … Sorry, they are not. Most of them are a **constant fight against corporate inertia, traditions, and habits, as well as people's unwillingness to change**. It can be done, but it is not easy. Now, go out, fight corporate inertia, and organize your industry!

19 On the Span of Control

Christoph Roser, May 07, 2019, Original at https://www.allaboutlean.com/span-of-control/

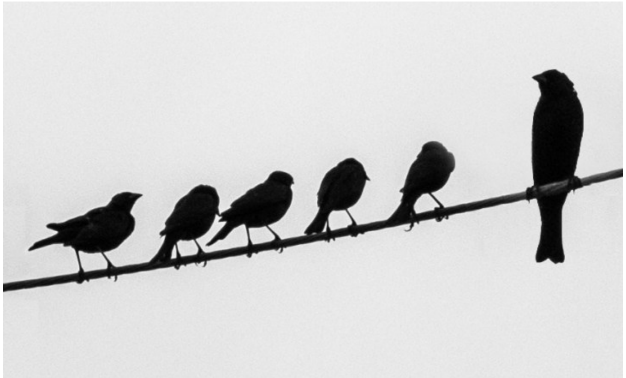

Figure 151: Little birds and big bird on wire (Image Tomascastelazo modified by Roser under the CC-BY-SA 3.0 license)

Pretty much all companies are based on a hierarchical structure. One superior manages multiple subordinates. The question is: How many subordinates should be managed by a superior? This is also called the span of control. This depends on a number of factors. Let's have a look at efficient group sizing.

19.1 Introduction

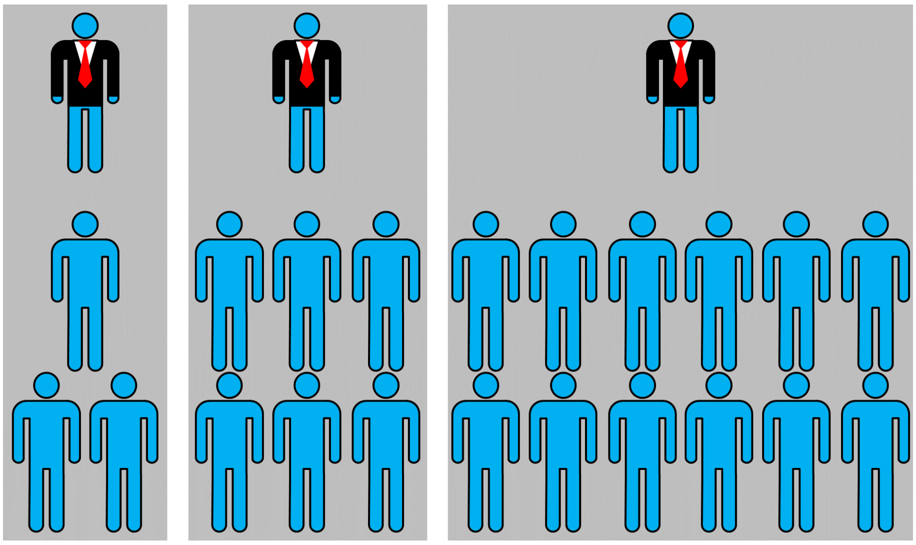

Figure 152: Team Hierarchy (Image Roser)

The leadership span depends in short on the workload of the supervisor. If it is a lot of work to supervise, then the supervisor can manage fewer people than if it would be an easy supervision job. This workload of managing people has to be combined with "other" work the supervisor has.

In the end, the time needed to both manage and do other jobs should be sufficient to do all tasks properly. Overloading the supervisor will have him out of necessity doing jobs sloppier or not at all to manage the workload within the available time. If the supervisor has too much work, it may be necessary to adjust the workload by adding hierarchical levels, changing group size, or assisting him with his other tasks.

When determining the right workload, unfortunately, looking at the available free time of the supervisor is not a good indicator. Humans (including supervisors) have the tendency to fill the available time with the available work. If the supervisor has too much time, he will fill up the time somehow. Complaining about too much work is also a common habit. Therefore finding the right level is a challenge.

In the worst case, the supervisor (or the worker in general) gets used to the situation of having too much or too little time. He may see wasting time (if too much time) or making sloppy jobs (if too little time) as the new standard. You would not only need to adjust the workload or time, but also the attitude of the supervisor toward work. Overall, this topic has lots of options to mess things up, which DOES happen regularly in industry.

19.2 Effects Impacting Leadership Span

Here is an overview of the most important factors impacting the leadership span.

19.2.1 Complexity of the Supervision

One aspect impacting group size is the complexity of the supervision. Some things are easy to supervise, and the supervisor can manage multiple people. Other jobs require a lot of close supervision. This depends on the skill of the subordinates (see next point), but also on the complexity of the task at hand.

Figure 153: A typical McKinsey Team (Image Traimak with permission)

For example, consulting requires a lot of close coordination between team members. At McKinsey, we usually worked in groups of 1 to 3 people below the first level of management, since a lot of coordination was required, including coordination with the people of the client even if they were not part of our team.

An assembly line where every part is the same can be supervised easier than a line where every product is different. The more similar the tasks, the easier to supervise.

19.2.2 Independence of the Subordinates

Figure 154: Pointing in a Factory (Image Tyler Olson with permission)

Another factor is the level of independence of the subordinates. Do you manage a team of beginners who need constant supervision and coaching? Or is it a group of experts who are best left alone after giving them the basic outline of the task? Naturally, it is more time consuming to handle newbies than seasoned experts.

19.2.3 Location of the Task (Ease of Communication)

Another factor is the location of the task. Are all people close at hand, ideally in the same room? In this case the manager can assist any needs quickly, and can switch between different employees quickly. Communication can be fast. Clues can be picked up non-verbally (employee getting nervous: Go and ask!). The supervisor can communicate with different people in short succession.

Figure 155: Where is everybody else? (Image Raymundo Valencia in public domain)

If the employees are more distributed, it is more difficult to supervise. The supervisor may have to do a lot more walking around, which consumes time. Getting information is also more time consuming for the same reason.Overall, the more distributed the employees are, the harder it is for the supervisor to actually supervise, since the communication requires more effort.

If they are working off site, it may be impossible to do supervision. The supervisor may be able to accompany one off-site worker, but cannot attend to any other workers at the same time. In this case it is common to pair an experienced worker with a newbie to teach him the ropes and be sort-of supervisor.

19.2.4 Continuous Improvement (Kaizen)?

One of the (often forgotten) side jobs of a supervisor is to work on continuous improvement, both on his own and with his team. Here I have data on an interesting correlation. I measured efficiency (as the percentage of value adding time at assembly lines) and compared it with the lowest-level team size on the shop floor. For the full data see my Grand Tour of Japanese Automotive. The efficiency (red line) and the team size (blue bars with orange for the range of team sizes) is shown below.

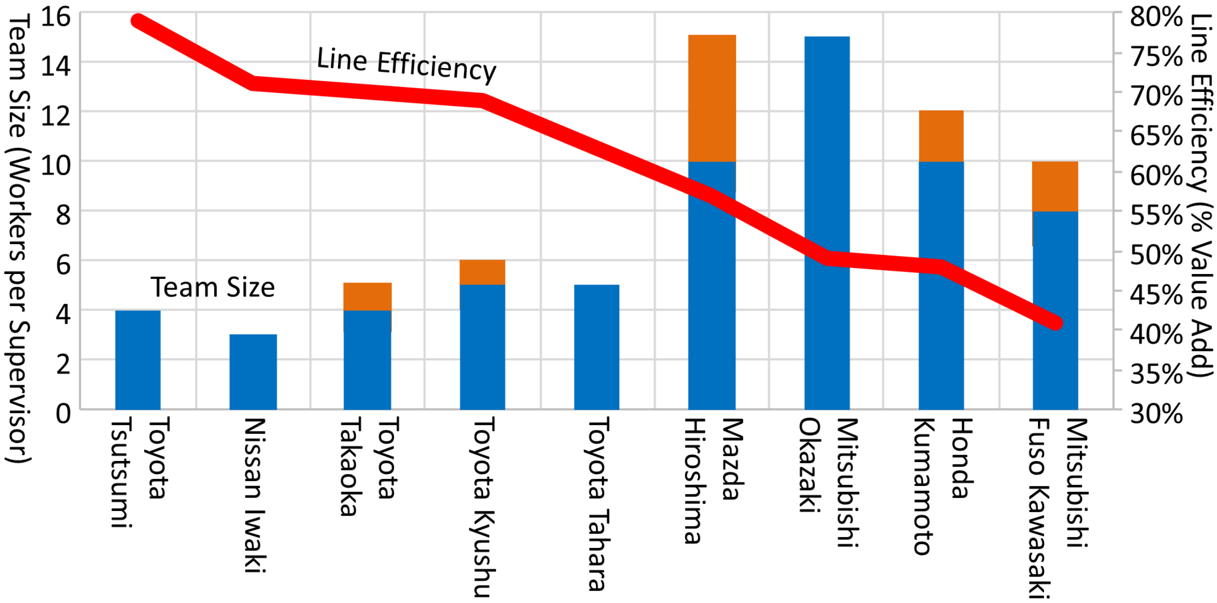

Figure 156: Team Size and Efficiency (Image Roser)

Most curiously, plants that worked with smaller teams were much more efficient (i.e., the workers used more of their time actually working on the product rather than walking, waiting, transporting, or talking). And this included the supervisors at the line. Of course, correlation does not mean causation, but I do like to think that there is some connection. I believe smaller teams have more time to work on improvements to make the process faster, better, and cheaper.

It also matches my experience that an overloaded supervisor will drop optional things, with improvement work often being considered optional. I strongly disagree with this, but if I would be a shop floor team supervisor, I too would under time pressure put more effort into keeping the line running rather than improving the line, since I would get yelled at much more for a stopped line or late deliveries than a possible improvement project.

19.2.5 Other Workload of the Supervisor

Figure 157: Overworked worker (Image Roser)

Also do not underestimate the other workload of the supervisor. To give you a few examples from the shop floor, he may be responsible for data entry into the ERP system. He may be the first contact for maintenance and repairs. He could be the one planning the work schedules and absences. He could have to cover unplanned absences. He is usually the one looking for missing material. And so on. While many of these tasks are often not recognized at the higher levels of management, they are essential for a properly functioning shop floor.

19.3 Some Examples

In Japan, many automotive plants have a shop floor supervisor in charge of 3 to 15 people. Smaller teams seem to indicate better performance, as the supervisor has more time to help, fix, and especially improve things. Information flows faster and better in both directions, since a manager managing 5 people actually **has time for his people!**

In German automotive, a lowest-level supervisor often covers 20 to 25 people – on top of data entry, material search, organizational stuff, maintenance, and many other things. He simply **has no time for his people** to manage them properly. Additionally, the workload does go up more than the number of people. If you double the number of people supervised, you more-than-double the work for the supervisor. This would yet be another reason for smaller teams.

In the past, there was often a span of control of around 4 workers per supervisor. Nowadays it is rare to find a supervisor in charge of less than 10 people, and often even more. It is argued that modern communication makes communication easier and allows for easier supervision, and there is definitely some truth in that. However, I argue that another factor is continuous pressure to be profitable. Since the value of a supervisor cannot realistically be determined by cost accounting, cutting supervisors was seen as an easy way to cut cost. Unfortunately, reducing supervisors not only cut cost, but may have cut efficiency and productivity even more.

Overall, I do not have you an easy rule for the span of supervision. There are some available in literature, but I am not sure how valid they are. However, if you are a Western company, in all likelihood your span of supervision is too big. But equally likely, you will not have the money for more supervisors. Nevertheless, **go out, try to reduce the span of control, and organize your industry!**

20 On Adjusting Supervisor Workload

Christoph Roser, May 14, 2019, Original at https://www.allaboutlean.com/supervisor-workload/

Figure 158: Drowning in Boxes (Image alphaspirit with permission)

In my last post I looked at the span of control. This is very related to the workload of the supervisor. Hence in this post I would like to discuss how to adjust the supervisor workload. Usually, this is to reduce the workload, as most shop-floor supervisors are in my opinion overworked and have no time left for improvement. In some cases, however, you may have a situation where you want to give the rare underworked supervisor more work. Most of the approaches presented will work in both directions. Let's look at some ideas:

20.1 Split/Merge Groups

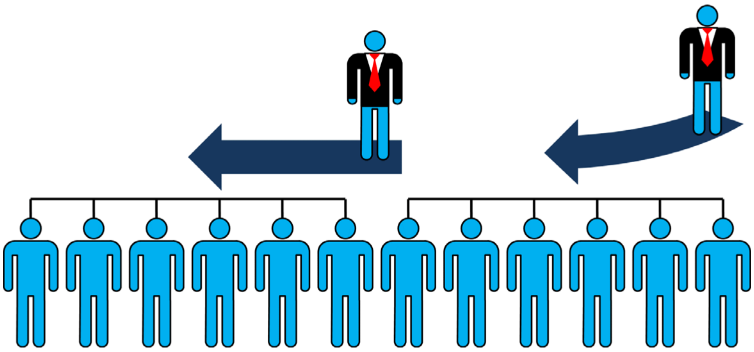

Figure 159: Split Team (Image Roser)

A common and easy solution to reduce the workload of supervisors is to make smaller groups. Toyota and Nissan use this approach on the shop floor, and their lowest-level-hierarchy span of control is much smaller than what Western industry uses. This seems to correlate with a much better performance in terms of speed, quality, and even cost on the shop floor (see my post [The Grand Tour of Japanese Automotive]() for details).

This could mean that you need more supervisors. In some cases this is the correct way to do it, and an additional person on the shop floor will reduce the work of the other supervisor(s).

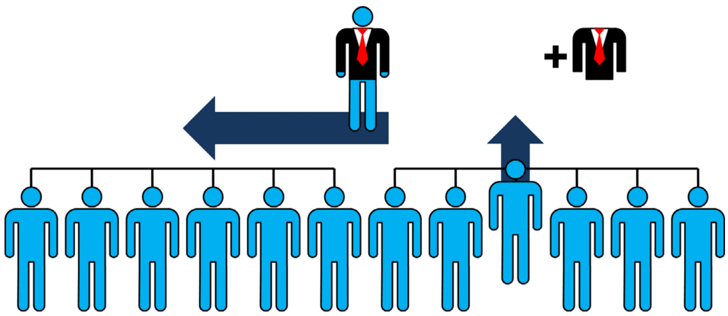

Figure 160: Promote Worker (Image Roser)

Another option would be to promote a qualified worker to supervisor. On paper this will not increase the available manpower. While you have one supervisor more, you have one worker less. However, the work of managing other people does not go up linearly. Managing twenty people is MORE than double the work than managing ten people. Managers make do by cutting corners, and the quality of the management suffers. At Toyota, the lowest level (team leader) also helps out with regular work, and takes over the job of a worker if the worker is absent.

20.2 Add/Remove a Level of Hierarchy

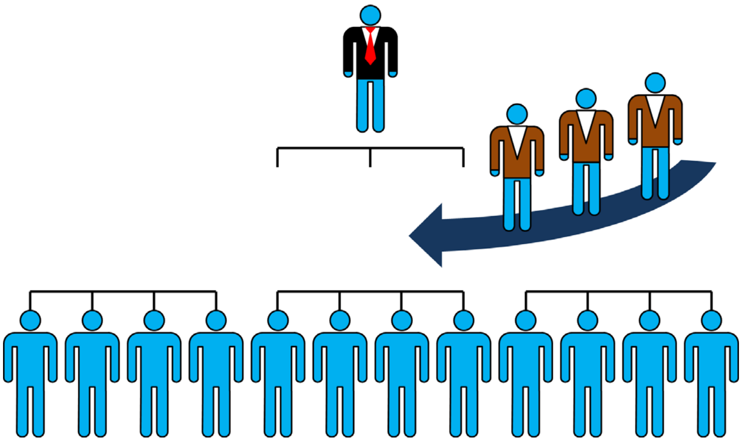

Figure 161: Add Hierarchy Level (Image Roser)

If there is the need to split a lot of groups, it may be of interest to consider adding a level of hierarchy. This option is only for larger changes, as you need many more supervisors (to be hired or promoted). Not so good for small changes. It also comes with a lot of advantages and disadvantages, like giving people more chances for promotion but also making upper-level people more remote.

20.3 Give an Assistant

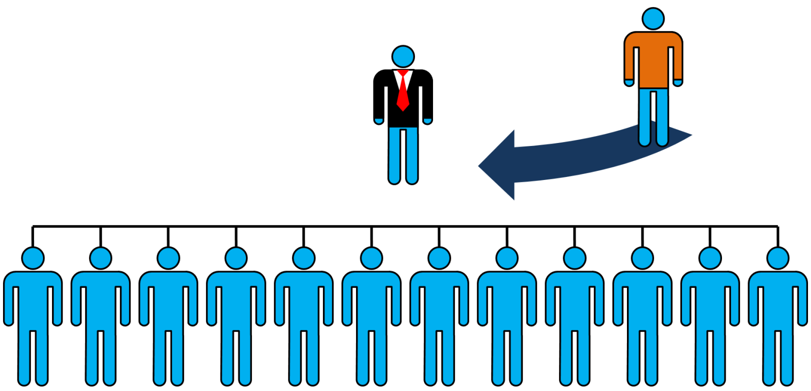

Figure 162: Add Assistant (Image Roser)

Another idea is to give the supervisor an assistant, or maybe an assistant for a group of supervisors.

On higher levels of hierarchy, this is very common. An assistant is often even seen as a status symbol, and on the highest levels people have multiple assistants for different tasks. In some cases such an assistant position may even be a fast-track for promotion.

On the lower levels of hierarchy, such assistants are rare. However, just because lower-level supervisors are paid (much, much) less, doesn't mean that their workload is less. The work may have less decision making under uncertainty, and they usually leave at the end of the shift rather than amassing overtime, but they may also be overworked.

For example, when working closely with shop-floor supervisors, I notice that they do a LOT of data entry, and nobody on the higher levels cares much about that work as long as it is done. However, it is a lot of work, and I'm often wondering if it would be possible to centralize this data entry with an assistant. This assistant would specialize in data entry and could probably do the work faster than the different supervisors.

20.4 Add/Remove Other Tasks

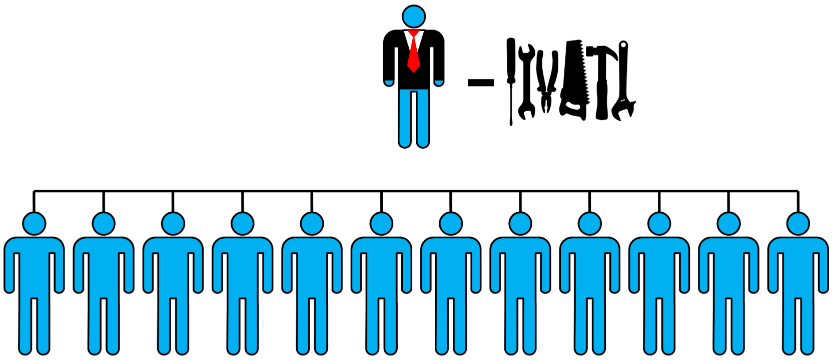

Figure 163: Reduce Workload (Image Roser)

Another option is to remove (or add) tasks from the workload of the supervisor. Are there certain repetitive tasks, possibly repetitive across multiple supervisors? Maybe you can get a person who specializes in this topic and subsequently can do the job better and faster than a supervisor?

Toyota, for example, has a mind-boggling number of people on the shop floor who solve problems. Western companies would be amazed to see how many people at Toyota exist solely to support the shop floor. In the Western world these people were often cut, with drastic results on quality and speed.

20.5 Optimize

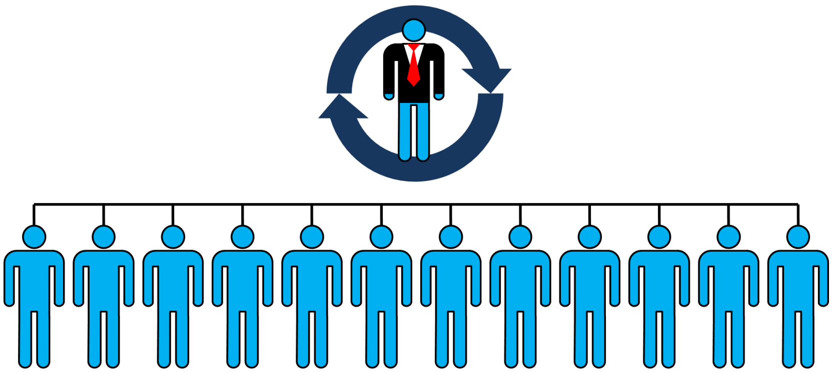

Figure 164: Optimize Supervisor Workload (Image Roser)

Yet another way to reduce the workload of the supervisor is to optimize and improve his work. Can you reduce walking distances? Does he have the tools he needs, and are these tools working well? Quite a few tools from the lean toolbox could be applied here, from 5S to the spaghetti diagram.

This, by the way, is one of the few ideas that work only in one direction. Optimization is there to make things easier or better. Nobody in their right mind would make things intentionally worse just so his people have more workload.

20.6 Don't Split Workers (Too Much)

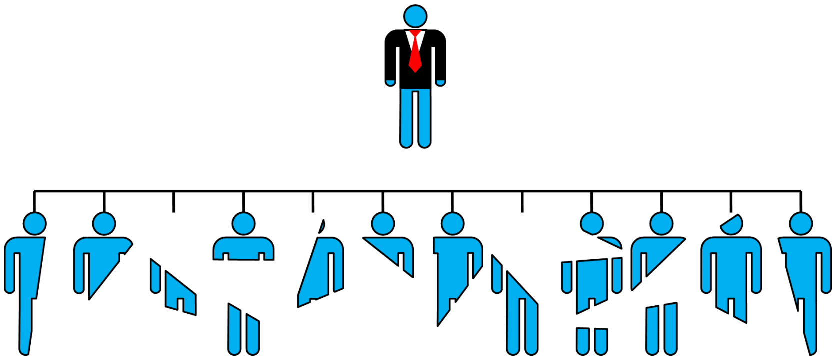

Figure 165: Manage Split Workers (Image Roser)

Another common mistake is to assign the same worker to multiple groups. This is not common on the shop floor, but very common in projects.

A worker is most effective if he has two or thee projects. Any more than that and the effort of just keeping up to date becomes too much.

Let me give you a real-world example I've observed. The leader in question was in charge of three projects. Each project had about forty people assigned to the project. However, the total number of people on all three projects was forty-seven! Even worse, the total number of man hours assigned to the project was the equivalent of thirteen people! If you are now questioning my math skills … well … **The average worker was assigned to ten different projects!** The average worker was split on ten projects, some of them with this supervisor, others with other supervisors, with roundabout 10% of his time assigned to one project. **This is madness!**

Just keeping up to date and attending the necessary meetings required more than 10% of the time. The workers did what every sensible worker would do: they picked one or two projects that they worked on, and pretty much ignored the rest.

On the shop floor, similar things happen to supervisors, who often have to do all the improvement projects. A plant I know had about forty-five major improvement projects, but almost all of them relied on a critical shop-floor manager who was also busy with keeping the plant running in the first place. Hence, there was not really any improvement going on. Please don't make the same mistake, and don't assign too many projects to a worker.

20.7 Don't Cut the Organization to the Bone

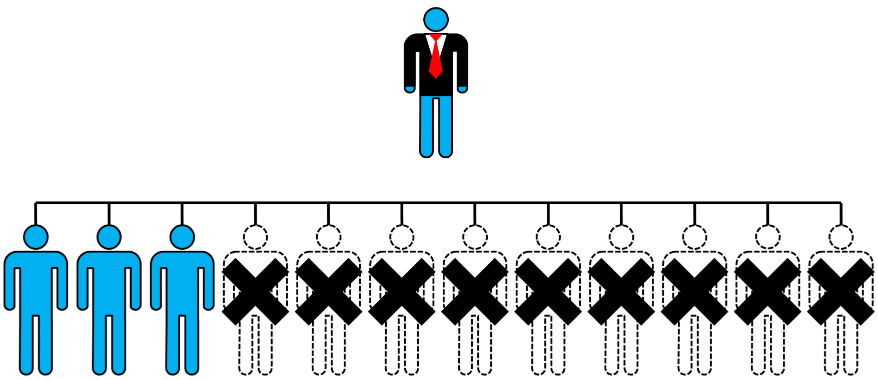

Figure 166: Crop to the Bone (Image Roser)

A lot of the problems of overloaded workers come from cost cutting. If you read some of my other blog posts, you may know that I often have a problem with cost accounting. A lot of the work of supervisors and support staff is hard to quantify, and the value is difficult to measure.

Traditional cost accounting often has the approach that if they cannot measure it, it must be zero. Hence the value of support staff and supervisors is zero. The cost of an employee, on the other hand, can be measured well. As a result, over the years more and more of this "cost without value"was cut, and many shop floors are running on fumes. (Toyota being again a notable exception).

20.8 Consequences of Excessive Workload

Overall, try to have a good workload for your supervisors (and also for your workers). It may not always be easy to determine, as few workers will complain about too little work. It may also be difficult to hire more people, as this will make you unpopular with your own boss (or your shareholders).

If your people are continuously overworked, people will start cutting corners and do sloppy jobs simply because they don't have time. Eventually, this workload will be seen as normal, and even if you increase the number of people, the "sloppy job is okay" attitude may prevail. Similar effects may happen with underwork (which is common in Japanese offices, where little gets done in an awful long workday).

Now, **go out, adjust the workload of your supervisors (and by that I mean reduce, not increase!), and organize your industry!**

21 Mixed Model Sequencing – Introduction

Christoph Roser, May 21, 2019, Original at *https://www.allaboutlean.com/mixed-model-sequencing-1/*

Figure 167: Ancient Egypt Necklace (Image metropolitan Museum of Art in public domain)

In a mixed model production line, different products may have different work content at different stations. Hence, some stations may need a longer or shorter time depending on the product. This requires careful planning of the assembly line. If this is not taken into account, it may cause significant idle time with all stations along the line. This is the first of a (very) long series of posts looking at Mixed Model Sequencing (i.e., the behavior of unbalanced workloads, and different ways to address these issues).

21.1 The Basic Problem

If your production line produces different types of products, chances are that the workload for these products is not the same for all stations. Let's take an example of an assembly line from the car industry, although similar situations can be found in almost any multi-product production line.

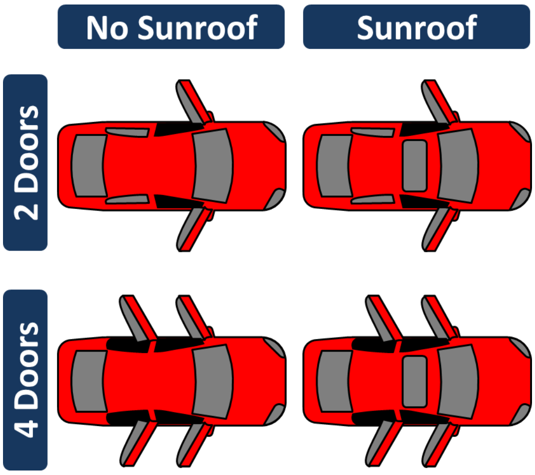

Figure 168: 2 Door 4 Door Sunroof (Image Roser)

You can get cars with two doors or four doors (or five if you count the trunk). You can also get cars with a sunroof or without a sunroof. Just with these two options you can configure four different cars as shown here.

Naturally, the station assembling the doors will have much more work with four doors than with two. The work content changes with the product that arrives.

Similarly, the workload of the sunroof assembly station also depends on the product, namely if there is a sunroof or not. If there is no sunroof, this station is completely idle. If there is a sunroof, the station is completely busy.

These are only two examples. There are many more options that can change the workload at different stations. Just go through the list of extras and options when buying a new car, and you will see how this can affect the work during assembly. Often, the actual work differences behind the scenes are even larger than what you can see in the catalog. For example, different car engines may come with different exhaust systems, different brake types, maybe changes to the gearbox, and so on.

Figure 169: Hyundai Veloster, which has one door on the driver side and two on the passenger side (Image loubeat under the CC-BY 2.0 license)

In a nicely balanced line, all stations have a similar amount of work. For more on this, see my extensive series of posts on line balancing. However, if the products are different, the balancing is not as simple as giving every station the same average work content. If you assume your car has, on average, three doors, you run the risk of overloading and underloading your door assembly depending on if it is a two-door or a four-door vehicle (unless it is a Hyundai Veloster, which actually has only three doors).

Using lean vocabulary, this uneven workload is a case of mura (斑 unevenness), one of the three evils in manufacturing (besides muda/waste and muri/overburden).

21.2 Worst-Case Scenario

For the sake of understanding, let's have a look at a worst-case scenario. Let's assume a pulsed line, where all parts move at the same time.

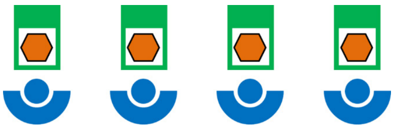

Figure 170: Pulsed Line. The animated image can be found at https://www.allaboutlean.com/timing-of-flow-lines-1/ (Image Roser)

Let's further assume all stations are perfectly balanced, and the only variable workload is the sunroof assembly station. If a car without a sunroof is produced, the sunroof assembly will be 100% idle. The time of the worker is wasted, as will be the money the company spent on his time, as shown below.

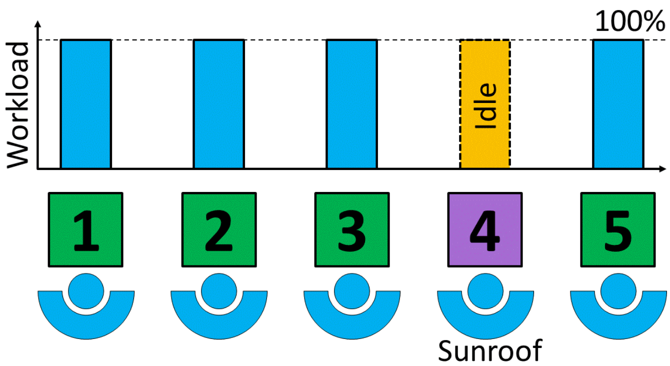

Figure 171: Pulsed Line Idle Sunroof (Image Roser)

This is an example of a station having less work. However, it is even worse if a station has more work than the others. Let's again assume all stations are perfectly balanced, and only the door assembly station has a variable workload. If a four-door car passes by, this station has much more work than the other stations along the line, and **all other stations have to wait** for the door assembly to finish. The wasted time and money is even more than when a single station has to wait for the others.

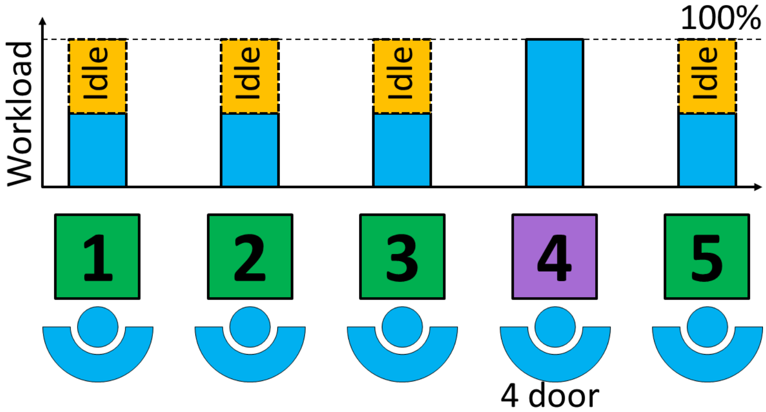

Figure 172: Pulsed Line Busy 4 Door (Image Roser)

If all of your other stations have to wait on one station, you are wasting lots of capacity. In this case it may be better to add capacity to the critical station, even if it is not used all the time. For example, you may add another worker to the door assembly, even if both workers are half-idle if a two-door vehicle comes along, as shown below. You may rather have two workers idle a bit every now and then, than have all other workers wait for the overloaded worker at the critical station. This solution is not pretty, and you may think about improving it eventually. But in manufacturing, you may not always have the time for a pretty solution, and this excess capacity will keep you going until you have time to improve it.

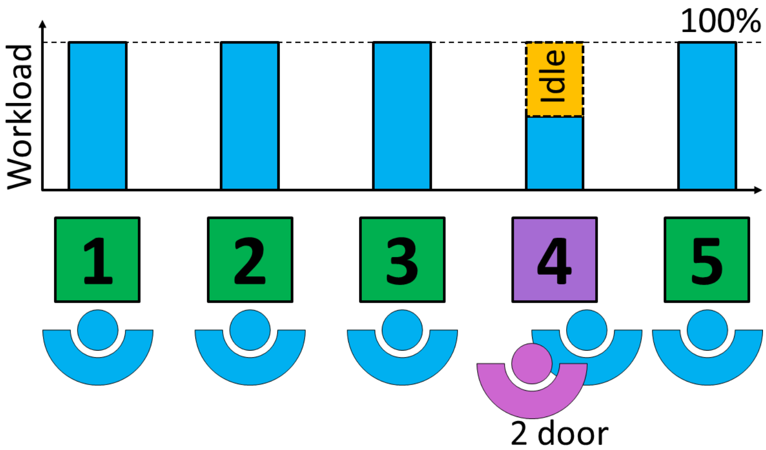

Figure 173: Pulsed Line Idle 2 Door (Image Roser)

Please note that we are worried here primarily about the **waiting time of the operators**. **Waiting time of machines** is usually cheap and nothing to worry about. This is especially true if the machine is faster than the line speed and hence not the bottleneck, as with the sunroof example above. Only if a bottleneck machine causes the entire line has to wait, like in the four-door example above, should you to consider alternative options. If one machine causes the rest of the line to wait, the line performance starts to suffer.

A pulsed line is usually the worst-case scenario for unbalanced workloads. A line where the parts can move anytime they are done is much easier. Even a continuously moving line can handle such situations much better.

21.3 Solution Approaches

Figure 174: Problem & Solution (Image Kenishirotie with permission)

You can see how in longer lines with a larger variety of products (for example, automotive assembly), this can quickly cause a lot of problems and expenses. There are a couple of different approaches to help with this issue. Please be aware that **not all of them work in all cases, and usually they are neither easy nor perfect!** Anyway, the three fundamental different approaches are:

- **Just Make the Problem Go Away**: You would just change your production system or product to remove the whole issue with varying workload. Wouldn't that be nice. But don't open the champagne yet; for most situations, this one does not really work.
- **Adjusting Capacity**: If there are different workloads, you add or remove capacity (i.e., add or remove workers). Usually works only if the increased or decreased workload persists over longer periods.
- **Adjusting Sequence**: Adjusting the sequence of products so that the workload does not fluctuate too much. Usually this means a station gets a "busy" product followed by a "non-busy" product so that on average they can manage. This would be, for example, a four-door car followed by a two-door car.

Also, all of the three approaches above can be combined, and are often combined. In the next few posts I will go into more detail on how to tackle the problem of unbalanced workload. Until then, **go out and organize your industry!**

22 Mixed Model Sequencing – Just Make the Problem Go Away

Christoph Roser, May 28, 2019, Original at https://www.allaboutlean.com/mixed-model-sequencing-2/

Figure 175: Ancient Sumerian Necklace (Image Metropolitan Museum of Art in public domain)

Your production line may have different workloads for different product variants. This unevenness causes waste and overburden. In this series of posts I will look at ways to address this unevenness. The first post was an introduction to the topic. This second post will look at ways to simply eliminate the problem – although this may not be feasible for many cases. In the next posts I will look at adjusting the capacity and finally at adjusting the product sequence through Mixed Model Sequencing.

As I said in my previous post, one option to resolve the issue with unbalanced workloads is to simply make the problem go away. While this sounds sweet, it is often unfortunately not an option. Nevertheless, you should check if it is possible, as it may make your life in production easier and your company more efficient.

22.1 Adjust the Product

One possibility is to adjust the product. For example, if you are attaching wheels to the car, some wheels may have four nuts and other five or even more. Obviously, the more nuts a wheel has, the longer it takes, and you have a product-dependent variable workload.

Figure 176: Three, four, and five lug nuts on a car wheel (Image Bindydad123 under the CC-BY-SA 4.0 license and (Image PXHere in public domain)

One possibility is to change the product to remove these differences and make the workload more consistent. If you can do that, try to move to the faster/better/cheaper options and away from the expensive ones.

Unfortunately, the product design is often out of your hands, and the designers will claim that you can sell so many more cars if they have an extra lug nut. Depending on the customer this may even be true. If your company offers only two-door vehicles, you miss out on a large market segment. The same if you do not have sunroofs. Hence, usually our options are limited here, but maybe you are one of the few lucky ones that can twiddle with the design to improve manufacturing.

22.2 Adjust the Process

If you cannot change the product, try to adjust the process. Can you make tools, jigs, gizmos, or other devices that simplify the different workloads so that the actual time is the same? Of course, here you should pretty much always go for a faster solution. There is no point in intentionally making the process **slower** for one product variant!

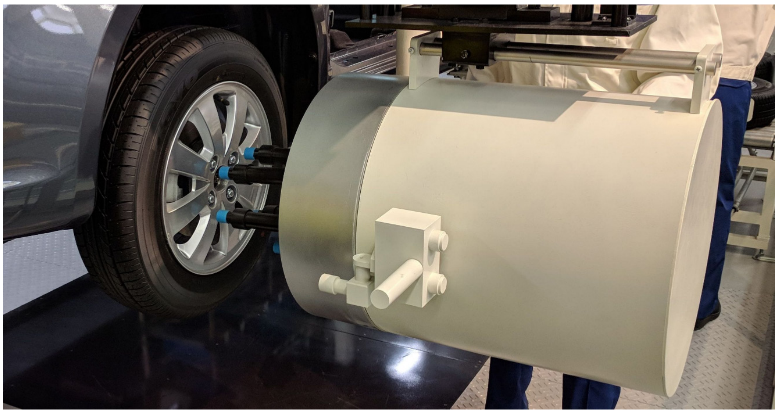

Figure 177: Suzuki Tire Assembly (Image Roser)

Let's stay with the while example. If you cannot change the number of nuts, can you simplify the tightening of the nuts? Many car companies use an automated tire assembly device as shown here in a photo from Suzuki. While this device in the photo is designed for four lug nuts, you probably can design a similar device that can handle three, four, five, or even more nuts automatically at the same time. In this case you have eliminated the uneven work content.

This example with the lug nuts is only an example, and probably won't fit your situation. But do think a bit if there are ways you could reduce the unevenness of the workload through better tools. It may not always work. For example, if you install a sunroof, there will be additional work no matter what you do. Yet even here you may have options.

Figure 178: Car Seat Assembly (Image Ford Motor Co. under the CC-BY-SA 2.0 license)

For example, you could build a fully automatic sunroof installer. True, the work content is still fluctuating, and the machine may have idle times if there is no sunroof to be installed.But it is a machine and not a human worker, and as long as the machine can satisfy the line takt, I don't care much about waiting times of machines. While not completely irrelevant, it is usually so far down my list of priorities that it is nothing to worry about.

The bigger issue is the cost-benefit analysis. Creating a fully automatic sunroof installer may not even be feasible. At least it is quite expensive, and the cost benefit of having less fluctuating workloads for the human workers may not be worth the cost of an expensive and potentially troublesome machine. Your choice.

22.3 Push it Elsewhere

Figure 179: Lego pushing wheel (Image Mirco Vacca with permission)

Sometimes it is possible to move the problem workload out of the main line and towards a secondary production that supplies the main line. This can reduce the problems at the main line, although it may introduce problems at the secondary system. However, sometimes the secondary system can be adapted easier to handle the issue. Even if it is still a problem, the secondary line may be shorter, fewer workers may be affected, and the problem may be overall smaller.

22.4 Separate Production Lines

Another option for changing the process is more radical, and for most situations useless. You could split the production line into two independent lines. One line makes only product A, while the other line makes product B. Each line could then be optimized and balanced for its own product, with no worries about the workload of the other product.

For example, instead of having a joint assembly line for two-door and four-door cars, you make one line exclusively for two doors and one line exclusively for four doors as shown below. This, however, is only a theoretical example, as it makes no sense for the car industry to split these lines. You will lose flexibility, you will lose efficiency, your machine expenses will go up, and so on. In most cases it makes no sense to split a line.

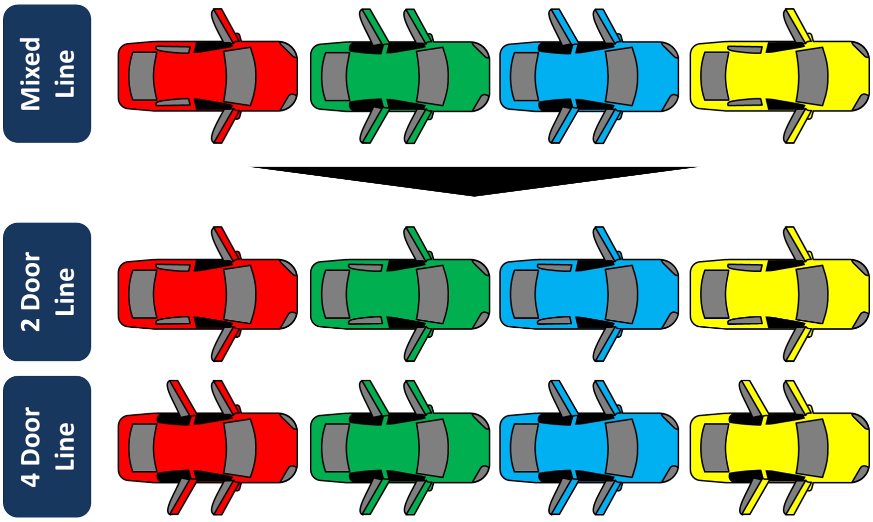

Figure 180: Split Assembly Line by product Type (Image Roser)

But maybe, just maybe, you have a situation where this is feasible, where you have only one or two product variants or product families, and inexpensive machines. For example, if it is a purely manual assembly, you could have a long line with ten people or two shorter line with five people, each line making only one product (family).

As you can see, there are some options on how to eliminate the problem of uneven workload due to different products – or at least push it on a machine where we don't really care, at least not too much. Yet they may be difficult in many cases. It is quite likely that you read through this post looking for a solution, only to find that none of these options will work for you. But don't despair. In my next few posts on the topic of product-dependent workload, I will look at adjusting the capacity, before finally getting to product sequencing to handle workload differences. This last one is usually feasible, but it is also more of an art than a science. Until then, go out, see if you can reduce unevenness in the workload (or elsewhere), and **organize your industry!**

23 Mixed Model Sequencing – Adjust Capacity

Christoph Roser, June 04, 2019, Original at https://www.allaboutlean.com/mixed-model-sequencing-3/

Figure 181: Ancient Egypt Animal Shaped Beads Necklace (Image Metropolitan Museum of Art in public domain)

Production lines with a product mix may have different workloads at different stations for different products. This can cause waste. In this third post in the series I will look at options on how to adjust the available capacity to ease this problem. In my next post I will look at Mixed Model Sequencing to adjust workload differences.

23.1 The Basic Idea

In its ideal form, adjusting capacity would mean to add one (or more) workers whenever a part with more work comes along, and remove them again once a part with less work comes along. For the example of an automotive assembly line with two-door and four-door vehicles, this could look like the animation below. A worker magically pops up whenever there is a four-door assembly, and disappears again if there is a two-door assembly.

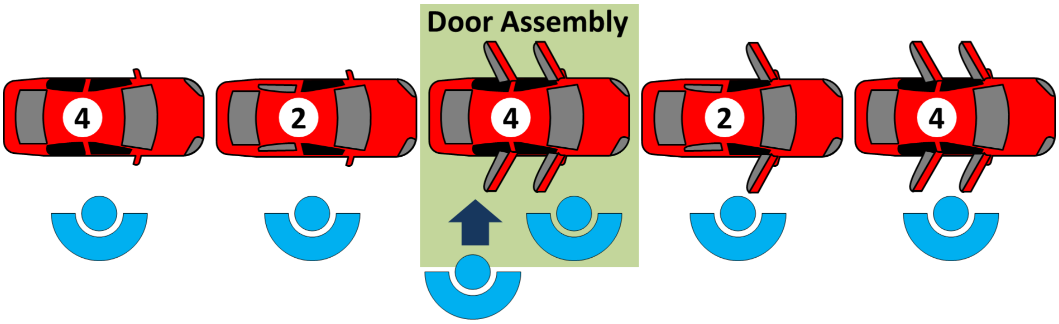

Figure 182: Animated Example for an Ultra Flexible Capacity Line. The original image can be found at https://www.allaboutlean.com/mixed-model-sequencing-3/ (Image Roser)

You probably can already see the problem here. It will be extremely difficult to manage the worker's workload. It will be very difficult to give the worker something else to manage. You (usually) just can't conjure some capacity out of thin air at a moment's notice. In the following I will show you some ideas how you still could manage.

23.2 Find Another Task to Fill the Capacity

It is possible to have a second worker at the critical station, as shown in the animation above. This worker can help if there is a part that requires more work (i.e., a four-door car). The challenge is to not waste this worker's time if a part requires less work (i.e., a two-door car).

You would have to find an useful task for the extra worker that he can do in the available time. This is usually not easy, or even impossible. If the worker is busy with four-door cars, the other task would have to wait. Hence, the other task must be a job that **can wait anytime**. The other job would also be something that **can be dropped anytime**, since the door assembly can **NOT** wait. Whenever a four-door car comes along, the worker must go to assemble doors. If you can find such a task, then this may be possible.

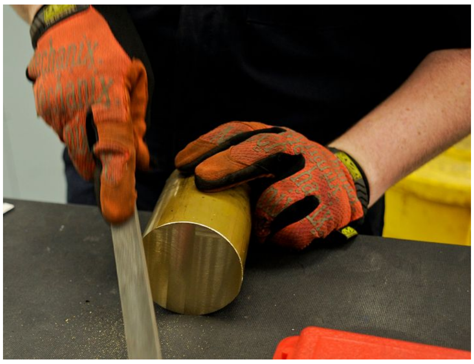

Figure 183: Deburring (Image Mariette M. Adams in public domain)

For example, you may have a box of sheet metal parts that need to be deburred. The additional worker deburs these items until he is needed at the assembly line, in which case he can just drop the part and continue later.

Even then, frequently interrupted work is a much less efficient work. In the worst case the worker may get confused and think he already did a secondary task, but in reality he did not yet do it. Subsequently a not-completed item from the second task is moved forward. For example, you can easily imagine a half-deburred part mistaken for a completely deburred part.

It will be easier if there is a longer time between parts, either due to a larger cycle time or due to infrequent larger products on the main line (i.e., four-door vehicles would be rare). In this case there will be less interruptions. Nevertheless, this solution is rarely feasible and usually not worth the effort.

23.3 Larger Batches

Another option is to make larger batches – and let me just say right at the beginning that **I don't like large batches!** A smaller batch size is usually much more efficient, and increasing batch sizes is usually the wrong way. The vision of lean is a lot size of one, or a one-piece flow. Never lose sight of that!

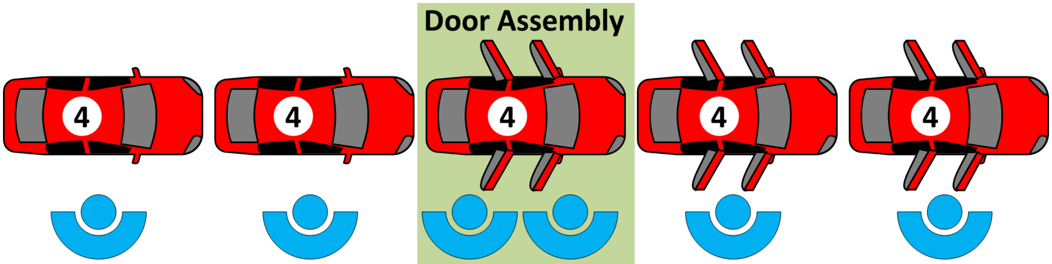

Figure 184: Four Door Assembly Only (Image Roser)

This being said, you could make larger batches, which gives you more time to move workers around. In the extreme this could be a different product type every shift. For example, you could do one shift of only two-door cars, and man the door assembly workstation with only one worker. In the next shift you produce only four-door cars, and man the workstation with two workers.

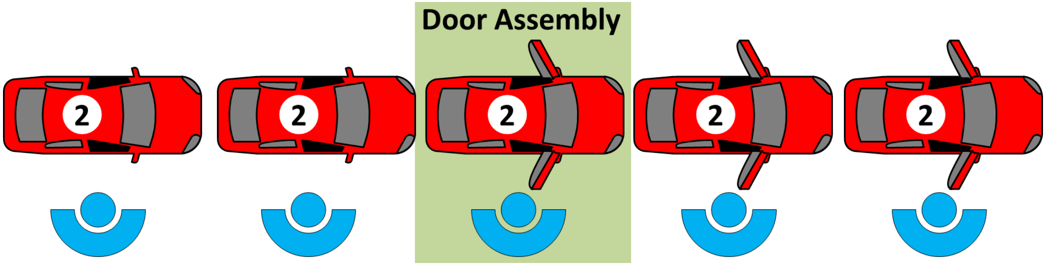

Figure 185: Two Door Assembly Only (Image Roser)

This would give you much more time to assign another worker. It is difficult to assign a worker a different task if he has only sixty seconds to spare. It is much easier to assign a different task if he has a few hours or even a whole shift to spare.

But again, if you choose this path you are walking in the wrong direction. For example, the supply lines that create your doors would be totally overloaded during a four-door shift, and would twiddle their thumbs during a two-door shift. You could improve this by having a larger buffer of doors waiting for assembly, but then increasing inventory is also not good. Overall, this option should be avoided due to all the negative consequences. It is quite likely to be penny wise but pound foolish.

23.4 Flexible Work Assignment

The whole problem of different workloads disappears if your workers have a flexible work assignment (i.e., the worker can take however long it takes to do a task before moving to the next task). Job shops would be one example, although job shops are usually inferior to flow lines. Other options are related to work cells.

23.4.1 One-Person Work Cell

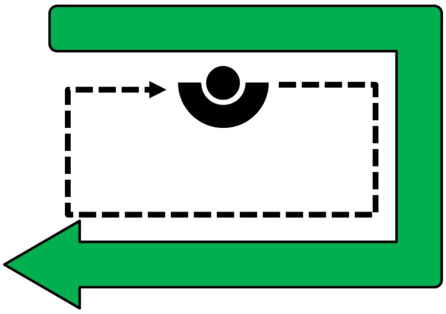

Figure 186: One Person in U Loop (Image Roser)

If one person does the entire job within a small manufacturing line (commonly called a work cell, although this term is used often also in different contexts), differences in workload content are not a problem. The worker just takes as much time as is needed, and the part is completed accordingly. A two-door vehicle without sunroof would be finished faster, whereas a four-door vehicle with a sunroof would be finished later. The differences in work content merely change the rate at which parts are completed.

23.4.2 Bucket Brigade

One such line layout is the bucket brigade (also known as *bump-back* or *bouncing line*). The worker moves along the line with his part, until he meets the next worker coming back without a part. He then gives the part to the next worker who moves the part forward along the process. The first worker then walks back until he meets the preceding worker to take the part or to get a new part at the beginning of the line. This way the capacity of the line is self-adjusting to the workload, and the workers utilization is easier to manage.

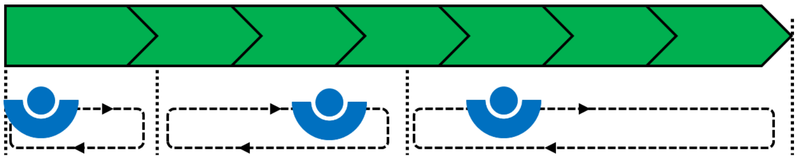

Figure 187: Animated Bucket Brigade Loops. The original image can be found at https://www.allaboutlean.com/bucket-brigade-1/ (Image Roser)

On the other hand, there are also some requirements for this to work. For example, all workers need to be trained at all stations, and it works best for shorter cycle times. For more details, check my series of two posts on the bucket brigade, with Part 1 here and Part 2 here.

23.4.3 Rabbit Chase

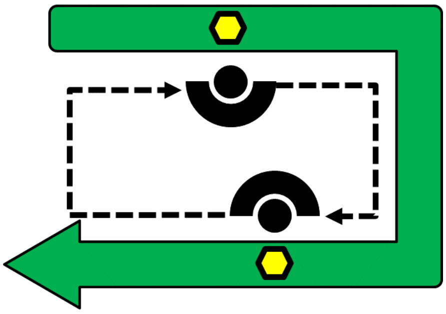

Figure 188: Animated Lean Rabbit Chase. Animated image can be found at https://www.allaboutlean.com/rabbit-chase/ (Image Roser)

Another option is the rabbit chase. A small number of workers, ideally in an U line, process an entire part from beginning to the end. Different work content and hence different speeds of the workers are buffered through the distance between the workers. Here, too, some restrictions apply similar to the bucket brigade. For more details, see my post on the rabbit chase.

Overall, the idea of adjusting capacity may sometimes work, but for many cases it is not so hot. There may also be the risk of "forgetting" a task that leads to an incomplete product. Especially for large and complex lines, a better way to manage different workloads is the sequencing of the products. I will talk more about this in the next post. Until then, stay tuned, **go out, and organize your industry!**

24 Mixed Model Sequencing – Basic Example Introduction

Christoph Roser, June 11, 2019, Original at
https://www.allaboutlean.com/mixed-model-sequencing-4/

Figure 189: Ancient Necklace of Wah (Image Metropolitan Museum of Art in public domain)

Your products may have different work content on your production line, which may make your line less efficient. One possible solution is Mixed Model Sequencing, a way to adjust the sequencing of your products to make the average work content stable. In previous posts I looked at the basics, at how to avoid the problem in the first place and how to play with capacity. However, especially for large complex lines (i.e., automotive), sequencing is often a suitable approach to manage different work contents.

24.1 Principle of Mixed Model Sequencing

Some of your products have a higher workload at a workstation (e.g., a four-door car at the door assembly). Hence, the other products have a lower workload (e.g., a two-door car at the door assembly). The basic idea is to alternate these products, and set up the critical station so that they can handle the average workload. For example, you alternate two-door and four-door cars, and the door assembly is designed to assemble three doors on average per cycle. I tried to illustrate this with the animation below, where two people on average assemble three doors each per cycle.

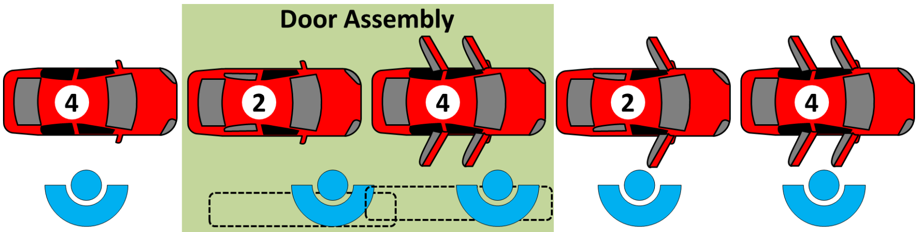

Figure 190: Animated image of a Double Door Assembly Station. The original image can be found at https://www.allaboutlean.com/mixed-model-sequencing-4/ (Image Roser)

In the following I will explain the fundamental approach to resolve a single unbalanced workstation. However, in reality you will have multiple workstations whose workload is affected by the product type. This makes everything more complex – but more about this later.

24.2 A Bit About Takt Times

Figure 191: Metronome (Image Vladimir Voronin with permission)

If you have a station with a mixed workload and hence a mixed cycle time, you try to set the average cycle time equal to the overall cycle time of the line. Or, to be more precise, you set the average takt time of the mixed workload station to the average takt time of the line. For more on the difference between cycle and takt time, see my post On the Different Ways to Measure Production Speed.

24.2.1 Two-Door and Four-Door Example

For example, your typical automotive assembly line completes a car every 60 seconds. Let's assume mounting four doors takes twice as long than two doors (even though this is technically not quite true; e.g., front doors often have more electronics).

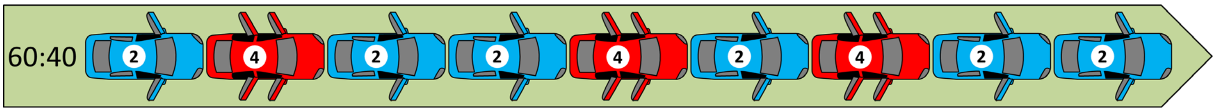

Figure 192: Two Products 60-40 Sequence (Image Roser)

The next information you need is the ratio of two doors to four doors (or, generally speaking, how many products of which type do you have). Here I will calculate for you the cycle times in dependence of the percentages of two-door and four-door vehicles. Assume the following variables:

- TT: Takt time of the entire line (the line takt)
- TT2: Takt time for the mounting of two doors
- TT4: Takt time for the mounting of four doors
- P2: Percentage of the cars that are two doors
- P4: Percentage of the cars that are four doors

In this case the following equation must hold true:

$$TT = TT_2 \cdot P_2 + TT_4 \cdot P_4$$

Where

$$P_2 + P_4 = 100\%$$

or

$$P_4 = 100\% - P_2$$

In our case we assumed that four doors take twice as long as two doors, hence also:

$$TT_4 = TT_2 \cdot 2$$

Using this additional information, we can simplify the first equation as follows:

$$TT = TT_2 \cdot P_2 + TT_2 \cdot 2 \cdot (100\% - P_2)$$

$$TT = TT_2(P_2 + 200\% - 2 \cdot P_2)$$

$$TT = TT_2(200\% - P_2)$$

$$TT_2 = \frac{TT}{200\% - P_2}$$

Depending on the share of two doors and four doors, you need to get the takt time to a certain value so the average takt time is consistent with the line takt. I have plotted this relation for you below (assuming four doors take twice as long as two doors, with a line takt of 60 seconds).

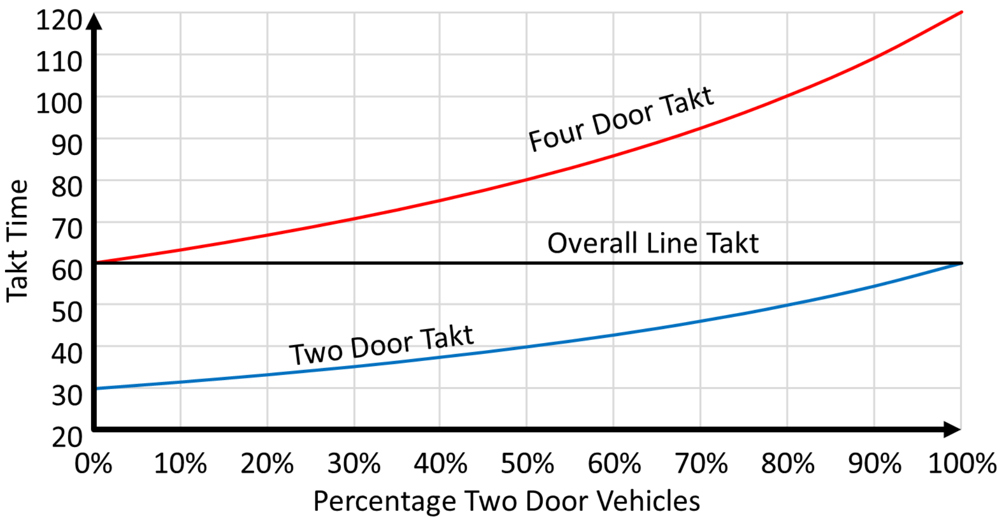

Figure 193: Takt Times Two and Four Door Example (Image Roser)

If you make only four-door vehicles (i.e., 0% two-door vehicles), then you need to get the takt time of four-door vehicles to 60 seconds. Similarly, if you make only two-door vehicles, you also need to get the takt time of two door vehicles to 60 seconds. If you make 50% of each, then you need a takt time of 40 seconds for the two doors and 80 seconds for the four doors for an average takt of 60 seconds.

24.2.2 General Equations

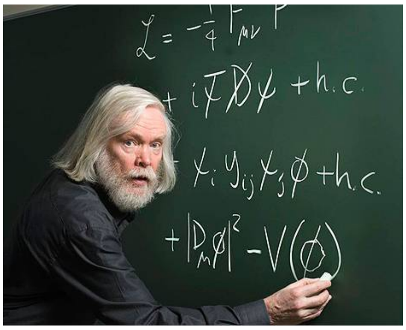

Figure 194: Scientist on Blackboard (Image Open Knowledge Foundation under the CC-BY 2.0 license)

The above example was for only two types of products (two and four doors), and we knew exactly the relation of the takt times between them (four doors take twice as long).

In reality, you may have many more different products. The relation of their takt times may also be not well defined. The takt time may even change between zero and a larger value (e.g., if you install a sunroof or if you do not). Generally speaking, if you have k different products with

- TT: Takt time of the entire line (the line takt)
- TT_n: Takt time for product n
- P_n: Percentage of the produced items that are of type n

Then the general formulas would be as follows:

$$TT = \sum_{n=1}^{k} TT_n \cdot P_n$$

$$\sum_{n=1}^{k} P_n = 100\%$$

The sum of the shares of the takt time has to be the line takt, and the sum of all shares has to be 100% (or 1 if you use mathematical notation). The challenge is to figure out this equation without knowing how much longer part type 25 takes in comparison to part type 33, and all in between. You probably also won't get hard numbers for the percentages, as they are only future predictions, and, hey, customers change their minds all the time, especially in the future.

For a mathematician this would be a disaster, having an equation where you don't know for sure what the values are. In lean we have this all the time, and hence you are probably familiar with the concept of guessing. If you don't have a number, just ask an expert for his opinion, or take the best guess you can get (it may even be your own guess).

Also keep in mind that you can change the takt times by getting better tools. Use jigs, rigs, and other gizmos to reduce a takt time that causes you problems. More on this later. You also don't need to be perfect with your math. The progress of timing a production line includes a lot of uncertainties, and just getting close to the targeted value is good enough. If the random chances turn against you and it takes too long, then you just have to put in an improvement effort to make it faster again.

So, this is the fourth post on product-dependent workload, and the first on sequencing. The sequencing part is actually a bit tricky, hence there will be a few more posts on this type of sequencing. Until then stay tuned, and **go out, and organize your industry!**

P.S. Many thanks to Mark Warren for his input.

25 Mixed Model Sequencing – Basic Example Workload and Buffering

Christoph Roser, June 18, 2019, Original at https://www.allaboutlean.com/mixed-model-sequencing-5/

Figure 195: Antique Frankish Beads (Image Metropolitan Museum of Art in public domain)

Adjusting the production sequence is a popular way of handling production lines with a product-dependent workload. This post is part of rather long series on Mixed Model Sequencing. In the last post I discussed the basics of sequencing and the calculation of the takt time. This post describes the basics of adjusting workload and buffering – but still for a simple case of only one imbalanced station. Subsequent posts will get more serious with multiple imbalances. But let's continue with our simple single imbalance example.

25.1 A Bit About Workload

Figure 196: Bottle Sizes (Image Andy Hay under the CC-BY 2.0 license)

In the last step in the previous post, you figured out what takt time you need for which product variant. Let's again take the example of a two-door and a four-door car, where mounting four doors takes twice as long as two doors, and the target takt time is 60 seconds per car. If you have a 50:50 mix, you need to get the four-door takt to 80 seconds and the two-door takt to 40 seconds to get an average of 60 seconds.

Hence you now would have to make sure that the workstation is able to do a two-door mount in 40 seconds and a four-door mount in 80 seconds on average, meaning the takt time has to be a bit faster, with the difference being the OEE. For this you can use all the usual tools to change takt times (and cycle times).

If the workload is too much, you could add another worker. Or you could split the work among different stations. Or you could add or optimize tools and machines to make the work faster.

If the workload is too little, you could remove a worker. Or you could add smaller tasks maybe from other stations. Overall, you have all the tools available to adjust your workload as you would have in normal line balancing. However, there are a few limitations specific to stations with mixed products that have different workloads or takt time:

- **You must have the same number of operators regardless of the product type**. Do not move operators around on very short notice depending on which product comes down the line. Additionally, they both should not have any significant idle time.

- **Tasks cannot move between stations depending on the product**. For a negative example, consider if you mount two doors you also attach a mirror, but if you mount four doors the mirror is done somewhere else. This is bad. This will create even more imbalance elsewhere, and you may end up with cars where the mirror was forgotten. I advise against such shifting around of tasks among stations during normal operation. All the tasks need to stay at this station. You may skip tasks (i.e., not mounting a sunroof if the car does not need a sunroof), but you cannot shift them elsewhere temporarily.
- **The average takt time must match the line takt**. As discussed before, the average takt time (or cycle time) across all product variants for this station needs to match the overall takt time of the entire line. On average, the workstation should neither be faster nor slower than the rest of the line.

Again, we don't care about machine waiting times. Hence, feel free to have one machine or tool exclusively for one product and another exclusively for another product if it helps you with your task.

25.2 A Bit About Buffering

So now you have created good takt times for your different products at this workstation. On average, your station takt matches the line takt. But never forget that the individual product types do not match the line takt. Some products may take longer (a four-door vehicle), others may be faster (a two-door vehicle), and again others may be just with the takt (a three-door vehicle?).

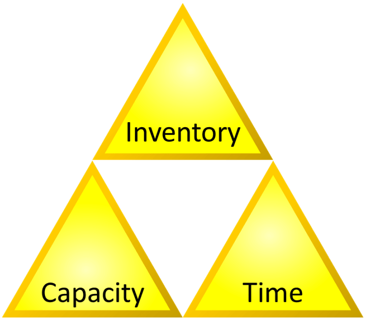

Figure 197: Triforce Inventory Capacity Time (Image Roser)

Hence, even if the average fits, the individual products won't. As a result you need to buffer these fluctuations. There are Three Fundamental Ways to Decouple Fluctuations: You create either an inventory buffer, a capacity buffer, or a time buffer. Most useful is the inventory buffer. But before I go into inventory buffers in more detail, let's get the other two out of the way.

25.2.1 Capacity Buffers

Figure 198: Male and Female Worker (Image Wayhome Studio with permission)

A capacity buffer would increase and decrease the available capacity depending on the demand. In my previous posts I already talked a lot about why it is a bad idea to have workers move in and out of the line on short notice, so let me just repeat that this is a bad idea. However, for smaller fluctuations we can use human nature to help us. Human workers can work at different speeds. On average, it should be a workload that keeps the worker busy without overloading, and that he can do every day for years without problems.

However, for short periods a worker can work a bit faster if he is able to rest a bit afterward. Hence, for product variants that have only a slightly longer takt time or a slightly smaller takt time, the worker may be able to work a bit faster during a busy part, and relax a bit during the not so busy part.

This may even happen automatically. If the worker sees that there is a larger bit of work coming, or that his colleagues may have to wait for him, he will speed up a bit. If there is little work and he would have to wait on other anyway, he will slow down a bit. A difference of ± 10% of the takt time is often not a problem for human workloads as long as the average is fine. Selling the idea to the unions, on the other hand, may be a bit more difficult.

25.2.2 Time Buffers

Figure 199: Time Spiral (Image mipan with permission)

Another option is time buffers, and these are usually the worst type. It means if the station does not get done in time, others will have to wait. If the station is faster than the others, the station has to wait. Please remember that we do this whole exercise solely to reduce waiting times in human workers and to avoid inefficiencies and waste.

This type of buffering does not need any particular planning, as it is the default way of decoupling fluctuations. If something goes wrong, no matter if it is with capacity (production) or inventory, someone has to wait (customer, workers, suppliers, etc.). Again: We want to avoid that!

25.2.3 Inventory Buffers

Figure 200: Warehouse worker checking the inventory (Image WavebreakMediaMicro with permission)

Often the best way of buffering is through inventory. You can create a buffer inventory before the workstation that fills up if the product variant takes longer, and empties again if the product variant is faster. Similarly, you can create a buffer inventory after the workstation that empties if the product variant takes longer, and fills up again if the product variant is faster. Psychologically, a buffer before often feels better for the workers at the workstation, but technically it makes little difference.

Here you also have to consider the type of line. If it is an unstructured timed line as shown below, you simply add the required buffer before.

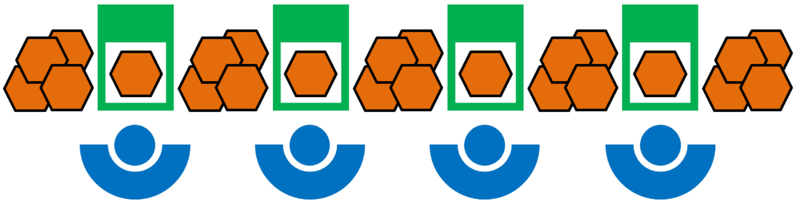

Figure 201: Animated Image of Untimed Line. The animated image can be found at https://www.allaboutlean.com/timing-of-flow-lines-1/ (Image Roser)

We already discussed that a pulsed line (shown below) is not well suited for this type of mixed model sequencing.

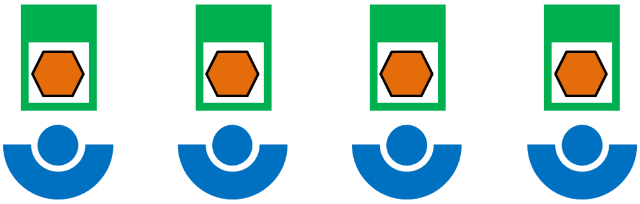

Figure 202: Animated image of Pulsed Line. The animated image can be found at https://www.allaboutlean.com/timing-of-flow-lines-1/ (Image Roser)

With a continuously moving line your buffer is simply a wider slot on the line. For example, if your line moves at 6 meters per minute, a station requiring 60 seconds would have a 6-meter-wide slot. If the station requires 80 seconds, the station would have a 8-meter-wide slot. If a part requires sometimes 40 seconds (two doors) and sometimes 80 seconds (four doors), you need a 8-meter-wide slot, even though you sometimes use only 4 meters if a two-door comes along. The difference is your buffer to manage the fluctuations in the work content.

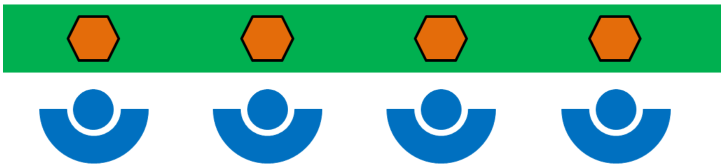

Figure 203: Animated image of Continuous Moving Line. The animated image can be found at https://www.allaboutlean.com/timing-of-flow-lines-2/ https://www.allaboutlean.com/timing-of-flow-lines-1/ (Image Roser)

This continuously moving assembly line has the advantage that you can add buffers that are not the equivalent to an integer number of parts.For the unstructured line, you either add a buffer capacity of one (or an integer number of parts), or you don't. You can't add the equivalent of half a part of buffer for an unstructured line. However, for a continuously moving line this is possible. We just did that. By making the slot 8 meters wide, we added the equivalent of one-third of a part as buffer capacity. Neat, isn't it?

I will continue this in my next post. This series on product dependent workload turns out to be rather long, so stay with me. Until then, **go out and organize your industry!**

P.S. Many thanks to Mark Warren for his input.

26 Mixed Model Sequencing – Basic Example Sequencing

Christoph Roser, June 25, 2019, Original at https://www.allaboutlean.com/mixed-model-sequencing-6/

Figure 204: Beads from Ancient Egypt (Image Metropolitan Museum of Art in public domain)

Mixed Model Sequencing is one way to handle products that have different workloads on a production line. In this post I continue to describe the basics if there is only one station where products vary in workload. This is part of a (rather long) series of posts on product-dependent workload. In my next posts I will go into situations where multiple stations are affected by product-dependent workload.

26.1 A Bit About Sequencing

Let's have a look at the sequence. Keep in mind we are still at the simple example of only one workstation with a part-dependent workload difference to explain the concept. It will get quite a bit more difficult later on.

The idea is to create a sequence with the product types spread as evenly as possible. If you have two products with a 50-50 mix (two door and four door), then you would alternate the two product types. If the mix is 75-25, then three of one type would be followed by one of another type. Below are a few options of sequencing two and four door cars depending on their share.

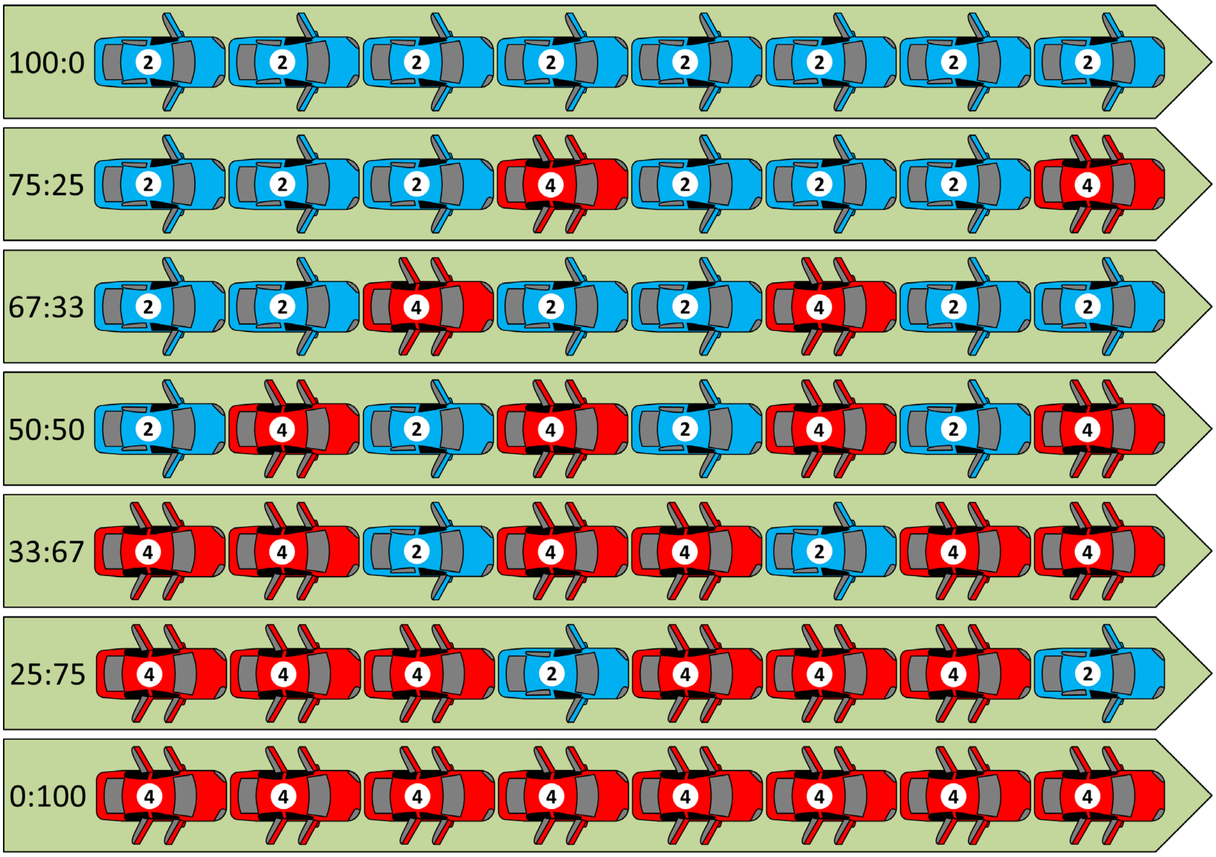

Figure 205: Two Product Sequence Options (Image Roser)

However, this may not always be that easy. The above examples just happen to be a ratio that results in nice sequences. A 60:40 ratio, however, would be a sequence where you sometimes have two 2-doors and sometimes one 2-door before having a 4-door vehicle as shown below.

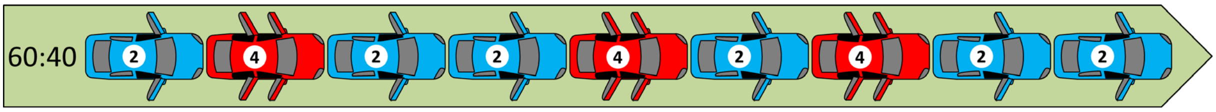

Figure 206: Two Products 60-40 Sequence (Image Roser)

This becomes even trickier if you have more than two product types. I have explained the basic approach already in my post [Introduction to One-Piece Flow Leveling](). However, for normal one piece flow leveling you would start with the largest quantity. Here you **start with the product that has the largest takt time difference to the target takt time**. For example a product that is 30 seconds faster (or slower) than the takt time takes priority over a product that is only 20 seconds faster (or slower) than the takt time.

You start with the first product, and divide it as evenly as possible on the slots. Mathematically you take the number of products of that type divided by the total number of products. In other words, if every 3.24th product is of this type, then every 3.24th slot is for that product, and you round to the nearest slot. Then you take the next part type, and do the same – except that if the slot is occupied, you take the nearest available slot. You continue with all product types in order of priority (takt time difference).

If your sequence does not distribute the product types evenly, you may need more buffer. Let's take again a 50:50 mix of two-door and four-door cars. But this time for demonstration, let's mess up the sequence with a lot size of two as shown below.

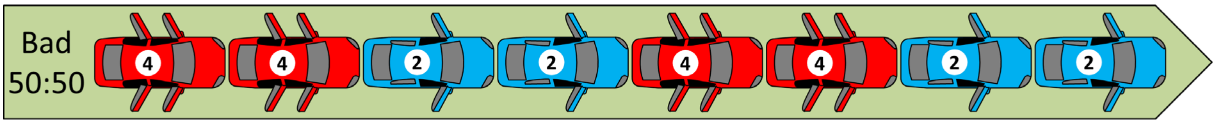

Figure 207: Two Products with a bad 50-50 Sequence (Image Roser)

Now instead of being 20 seconds over and 20 seconds under the average, you will be 40 seconds over (after two 4-door cars), followed by 40 seconds under (after two 2-door cars). Your fluctuation just got bigger, and you need more buffer to handle these larger fluctuations. Yet one of the main goals of this type of sequencing (besides reducing waiting time) is to reduce buffer space. Remember, the space directly at the assembly line is your most valuable space, and the more space you use simply to buffer parts the more waste you create by walking, bad information flow, longer transport, and so on. Hence, it is important to **spread your takt time fluctuations as evenly across the sequence as possible!**

26.2 Unfortunately the World Is More Complex

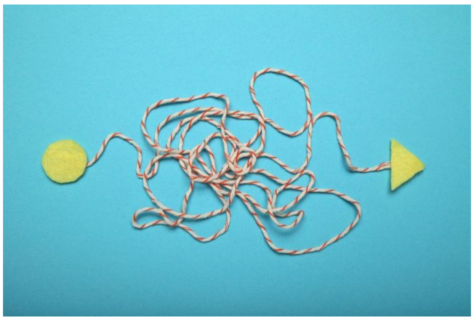

Figure 208: Complex Problems (Image Andrii Zastrozhnov with permission)

The above examples already look complex, with multiple products in varying quantities. However, having only one workstation with products with different workload is still easy. It becomes really tricky if multiple stations have workloads that depend on the product type.

This is usually the norm in the automotive industry. I already mentioned the number of doors and the sunroof. A car can have regular seats, or seats with heating, or with massage functions, or electric adjustments, all of which may have different connectors. There may be an air condition, or maybe not. The customer may have spent extra on a chrome package, illuminated door handles, color applications, door and trunk protectors, or many more options. For example, I read that BMW on average makes two identical cars per variant and year. Only once per year the very same car may come down the line.

If you now optimize the sequence for one workstation, it may be a very bad sequence for the other workstation. If you optimize for the other workstation, it may in turn be a bad sequence for the first workstation.

This may be complicated even more if the station also would have to do rework to fix issues that where done wrongly beforehand. I know, it should not have happened in the first place, but sometimes it does.

The overall goal is again to reduce idle times of workers while increasing inventory as little as possible.

26.3 How Many Option Are There?

Figure 209: Businessman and many doors (Image alphaspirit with permission)

Just for kicks I looked into the number of possible ways to sequence a production. This depends on the number of product variants and the total number produced. Generally speaking, for n variants and m number of products in total, the possible number of different sequences S are

$$S = m^{n-1}$$

Unfortunately, this number will go up rapidly. If you make 1,000 parts in one day with 30 different product types, the number of possible sequences 10^{87} already exceeds the number of atoms in the universe (10^{78} to 10^{82}, there is a bit of uncertainty here). If you add the restriction of making only a certain quantity of each type, the number of sequences will go down a bit, but it will still be too much to have a computer compare all possible sequences. Hence, a good sequence is often more of an art than an algorithm. Due to the number of possible options, you will probably not achieve a perfect sequence, but hopefully one that is good enough. In the following posts I will describe you an approach that helps you to find a good sequence. Until then, **go out, consider where you have product dependent workloads, and organize your Industry!**

P.S. Many thanks to Mark Warren for his input.

27 Mixed Model Sequencing – Complex Example Introduction

Christoph Roser, July 02, 2019, Original at https://www.allaboutlean.com/mixed-model-sequencing-7/

Figure 210: Animal shaped Beads on a String (Image Metropolitan Museum of Art in public domain)

In this rather long series of posts on Mixed Model Sequencing to handle product-dependent workload, we are finally getting to the most complex part: Sequencing of products when multiple stations have product-dependent workloads. Read on:

27.1 How To Address Multiple Workload Imbalances

In the following I will describe an approach on how to create a sequence with the goal to reduce the idle time of your workers while keeping the buffer inventories small. This is all with production where the workload depends on the type of product.

But first a small warning: There is probably not a perfect solution. We hope to find a solution that is not too bad. This is a somewhat messy process, and you will quite likely have to do quite a few iterations until you are satisfied (or running out of time).

Also, one of the requirements for this to work is a lot size one. If your system can work only in larger batches, then it is usually not possible to have a batch with a lot of work at one station to be canceled out by another batch with little work at this station. If you have a lot of four-door vehicles, excess work will accumulate. While this overload may be reduced by a corresponding underload (e.g., a two-door vehicle), you would need a large batch before the workstation to achieve this. This large batch could cause more waste than the over/underload.

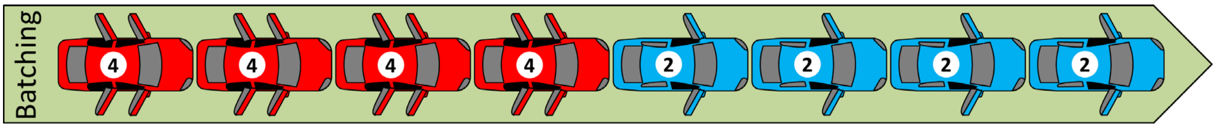

Figure 211: Batching 2 Door 4 Door Vehicles (Image Roser)

If you are worried about this being difficult, don't worry, you will get lots of practice by redoing the sequence every few weeks due to a change in customer demand, new machines, new products, and so on. In a year you will be good at this!

27.2 Determine Customer Takt, Line Takt, Cycle Time

Figure 212: Metronome (Image Vladimir Voronin with permission)

Before even thinking about the sequence, we would need to know the customer takt. What is the average time between one customer's demands? This customer takt defines the overall requirement on your production line. For more on customer takt, including how to calculate it, see my post How to determine Takt Times.

Based on the customer takt, you get the line takt (which is pretty much the same thing, ideally). The line takt multiplied with the OEE (or, at this point an estimate of the OEE) gives the target cycle time. I have described the whole process in my series on Line Balancing.

At the end, you need to know how fast your production line has to be on average, either as a line takt (including losses) or as a target cycle time (excluding losses).

27.3 Get All the Work Contents

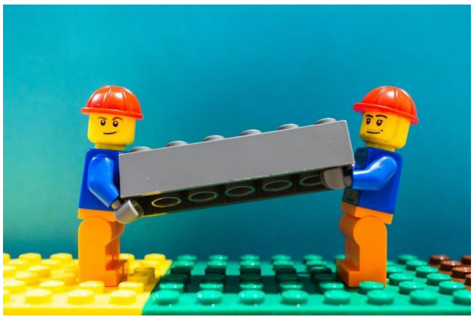

Figure 213: Lego Workers (Image Mirco Vacca with permission)

Now we need to get the details on the work content of the different products. This will be time consuming.

You would need to look at all the different tasks that need to be done for each product variant, and how long they take. If product A gets something attached with 4 screws, how long will it take? If product B has only three screws, how long will this take?

If you have no line yet to look at, a system of predetermined times like MTM may help you to get an estimate of the durations. If you do have a line, you may consider measuring it directly at the line – but make sure that the people and the union are informed about this. Depending on the legal situation, they do have veto rights on this in some countries.

If you measure by hand, also keep an eye out for tasks where the workload fluctuates a lot. This is an indication for a bad standard or insufficient training. Getting these fluctuating durations for the same task at the same part under control will help you later on.

Also keep an eye out if the stations have to do rework. This would also have to be included in the data collection. It also helps if you list the tasks in the sequence that they have to be done. While this sequence is not final, it makes things easier later on compared to a randomly ordered list of tasks.

Luckily, a lot of the tasks and their durations are usually identical. But there will be enough tasks left that are different, and these will make life tricky for you.

Below is a (fictional) example of installing car seats. Some steps are the same across all models. In the example below step 4, plugging in the sensor takes the same duration across all product types. Other steps are complementary, as for example steps 1 and 2. You attach a seat either with 4 screws or with 6 screws, but you do either one or the other. You do not attach two driver seats, nor do you attach zero driver seats. Similarly, you add only one floor mat, although the durations may depend on the model as seen in steps 7, 8, and 9.

	A	B	C	D	E	F	G	H	I	J	K
1		Quantity	2355	2766	121	5400	5500	790	670	130	
2		Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%	
3		Product	2433	2434	2435	2436	2437	2438	2439	2440	⋮
4	Task	Description									
5	1	Attach Driver Seat with X screws	28	28	28	28	28	35	35	35	
6	3	Plug in Seat heating	5	5				5		5	
7	4	Plug in Sensor	5	5	5	5	5	5	5	5	
8	5	Plug in Massage			7			7		7	
9	6	Plug In Seat Adjustment	6	6	6					6	
10	7	Add floor mat	15	15	15	15	15	15	15	15	
11	10	...									

Figure 214: Excel Sample Task List (Image Roser)

Other tasks are independent of each other. The seat may have seat heating, electric seat adjustments, and even a massage function – and there may be any combination that depends only on the product type. You could have all of them (as with product type 2440), or none (as with product type 2436, 2437, and 2439).

Overall, getting this data will take time. Some of the data may not even be available, or will be subject to a lot of guesswork.

27.4 Reduce or Eliminate Overloads, Spreads, and Fluctuations When Possible

Now would be a good time to see if you can eliminate some of these fluctuations. Can you make tools that make the tasks more similar? Can you change the training, improve the standards? In the example from above, the floor mat installation has been standardized and now takes only 15 seconds regardless of the type of the mat (and the type of the vehicle). The more differences you can get away, the easier it will be later on.

	A	B	C	D	E	F	G	H	I	J	K
1		Product	2433	2434	2435	2436	2437	2438	2439	2440	⋮
2	Task	Description									
3	1	Attach Driver Seat with X screws	28	28	28	28	28	35	35	35	
4	3	Plug in Seat heating	5	5				5		5	
5	4	Plug in Sensor	5	5	5	5	5	5	5	5	
6	5	Plug in Massage			7			7		7	
7	6	Plug In Seat Adjustment	6	6	6					6	
8	7	Add floor mat	15	15	15	15	15	15	15	15	
9	10	...									

Figure 215: Excel Sample Task List Condensed (Image Roser)

It may also be a good time to condense complementary tasks into one line. In the example above I merged the seat attachment into one line. The times are still different with 4 and 6 screws, but having only one line makes it slightly easier later on.

Figure 216: 3, 4, and 5 lug nuts (Image Bindydad123 under the CC-BY-SA 4.0 license and PXHere in public domain)

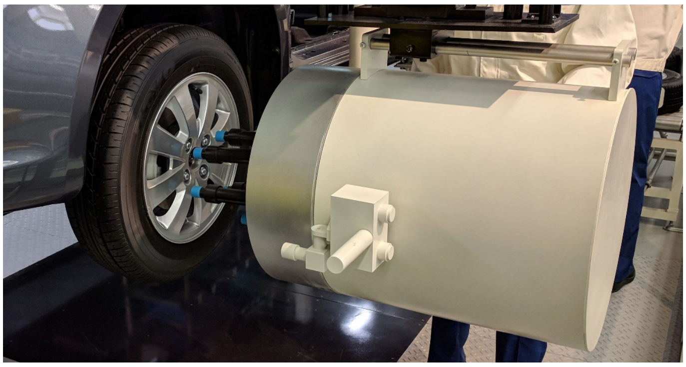

Figure 217: Suzuki Tire Assembly (Image Roser)

In one of the previous posts of this series, I had the example of a wheel with different numbers of lug nuts, which could be avoided using a specialized tool that can attach various lug nuts in the same time.

In my next post we will look at the average work durations for the different tasks, how to determine them, and how to combine them into workstation-sized bits. In the meantime, **go out and organize your industry!**

P.S. Many thanks to Mark Warren for his input.

PSS: The Sequencing Example Excel File for posts 7 to 11 with the complex example is available at https://www.allaboutlean.com/wp-content/uploads/2019/05/Sequencing-Example-Excel-File.xlsx. Please note that this is not a tool, but merely some of my calculations for your information.

28 Mixed Model Sequencing – Complex Example Data Basis

Christoph Roser, July 09, 2019, Original at https://www.allaboutlean.com/mixed-model-sequencing-8/

Figure 218: Ancient Iranian Necklace (Image Metropolitan Museum of Art in public domain)

Having a product mix with different workloads at different stations is challenging. Hence this is getting to be a pretty long series of blog posts on Mixed Model Sequencing. Let’s continue:

28.1 Get the Quantities to Be Produced

The next step is a smaller one. You need to figure out the quantities that have to be produced of the different products. The time frame in question is the time frame for which the sequence should hold. Since you don’t known the future, this will be an estimate. As such, it will be flawed. Such is life. Just take the best guess you can and move forward with that number.

For the Excel example from the last post I added two rows that have the (expected) quantity, for which I also calculated the percentage thereof. Our high runner is model 2437, and our rare exotic models are 2435 and 2440.

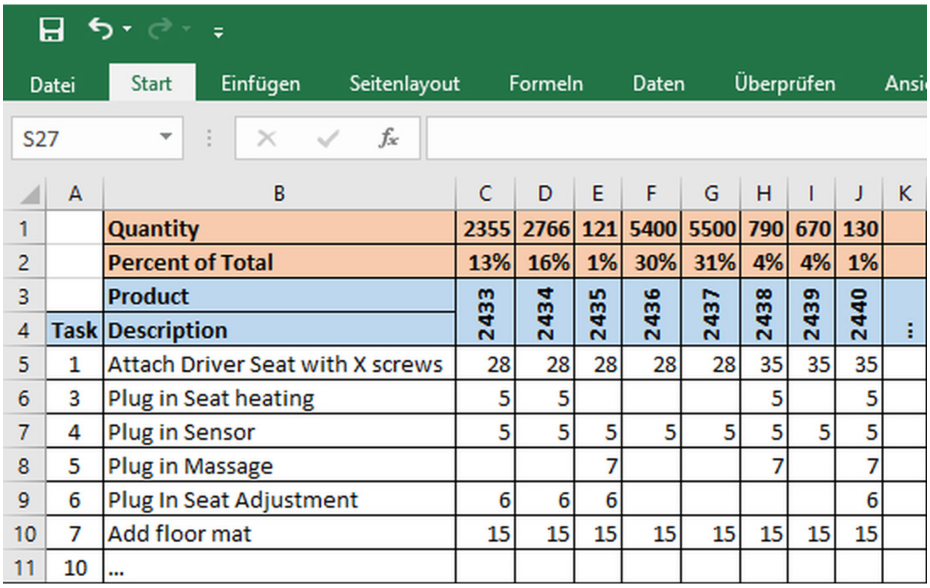

	A	B	C	D	E	F	G	H	I	J	K
1		**Quantity**	**2355**	**2766**	**121**	**5400**	**5500**	**790**	**670**	**130**	
2		**Percent of Total**	**13%**	**16%**	**1%**	**30%**	**31%**	**4%**	**4%**	**1%**	
3		**Product**	**2433**	**2434**	**2435**	**2436**	**2437**	**2438**	**2439**	**2440**	
4	**Task**	**Description**									⋮
5	1	Attach Driver Seat with X screws	28	28	28	28	28	35	35	35	
6	3	Plug in Seat heating	5	5				5		5	
7	4	Plug in Sensor	5	5	5	5	5	5	5	5	
8	5	Plug in Massage			7			7		7	
9	6	Plug In Seat Adjustment	6	6	6					6	
10	7	Add floor mat	15	15	15	15	15	15	15	15	
11	10	...									

Figure 219: Excel Sample Task List Condensed Quantity (Image Roser)

Please note that this quantity is the quantity you want to produce in the period you are sequencing. This depends on the cycle time. If it is a fast cycle time (e.g. less than three minutes), you may sequence only for one shift, and create a new sequence again for the next shift. If you have longer cycle times but below an hour you may sequence a day or a week. If your cycle time exceeds hours you may sequence a month. If your cycle time is multiple hours … then you probably don’t need to sequence at all but adjust the capacity by adding or removing workers for one shift to manage the excess workload or excess idle times.

28.2 Determine Average Contribution to Cycle Time

The next step is simple math. We would need to figure out how long each step takes on average. The table above gave us the time for each individual step and product type. We now multiply these times with the percentage share of the product type, and sum this up across all products. This gives us the weighted arithmetic mean of the task time. The result could look like this example below.

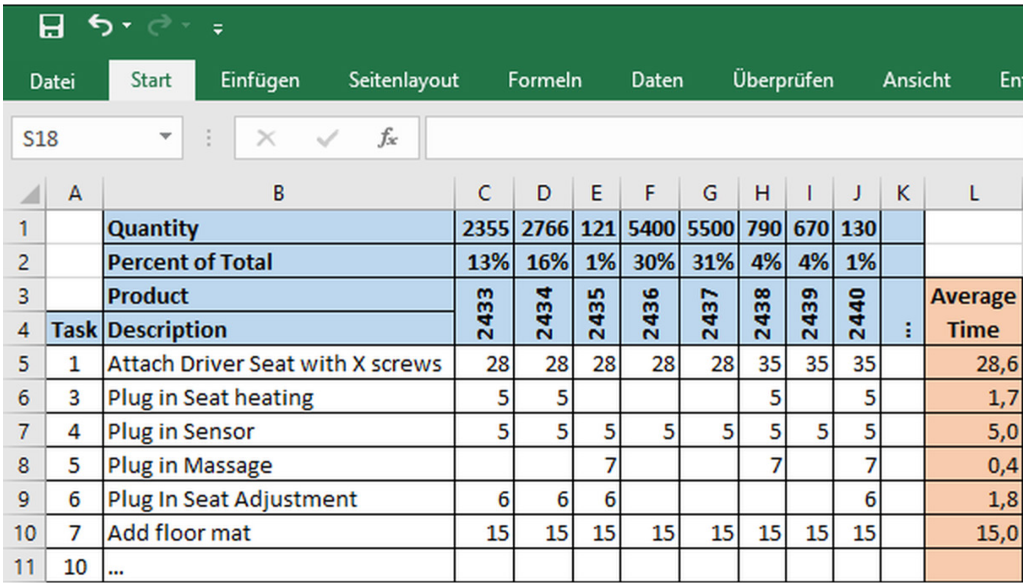

	A	B	C	D	E	F	G	H	I	J	K	L
1		Quantity	2355	2766	121	5400	5500	790	670	130		
2		Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%		
3		Product	2433	2434	2435	2436	2437	2438	2439	2440	...	Average
4	Task	Description										Time
5	1	Attach Driver Seat with X screws	28	28	28	28	28	35	35	35		28,6
6	3	Plug in Seat heating	5	5				5		5		1,7
7	4	Plug in Sensor	5	5	5	5	5	5	5	5		5,0
8	5	Plug in Massage			7			7		7		0,4
9	6	Plug In Seat Adjustment	6	6	6					6		1,8
10	7	Add floor mat	15	15	15	15	15	15	15	15		15,0
11	10	...										

Figure 220: Excel Sample Task List with Times (Image Roser)

You could also calculate the maximum, minimum, spread, and standard deviation of the average task time. However, at this time these details are not yet that useful.

28.3 Group into Workstation-Sized Workloads

Now comes a more challenging task. You need to group these tasks into groups with the workload for different stations. For now we ignore the times for the individual products and consider only the weighted average time across all products. The total average cycle time across the tasks for one station should be close to the target cycle time (or the average takt time should be close to the target line takt, whichever way you prefer it). If you are a tiny bit above or below the target cycle time, don't sweat it.

Also keep in mind that you usually can change the sequence somewhat. If I plug in the seat sensor first or the heating first probably does not matter, but it may help you to get a nice average workload for a station. Just make sure the sequence is possible, and that this sequence does not create other problems (i.e., if the worker has to walk from one end of the car to the other end and back).

At the end, all tasks have to be assigned to a workstation, and all workstations have cycle times matching the target cycle time. (If at the end there is work for half a workstation left, try not to spread this time across all workstations, but instead pool it into one workstation. This makes it easier with some improvements to eliminate the workstation altogether for a more efficient line.)

I have written a series of blog posts on this too. Line Balancing Part 5 – Balancing Using Paper and Line Balancing Part 6 – Tips and Tricks for Balancing are the two parts that are most relevant here.

Below would be the example of the seat mounting workstation. The cycle time ranges from 48 seconds (for product numbers 2436 and 2437) to 73 seconds (for product number 2440), with a weighted average of 52.6 seconds. Emphasis on weighted average, as this is not just the average of the individual product times but has to be combined with the percentage of this particular product.

Task	Description	2433	2434	2435	2436	2437	2438	2439	2440	⋮	Average Time
	Quantity	2355	2766	121	5400	5500	790	670	130		
	Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%		
1	Attach Driver Seat with X screws	28	28	28	28	28	35	35	35		28,6
3	Plug in Seat heating	5	5				5		5		1,7
4	Plug in Sensor	5	5	5	5	5	5	5	5		5,0
5	Plug in Massage			7			7		7		0,4
6	Plug In Seat Adjustment	6	6	6					6		1,8
7	Add floor mat	15	15	15	15	15	15	15	15		15,0
	Seat Mounting Total	59	59	61	48	48	67	55	73		52,6

Figure 221: Excel Sample Task List for one Station (Image Roser)

28.4 Determine Spread of Workload

Finally you end up with a list of all workstations, how long each product variant takes at each station, and how long it will be in average based on the quantities.

Now we look at the spread of the workloads. You already have the average time for each station. Next we calculate the longest time at each station (i.e., how long does the model with the longest time take at this station). For the dashboard mounting station below, this would be products 2435 and 2440 with 70 seconds cycle time each.

Task	Workstation	2433	2434	2435	2436	2437	2438	2439	2440	⋮	Average Time	Min Time	Max Time	Spread	Standard Deviation
	Quantity	2355	2766	121	5400	5500	790	670	130						
	Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%						
1	Dashboard	41	61	70	40	66	40	40	70		51,9	40	70	30	13,5
3	Sunroof Mounting	0	111	0	0	112	0	0	0		52,1	0	112	112	48,3
4	Rear Seats	50	50	50	50	50	70	66	72		51,7	50	72	22	9,5
5	Passenger Seat & Glove Box	50	49	47	55	53	59	50	50		52,7	47	59	12	3,6
6	Driver Seat	59	59	61	48	48	67	55	73		52,6	48	73	25	8,1
7	Door Mounting	75	30	76	30	77	34	30	75		51,4	30	77	47	22,4

Figure 222: Excel Sample Station List (Image Roser)

We do the same with the shortest cycle time. For the dashboard station this would be product 2436 and 2439 with 40 seconds each. We also calculate the spread (i.e., the difference between the cycle time of the fastest product and the cycle time of the slowest product). For the dashboard this would be a spread of 30 seconds between the fastest and the slowest product variant. The larger the spread, the more challenging it will be to fit into the average cycle time.

Another value I find useful is the total and average work per product type rather than per workstation. This is shown in the example below. Here product 2437 sticks out with a total workload of 406 seconds total or 67.7 in average, almost twice that of the station with the smallest workload for product 2436 with 37.2 seconds.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1		Quantity	2355	2766	121	5400	5500	790	670	130						
2		Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%						
3		Product	2433	2434	2435	2436	2437	2438	2439	2440		Average	Min	Max		Standard
4	Task	Workstation									:	Time	Time	Time	Spread	Deviation
5	1	Dashboard	41	61	70	40	66	40	40	70		51,9	40	70	30	13,5
6	3	Sunroof Mounting	0	111	0	0	112	0	0	0		52,1	0	112	112	48,3
7	4	Rear Seats	50	50	50	50	50	70	66	72		51,7	50	72	22	9,5
8	5	Passenger Seat & Glove Box	50	49	47	55	53	59	50	50		52,7	47	59	12	3,6
9	6	Driver Seat	59	59	61	48	48	67	55	73		52,6	48	73	25	8,1
10	7	Door Mounting	75	30	76	30	77	34	30	75		51,4	30	77	47	22,4
11		Total	275	360	304	223	406	270	241	340		312,2	223	406	183	
12		Average	45,8	60	51	37,2	67,7	45	40	57		52,04	37,2	67,7	30,5	

Figure 223: Excel Sample Station List Product Total (Image Roser)

Yet another way to analyze the problem is to look a the largest and smallest cycle times for each product. In the example below, I marked the largest and smallest cycle time for each product, with product 2437 and 2434 having the largest cycle times at the sunroof mounting. All other products have the smallest cycle times at the sunroof mounting of zero.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1		Quantity	2355	2766	121	5400	5500	790	670	130						
2		Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%						
3		Product	2433	2434	2435	2436	2437	2438	2439	2440		Average	Min	Max		Standard
4	Task	Workstation									:	Time	Time	Time	Spread	Deviation
5	1	Dashboard	41	61	70	40	66	40	40	70		51,9	40	70	30	13,5
6	3	Sunroof Mounting	0	111	0	0	112	0	0	0		52,1	0	112	112	48,3
7	4	Rear Seats	50	50	50	50	50	70	66	72		51,7	50	72	22	9,5
8	5	Passenger Seat & Glove Box	50	49	47	55	53	59	50	50		52,7	47	59	12	3,6
9	6	Driver Seat	59	59	61	48	48	67	55	73		52,6	48	73	25	8,1
10	7	Door Mounting	75	30	76	30	77	34	30	75		51,4	30	77	47	22,4
11		Total	275	360	304	223	406	270	241	340		312,2	223	406	183	
12		Average	45,8	60	51	37,2	67,7	45	40	57		52,04	37,2	67,7	30,5	
13		Min	0	30	0	0	48	0	0	0	0					
14		Max	75	111	76	55	112	70	66	75	0					

Figure 224: Excel Sample Station MinMax Cycle Time (Image Roser)

28.5 Focus on Your Biggest Rocks first

Figure 225: Rock Stacking Iceland (Image Roser)

Now we have our data together. But before we start sequencing, we have to consider which parts to sequence first. For a good sequencing, we have to start "with the big rocks." But what do we mean by *big*? There are many different aspects that come into play:

- **Longest Cycle Times**: The longer the cycle time, the bigger the problem to fit it into our sequence. Hence, products that have an excessively long cycle time are probably among the first to be sequenced. In our example, products 2437 and 2434 with a sunroof mounting cycle time of 112 and 111 seconds are the largest rocks we have. This is followed by different door mountings, the rear seats, and the dashboard.
- **Shortest Cycle Times**: These are also relevant, as they may cause excess idle time. They can also be used to counteract the longest cycle times.
- **Largest Spread**: The difference between the largest cycle time and the shortest cycle time for a station is also relevant, although this is often similar to the parts with the longest and shortest cycle times.
- **Largest Average Work per Product**: Another way to view this is by looking at the largest average work per product.
- **Smallest Average Work per Product**: The smallest average work per product can be relevant, similar to the shortest cycle times.
- **Largest Cycle Time per Product**: And yet another item to consider is the largest cycle time per product
- **Smallest Cycle time per Product**: Similarly, the smallest cycle time per product may also be relevant to offset the largest ones.
- **Largest Quantities Produced**: Starting the sequencing with the largest quantity products will help to spread most of the material evenly. If there is no product-dependent workload, then this is the standard approach for sequencing.

Now you have a lot of options for where to start. Unfortunately, a lot of these may be giving you different priorities. A product that has a longest cycle time may also have the shortest at another station, or a perfectly average one at the third. As a small suggestion, you probably can ignore anything that falls within 10% or maybe even 20% of the mean. These will sort itself out by people automatically working a tick faster or slower depending on what is needed.

In the next post I will show you how to make a sequence out of this mess of somewhat conflicting priorities. However, this is not foolproof, and you will need lots of iterations to get to a good solution. Until then, **go out, ponder your largest fluctuation, and organize your industry!**

P.S. Many thanks to Mark Warren for his input.

PSS: The Sequencing Example Excel File for posts 7 to 11 with the complex example is available at https://www.allaboutlean.com/wp-content/uploads/2019/05/Sequencing-Example-Excel-File.xlsx. Please note that this is not a tool, but merely some of my calculations for your information.

29 Mixed Model Sequencing – Complex Example Sequencing 1

Christoph Roser, July 16, 2019, Original at [*https://www.allaboutlean.com/mixed-model-sequencing-9/*](https://www.allaboutlean.com/mixed-model-sequencing-9/)

Figure 226: Beads from Chlorite Necklace (Image Metropolitan Museum of Art in public domain)

Sequencing products due to different workloads of different products at different workstations is tricky. This is now the sixth post on Mixed Model Sequencing, and we finally start with our sequence! Wohooo!

29.1 Sequence the First Product

The first product to be sequenced is usually easy; it is the biggest rock you have (i.e., the part that has the largest cycle time at any station). In the example we used so far, this would be product number 2437, since it has the largest cycle time for the mounting of the sunroof of 112 seconds. This is our biggest rock. I have marked this in red in the table below.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1		Quantity	2355	2766	121	5400	5500	790	670	130						
2		Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%						
3		Product	2433	2434	2435	2436	2437	2438	2439	2440	⋮	Average	Min	Max		Standard
4	Task	Workstation										Time	Time	Time	Spread	Deviation
5	1	Dashboard	41	61	70	40	66	40	40	70		51,9	40	70	30	13,5
6	3	Sunroof Mounting	0	111	0	0	112	0	0	0		52,1	0	112	112	48,3
7	4	Rear Seats	50	50	50	50	50	70	66	72		51,7	50	72	22	9,5
8	5	Passenger Seat & Glove Box	50	49	47	55	53	59	50	50		52,7	47	59	12	3,6
9	6	Driver Seat	59	59	61	48	48	67	55	73		52,6	48	73	25	8,1
10	7	Door Mounting	75	30	76	30	77	34	30	75		51,4	30	77	47	22,4
11		Total	275	360	304	223	406	270	241	340		312,2	223	406	183	
12		Average	45,8	60	51	37,2	67,7	45	40	57		52,04	37,2	67,7	30,5	
13		Min	0	30	0	0	48	0	0	0	0					
14		Max	75	111	76	55	112	70	66	75	0					

Figure 227: Excel Sample Station Biggest Rock (Image Roser)

The sequence for this part 2437 is easy. We want to produce 5,500 of this type out of 17,732 parts in total. Hence our sequencing interval S for our product k with a quantity Q_k out of n products in total is:

$$S_k = \frac{\sum_{m=1}^{n} Q_m}{Q_k}$$

or for our case:

$$S_{2437} = \frac{17732}{5500} = 3.22$$

Hence, every 3.22 items in the sequence should be this product 2437. If we start at position 1 with this type, then the next one should be at 1+3.22 = 4.22. The next one afterward should be at 4.22+3.22 = 7.45, then at 10.67, at 13.90, and so on. Since this is the first product, all slots are still available, and we can simply round to the nearest integer, giving us this product 2437 at positions 1, 4, 7, 11, 14, etc. in our production sequence. The first 30 slots are visualized below.

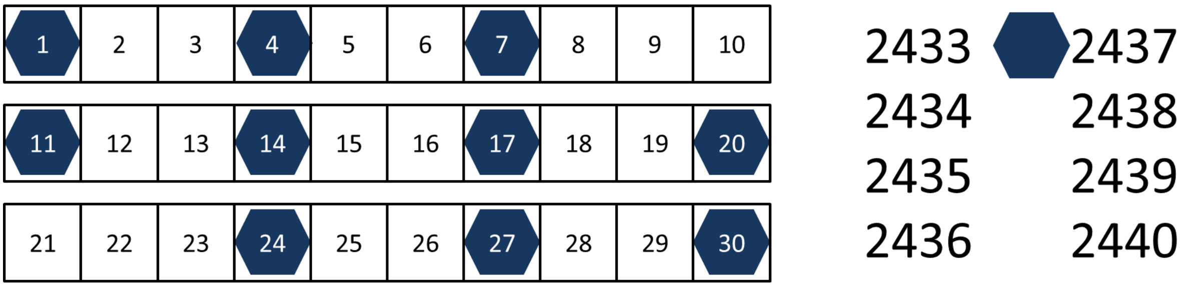

Figure 228: Sequencing Example Product 1 Sequenced (Image Roser)

29.2 Sequence the Second Product

Now comes the ambiguous part: Which product should we sequence next? Should we take product 2434 because it has the second-biggest cycle time, also at the sunroof mounting? Or should we take some of the products that have no sunroof to offset the excessive sunroof time of 2437? If so, it could be product 2436, which also has a similar quantity to product 2437, which we just sequenced (grayed out in the table below). Or should we go for something completely different by sequencing 2435, because it has the largest cycle time for the door mounting?

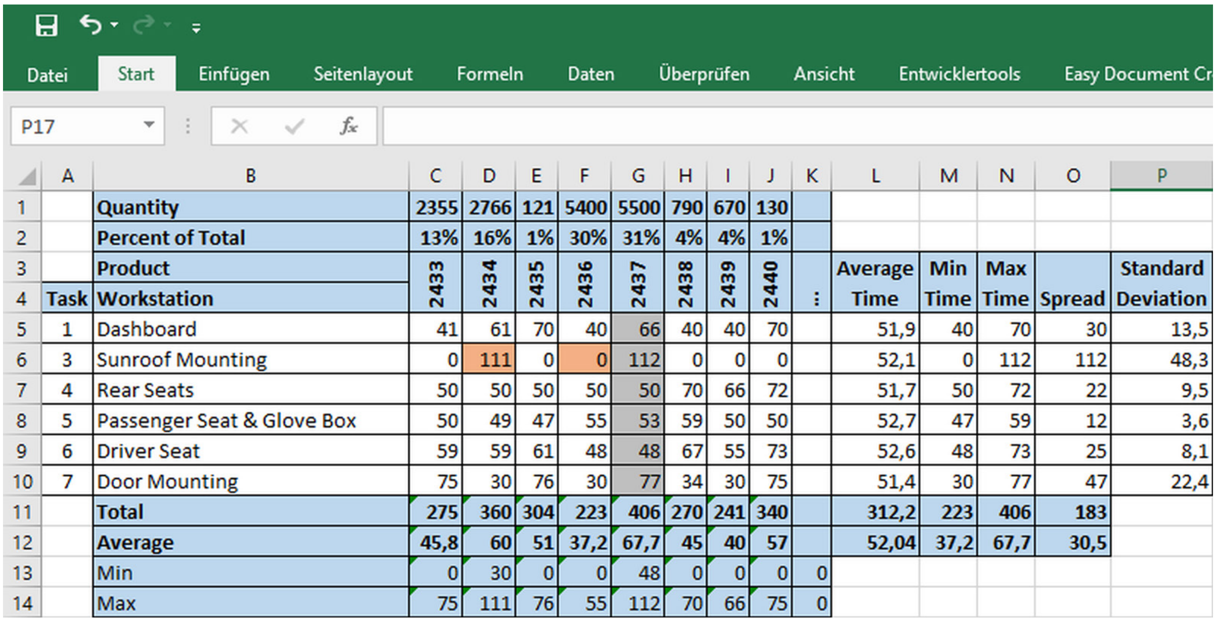

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1		Quantity	2355	2766	121	5400	5500	790	670	130						
2		Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%						
3		Product	2433	2434	2435	2436	2437	2438	2439	2440		Average	Min	Max		Standard
4	Task	Workstation									:	Time	Time	Time	Spread	Deviation
5	1	Dashboard	41	61	70	40	66	40	40	70		51,9	40	70	30	13,5
6	3	Sunroof Mounting	0	111	0	0	112	0	0	0		52,1	0	112	112	48,3
7	4	Rear Seats	50	50	50	50	50	70	66	72		51,7	50	72	22	9,5
8	5	Passenger Seat & Glove Box	50	49	47	55	53	59	50	50		52,7	47	59	12	3,6
9	6	Driver Seat	59	59	61	48	48	67	55	73		52,6	48	73	25	8,1
10	7	Door Mounting	75	30	76	30	77	34	30	75		51,4	30	77	47	22,4
11		Total	275	360	304	223	406	270	241	340		312,2	223	406	183	
12		Average	45,8	60	51	37,2	67,7	45	40	57		52,04	37,2	67,7	30,5	
13		Min	0	30	0	0	48	0	0	0	0					
14		Max	75	111	76	55	112	70	66	75	0					

Figure 229: Excel Sample Station Second Biggest Rock (Image Roser)

There is no single right answer here. Depending on which aspects you look at, you may pick a different product for sequencing. I go for product 2434 here, since it has the second-largest cycle time overall due to the 111 seconds for the sunroof mounting, but this is just a hunch. The sequencing interval is 6.41, hence every 6.41 slots I would like to have this product.

$$S_{2434} = \frac{17732}{2766} = 6.41$$

The first slot is already occupied, hence we cannot start this sequence with the first slot. The second slot would not be good either, because then we have two sunroof mountings directly adjacent to each other in slot 1 (product 2436) and slot 2 (product 2434). This would require even more buffer before the sunroof mounting to decouple **two** excessively long cycle times in sequence. So we move to the third slot … and have the same problem again since the fourth slot is another sunroof again.

Damn!

This is a tricky situation. We seem to have always two sunroofs adjacent to each other in the sequence. I picked the second slot to start the sequence, giving me product 2434 at slots 2, 8, 15, 21, and so on. If the sequence arrives at a slot that is already occupied, then you simply take the nearest available slot. The result is visualized below, and we have lots of sunroof mountings adjacent to each other in the sequence. This is a problem!

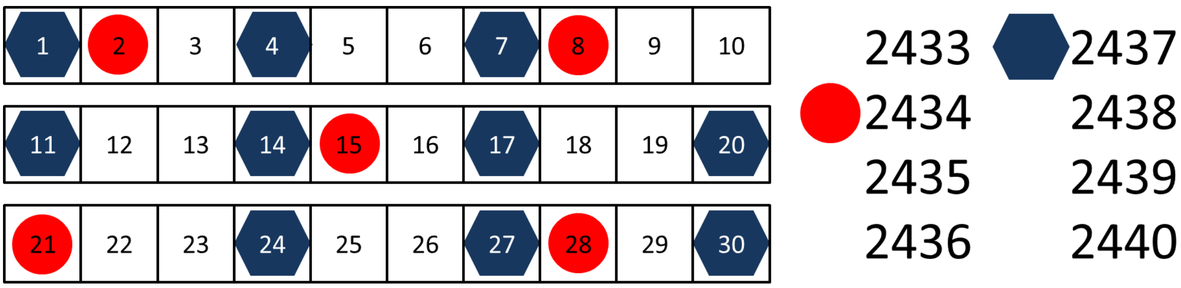

Figure 230: Sequencing Example Product 2 Sequenced (Image Roser)

However, this is not yet the final sequence. I can shuffle both the initial product 2437 (our biggest rock) and the second product 2434 (our second-biggest rock) around. I cannot remove any product, but I can move them within the sequence. And, lo and behold, we can get a sequence where there is always at least one slot free between sunroof mountings!

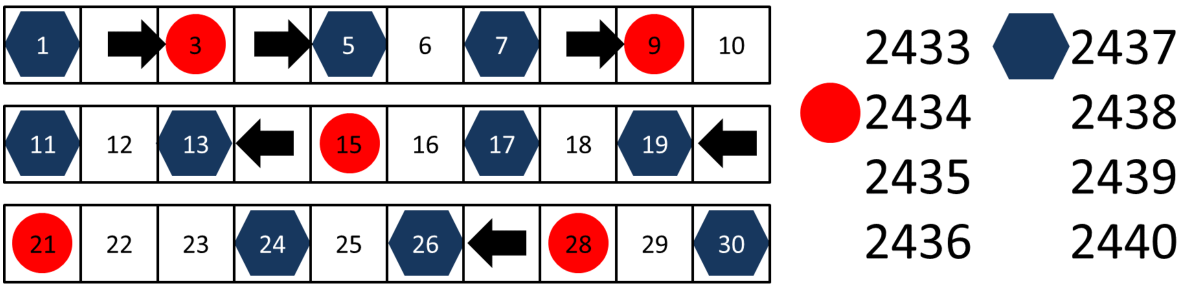

Figure 231: Sequencing Example Product 2 Shuffling (Image Roser)

While this kind of shuffling is time consuming and hard to automate, it can save you a lot of problems later. The updated sequence is shown below.

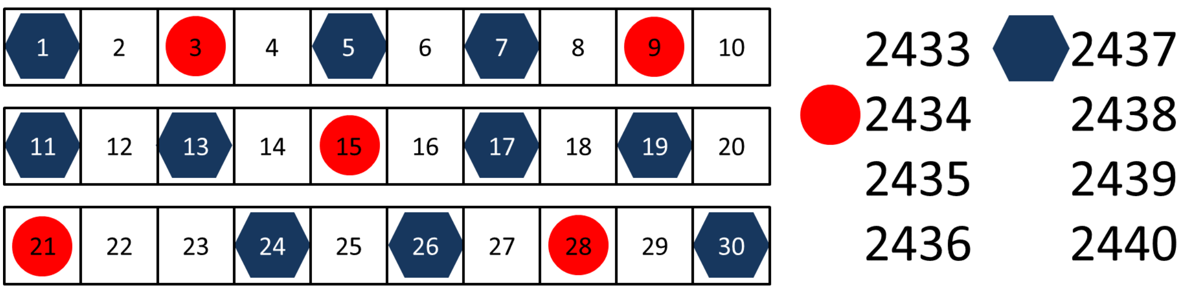

Figure 232: Sequencing Example Product 2 Re-Sequenced (Image Roser)

Generally speaking, if you have no more than 50% of the products with one certain aspect, you can always keep a slot free between them. In our example, we had a total of 8,266 products with sunroof (across two product variants), representing 47% of all products. Hence we can make a series of alternating *sunroof – no sunroof*, sometimes even with two empty slots in between.

However, this can only be guaranteed if you start with an empty sequence. If the sequence has been already filled with other products, our options are limited, and we may not be able to keep a nice alternating sequence without breaking some other sequence benefit. This is the idea of always putting in the biggest rock firsts, so if you run into conflicting sequences later, you already have the biggest problems out of the way.

29.3 Sequence the Third Product

So, what should be our third-biggest rock? The largest deviation from the average of the not-yet-sorted products are now all models without a sunroof. However, since this means all remaining available slots will be filled with "no-sunroof" vehicles, there is no priority to assign them. Instead, we would go for the next biggest deviation which we have in the door-mounting process with product 2435.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1		Quantity	2355	2766	121	5400	5500	790	670	130						
2		Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%						
3		Product	2433	2434	2435	2436	2437	2438	2439	2440	:	Average	Min	Max		Standard
4	Task	Workstation										Time	Time	Time	Spread	Deviation
5	1	Dashboard	41	61	70	40	66	40	40	70		51,9	40	70	30	13,5
6	3	Sunroof Mounting	0	111	0	0	112	0	0	0		52,1	0	112	112	48,3
7	4	Rear Seats	50	50	50	50	50	70	66	72		51,7	50	72	22	9,5
8	5	Passenger Seat & Glove Box	50	49	47	55	53	59	50	50		52,7	47	59	12	3,6
9	6	Driver Seat	59	59	61	48	48	67	55	73		52,6	48	73	25	8,1
10	7	Door Mounting	75	30	76	30	77	34	30	75		51,4	30	77	47	22,4
11		Total	275	360	304	223	406	270	241	340		312,2	223	406	183	
12		Average	45,8	60	51	37,2	67,7	45	40	57		52,04	37,2	67,7	30,5	
13		Min	0	30	0	0	48	0	0	0	0					
14		Max	75	111	76	55	112	70	66	75	0					

Figure 233: Excel Sample Station Third Biggest Rock (Image Roser)

Since we make only 121 products of this type, the sequencing interval is 146.55 hence only every 146.55 slots, I would like to have this product.

$$S_{2435} = \frac{17732}{121} = 146.55$$

Now, I could start the sequence at the first free slot, slot number 2 as shown below.

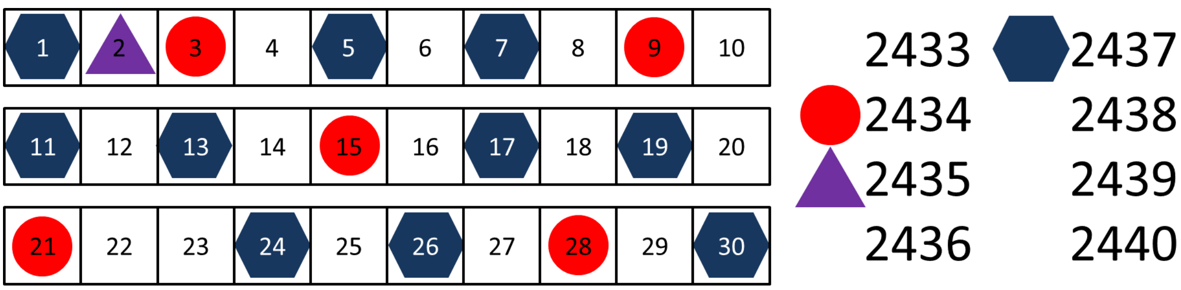

Figure 234: Sequencing Example Product 3 Sequenced (Image Roser)

However, looking into the data in more detail, this would give me a four-door vehicle 2435 directly adjacent to another four-door vehicle 2437. Again, I would have two longer cycle times in sequence for the door-mounting station. There are not many options left, but I could move this product from slot 2 to slot 22, where it is adjacent only to a two door-vehicle 2434.

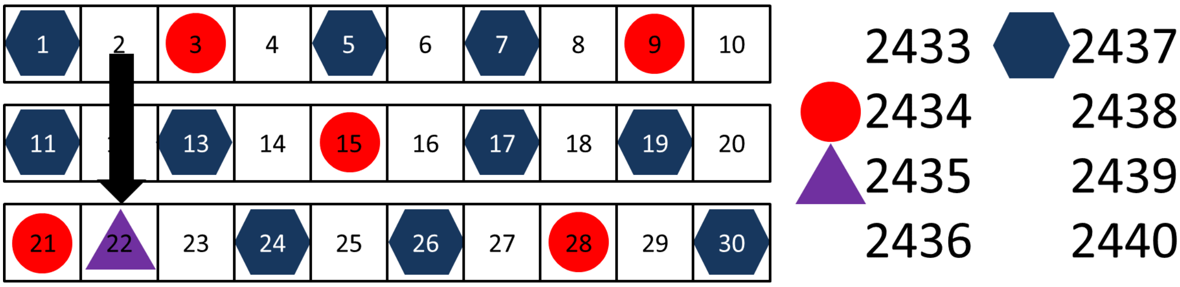

Figure 235: Sequencing Example Product 3 Shuffling (Image Roser)

The resulting sequence is visualized below:

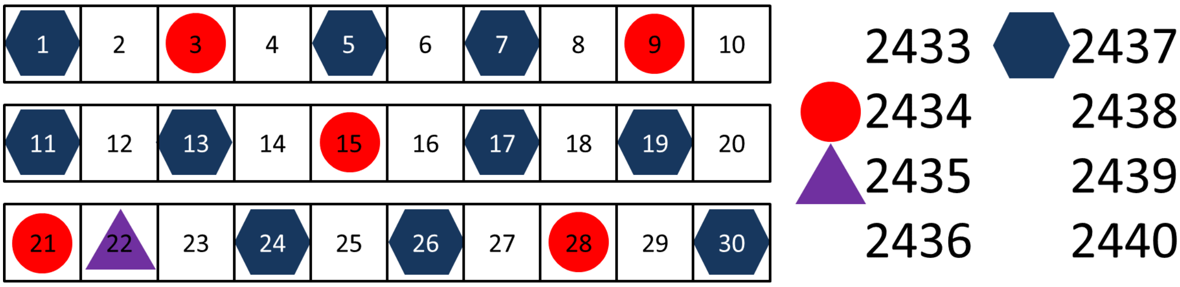

Figure 236: Sequencing Example Product 3 Re-Sequenced (Image Roser)

So far we have taken care of our thee biggest rocks, and so far we have been lucky. We were always able to find a nice spot somewhere for all of our products, where we did not have an excessive accumulation of sequential cycle times.

However, as you can see in the sequence above, now we have only slots left that are adjacent on both sides. It is going to be tight! But this has to wait for the next post. Now **go out, look again for your biggest fluctuations, and organize your industry!**

P.S. Many thanks to Mark Warren for his input.

PSS: The Sequencing Example Excel File for posts 7 to 11 with the complex example is available at https://www.allaboutlean.com/wp-content/uploads/2019/05/Sequencing-Example-Excel-File.xlsx. Please note that this is not a tool, but merely some of my calculations for your information.

30 Mixed Model Sequencing – Complex Example Sequencing 2

Christoph Roser, July 23, 2019, Original at https://www.allaboutlean.com/mixed-model-sequencing-10/

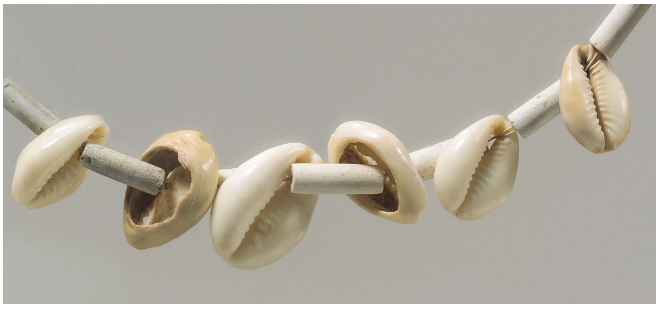

Figure 237: Shell Beads Necklace (Image Metropolitan Museum of Art in public domain)

In this seventh post on Mixed Model Sequencing, I will finish the sequencing of the more complex example with Product-Dependent Workload and Mixed Model Sequencing. This is now the tenth post in this series. I knew this sequencing topic was demanding, but even I am surprised how much there is to cover. Thanks for staying with me, and read on.

30.1 Sequencing the Fourth Product

So, what product should we sequence next? I went for the next largest cycle time of product 2433 at the door-mounting station.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1		Quantity	2355	2766	121	5400	5500	790	670	130						
2		Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%						
3		Product	2433	2434	2435	2436	2437	2438	2439	2440		Average	Min	Max		Standard
4	Task	Workstation									:	Time	Time	Time	Spread	Deviation
5	1	Dashboard	41	61	70	40	66	40	40	70		51,9	40	70	30	13,5
6	3	Sunroof Mounting	0	111	0	0	112	0	0	0		52,1	0	112	112	48,3
7	4	Rear Seats	50	50	50	50	50	70	66	72		51,7	50	72	22	9,5
8	5	Passenger Seat & Glove Box	50	49	47	55	53	59	50	50		52,7	47	59	12	3,6
9	6	Driver Seat	59	59	61	48	48	67	55	73		52,6	48	73	25	8,1
10	7	Door Mounting	75	30	76	30	77	34	30	75		51,4	30	77	47	22,4
11		Total	275	360	304	223	406	270	241	340		312,2	223	406	183	
12		Average	45,8	60	51	37,2	67,7	45	40	57		52,04	37,2	67,7	30,5	
13		Min	0	30	0	0	48	0	0	0	0					
14		Max	75	111	76	55	112	70	66	75	0					

Figure 238: Excel Sample Station Fourth Biggest Rock (Image Roser)

Since we make 2,355 products of this type, the sequencing interval is 7.53.

$$S_{2437} = \frac{17732}{2355} = 7.53$$

Starting at slot 2, this product would be in slot 2, 10, 17 … wait, this one is already occupied, the next free one is 16, and 25 in the sequence. The resulting sequence is shown below.

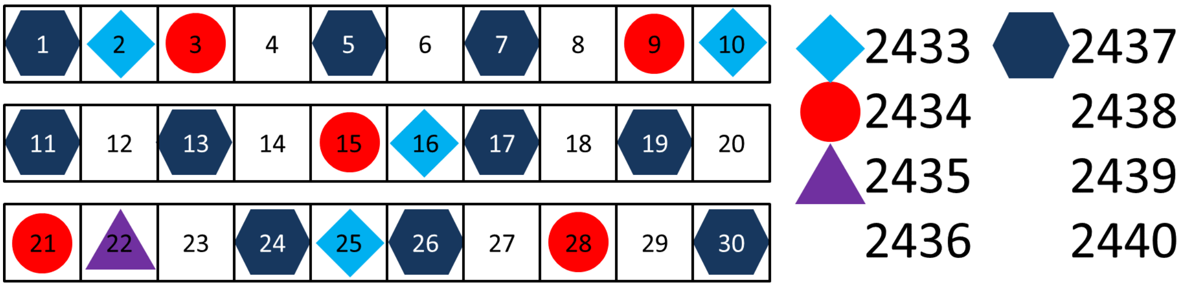

Figure 239: Sequencing Example Product 4 Sequenced (Image Roser)

Let's have a look if we want to shuffle something around to make an even better sequence. Product 2433 in slot 2 is adjacent to 2434 in slot 3, which is not a problem. Both the door mount and the sunroof mounting have cycle times in the opposite direction and cancel each other out. The adjacent product 2437 in slot 1 is a bigger problem, since both are four-door vehicles with a rather long door mounting cycle time. Since they are next to each other, this will accumulate, and we need a bigger buffer.

However, there are no other free slots that are NOT adjacent to 2437, so we cannot move this new product around. If we would move another part again, then somewhere in the other sequence it would get even worse. Therefore, while not ideal, part 2433 stays where it is in the sequence.

30.2 Sequencing the Fifth Product

Next, we go for product 2440, since it has the longest remaining excess cycle time of 75 seconds for the door mount (and also pretty large cycle times for the rear seats, driver seats, and the dashboard). With only 130 parts, this gives us a sequence interval of 136.40. Starting with the first available free slot this sequence is shown below.

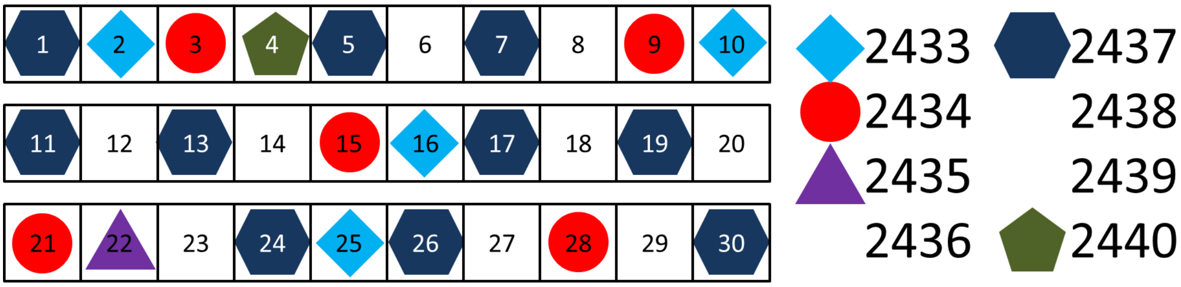

Figure 240: Sequencing Example Product 5 Sequenced (Image Roser)

This is also not ideal, since product 2440 in slot 4 clashes with 2434 on the left by both having longer dashboard mounting times, and with 2437 on the right by both having longer dashboard and door mounting times. But I cannot find a better slot, so it stays where it is in the sequence.

30.3 Sequencing the Sixth Product

Not many products left. The biggest cycle time is product 2438 with 70 seconds cycle time for the rear seat assembly. With 790 parts this gives us a sequence interval of 22.45. Starting at the first empty slot, this gives us position 6 and 28 … which is already occupied, hence I take an adjacent one with slot 27. The sequence is shown below.

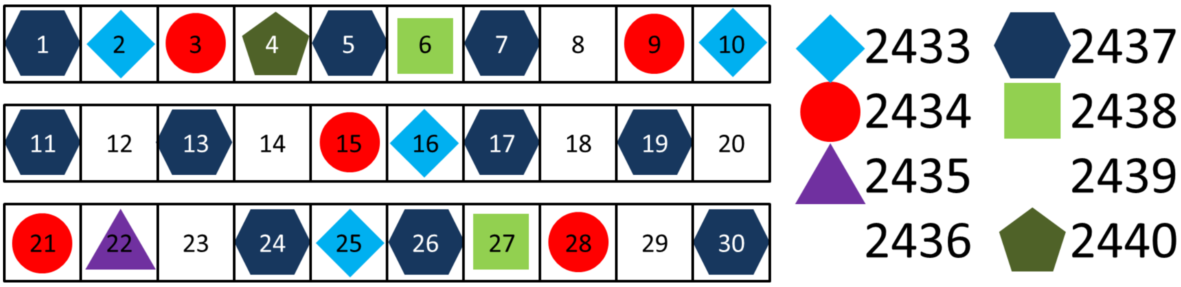

Figure 241: Sequencing Example Product 6 Sequenced (Image Roser)

Again checking for effects with adjacent products 2437 and 2434, there seem to be no bigger issues. If anything the door mounting of 2438 and 2434 are both short (i.e., both have not enough work for the cycle time). Having this in sequence accumulates available time that has to be used for another product, requiring again larger buffers. However, this is not a big issue, and I am not destroying the rest of the sequence by shuffling things around her.

30.4 Sequencing the Seventh Product

Only two products left! These two are not even too different with their cycle times. Both have two very short cycle times at the door mount, but product 2439 is a tick longer at the rear-seat mounting, hence I pick this one. With 670 parts we get a sequence interval of 26.47. Starting at the first available slot, I get position 8, 34, and so on. Always check if these positions are already occupied. The resulting sequence is shown below.

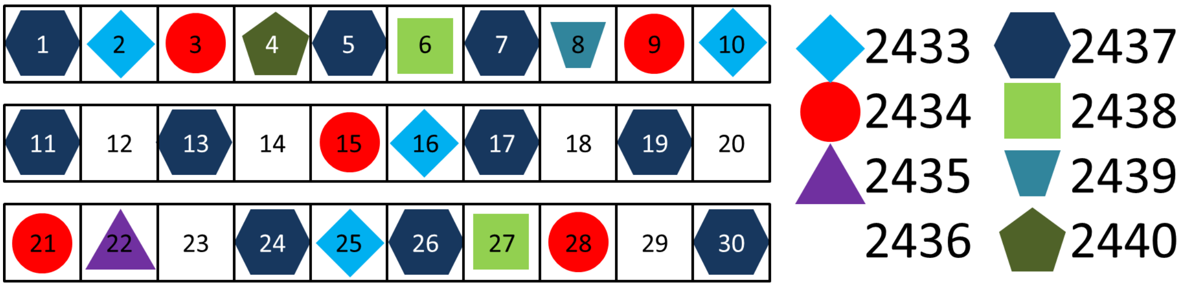

Figure 242: Sequencing Example Product 7 Sequenced (Image Roser)

There is not much conflict with the adjacent slots. Even if there is, not much I could do about this anyway at this point in the sequencing.

30.5 Sequencing the Eighth and Last Product

Only one product left, number 2436. This one is easy. You do not need to pick a product since there is only one left. You do not need to calculate a sequence interval (which would be 3.28 for 5400 parts). You do not need to worry about placement. Since it is the last product, all free slots must be this product type. The sequence is shown below.

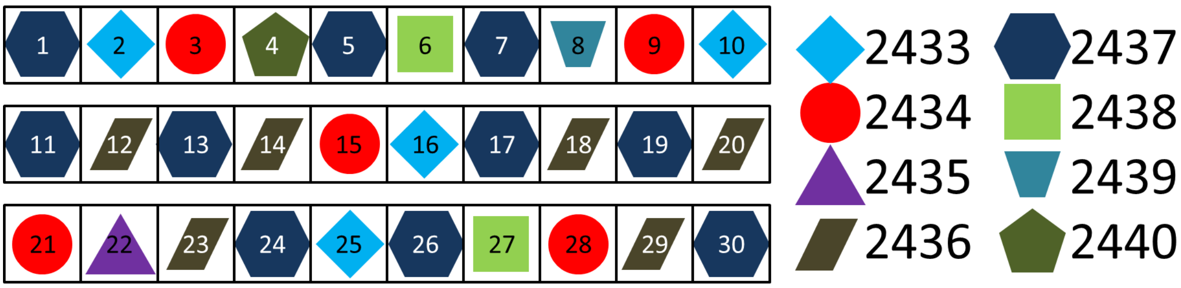

Figure 243: Sequencing Example Product 8 Sequenced (Image Roser)

Is it a good sequence? Maybe. Is it perfect? Probably not. Juts by looking at it, you can see that the last product 2436 is somewhat unevenly distributed. It is not present at all in the first row, is four times in the second row, and twice in the third row of this sequence. So it is probably possible to make better sequences than the one above. As I said above, sequencing is more of an art than a simple process. You could have sequenced the parts in a different order. You could have shuffled them differently than I did. There are many ways to do this differently. If you don't like the sequence, do it again, differently and better. It is a very iterative process.

30.6 A Bit on the Terminology

I was looking quite a bit for a proper English name for this method. Most articles just describe it as Mixed Model Sequencing, Mixed Model Balancing, or confusingly also just Line Sequencing (even though there are many other reasons to sequence like changeover sequencing, Just in Sequence, prioritization sequencing, EPEI, one piece flow, or just simple production sequences like FIFO, EDD, SPT and so on)

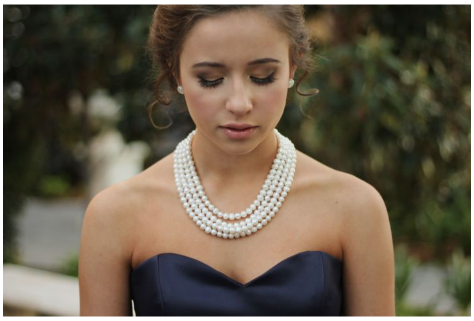

Figure 244: Pearl Necklace (Image MaxPixel in public domain)

In German it is much easier, and the word *Perlenkette* (pearl necklace) is used for this type of sequencing. If you mention *Perlenkette* to a German lean expert, he instantly knows what it is all about. This is by the way also the reason why all the posts on this topic start with an image of a necklace, all of them ancient with the images from the Metropolitan Museum of Art in New York. Only the lady here is more recent.

While the English word *pearl necklace* would be nice to use here in lean too, it unfortunately has a second definitely NSFW meaning. So, if you know a good, short, and generally accepted term for Mixed Model Sequencing, let me know.

In my next post I will show you how to check the quality of the sequence and how to determine the required buffer size for the sequence. Until then **go out, think about a better way to sequence your products for a mixed model assembly, and organize your industry!**

P.S. Many thanks to Mark Warren for his input.

PSS: The Sequencing Example Excel File for posts 7 to 11 with the complex example is available at https://www.allaboutlean.com/wp-content/uploads/2019/05/Sequencing-Example-Excel-File.xlsx. Please note that this is not a tool, but merely some of my calculations for your information.

31 Mixed Model Sequencing – Complex Example Verification

Christoph Roser, July 30, 2019, Original at *https://www.allaboutlean.com/mixed-model-sequencing-11/*

Figure 245: Ancient Egyptian Eye Beads (Image Metropolitan Museum of Art in public domain)

As part of a much larger series on Mixed Model Sequencing, this post describes how to verify the sequence quality. It also describes how to determine the required buffer spaces to buffer against these fluctuations in workload. There may be some wiggle room here. Read on:

31.1 Verify Sequence Quality

Next, we have to verify the quality of the sequence. Let's demonstrate it with the sunroof-mounting station. The average cycle time at this station is 52.1 seconds. Whenever a model 2434 with a sunroof comes along, it takes 111 seconds and hence 58.9 seconds longer than average. For a model 2437, it takes 112 seconds and hence 59.9 longer than average. All other models have no sunroof, and they have a cycle time of 0 seconds and hence 52.1 seconds less than average. Please note that the average is based on the average cycle time of this station and not the line takt or target cycle time. The table below shows all the data, with the sunroof mounting over/under times marked in red.

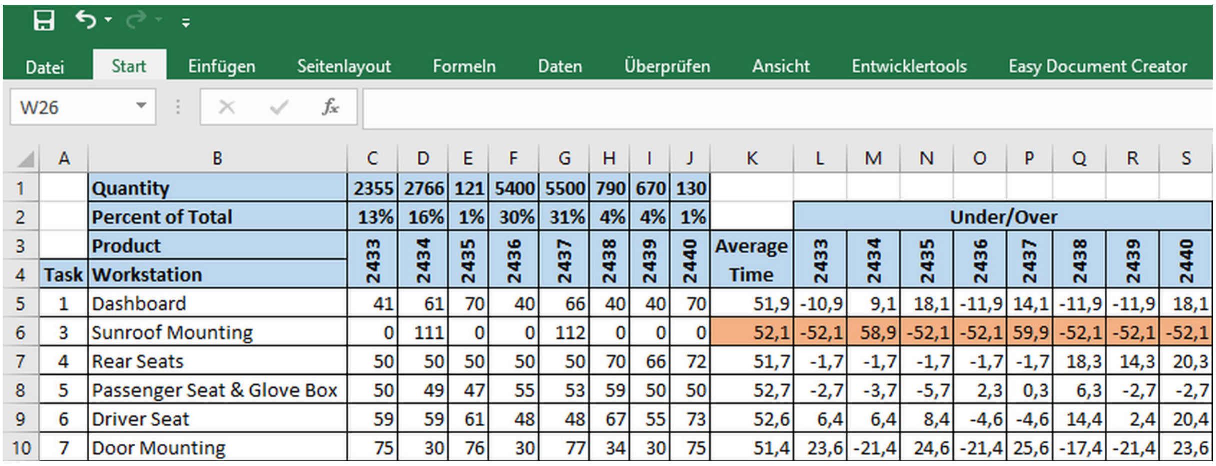

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S
1		Quantity	2355	2766	121	5400	5500	790	670	130									
2		Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%		Under/Over							
3		Product	2433	2434	2435	2436	2437	2438	2439	2440	Average	2433	2434	2435	2436	2437	2438	2439	2440
4	Task	Workstation									Time								
5	1	Dashboard	41	61	70	40	66	40	40	70	51,9	-10,9	9,1	18,1	-11,9	14,1	-11,9	-11,9	18,1
6	3	Sunroof Mounting	0	111	0	0	112	0	0	0	52,1	-52,1	58,9	-52,1	-52,1	59,9	-52,1	-52,1	-52,1
7	4	Rear Seats	50	50	50	50	50	70	66	72	51,7	-1,7	-1,7	-1,7	-1,7	-1,7	18,3	14,3	20,3
8	5	Passenger Seat & Glove Box	50	49	47	55	53	59	50	50	52,7	-2,7	-3,7	-5,7	2,3	0,3	6,3	-2,7	-2,7
9	6	Driver Seat	59	59	61	48	48	67	55	73	52,6	6,4	6,4	8,4	-4,6	-4,6	14,4	2,4	20,4
10	7	Door Mounting	75	30	76	30	77	34	30	75	51,4	23,6	-21,4	24,6	-21,4	25,6	-17,4	-21,4	23,6

Figure 246: Excel Sample Station Average Over Under (Image Roser)

Now we simply go through the sequence and sum up the times that the process is over or under average speed. Across all parts this should give a value of zero again. Below is once more the sequence we had from the last post.

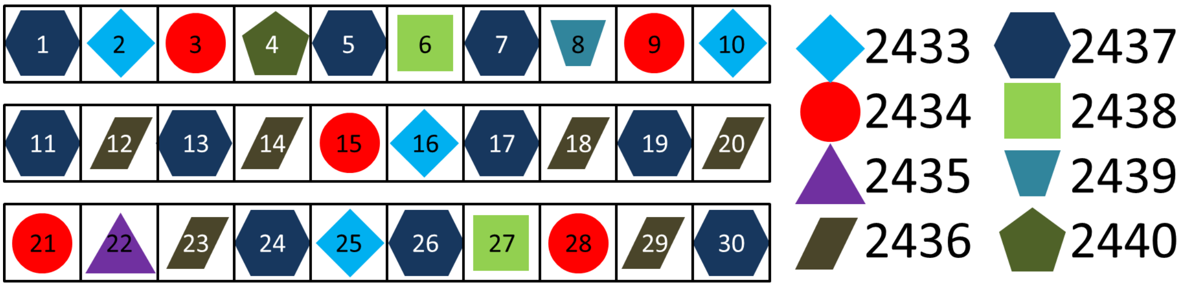

Figure 247: Sequencing Example Product 8 Sequenced (Image Roser)

Starting at zero, the first part has a sunroof, adding 59.9 to the accumulated times. The second part had no sunroof,reducing it by 52.1 seconds to 7.9 seconds. The third part added a roof again adding 58.9 seconds to an accumulated total of 66.8 seconds at the third slot. The graph below shows the first 30 slots. At slot 22 we had two "non-sunroof" vehicles in a row. In the long run the values will go up and down, they can also go negative, to eventually reach zero again.

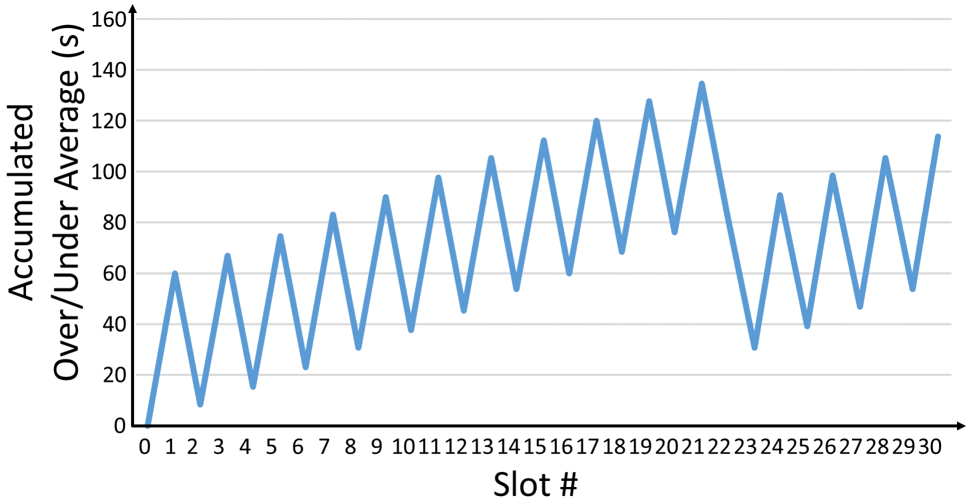

Figure 248: Sequencing Example Accumulated Sunroof Time (Image Roser)

The graph below shows the first 30 slots for all stations. Notice how the sunroof station has the most regular zigzag curve? These were the first two product types we sequenced.

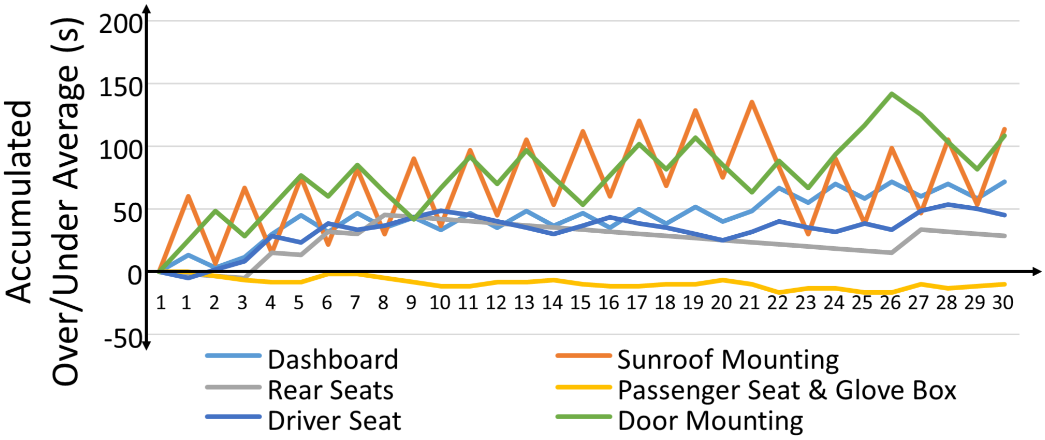

Figure 249: Sequencing Example Accumulated All Time (Image Roser)

Overall, you will notice an upward trend at most stations. While eventually things will go down again, larger deviations from the average are not good, as this requires larger buffers. Hence, this is not such a good sequence.

The reason for this is part type 2436. You remember how in the last post this part was the last one we sequenced, and it had a pretty bad sequence. While there should have been a part every 3.28 slots, due to other parts being sequenced earlier this part had to do with the slots that were left at the end. Hence, during the first 30 slots there were only 6 of this part type, where there should have been 9.

This is a problem because across the board this part type had a very low work content. The average work content per station for this part type was only 37.2 seconds, much less than the average of 52.04. The data is once again shown below, with the average work contents for all parts shown in red.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1		Quantity	2355	2766	121	5400	5500	790	670	130						
2		Percent of Total	13%	16%	1%	30%	31%	4%	4%	1%						
3		Product	2433	2434	2435	2436	2437	2438	2439	2440		Average	Min	Max		Standard
4	Task	Workstation									⋮	Time	Time	Time	Spread	Deviation
5	1	Dashboard	41	61	70	40	66	40	40	70		51,9	40	70	30	13,5
6	3	Sunroof Mounting	0	111	0	0	112	0	0	0		52,1	0	112	112	48,3
7	4	Rear Seats	50	50	50	50	50	70	66	72		51,7	50	72	22	9,5
8	5	Passenger Seat & Glove Box	50	49	47	55	53	59	50	50		52,7	47	59	12	3,6
9	6	Driver Seat	59	59	61	48	48	67	55	73		52,6	48	73	25	8,1
10	7	Door Mounting	75	30	76	30	77	34	30	75		51,4	30	77	47	22,4
11		Total	275	360	304	223	406	270	241	340		312,2	223	406	183	
12		Average	45,8	60	51	37,2	67,7	45	40	57		52,04	37,2	67,7	30,5	

Figure 250: Excel Sample Station List Product Total (Image Roser)

Hence, it would have been a much-needed part type to reduce the workload again. Unfortunately, this part type is now very unevenly distributed, and our sequence is not that good. Remember how I mentioned Mixed Model Sequencing being an iterative process? That's right, we should go back and do the sequence again. Maybe this time we will sequence part 2436 earlier, possibly directly after part types 2437 and 2434 (the two part types with a sunroof).

Also remember how I mentioned in an earlier post that there are many different aspects that can influence sequencing? In our example we used mostly the largest cycle times for any product at at any station, but maybe we should have also been looking at the work content per product variant. Hence, for a real-world sequence I would go back and do it again, hopefully better – unless I find another aspect I should have been considering too. Here, however, I continue with the sequence we already have. Again, it is iterative and more of an art than a science!

31.2 Buffer Size

	A	B	C	D
1	Accumulated	Min	Max	Spread
2	Dashboard	0,0	72,7	72,7
3	Sunroof Mounting	0,0	134,9	134,9
4	Rear Seats	-5,0	44,7	49,7
5	Passenger Seat & Glove Box	-15,9	0,3	16,3
6	Driver Seat	-4,6	54,3	58,9
7	Door Mounting	0,0	142,2	142,2

Figure 251: Excel Sample Station Peak Under Over and Spread (Image Roser)

Anyway, let's have a look at the buffer size. The table here shows the largest positive and negative accumulation of work content (i.e., the value of the top and bottom most peaks).

We are interested in the spread (i.e., the difference between the smallest and the largest peak). This is the fluctuation that we have to cover. This is the amount of buffer that we would need to provide for a smooth operation.

For a continuously moving assembly line, this would be represented by the width of the space allocated to this workstation. If your line moves 6 meters per minute, then a buffer of 72.7 seconds is an additional 7 meters of slot space.

For a line with buffer slots between stations, this spread would have to be divided by the cycle time (or takt time, depending what you used to calculate the spread) to determine the number of buffer spaces.

Looking at the data table above this seems like a lot, but there is some wiggle room. This peak buffer space is needed only a very few times during the production. Here are a few options:

- You probably plan to include some buffer anyway. This buffer here does not have to be on top of the regular buffer, but can be combined with the regular buffer. For example, your average cycle time for the rear seats is 51.7 seconds. If the whole line moves at a line takt of 60 seconds, then you automatically have 8.3 seconds buffer here. If the work content fluctuation in you rear seat station, ask for a buffer of 49.7 seconds, then you do not have to add them, but merely take the larger one.
- Humans work faster if there is a lot of work, and slower if there is less. Hence, a part of the fluctuation can be buffered with a changing human work speed. This can be 10% of the buffer, or 20% if you're daring, but probably not 30%. The tighter the buffer, the more likely that the workers will not be able to make it even with a (short time) extra effort.
- But even this is not a huge problem. If the buffer is not enough, then the station cannot make it in time, and the rest of the line has to wait. If this happens once per shift for 10 seconds, so what. While not ideal, a delay for the rest of the line of 10 seconds in a shift may be preferable to an extra 3-meter floor space for buffer **all the time!**

So you see, there is some wiggle room for the buffer. If it is not wiggly enough for you, you could look for an even better sequence. Or you can also change the product design, add better machines, and do a lot of other things as discussed in a previous post to make these problems simply go away. Or you bite the bullet and add the buffer. Or you try if you can get away without that much buffer, and change it back if it does not work. The possibilities here are endless. **Now, go out, make your buffer fit your needs, and organize your industry!**

P.S. Many thanks to Mark Warren for his input.

PSS: The Sequencing Example Excel File for posts 7 to 11 with the complex example is available at https://www.allaboutlean.com/wp-content/uploads/2019/05/Sequencing-Example-Excel-File.xlsx. Please note that this is not a tool, but merely some of my calculations for your information.

32 Mixed Model Sequencing – Summary

Christoph Roser, August 06, 2019, Original at
https://www.allaboutlean.com/mixed-model-sequencing-12/

Figure 252: White Beads Necklace (Image Metropolitan Museum of Art in public domain)

Mixed Model Sequencing to manage different product types with different work content is tricky. This is now the twelfth post of this series. I knew it would be long, but I never guessed that it would be that many posts. This is almost a book (and will probably be part of a book in the future).

32.1 Summary of Steps for Mixed Model Sequencing

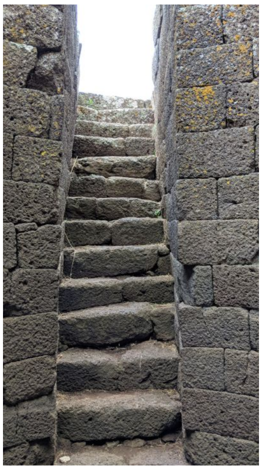

Figure 253: Ancient Stairs in Sardinia (Image Roser)

Since this is a pretty long series, let me give you a recap and overview of the different steps.

- **Big Picture:** The first step in sequencing products for a line with a product-dependent workload is to get an overview of the big picture: your customer takt, line takt, and cycle time (see post #7 for details).
- **Work Content:** Next you need to understand the work content: what tasks need to be done for which product and how long it will take. If you already have a production line, you need to know the work content at each station. If you don't yet have a production line, you need to know the work content for each individual task and need to arrange them into workstation-sized chunks (see post #7 and post #8 for details).
- **Eliminate Fluctuations** (where possible): In some cases it may be possible to reduce or eliminate these differences in work content. Now would be the time to do this, as this may make your sequencing easier (see post #2 and post #7 for details). You can also check if you can **adjust the capacity**, although this is often not possible or is inferior to a good sequence (see post #3 for details).

- **Production Quantities**: You also need (an estimate of) the production quantities: how much of which product you have to produce (see post #8 for details).

Figure 254: Rock Stacking Iceland (Image Roser)

- **Big Rocks First**: Find out which products to sequence first. Possible criteria are: *Longest Cycle Times, Shortest Cycle Times, Largest Quantities Produced*, Largest Spread, Largest Average Work per Product, Smallest Average Work per Product, Largest Cycle Time per Product, or Smallest Cycle time per Product. I've often found the first three especially helpful, but this depends on your products and their workload (see post #8 for details).
- **Sequence Products One by One**: Now you sequence products into your production sequence. Starting with the biggest rock, you try to distribute it as evenly across the schedule as possible. Try to avoid accumulating excess workload or excess idle time at workstations – although this may not always be possible (see post #9 and post #10 for details). Check the sequence afterward for excess workload or excess idle time at workstations, and consider adjusting or re-sequencing (see post #11 for details). This step may have to be repeated multiple times until a suitable solution is found.
- **Define Buffer Size**: Create buffers before and after the critical stations that can handle the excess workload or excess idle time caused by the product-dependent workload. This does not need extra safety margins, and you may be able to get away with a smaller buffer (see post #11 for details).

32.2 Rinse and Repeat

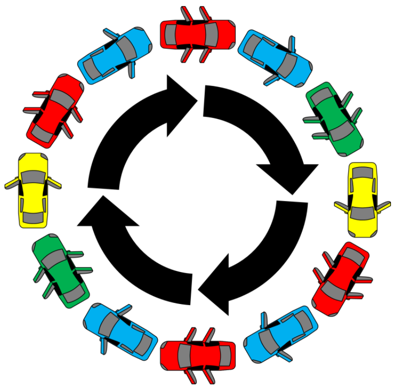

Figure 255: Cars going in Circle (Image Roser)

These steps above help you create a sequence for one production line. This may be challenging. But don't worry, you will have to do it over and over again. First of all, your first sequence is unlikely to be very good, and multiple tries may be needed to get a good sequence.

Second, your system will change all the time. The sequence you create will be valid only for a short period of time. The next period may have a different demand, and hence may need a different sequence. Sooner or later you will also have changes to the machines and processes, which will affect your cycle times. Overall, you will have to do the sequence again and again every few days or weeks.

The period for which you sequence your products depends on the cycle time. If it is a fast cycle time (e.g., less than three minutes), you may sequence only for one shift, and create a new sequence again for the next shift. If you have longer cycle times but less than an hour, you may sequence a day or a week. If your cycle time exceeds hours, you may sequence a month. If your cycle time is multiple hours … then you probably don't need to sequence at all but adjust the capacity by adding or removing workers for one shift to manage the excess workload or excess idle times. Lucky, ain't it.

32.3 Software Tools?

Figure 256: Beep-bop … Here's your sequence… (Image Boffy b under the CC-BY-SA 3.0 license)

Soon you will figure out that sequencing by hand is a time-consuming and not-very-entertaining task. You may wonder if you should use software tools to do the sequencing. **By all means, please do!** This kind of optimization is well suited for software algorithms.

If you have all the data, a software tool can run many different sequences for comparison. Please keep in mind that due to the length of the sequence and the number of products, this will not be an exhaustive search, as the number of combinations will quickly exceed the number of atoms in the universe. But even if it is not the absolute perfect solution, pretty good is good enough for us.

The algorithm may create a first sequence and then try to shift products around to reduce excess workload and excess idle time. With luck, you can swap a part in the sequence and neutralize an excess workload at one time with an excess idle time at another time. Even if the program is not as smart as you, it can try out things much faster than you can, and may end up with a better solution than you. Most importantly, it saves you time!

The question now is: Which software tool? Here I'm not an expert, and I don't have an overview of the available software packages. For convenience sake I would recommend checking if your ERP program has such a package. If the software tool is integrated with all the data in your system, it will save you a lot of time. Otherwise you may have to shift data around from your ERP system to the sequencing package and then the sequence back again into the ERP tool. Probably most sequencing packages out there manage a reasonable sequence – as long as the data is correct.

Figure 257: Hopefully not your data… (Image RitaE in public domain)

This is my one caveat: **Make sure the data in the system is correct**. As with any computer system, garbage in – garbage out. Occasionally check if the data is still good, or if the production line or the products have changed and the digital twin lags behind.

This now concludes this series of twelve (!!!) posts on Mixed Model Sequencing. But, what can I say, it's a difficult and demanding topic. Thanks for staying with me through this long series, and I hope it was helpful for you. Now go out, sequence your products to get the last bit of efficiency out of it, and **organize your industry!**

PS: Many thanks to Mark Warren for his input.

PSS: The Sequencing Example Excel File for posts 7 to 11 with the complex example is available at https://www.allaboutlean.com/wp-content/uploads/2019/05/Sequencing-Example-Excel-File.xlsx. Please note that this is not a tool, but merely some of my calculations for your information.

33 Cardboard Engineering – Preparation

Christoph Roser, August 13, 2019, Original at https://www.allaboutlean.com/cardboard-engineering-preparation/

Figure 258: Kid in a cardboard Paper Plane (Image Tverdokhlib with permission)

Cardboard Engineering (CBE, sometimes also Cardboard Modeling) is in general the building of models from cardboard. These models are usually quick and inexpensive to build, but often not very durable. In lean manufacturing, these cardboard models are often workstations or entire assembly lines to test different concepts before building the whole thing in more expensive and time-consuming aluminum and steel. This allows faster and easier experimentation with different concepts to improve your production system.

33.1 Introduction

Figure 259: Nintendo Labo Piano (Image Roser)

Cardboard is often used for the packaging of goods. Afterward it is disposed of (and hopefully recycled). Hence, it is a readily and freely available material. Even if you buy it (because you need more, or you want higher grade), it is not that expensive. Hence it fits one requirement for models: it is cheap!

Second, it is easy to modify. Cutting and gluing cardboard is quick and easy. Hence it is also quick and easy to build models.

Its disadvantage is that it is not very durable. Cardboard models don't last very long if you have to handle them.

Figure 260: Cardboard Boat (Image Roadell Hickman under the CC-BY-SA 4.0 license)

Cardboard models have many uses. The image above shows a Nintendo Labo cardboard piano interacting with the Switch gaming console. Cardboard boat races or bridge constructions are a fun and easy way to teach engineering skills. Cardboard toy planes are also common

33.2 What Is It Used For in Lean?

We are interested in the use of cardboard engineering for lean manufacturing. These models are usually on a 1:1 scale (i.e., full sized). They are to simulate a subsequent real-world situation in order to identify problems and improvement potentials. The model is then adjusted to eliminate the problem and to improve the performance of the system. And, oh, it is also usually great fun and a wonderful way to create team building, employee interaction, and to improve motivation!

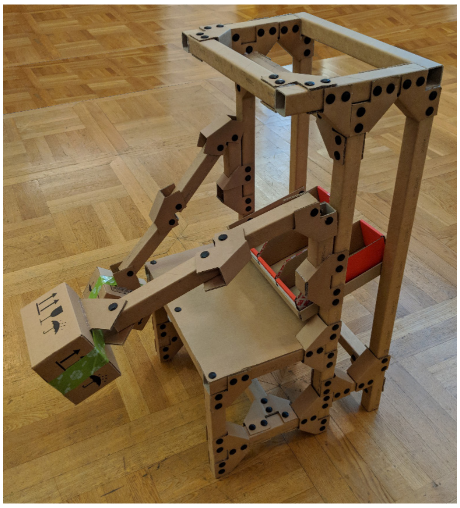

Figure 261: ASSTEC Cardboard Robot (Image Roser)

There are a couple of uses for these models. I usually find that there are two main groups of models: rough models for large systems, and detailed models for smaller systems. The rough models are for the big picture. In other words, "what goes where" to optimize the floor space. The detailed models are to optimize the actual processes like ergonomics, performance, material flow, and so on. It is rare to have a model that is both large and detailed, since it simply takes too much time to build. A possible combination, however, is a large model that has some detailed parts for optimization. Anyway, the two main directions are as follows:

Optimize Layout: Figure out if the machines and equipment fit in the space you have on the shop floor, and how to best fit it. Can you still access it for maintenance? Are the emergency exits blocked? Is there enough space for the forklift to pass through? If your goal is to merely find out if it fits, then you do not need to model the system in great detail. A rough model is easy enough here.

Optimize for Ergonomics, Performance, and Material Flow: Where should be what? Is it within easy reach, even for the not-so-tall employees? How long do the tasks take? Where does the material go into the system? Where does it go out? How is it handled? Do you have enough space? Do you have enough material? For these kind of questions you need a more detailed model.

33.3 What Do You Need?

Figure 262: Empty Cardboard Boxes (Image Vera Kratochvil in public domain)

Cardboard. Lots of cardboard. This should usually not be a problem. Pretty much any plant I know throws out piles of cardboard on a daily basis. Just collect these boxes for a few days and you should be all set. Boxes that are still re-foldable into a box rather than flat may make your life easier later on. You can also use flat cardboard, but in this case you have a bit more work to create 3D models out of 2D cardboard. If you happen to come across some cardboard tubes or similar items, you may put them aside too, since they may be useful later on. If for some reason you do not have cardboard, you can simply buy it. Also, while it is called Cardboard Engineering, you are not limited to cardboard. Feel free to use whatever is at your disposal. If you have some wood or plastic tubing lying around, feel free to use them too. If you have a work table that won't change, just build the gadgets on top of the table. Keep in mind, however, that cardboard is easiest, and you may need additional tools like saws, drills, and screws if you use other materials.

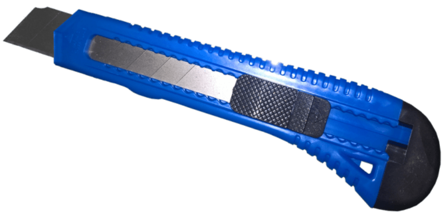

Figure 263: Carpet Knife (Image Maxpixel in public domain)

Something to Cut. You will have to cut the cardboard down to size. A pair of scissors, however, is not the perfect tool for cardboard (although it is still good to have one around). Much better is a box cutter, also called a carpet knife. These are also much safer than a normal knife. Make sure you have enough knives. If you have sturdier material like more rigid cardboard profiles, you may also need a saw. A simple hand saw will usually be good enough, although you may also use electric saws.

Figure 264: Hot Glue Gun (Image Roser)

Something to Join. You will also have to attach cardboard pieces to other cardboard pieces. Here you have multiple options. You can get adhesive packing tape. You probably should have these rolls somewhere in your plant too. You can also get a stapler, or even a specialized box stapler if you have one. You can also use glue. Here I recommend a hot glue gun; the glue hardens much faster and holds much better than a regular glue. It also makes a sturdier connection than tape. If you plan to use a glue gun, make sure you have an extension cord. The cable of the glue gun is usually not long enough for larger projects. If you have multiple glue guns, plan a separate extension cord for each of them, rather than a multi-outlet extension cord. Otherwise you are again limited in range by the distance to the other glue guns.

Figure 265: Flipchart Markers (Image Jan Hagelskamp1 under the CC-BY 4.0 license)

Something to Mark. Finally, you will probably want to mark the model, add signs and directions, and in general fiddle with visual management. This is especially important for detailed models. hence make sure to bring some flip chart markers in different colors. You can also bring colored paper for fancier markings. Adhesive dot labels or similar may also be used. In a pinch, adhesive packing or painter's tape can also be used to mark the cardboard.

Something to Document. The model may not be there for long. Hence, take pictures while it is still looking good. Make sure you have permission to take photos. You may even take videos. Again, make sure you have permission and the union representative is informed (or even involved in the workshop).

So now you have all the material you need to make a cardboard engineering workshop. This I will describe you in the next post. Until then **go out, get your glue guns ready, and organize your industry!**

34 Cardboard Engineering – Workshop

Christoph Roser, August 20, 2019, Original at https://www.allaboutlean.com/cardboard-engineering-workshop/

Figure 266: A Cardboard Boat Race (Image Roadell Hickman under the CC-BY-SA 4.0 license)

In my last post I talked about what you need for Cardboard Engineering. In this post I will show you how to do a Cardboard Engineering workshop. Spoiler: Keep in mind that **the goal is not to just put something together but to try out different options** (and I will repeat this a few times in this post). It is very easy to have fun with cardboard while learning very little about the problem you want to investigate!

34.1 The Goal of the Workshop

Figure 267: Archery Target with Arrows (Image Casito under the CC-BY-SA 3.0 license)

Before you start, you should be clear on what you want to achieve in the workshop. As per my last post, the two most common goals are 1) to create a layout for a production system on an already existing shop floor; and 2) to optimize a single workstation for ergonomics, layout, and material flow. Make sure you know which one you want. If you want both the big picture layout and the detailed set up of every workstation, you are biting off more than you can chew! While you may be able to set something up both big and small, keep in mind: **The goal is not to just put something together but to try out different options!**

Also understand the scope of your workshop. What are the boundaries of the material and information flow you want to analyze. The better you define your problem, the better will be your solution.

34.2 Team

Figure 268: Construction Worker Team (Image Vittaya_25 with permission)

As usual, the team should be cross-functional. As this is a very hands-on workshop, you should include more than just one operator from the line or workstation (or a similar line or workstation if you are creating a new one). You also need a moderator, which may also be the manager. It also helps to include an engineer who will later turn our cardboard models into reality. Ask if a member of the unions wants to participate. For layout optimization it may also help to have someone familiar with the infrastructure and the requirements on emergency exits and the minimum width of paths and roads and other stuff. Altogether there could be between four and six people involved. While I often split larger groups into sub-teams, it may not work well, especially if you want to set up the layout.

34.3 Duration

Figure 269: Time on Hand (Image geralt in public domain)

How long should you plan for the workshop? It depends. Usually it is difficult to do it properly in less than one day. Sure, you can put something together in half a day, but again **the goal is not to just put something together but to try out different options!** Such workshops can last between one day and a whole week, depending on the complexity of the problem. I will talk a bit more about this farther down at the Agenda.

34.4 Final Preparations

Besides organizing all the tools for the workshop (cardboard, cardboard cutters, glue gun, … see my last post), you also will have to make sure you have proper space for the cardboard designs. If it is a layout optimization of a new line, it would be best to do it right on the shop floor – which of course requires this space to be free. If the shop floor is not available for a cardboard layout, it may be very difficult to do it somewhere else. The idea of the cardboard model for the layout is to see how it fits on the shop floor, and hence makes only sense if it is done on the shop floor. If you have no access, you may be better off doing it on a paper layout or a computer.

If you optimize only one workstation, it may be useful to do it right on the shop floor. However, this is often not possible (e.g., since the line is running or the space is occupied). In this case feel free to use any kind of meeting room that is big enough. It may help to have a few samples of the product components and the completed product to see if the model fits the product.

34.5 Agenda

Figure 270: Checklist (Image Clker-Free-Vector-Images in public domain)

Now lets have a look at the workshop itself. Below are the different agenda points. This is loosely based on [My Workshop Structure for Creative Problem Solving](#).

34.5.1 Introduction

Welcome, warm-up, introduction of problem, introduction of people … if you have done improvements workshops before, you know the drill.

34.5.2 Creating Ideas and Narrowing Down the Solution Space

Figure 271: Hanging Light Bulbs (Image Ross Dunn under the CC-BY-SA 2.0 license)

Even though it is tempting, you should not start with the cardboard construction right away. If you would do so, you will get a cardboard gadget, but it will be far from an improved solution. Hence, before touching the cardboard, you should first narrow down the solution space to a few feasible solutions. This is almost like a mini-workshop, and this would be the workshop you would be doing if you would not use cardboard engineering. Again, the goal is to come up with an useful solution that you think is superior to the other concepts and ideas you had. This can take around half a day, although it could be less for simple problems, or more for complex ones.

34.5.3 Building Cardboard Models (More Than Once)

Figure 272: Bean Bowling (Image Susan L. Davis in public domain)

Finally, you get to build the cardboard. Have the team create a cardboard model. This is usually quite fun, and can also be an excellent team-building exercise. This is also usually not challenging for the moderator, as the team pretty much runs by itself.

However, this is no time to idle. The moderator needs to think about all the possible problems, shortcomings, and flaws of the model. You may have to force yourself, but rip the model apart mentally – because the team may have to do just that later on in reality. After all, **the goal is not to just put something together but to try out different options**, and the current solution can always be improved.

So, after the model is built, have your team look for flaws and things that are less than optimal. I do not mean the quality of the model – it is temporary anyway – but for things that are not perfect or that can be done better. This may be difficult for the team, since the cardboard model is now their baby, but we are looking for a prettier baby now! If it is the optimization of a work station, have a worker mock-assemble a product a few times to see what is good what not.

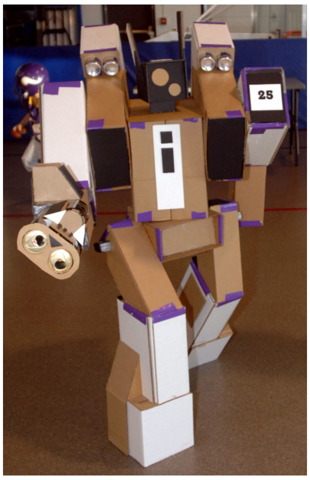

Figure 273: Cardboard Robot US Air Force (Image Susan L. Davis in public domain)

After this discussion, the new ideas are put into cardboard. Do not hesitate to build a completely new model. Especially for workstation optimization it is not a problem to build a second model next to the first one. For smaller changes you may also simply fiddle with the existing model. For a layout optimization you may have to rip out entire cardboard machines and build them again. If you change or destroy a model, don't forget to document it before you change it!

As before, while the team is building, think about all the sub-optimalities of the second model. While it is (hopefully) improved, it still can be done better. We are still looking for a prettier baby. Unless you are running out of time, you can now discuss this with the team, and build a third model, similar to the process leading to the second model. Sometimes one model is good enough, but usually two or three models give better solution. Rarely do you need more than four models. Because, again, **the goal is not to just put something together but to try out different options!**

As for the time required, I estimate how long it takes to make one cardboard model and multiply by three. Granted, the second and third model will probably go a lot faster, but you also want to have discussion time in between to consider new ideas.

34.5.4 Wrap-Up

Now you have multiple models. Pick the one you and your team consider to be the best. If you have more final ideas, feel free to create a final tweak of the model. Sometimes this may also be a zombie consisting of part of one model and another part from another model. Whatever works best is the way forward. You also need probably at least one hour to document the model and to transfer the information to the people who do the real construction, possibly quite a bit more.

34.6 Summary

So this is the approach to use cardboard engineering in lean. Did I mention that **the goal is not to just put something together but to try out different options?** This is probably the most common pitfall of cardboard models, where the joy of having a nice model is valued more than having the best possible solution. We don’t care about the model; it is going to be ripped apart and trashed soon anyway. The solution, however, will stay with us after it is implemented. Hence, **the goal is not to just put something together but to try out different options!**

In my next post I will show you a few alternatives and variants with different materials, from ultra-low-cost to fancy, shiny stuff. Until then, **go out, try out different solutions, and organize your industry!**

35 Cardboard Engineering – Alternatives

Christoph Roser, August 27, 2019, Original at [*https://www.allaboutlean.com/cardboard-engineering-alternatives/*](https://www.allaboutlean.com/cardboard-engineering-alternatives/)

Figure 274: Children with Cardboard Box Toys (Image CHOReograPH with permission)

Cardboard Engineering is a quick way to try out different configurations in reality. As the name says, this is done using cardboard. However, there are a few alternatives. Let me show you a portfolio of different ways to make cardboard models with (not only) cardboard, from ultra-cheap to very fancy. Please note that the fancier methods are usually not so well suited for layout optimization, but more for workstation optimization.

35.1 Sidewalk Chalk

If you are doing a layout workshop, a very cheap and quick alternative is to mark on the floor using chalk. Normal blackboard chalk is too brittle here, but for very little money you can get sidewalk chalk (or you borrow it from your kids).

Rather than building a cardboard model, you simply draw the future layout directly on the shop floor. This is a quick and easy way to visualize the future shop floor layout. Use different colors, and label the items. Also, for obvious reasons do not use it in a clean room or generally on clean floors! In this case painters tape maybe an alternative.

35.2 Lumber

Figure 275: Hermann Ultraschall Wood Mock Up (Image Hermann Ultraschall with permission)

A very cheap and readily available material is lumber, boards, and other wood products that can be assembled into different shapes. The images here are courtesy of [Hermann Ultraschall](https://www.hermannultraschall.com).

This material is readily available in pretty much all hardware stores, and it is likely that at least some of your team has some experience assembling wood. Hence, it permits you to create quick mock-ups of your workstations.

Of course, you need some more tools like a wood saw (hand or electric) and a screwdriver with bits and screws (and I heavily recommend a cordless one over non-cordless or even manual screwdrivers). For corner connectors you can either improvise with screws or also get some cheap wood connectors made out of perforated sheet metal. It is also possible to use a nail gun, but please be aware that nail guns are a major accident hot spot, especially for people who use one only occasionally.

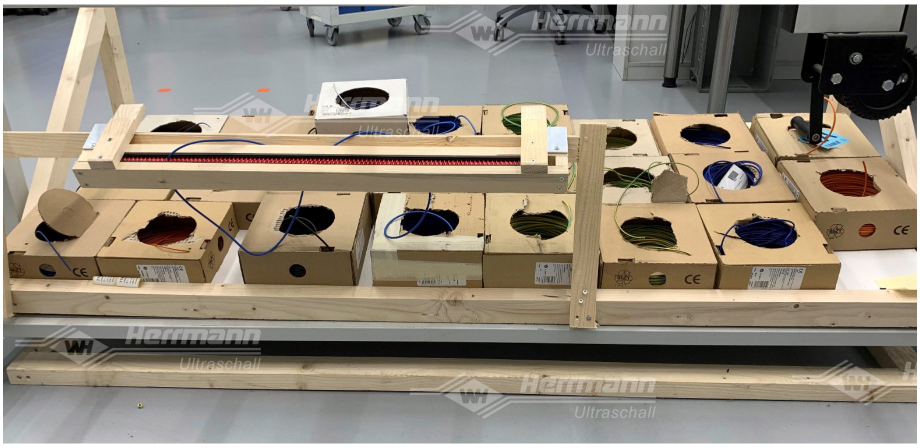

Figure 276: Hermann Ultraschall Wood Mock Up (Image Hermann Ultraschall with permission)

35.3 ASSTEC Cardboard Engineering Kits

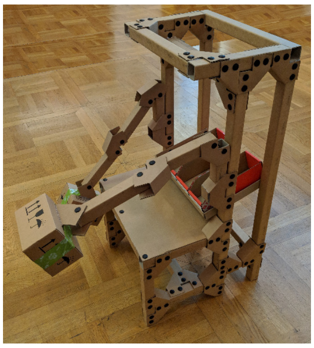

Figure 277: ASSTEC Cardboard Robot (Image Roser)

You can also buy quite nice pre-prepared cardboard items designed especially for cardboard engineering. I got a free sample of such a product from ASSTEC in Germany. Their kits consist of 45x45mm rigid cardboard tubes, corner connectors, plugs, and flat cardboard plates. The photo here shows a cardboard robot that was presented by ASM at the Lean Conference in Frankfurt 2018. On the ASSTEC website you can see photos of much larger installations. The result is actually quite sturdy. On one of their images I saw a young lady sitting on such a structure – although this lady probably weighs significantly less than me.

The photo below shows the details of these black plugs in open and closed positions, both on their own and from the inside of a tube. These plugs can be reused for the next project.

Figure 278: ASSTEC Plug System (Image Roser)

Figure 279: ASSTEC Joints (Image Roser)

The photo here shows the unassembled and assembled corner pieces. These cardboard pieces are pre-cut and notched for easy folding and assembly. You need a **drill** with a **wood drill bit** to drill the holes into the cardboard tubes. You also need a **wood saw** to cut the profiles to size. The flat cardboard sheets were also notched to make folding easier. As the robot above shows, you can also make movable joint connections.

The end result looks quite impressive, and has the potential to impress your boss and your customers. They are almost too nice to throw out. It does take a little bit longer to build, and it requires also a bit more engineering skills, as you have to measure and cut to size, whereas with normal cardboard you often just "wing it." However, it looks much nicer. Below is their product demonstration video.

35.4 Plastic Tubes, Aluminum Profiles

The hardware store also sells other options for a quick mock-up. For example, you can buy cheap aluminum profiles with corner connectors to also create a mock-up quickly. Price-wise they are comparable to the cardboard profiles. They do require a bit more engineering thought for the lengths of the profiles to match, and sawing and drilling them is also a bit more difficult than the cardboard profiles from above. Another flaw is that they may look too permanent. You have to consider the reaction of management and workers to a permanent looking mock-up. It may be not a problem, or it may raise eyebrows about the perceived cost.

Figure 280: Hardware Store Profile and Connectors (Image Roser)

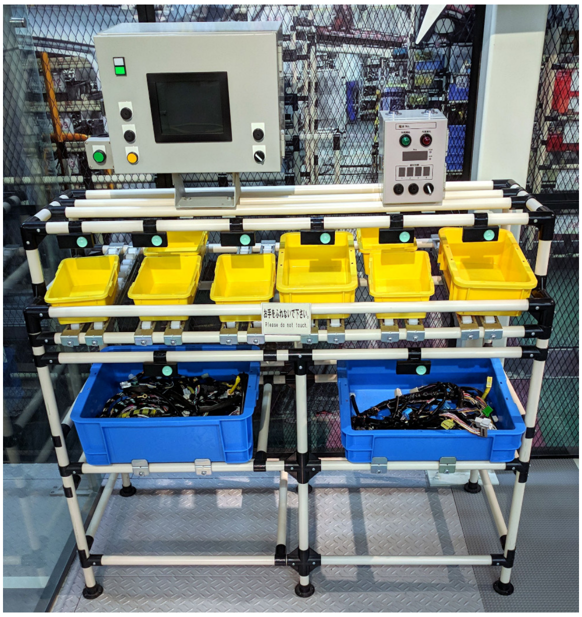

Figure 281: Suzuki Pipe Structure (Image Roser)

Similarly, there are also systems using inexpensive tubes and matching connectors. These are also sturdy enough for permanent structures, and fall somewhere between a mock-up and a permanent installation. The image on the left shows a permanent testing station from Suzuki in Hamamatsu, Japan.

Figure 282: Rexroth Profiles (Image Roser)

The gold-standard for permanent construction is of course dedicated aluminum profiles. One common maker is for example Bosch Rexroth. The image on the right is from the Bosch section in Arena 2036. However, since these special profiles cost three times as much as a simple aluminum or cardboard tube, it would be an expensive mock-up. It is best used for permanent designs – and of course you can make a similar workshop to directly create a final design rather than creating mock-ups. These profiles are widely used in Europe, but Japanese companies often go for cheaper round profiles.

35.5 EverBlocks

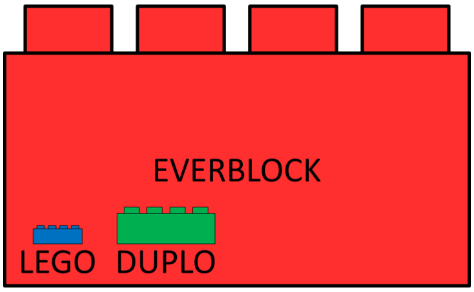

Figure 283: Lego, Duplo, and Everblock Size Comparison (Image Roser)

A new concept I recently came across is a Lego-like structure. You probably know LEGO toys, and can imagine that building a real-life structure from LEGO bricks is cumbersome. Even the next size up DUPLO would require still too many blocks. However, EverBlock makes LEGO-style blocks in 12x6x6 inches (roughly 30x15x15 cm).

Figure 284: Photo EverBlock (Image Everblock with permission)

While I haven't tried it myself yet, these blocks are designed for assembly of usable structures. Their portfolio pictures show desks, counters, room dividers, and many other examples. Hence I imagine they can also be used to create engineering mock-ups – which are also sturdy enough to be used. The system also includes flat surfaces and shelves. The blocks also have through holes that can be used for electric wiring or stability enhancing dowels. If necessary they can also be drilled or screwed into. Construction is – literally – child's play and very easy. After use they can also easily be disassembled. Long-term use may depend on your workshop environment, and I would not use this for example in a welding shop. But then, I wouldn't really use cardboard either.

A single 2×4 block costs around $7, or $5 for recycled plastic if color is not that important to you. A comparable structure made from EverBlocks is by my estimation probably a tad more expensive than the cheap cardboard or aluminum profiles, but cheaper than a Rexroth profile construction. You may need a few blocks for larger projects, but they sell anything from single blocks to 8,000-block truckloads.

35.6 Computer, CAD, and Virtual Reality

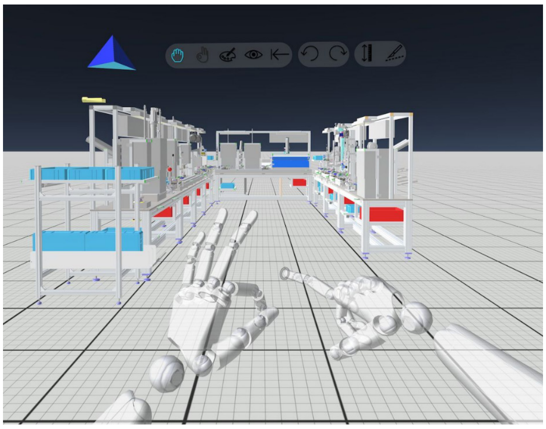

Figure 285: Screenshot of the R3DT Virtual Reality Environment (Image R3DT with permission)

Finally, there is the option of computer and virtual reality. You can create and change models in CAD – although this requires people skilled in the use of your program. Another option I tried out recently (and on which I will write a bit more soon) is a virtual-reality tool for shop floor design reviews by R3DT. Their product creates a virtual 3D environment using VR glasses like Oculus Rift and HTC Vive. Merely by moving your hands without any controllers you can experience, interact, and modify a virtual shop floor.

All modifications with computer models have the drawback of requiring a computer model, which may take a lot of time up-front. However, in many cases such a model already exists, which makes it a lot easier.

35.7 Which One to Use?

Which one you should use depends on your situation. If you have lots of lumber lying around anyway and your people know how to handle it, use lumber. If it should look extra nice, use Asstec cardboard profiles or aluminum profiles from the home depot. If your people feel comfortable with computers, you may use CAD or a virtual reality environment. And don't forget that you can **combine** these approaches. Make the frame from lumber, and use cardboard for the gadgets attached to the frame. Take whatever feels best for you.

Regarding cost, below is a rough estimate of the cost for a single 2m-long profile or a similar equivalent to give you an idea of the magnitude of the cost. This does not include any connectors or joints. For EverBlock I did an estimate for a 60×90 wall vs. a similar wall made from other materials to place it roughly on this chart. Of course, for all these there may be volume discounts or bulk orders possible. The CAD and Virtual Reality model is not on the chart, since you usually pay per seat, and the cost per model depends on how much you use it.

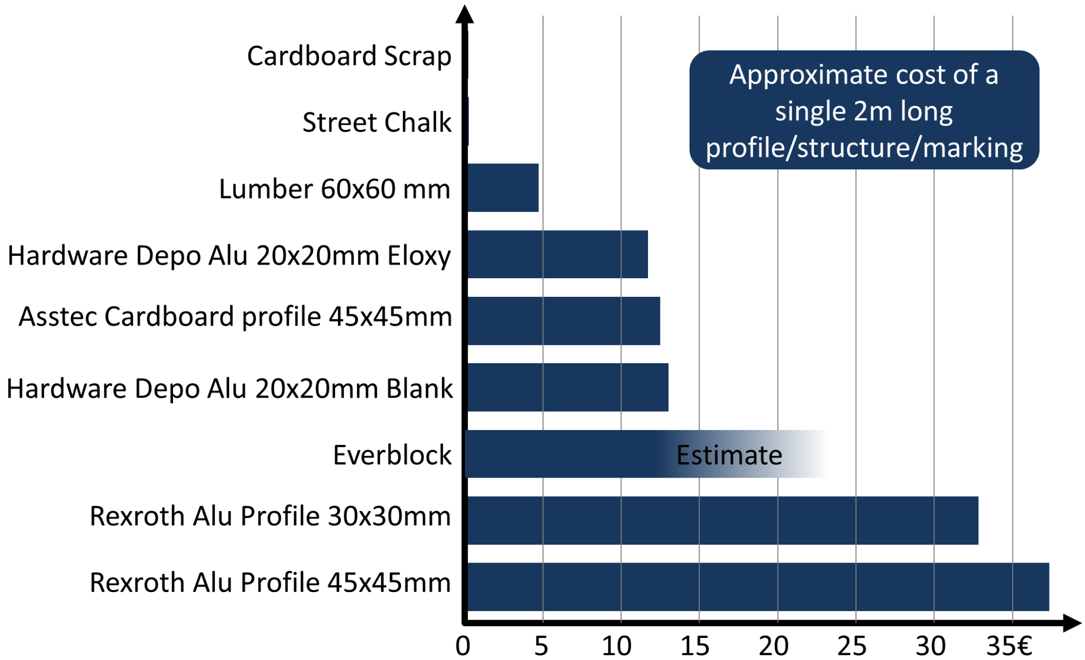

Figure 286: Cardboard Engineering Cost Comparison (Image Roser)

Overall, there are many different alternatives to cardboard engineering, and most of them can be combined. Hopefully this includes something that you like and can use for your work. Now, **go out, cut up some cardboard, or lumber, or aluminum profiles, or fire up the computer, and organize your industry!**

36 Happy 6th Birthday, AllAboutLean.com

Christoph Roser, September 01, 2019, Original at https://www.allaboutlean.com/6th-birthday/

Figure 287: 6th Birthday Cake (Image Jane Biriukova with permission)

Oh my gosh, it is six years already and 317 blog posts! Like clockwork every week, one blog post with at least 1,000 words. Time to celebrate again! Many thanks to all for reading and commenting. I am looking forward to keep this up for many more years to come.

36.1 Most Popular Posts

Figure 288: Top 10 (Image Roser)

The top ten blog posts of this year were in ascending order:

- How Many Kanbans? – The Kanban Formula Part 1 with 9,155 clicks.
- All About Spaghetti Diagrams with 10,066 clicks.
- All About Swim Lane Diagrams with 10,451 clicks.
- Introduction to Karakuri Kaizen with 11,159 clicks.
- Line Layout Strategies – Part 2: I-, U-, S-, and L-Lines with 12,828 clicks.
- How to Measure Cycle Times – Part 1 with 13,138 clicks.
- What Is Your Production Capacity? with 16,195 clicks.
- Visual Management with 16,336 clicks.
- Glossary of Lean Production Related Terms with 17,523 clicks – this one now has over 430 entries, and more are in the pipeline.
- The (True) Difference Between Push and Pull with 22,776 clicks – this one is quite popular and made the Top 10 list pretty much every year since I wrote it in 2015.

The popularity of my blog was also steadily increasing, and I was approaching almost 3,000 clicks per day in February, being among the top 350,000 websites in the United States according to Alexa.

Another blogger somewhere wrote that in order to be a successful blogger, you should spend 20% of your effort on creating content, and 80% of your effort on promoting your content. However, I don't necessarily want to be a successful blogger but primarily a good blogger, and for me it is almost in reverse. I spend 80% of the time creating content, and 20% of the time promoting it. Hence, I will just keep on writing, because I really enjoy writing about lean.

36.2 Upcoming Stuff

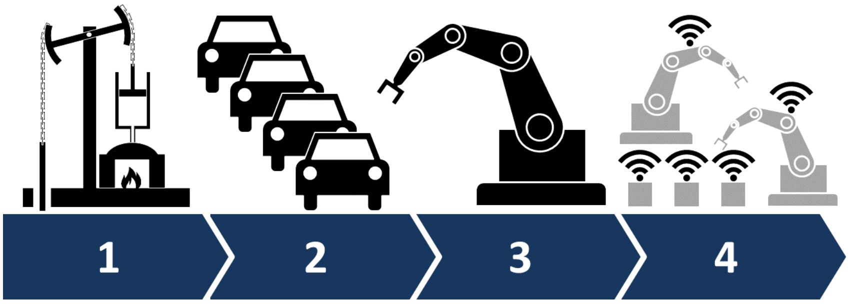

Figure 289: Industry 4.0 (Image Roser)

At the beginning of July, a couple of friends and I organized a non-profit tour through south Germany to study Industry 4.0. This of course resulted in a series of blog posts on the state of Industry 4.0 in Germany. This series with six or more posts will kick off next week on Tuesday, September 10, and take you through Bosch, Kärcher, Trumpf, Siemens, ABB Stotz-Kontakt, Audi, and many more. I hope you will like it.

AllAboutLean.com

Figure 290: All About Lean Logo (Image Roser)

I am also working on a series of collected volumes based on my blog. These are nothing else but my blog posts collected in the form of a book, with one volume for each year. Since my blog is already available for free, I will also provide a free download of the PDF and Ebook files on my blog. If you want a paper version, it will be available by print-on-demand through Amazon, although this won't be free. The idea is that some of you may want to download the content of my blog to have it on your own hard disk.

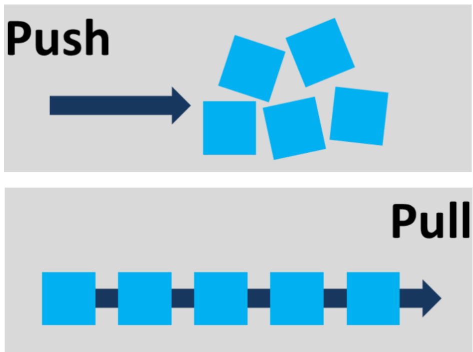

Figure 291: Somewhat misleading push pull illustration (Image Roser)

I am also working on a second series of books on different lean topics. While they are also still based on my blog posts, they are heavily edited, lots of content added, and put into a consistent structure so that it makes a consistent reading. The first book will be on pull production. Hopefully it will be ready next year (but then, I said this already last year …). Anyway, it is in the works.

36.3 Awards and Praise

Figure 292: Top 20 Lean Manufacturing Blogs Award (Image Feedspot with permission)

In February I received a very nice email from Anuj Agarwal, founder of Feedspot. They curate a list of Lean Manufacturing blogs, and AllAboutLean made it into the Top Twenty. Check out their list of the top 20 manufacturing blogs for other great lean blogs!

Figure 293: Kanbanize Logo (Image Kanbanize for editorial use)

Kanbanize selected 19 lean blogs worth reading, and, guess what, I made the list :). They are a software provider for lean and agile project management.

I also received lots of positive comments from you on the blog in general or articles in particular. Below is just a small selection

> *Excellent article! The writer walks the talk. He did a good job in putting credible details for the reader's appreciation. Thanks! (Rey Elbo at Pay Attention to Details – Operator Training at Toyota and Scania)*
>
> *Great reading and back ground (Terry Brook on TWI)*
>
> *Thank you for this excellent article. I learned something new with each job that you published. Congratulations. (Pedro Chipana at Hoshin Kanri)*
>
> *I want to praise everything you write on about on your blog, I am constant reader and always looking forward to new posts (Dimitrije on Maintaining Weak FIFO in Parallel FIFO Lanes)*

I could go on and cite many more comments. Your comments definitely motivate me to keep up the writing. Many thanks to all for commenting here, on LinkedIn, or elsewhere.

36.4 Odd Requests…

On the other end, popularity comes at a price, and I am getting repeated requests from others to write articles for my blog – although I don't really see how e.g. a travel report would fit into lean manufacturing?!?! In any case, I decline all such requests to keep the quality high (or at least I hope to do so). Everything written here (for better or for worse) is by me. Only the spelling is checked externally, because this is definitely not my strong point.

Figure 294: Toyoda Model G Automatic Loom (Image Roser)

A while back I also wrote on the historic Toyoda Model G loom, a museum piece that was groundbreaking in 1920, but has been long since overtaken by better products. Since then I have been getting repeated requests from India if they can buy it from me, that I should send a catalog with loom models and spare parts, and prices in rupees, and apparently they seriously want to buy an almost 100-year-old loom to start production. Highly interesting. In any case, **I have no looms to sell, please stop asking!**

36.5 Summary

Overall, another successful year of blogging comes to an end, and I am looking forward for the next year. As of now I am in no danger of running out of topics. My "idea for a blog post" list has way over 200 entries. This should last another four years, and the list is usually getting longer rather than shorter. Anyway, many thanks for reading and commenting, keep on reading and commenting! Now, **go out, use lean in whichever way fits your situation, and organize your industry!**

37 Industry 4.0 Tour in Germany – A Van Full of Nerds – Overview and Audi

Christoph Roser, September 10, 2019, Original at https://www.allaboutlean.com/i4-0-tour-germany-1/

Figure 295: Van of Nerds at Bosch Reutlingen (Image Roser)

Recently I organized a non-commercial **Industry 4.0 tour** for some friends through my university, the Karlsruhe University of Applied Science. For the first week in July 2019, we rented a van and toured southern Germany. We visited fourteen different locations in five days to understand the current state of Industry 4.0 in Germany. Almost all of these locations were Industry 4.0 award-winning enterprises. However, our assessment of Industry 4.0 often differed from these awards. Since we all come from the lean corner, we often have a different outlook on things than people who specialize in Industry 4.0. Let me give you an overview of our tour:

37.1 The Van Full of Nerds

37.1.1 The Nerds

Figure 296: Van of Nerds at Dinner in the Alte Brauerei Weingarten (Image Roser)

The group consisted of eight people from all over the world coming together for one week of studying Industry 4.0 and having fun. We called it a “van full of geeks” until one member pointed out that geeks are wanna-bes and nerds are the real thing, hence we renamed it to a “van full of nerds.” If you work in lean, you surely will recognize some of these names. We are (in alphabetical order):

- Michel Baudin (Lean Expert and Consultant, USA)
- Prof. Hironori Hibino (Tokyo University of Science, head of Japan Industry 4.0 government group, Japan)
- Dr. Kai Lorentzen (Senior Product Manager I4.0 Bosch, Germany)
- Prof. Torbjörn Netland (Production Professor at ETH Zürich; Switzerland)

- Dr. Ralph Richter (Retired Plant Manager of Bosch, Production Researcher, Germany)
- Franck Vermet (Production Systems Mentor)
- Mark Warren (Lean Expert, Researcher, and Historian, USA)
- Prof. Christoph Roser (Lean Professor, Author of this blog, Germany)

37.1.2 The Plants

Our week was quite packed with stops. Altogether during this fun week, we had fourteen stops in five days on our itinerary. Seven of them were plant visits, and another seven were presentations, try-outs, and demonstrations. Hence we had almost three study events per day! Here is the list of plants we visited in southern Germany to study Industry 4.0 (in order of the visits):

- Bosch Wafer Fab, Reutlingen: Wafer factory producing sensors, won the 2017 Industry 4.0 award for consistently networked factory
- Kärcher, Winnenden: Produces floor cleaning machines of all sizes and types
- Bosch, Feuerbach: Huge plant, we looked at the Bosch Connected Industry and the Nexeed Transparency Kit
- Trumpf, Gerlingen: Smaller plant of Trumpf, making punching tools for sheet metal processing
- Siemens, Amberg: Famous plant making programmable logic controllers, won the 2018 Industry 4.0 award for smart factory, and many more
- ABB Stotz-Kontakt, Heidelberg: Making fuses, won the 2016 Industry 4.0 award for automation and networking of production to control the increasing number of variants
- Audi, Neckarsulm: Producing the higher end cars of the Audi brand (e.g., A7, A8)

Most of these were near Stuttgart, but we drove up to three hours to see Siemens, the plant farthest away from our hotel in Stuttgart.

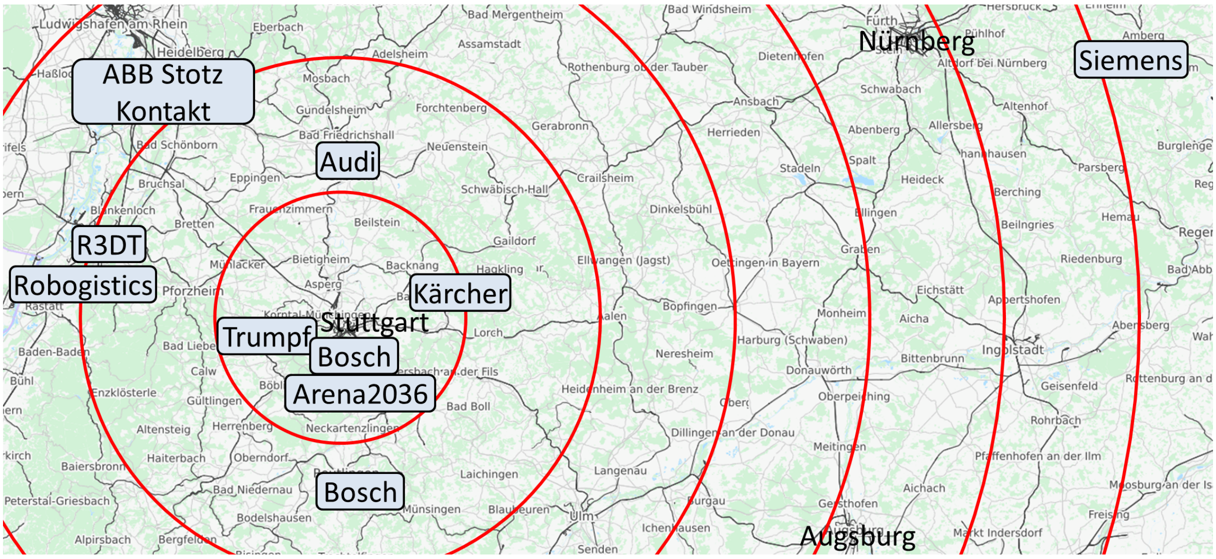

Figure 297: I4.0 Tour Map – A Van full of Nerds (Image OpenStreetMap contributors under the CC-BY-SA 2.0 license)

We especially liked the ABB-Stotz Kontakt plant, which we thought had the best Industry 4.0 approach, with a particular focus on finding out what really is useful for their factory. But overall it was interesting to see the different approaches used by the different factories. All of them had their strengths and weaknesses. In this series of posts I will focus on the positives, and glance over the weaknesses, since it is impossible to make a complete assessment of the plants.

37.1.3 Other Events

We also had a few presentations, demonstrations, and try-outs. Below are these visits (in order of the visits):

- Bosch at the Arena 2036, showing us their vision of the factory of the future
- Drag&Bot at the Arena 2036 for programming robots
- NAiSE at the Arena 2036 providing sensors for tracking intra-logistics
- ThingOS at the Arena 2036 providing connectivity for Industry 4.0, smart home, and smart retail
- Robogistics Laboratory at the Karlsruhe University of Applied Science, researching at the cutting edge between Robotics and Logistics
- R3DT providing virtual reality for work and assembly planning including ergonomics
- Klingelnberg: Maker of machines for gear cutting, won the 2016 Industry 4.0 award for the introduction of a cyber-physical production system in bevel gear production

Figure 298: Rothenburg ob der Tauber (Image Berthold Werner in public domain)

As you can see, it was quite a packed week with fourteen study events – and we even fit in a visit to the quaint medieval town of Rothenburg ob der Tauber. Overall it was a very educational and fun week. In the following I will look at some locations in more detail and show you what we saw, what we liked, what we didn't like, and why.

This starts a series of blog posts about what we saw and what we learned. While I was influenced by and received input from the other tour members, these blog posts are written by me and hence may reflect my opinion more than the others, although on many points we agreed. In this first post, besides this introduction, I will give you a (very brief) look into Audi in Neckarsulm.

37.2 Audi Plant, Neckarsulm

Figure 299: Audi Forum Neckarsulm (Image Joachim Köhler under the CC-BY 3.0 license)

Our last stop on the tour was the Audi plant in Neckarsulm. This tour was one of the normal tours for tourists and customers picking up their new Audis, so it was not an Industry-4.0-specific tour. Hence there won't be much detail. Subsequent blog posts on other plants will have much more details and information.

Anyway, the plant was very clean and highly automated, with 2,200 industrial robots. The mobile robots usually had nicknames to help their integration with the human workers. A car is made every 162 seconds in two shifts per day. The teams had around six to eight team members for every team leader, hence nicely small teams. As the City of Neckarsulm grew around it, the plant suffered from a lack of space, so they built upward with production on three levels. Most of us had a good feeling about the plant. The presentation was very good by a highly enthusiastic presenter, but since it was geared to a general audience (mostly people picking up their new Audis), it did not contain much information on Industry 4.0.

37.3 Summary

So this is the first post on our van full of nerds. In my next post I will present ABB Stotz-Kontakt in Heidelberg. While all plants had strengths and weaknesses, we liked this plant best. Until then, stay tuned, and **go out and organize your industry!**

P.S.: Thank you very much to everybody who hosted us and showed us their plants and products!

38 Industry 4.0 Tour in Germany – A Van Full of Nerds – ABB Stotz-Kontakt

Christoph Roser, September 17, 2019, Original at https://www.allaboutlean.com/i4-0-tour-germany-2/

Figure 300: ABB Logo (Image ABB for editorial use)

In this second post of the series on our van full of nerds touring southern Germany to study Industry 4.0, we will look at ABB Stotz-Kontakt in Heidelberg. Quite an impressive plant! Let me show you:

38.1 ABB Stotz-Kontakt Plant in Heidelberg

38.1.1 The Plant

Figure 301: Van of Nerds at ABB Stotz (Image Roser)

Out of all seven plants we visited, I liked the Industry 4.0 aspects of the ABB Stotz-Kontakt plant in Heidelberg the most. While some plants go overboard with Industry 4.0, at ABB Stotz it seemed to be a good and productive discussion about when and where to use computer. For example, the plant manager prefers digital dashboards, whereas the production manager prefers paper-based dashboards. Rather than just overruling the production manager, there seems to be a constructive discussion. While I prefer paper, my hope is not that "one side wins," but rather both sides come to an agreement.

Figure 302: ABB Circuit Breaker (Image ABB under the CC-BY-SA 3.0 license)

The ABB Stotz-Kontakt plant dates back to the *Stotz und Cie* company from 1891. While they are part of ABB since 1988, the people still feel that they are Stotz first and ABB second, and keep their distinct identity.

It is a plant producing circuit breakers near Heidelberg. They have around 1,200 employees. The first good sign of correct priorities was their safety instructions. While most plants we visited glossed over them, this company spent three to five minutes actually covering the important parts (including "hands on railing" on stairs, and "no mobile phone use while walking"). Also, all employees have a minimum of a Lean Six Sigma white belt. In addition, most of the employees have yellow, green, and some a black belt. While I am not a fan of Six Sigma, it is basically a quick fundamental lean training. They have two main assembly lines. The lines run twenty-four hours a day six days per week.

38.1.2 Lean Aspects

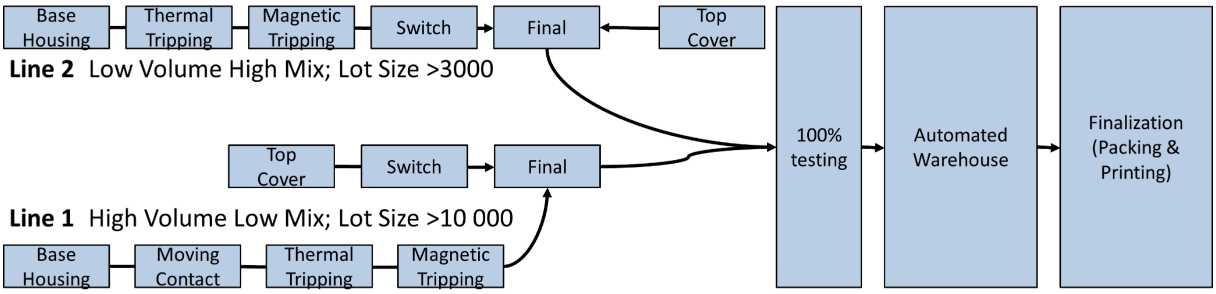

Figure 303: ABB Stotz-Kontakt Layout (Image Roser)

Line 1 is a fully automated one for high-volume, low-mix production. About 160 variants are produced on this line. Production is economic for lot sizes above 10,000 circuit breakers. Line 2 is a mix of automated and manual line for low-volume, high-mix production. This line handles 450 different variants and is economic for lot size 3,000 and up. The total changeover time of line 2 is about 90 minutes, but due to a running changeover, the lost production time is only 4 minutes. Some labor-intensive parts are received from a second plant in Bulgaria (low labor cost). Line 2 has a cycle time of 1.2 seconds, while line 1 has stations doubled resulting in a cycle time of 0.6 seconds.

They also use a 3D computer tomography machine with a resolution of 0.1µm resolution to analyze circuit breakers without opening them.

Figure 304: Stack Light (Image Ktm250-1150gs under the CC-BY-SA 3.0 license)

Compared to some other automated lines in other plant, the workload of the employees was appropriate (another company we visited had employees wasting a lot of time with idling, waiting, and other wasteful things). The material flow was very well visualized. They used stack lights extensively, but did not go overboard with too many stacks. Most machines had a stack light with only one color, some two, and very few three.

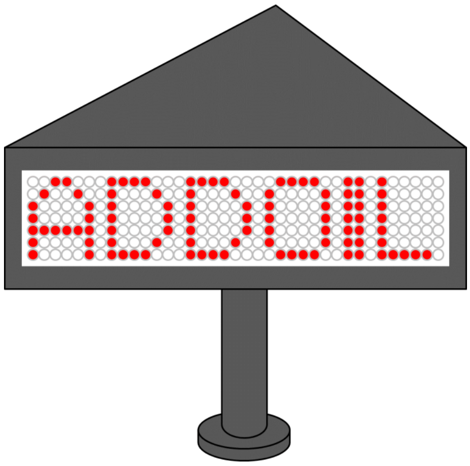

Figure 305: ABB Stotz-Kontakt LED Text Band (Image Roser)

For additional visualization they also had a triangular LED running text that gave more details on the error through an error code and an error text. This way it was already visible from far away what the problem at the machine was.

Figure 306: ABB YuMi Robot (photo not at Stotz-Kontakt Heidelberg) (Image Anthony O'Neil under the CC-BY-SA 2.0 license)

They also used some ABB YuMi collaborative robots (Cobots) to pack breakers into blister packs for retail. These robots were of an unusual design, having two arms per robot. Even though these are designed for collaboration and stopped when touched, they were partially shielded at head height to prevent eye injuries, where even a slight touch can hurt quite a bit. Nevertheless, these Cobots managed to keep work in Germany that was already scheduled for Bulgaria.

38.1.3 Industry 4.0 Aspects

Their approach to industry 4.0 was very structured. Every project that they do needs a business case. They had quite a few projects, but structured them into Focus Projects; In Progress; Quick Wins, and IT Standards. These projects are also grouped into four areas: Energy Management; Product Improvement; Smart Devices; and Quality Improvement. Below is a selection of these projects.

- Focus Projects
- Smart Wall for Shop Floor Management: For daily shop floor production meetings
- Workforce Management Systems: Allocation of work, skill matrix
- 3D printing of Molds: Faster prototyping, cost reduction
- In Progress
- Energy Monitoring (Power): Monitoring, cost calculation
- Energy Monitoring (Air Pressure): Monitoring, cost calculation
- Standard common data layer and HMI for line/plant production data: Common data structure for easy analytics
- Augmented Reality for Remote Support: Internal maintenance, external support, training, etc.
- Quick Wins
- Swift Insight IoT Production Monitoring System: Analysis of production data
- Wearable Devices: Receive downtime information, support
- IT Standards
- Standard for IoT devices: Standard for IoT at ABB
- Network Infrastructure (WIFI): Stable WIFI in all locations
- Cyber Security: Reduce risk, reduce exposure

All of these had a clear structure and a business case somewhat similar to an A3, and also a rough timeline. The project was described; the scope defined; current situation was presented; the expected outcome was given; a cost/benefit decision matrix with the main cost drivers was made; and an overview of the benefits was provided. Since the business case for such projects can often not be calculated numerically, it was often only a tendency (e.g., transparency up; cost down…).

To show just a few examples: They are experimenting with a digital assistant for manual assembly of the most low-volume products. The product is by a German company Ulixes, and provides a projection on the work surface to assist with assembly. Their demonstration video is below.

However, the implementation is far from certain. The risk of this system is that the worker feels belittled and micromanaged, and his skills not valued. They are discussing this, working together with the employees to see if it is really a help or just an expensive nuisance.

Figure 307: HoloLens Demo (Image Hoshinim under the CC-BY-SA 4.0 license)

They are also experimenting with Microsoft HoloLens for remote support and maintenance, as for example with their plant in Bulgaria. HoloLens are augmented reality glasses where you can add visualization of digital data on top of the real world environment. As Ulixes, this is in testing at ABB to see if it is truly useful or not (yet).

Again, the strong point we saw at ABB was this hard work to find out if a technology is really useful before it is rolled out on the shop floor. If you skip that, then you will waste time and energy. Overall, we all liked the plant very much, and many of us considered this to be the best plant in the tour.

In my next post we will visit Trumpf. I have posted about Trumpf before, and am quite impressed with their approach to business. My previous posts were about the main plant in Ditzingen, and this time we visited a much smaller plant in Gerlingen making stamping tools. Until then, **stay tuned, go out, and organize your industry!**

P.S.: Thank you very much to everybody who hosted us and showed us their plants and products!

39 Industry 4.0 Tour in Germany – A Van Full of Nerds – Trumpf Gerlingen

Christoph Roser, September 24, 2019, Original at https://www.allaboutlean.com/i4-0-tour-germany-3/

Figure 308: Trumpf Logo (Image Trumpf for editorial use)

As one of the stops on our van full of nerds touring southern Germany to study Industry 4.0, we also visited Trumpf. However, we didn't go to the main plant in Ditzingen but to a much-smaller plant in Gerlingen that makes stamping tools. I have heard a lot of good things about Trumpf Gerlingen, and wanted to see if it is true. Let's find out:

39.1 Trumpf Plant in Gerlingen

39.1.1 The Plant

Figure 309: Van of Nerds at Trumpf (Image Roser)

Trumpf is a famous family-owned German machine tool maker. I like their lean approach (called Synchro) quite a lot, and also wrote about it on this blog. However, the focus of this visit was not lean but Industry 4.0. We visited the much-smaller plant in Gerlingen, which has only 75 employees (out of 13,000 worldwide). This plant has a good reputation for its highly automated manufacturing system.

Below you see two examples of their punching tools, with a multi-tool in the second row. The first column is the top (the punch), the second the bottom (the die), and the third a punched sample of sheet metal. They also have a carrier that is not pictured here.

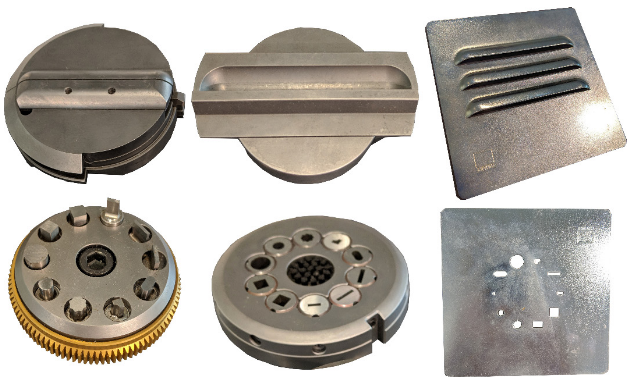

Figure 310: Trumpf Punching Tools (Image Roser)

What impressed us most about Trumpf Gerlingen was their clear purpose. They really asked thoroughly what they want to achieve, and then worked on achieving this goal. Okay, you may think this sounds easy, but it is not! Many of the other locations we visited rushed this step, and it often felt like they started with the solution (something with computers, obviously) and then looked for the matching problem. This is terribly wrong! At Trumpf Gerlingen, they asked, “Why do we want to do this?” and we as lean guys really like this approach. On a side note, I also saw this in their much-larger main plant in nearby Ditzingen, so it indeed seems to be part of the corporate culture!

39.1.2 Industry 4.0 Aspects

Anyway, their Industry 4.0 journey in Gerlingen started in 2009. Their clear goal was to get the part to the customer faster in order to stay competitive. Their motto was “The Customer starts the Machine.” In 2015 they moved to paperless manufacturing, adding a data matrix code to every part with the motto “The Workpiece as Information Carrier.” In 2017 they implemented the option to send in a photo of the part (the data matrix code) to reorder the same part again, with the motto “Photo Starts Machine” (although the customers don’t really use this much since the purchasing at the customer is usually far away from the part on the shop floor). This automated system works for the high runners and less-complex products (70% of volume). Their next goal is to extend the system to include more of the complex products, as for example multi-tool punches.

Figure 311: Selection of standard punch shapes (Image Roser)

Their system is quite nice. It all starts with their online shop for punching tools. With only a few clicks, registered Trumpf customers can select the shape, the dimensions, and other parameters of the tool. Only for more-complex tools do they need to interact with a Trumpf salesperson. Overall, this system allows for about 31 million possible punches, way too many for them to keep in stock. Once the customer clicks to order, the whole system goes into action. The order is added seamlessly to the Trumpf computer system. A punching tool consists of a punch at the top, a die at the bottom, and a stripper in between. These themselves can consist of multiple parts.

Figure 312: Trumpf Punching Tool Data Code (Image Roser)

Trumpf has 176 blanks for punches and 4 blanks for dies. Let's first look at the punches. Based on the order, a matching blank is selected. The first step is to laser engrave a data matrix code on the top of the blank. This code is unique and directly linked to the order. At every subsequent machine, the machine scans this code and automatically uses the correct program for this customer order, all of which in lot size one. No human interaction is needed except for bringing the parts to the machine. After the laser engraving, the punch goes to a grinding machine, which scans the code, selects the correct grinding wheel, and grinds the shape of the blank punch into the customer's desired shape. Another laser-engraving process adds text or an image, as chosen by the customer, to the side of the shaft.

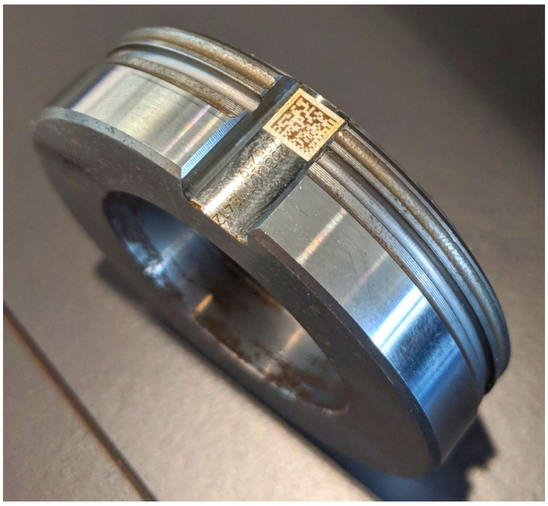

Figure 313: Trumpf Die Data Code (Image Roser)

There are much fewer variants of blank punches (around 4), and they are already engraved with a unique data matrix code. But here, too, a robot automatically selects the correct blank and adds it into an electric discharge machine (EDM). At the EDM, the data code is uniquely linked to the customer order, and again all subsequent machines simply scan the code and select the correct program. Overall, for standard tools the only human interaction between the customer order and the packaging is the testing and transport; everything else is fully automatic based on the data matrix code.

The results are impressive. Trumpf Gerlingen manages mass customization of their tools. If a customer orders a (non-complex) tool before 2:00 PM, the order is shipped on the same day. This is despite a huge fluctuation of orders, between 200 and 1500 per day. Including custom tools, 45% of the products arrive at the customer on the same or next day. Before 2009 it took them, on average, three days to produce an order; now it is down to less than one day. Personally, I was also surprised at the price, where punching tools can cost as little as €30. This gives them a market share of 70% for punching tools in Europe. In China they have only 5%–10%, since the Chinese market values price more than quality. They have around three days' worth of inventory of blanks in their warehouse.

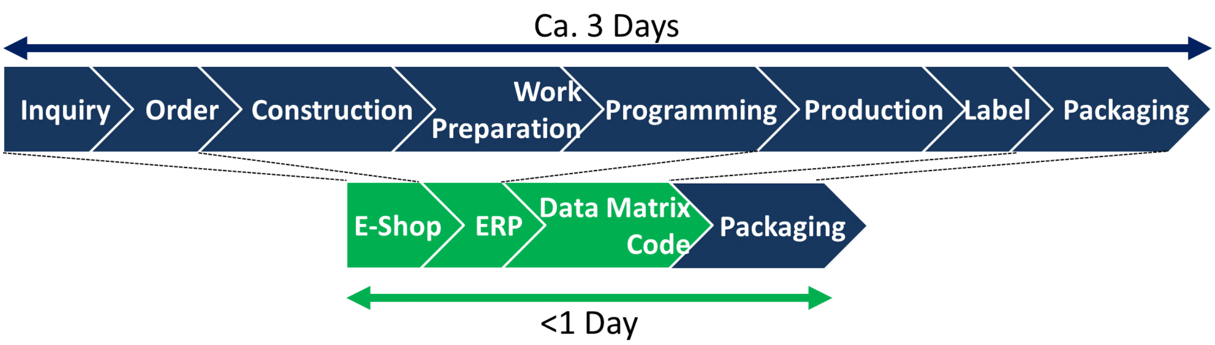

Figure 314: Trumpf Order Sequence (Image Roser)

39.1.3 Lean Aspects

Overall, we were quite impressed with their system, and also with their focus on performance rather than appearance. However, while the digital side works excellently, the flow of parts in the plant seems less structured, and cleanliness on a lower level compared to the other plants. Both material flow and cleanliness are much better in the main plant, Trumpf Ditzingen. Overall, many of us, including me, liked the Industry 4.0 aspects of Trumpf Gerlingen, especially their clear direction on what they want to achieve. They aimed to reduce the barriers between the customer and production, and made a quite seamless integration of the ordering process.

In my next post we will visit another impressive plant, the Bosch wafer fab in Reutlingen. Until then, **stay tuned, go out, and organize your industry!**

P.S.: Thank you very much to everybody who hosted us and showed us their plants and products!

40 Industry 4.0 Tour in Germany – A Van Full of Nerds – Bosch Reutlingen Wafer Fab

Christoph Roser, October 01, 2019, Original at https://www.allaboutlean.com/i4-0-tour-germany-4/

Figure 315: Bosch Logo (Image Bosch for editorial use)

The first plant we visited as our van full of nerds touring southern Germany to study Industry 4.0 was the Bosch wafer fab RtP1 in Reutlingen. Quite impressive. Let me show you what we found:

40.1 Bosch Wafer Fab Reutlingen

Figure 316: Bosch RtP1 Reutlingen (Image Roser)

40.1.1 The Plant

The Bosch wafer fab in Reutlingen produces – as the name implies – wafers. These are used on chip sensors. Probably 90% of all mobile phones in the world have a sensor from this plant. As with all wafer fabs, it is a highly complex material flow with around 800 steps to produce one waver, and the material flow doubling back on itself multiple times, including a loop over the Near East. There are multiple wafer lines in the plant, and we looked at the 200mm-diameter wafer line. They also have a 150mm line and a MEMS (micro-electro-mechanical-systems) line. The chips have between 7 and 34 layers. The smallest structures on the wafers are around 240nm, much larger than high-end computer chips with down to 7nm, but quite sufficient for sensors. Due to the larger resolution, the yield is also much better, around 94%–99.9% good chips. They also create micro-machines as part of their sensors.

One of the biggest challenges they faced was when they built the new 200mm wafer fab AND at the same time expanded into smartphone sensors AND at the same time started to use many external partners. Changing their business model in three directions at the same time (new plant, new suppliers, new customers) was very hard to do.

Like most wafer fabs, precision is crucial, and even small factors, like the rumbling of a truck on a nearby road or even the wind, can distort the results. Hence, the factory has a large foundation, and is practically a building in a building, where the outer building protects the inner one from the wind.

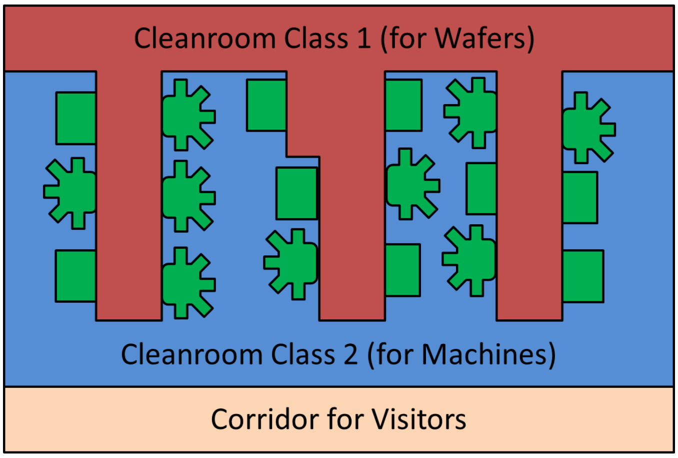

Figure 317: Bosch RtP1 Cleanroom (Image Roser)

The cleanrooms are actually split into two zones. The inner core, where the wafers are handled, is a cleanroom of class 1 with <10 particles of <0.1µm per cubic meter. Around this is the second zone where the machines are. This is a cleanroom of class 2 with less than 100 particles per cubic meter. The wafers are moved through openings between the two zones directly into the machines. This is shown schematically in the image here.

As far as lean goes, it was difficult to get a good understanding by just looking in from the outside corridor. Our presenter mentioned that lean and Industry 4.0 work well with each other.

40.1.2 Industry 4.0 Aspects

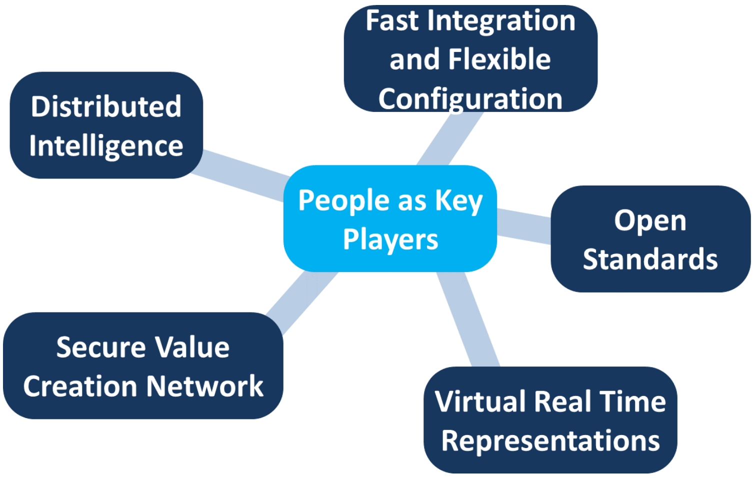

Figure 318: Bosch I4.0 Structure (Image Roser)

The Industry 4.0 model at Bosch has seven key elements as shown here. Six of them are arranged around the "People as Key Players" in the center, which was very important to the unions.

Obviously, an RFID chip is attached to pretty much everything, and at every step it is checked to be sure the products are what they should be.

The wafer fab has a high level of traceability, where every single chip can be traced back to where and when it was made. The decision to go for full traceability was made in 2017 to successfully to improve quality. They gave us one example where they noticed a slight increase in defects for chips. Based on their position of the wafer and the nature of the defect, they could track a likely source of a machine at a supplier in Near East, where a particular screw was probably misaligned. When they told the supplier, the supplier thought they are crazy … but they were right. This system generates around 50GB of data every day. Their philosophy is "As much information in the virtual world as necessary, as little as possible."

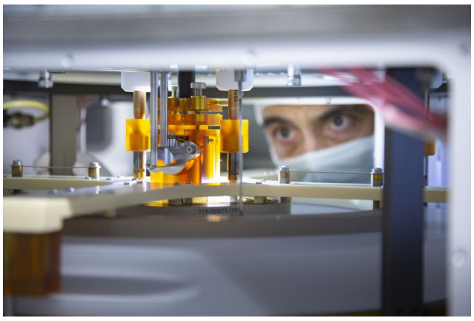

Figure 319: Reutlingen Wafer Fab (Image Bosch with permission)

They have an R2R (Run to Run) system where the machines communicate with each other to optimize performance. For example, if a machine detects that a coating is slightly to the thicker end of the tolerance limit, information is sent forward to the next machine to etch 1 second longer, and also relayed back to the previous coating machine to reduce the coating time by 2 seconds to adjust for this behavior. Implementing this was a lot of work, and the algorithms are hard-coded by a team of 4 people (not artificial intelligence … yet). Hence only the most relevant machines are included in this R2R system. Yet this greatly reduced the rework within the wafer fab.

They also use robotics extensively in the plant, with robots doing about 50% of the wafer handling. This was also a gradual change. They started with robots that were permanently mounted in one location. Then they introduced robots on sliding rails in 2014 ("HEROrails"). Finally they started to use fully autonomous robots in 2017 ("HEROfab") and self-navigating autonomous robots ("SCOUT"). To get the people used to the robots, they initially just had them move around on the main corridor. This allowed the human workers to build trust in the ability of the robots to stop and not run the human workers over. To further make them part of the family, they gave them human names (Margarete, Paula, Robert, Eva … the children of Robert Bosch).

Figure 320: Reutlingen Wafer Fab (Image Bosch with permission)

They also use RTD (Real Time Dispatcher), a mixture of rule-based and artificial-intelligence scheduling to optimize and prioritize the production schedule in real time. Goals include line balancing, customer priority, delivery dates, bottleneck management, and quality. This is quite tricky for the complex and iterating material flow in a wafer fab. However, they measured that human workers followed this sequence only 80% of the time and often used a different sequence than the one suggested by the AI. As it turned out, humans were able to see problems and understand the system sometimes better than the AI, and by changing the sequence outperformed the AI. Guess we humans are still able to compete with AI. They also need human ingenuity to drive continuous improvement, which computers also still cannot do.

Overall a quite nice plant with a lot of cutting edge technology. My next post will look at Kärcher, a German maker of home and industrial cleaning equipment, and Siemens. Until then, **stay tuned, go out, and organize your industry!**

P.S.: Thank you very much to everybody who hosted us and showed us their plants and products!

41 Industry 4.0 Tour in Germany – A Van Full of Nerds – Kärcher and Siemens

Christoph Roser, October 08, 2019, Original at https://www.allaboutlean.com/i4-0-tour-germany-5/

Figure 321: Logo Kärcher and Siemens (Image Kärcher and Siemens for editorial use)

As part of our van full of nerds tour through southern Germany to study Industry 4.0, we also visited two companies, Kärcher and Siemens. Siemens is probably well known to all of you. We went to their Amberg plant where they make programmable logic controllers. Kärcher is a smaller company that makes pressure washers and industrial cleaning machines. Let me show you what we found.

41.1 Kärcher Plant in Winnenden

41.1.1 The Plant

Figure 322: Van of Nerds at Kärcher (Image Kärcher with permission)

Kärcher is THE brand in Germany if you need any type of pressure washer or industrial cleaning device. From small handheld pressure washers for home use, to powered carts for industrial indoor and outdoor use, to specialized tools like ultra-high-pressure 4000psi washers to “wash the concrete off the rebar” during deconstruction and maintenance. I personally own their Window Vac to clean my windows. The family-owned company has around 13,000 people worldwide, of which 2,000 work in the headquarters in Winnenden, where they assemble mostly floor scrubbers and scrubber dryers, as well as do testing.

41.1.2 Lean Aspects

Figure 323: Kärcher Floor Scrubbers (Image Kärcher with permission)

The factory itself had, in my opinion, a good and clear material flow and structure, and I considered this to be a not outstanding but good lean factory (although one of the nerds in the tour disagreed with that). As far as Industry 4.0 goes, it is much less so. The high-volume, low-mix line is considered their "Industry 4.0" line and uses Pick by Light and digital instructions. I estimated the value-added time of their workers at around 50%–55%, which is pretty good considering the low cycle time and high variety. The lot size is one. Even if there are larger orders, they found it better for the system to distribute them into lot sizes of one. While mostly running in two full shifts, if demand is lower, a half-sized team at the night shift can use pre-assembled parts from the feeder-line buffers. The feeder lines work faster, and produce in one shift what the main line consumes in two shifts.

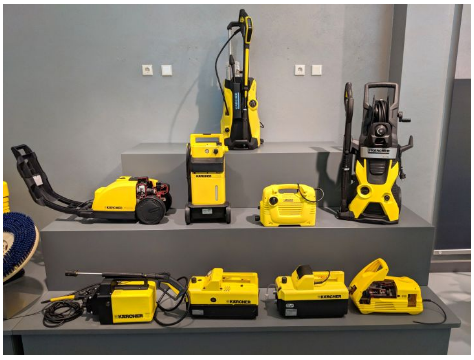

Figure 324: Kärcher Product Line Up (Image Roser)

The lower-volume, high-mix lines don't really use Pick by Light or digital instructions. All lines used an Andon (although mounted in a way so that visitors can't see it as per agreement with the unions). The plant has the ability to produce many thousands of variants, although the Industry 4.0 line is assigned to only around 85,000 variants. The team structure is nice and small, with around 6 people for every shift leader. They are trying to establish a central database to collect production data, but are handicapped by different MES systems at different machines. Unsurprisingly for a company specializing in cleaning products, the plant looked clean.

Overall nice but not very computerized … but this is fine by me. It is completely justified NOT to use Pick By Line if you have a high-variety, low-volume production and each "pick" box costs 50€ just for the gear, not to mention the effort of setting it up. Don't use computers just because it is possible to do so; use them when it is the best solution for your case.

41.1.3 Industry 4.0 Aspects

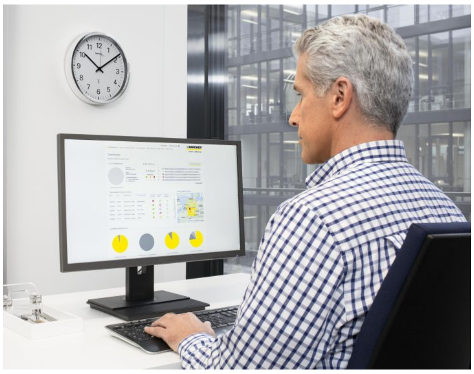

Figure 325: Kärcher Connected Industry (Image Kärcher with permission)

What I liked much more was the use of modern networking technology in their products. Most of their larger machines include a worldwide data SIM card, allowing them to do all kinds of tricks like tracking (Is you employee using your machines to clean somewhere else for his own profit on the side?); prevent this through Geo-Fencing (won't start outside of a permitted area); remote maintenance and diagnostics (Time for new brushes?); battery information (Do you need to charge?); and even location detection and prevention of starting the machine (one machine stolen in Europe was located in Morocco, shut down, and with the help of the local police brought back to its rightful owner). Their machines with a SIM card can do pretty much everything a mobile phone can do.

Overall, most of us (including me) liked the plant for having a level head and good common sense on what to do and what not to do.

41.2 Siemens Plant in Amberg

Figure 326: Siemens Plant Amberg (Image Roser)

The Siemens plant in Amberg is well known for its Industry 4.0 implementation and has won multiple awards, including the 2018 Industry 4.0 award for smart factory. It produces programmable logic controllers (e.g., the SPS7) and is widely known for its Industry 4.0 approach. They make 60,000 products per day in 120 variants, having around 350 changeovers per day. Their defect rates are around 10 dpmo (defects per million opportunities) (not to be confused with the more stringent measure of parts per million PPM).

Figure 327: Siemens Amberg (Image Siemens with permission)

The special feature of this plant is their underground material supply system. The material transport from the warehouse is in the basement. On the shop floor there are 9 large towers, which are an intermediate storage and also the connection between the automated material flow in the basement and the production on the shop floor (four of them are visible here in the press photo). Material can be moved automatically from any tower to any other tower or the warehouse, as decided by the order management system. On the shop floor the material is moved from and to the towers mostly by hand. It takes in average only ten minutes for requested material to arrive, and no more than twenty-five minutes in the worst case. In case of short-term changes in the production program, the system can deliver material for a new production job in very short notice.

These towers are also decoupling points between the SMD soldering lines and the final assembly. This allows focus on utilization for the SMD soldering and focus on customer orders on the final assembly, improving the flexibility of the production system. This results in a 99.5% delivery reliability.

The machines themselves can also be moved flexibly, and there are many prepared covered openings in the floor that can be opened up quickly to provide the machine with data, power, air, gas, and whatever it needs. About 25% of the machines are relocated every year.

Siemens in Amberg also has a very good traceability, with data going back decades. The components of every product can be traced back to the supplier. This allows them, in combination with a "SCOUT" quality system, to analyze the causes of defects and eliminate these causes. It definitely was an impressive technology in a huge plant.

This concludes the list of plant visits. In the next and last post of the series of this tour of a van full of nerds, I will show you an overview of the presentations and tryouts we did to understand Industry 4.0. Until then, **stay tuned, go out, and organize your industry!**

P.S.: Thank you very much to everybody who hosted us and showed us their plants and products!

42 Industry 4.0 Tour in Germany – A Van Full of Nerds – Presentations and Tryouts

Christoph Roser, October 15, 2019, Original at https://www.allaboutlean.com/i4-0-tour-germany-6/

Figure 328: Arena 2036 Building (Image Pjt56 under the CC-BY-SA 4.0 license)

To round up our tour of a van full of nerds to study Industry 4.0 in Germany, here is the report on different presentations and tryouts. These were not plant visits, but different demonstrations by some smaller and one not-so-small (Bosch) companies. The first four were at the Arena 2036, a research collaboration to explore the future of the automobile. The other three were at the respective companies locations. Also quite insightful.

42.1 Presentations and Tryouts

Below is a very brief summary of some of the presentations and tryouts we had during the tour. Keep in mind that these were all demonstrations, and not "normal" use in production.

42.1.1 Bosch at Arena 2036

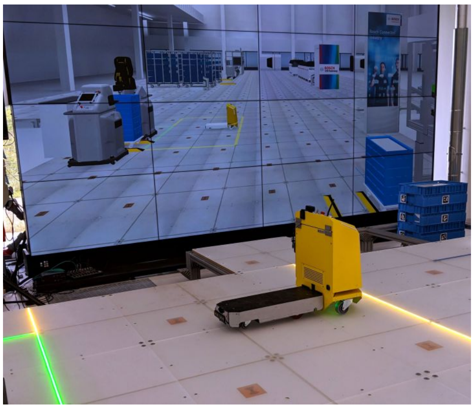

Figure 329: Arena 2036 Bosch Virtual World (Image Roser)

The Arena 2036 is a research collaboration to explore the future of the automobile. This includes its production. Bosch had an area where they researched on flexible planning of assembly lines. Their goal is not primarily the research of technologies, but the interaction of these technologies in the factory of the future. They want to explore which technologies will be useful in the future and what applications they can be used in.

Using NAiSE sensors (see below) they could determine the location of their machines, parts, and stations with high accuracy, and could create a real-time digital twin. If they moved the workstation, the station also moved in the digital world. They also used Pick by Light from ThingOS (also more below). Overall, it was a nice demonstration, although it is unclear when if ever this technology will be used for real shop floors. But then, their goal is to research the future possibilities.

42.1.2 Drag&Bot at the Arena 2036

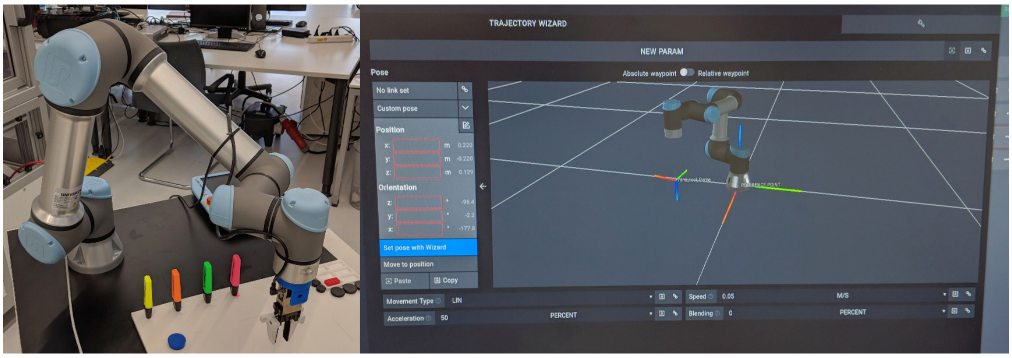

Figure 330: Arena 2036 Drag&Bot (Image Roser)

Drag&Bot is a small startup that aims to simplify robot programming. The goal is to use the same software interface to program robots of all kinds of brands (ABB, Denso, Kuka, Universal Robot …). The programming is a teach-in approach where the robots are guided by hand. It looked useful, although it is not my area of expertise.

42.1.3 NAiSE at the Arena 2036

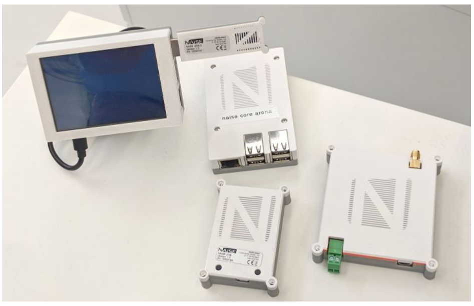

Figure 331: Arena 2036 NAiSE (Image Roser)

NAiSE provides an integrated system of hardware and software. The hardware (sensors and base stations) determine the exact position of items on the shop floor. The software can optimize routes. It is something similar to Google Maps for industrial and commercial buildings. Compared to other systems, it needs relatively few base stations. For the 7,000-square-meter arena they needed only 6 bases to achieve an accuracy of up to ±1 cm for immovable objects and ±7cm for moving objects. It does not need a line of sight, but like all similar systems, metal blocks the signal, and steel shelving would cause problems (or require more sensors).

42.1.4 ThingOS at the Arena 2036

Figure 332: Arena 2036 ThingOS (Image Roser)

ThingOS provides a platform to combine the numerous smart sensors that are part of modern Industry 4.0 and the smart factory, but also smart home and smart retail.

What I liked was their Pick by Light system. Rather than using expensive motion sensors, lidar, infrared, or light barriers to detect a pick, they simply stick a cheap RFID chip on every box and put a RFID reader on the wrist of the picker. The photo here shows the RFID sensor on the dominant right hand, and a small digital interface on the left hand.

Figure 333: Arena 2036 ThingOS Picking (Image Roser)

This interface also tells the user "three pieces" … "one piece" and so on for each pick. The second photo shows a "pick." They claim that their technology is 100 times cheaper than their competitors'. I am not sure about the 100 times, but it does look significantly cheaper than the other Pick by Light technologies that I know.

42.1.5 Robogistics Laboratory at the Karlsruhe University of Applied Science

Figure 334: Karlsruhe Robogistics Lab (Image Tobias Schwerdt with permission)

The Robogistics Laboratory at my university, the Karlsruhe University of Applied Science, is researching on the use of robots in logistics. My colleagues there are Prof. Christian Wurll and Prof. Björn Hein. Among the many things they do, we looked a bit more in detail on how to pick up shoe boxes from a bin – which is not as easy as it sounds. The boxes may be on top of each other, and also they are not firmly closed. This risk is that the robot grabs the box at the wrong end and the shoes fall out. A change in the box design apparently is not possible, since – surprisingly for me – these empty shoe boxes are also a collectors item.

42.1.6 R3DT Virtual Reality Tool

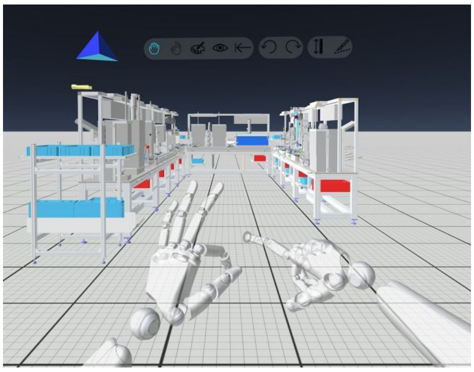

Figure 335: Screenshot of the R3DT environment (Image R3DT with permission)

R3DT provides a virtual-reality tool for shop-floor design reviews. I tried their product two years ago and found it a bit "clunky" and the hand sensors odd. Since then it has come a long way, and you no longer need to hold anything in your hand. The Leap Motion sensor attached to attached to VR glasses, like Oculus Rift and HTC Vive, can detect your bare hands in real time and add them to the virtual environment. Hence, moving your hands and controlling the software felt very natural.

Figure 336: R3DT Oculus Rift and Leap (Image Roser)

The latest version of the Oculus Rift, the Oculus Rift S, also no longer needs external trackers to determine your position. Overall it was very immersive – except for the grabbing and handling items. Picking and placing of items is possible, but not (yet) flawless.

What I liked, however, was the option "Ergo Check" to "scale you down." With a height of 187cm, I am above the 95th percentile male, meaning less than 5% of the men in Germany are taller than I am. Using the R3DT software, I could shrink down to a 5th percentile female, meaning I was now only 153cm tall. **This was a completely new and surprising perspective!** I no longer could look in or even reach boxes. The work surface was suddenly at chest height. Working at this station would have been much more difficult. Knowing this is one thing, experiencing it a completely different thing. According to R3DT, this is a world first in this simple form.

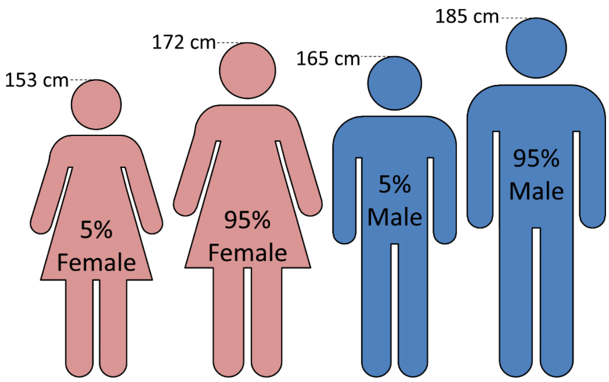

Figure 337: 95 Percentile and 5 Percentile Male and Female (Image Roser)

R3DT also has some other features like measuring distances, creating cross-sections, analyzing the reach and the visual field, and more. It was an interesting experience, although I think it makes only sense if you have a digital model already available. Creating a model from scratch is very time consuming, but the tool is intended primarily if you already have a computer model of your manufacturing system. It is already used by quite a few big-name companies. They also provide trial versions. Below is a video of one of their use cases where they check the reach of an operator at a workstation.

42.1.7 Bosch Connected Industry in Feuerbach

Figure 338: Bosch Nexeed (Image Bosch with permission)

Our visit to the Bosch Feuerbach plant was not so much to look at the plant, but at the Bosch Nexeed Transparency Kit and the Bosch Connected Industry (BCI). This kit aims to create more transparency for machines on the shop floor. This is done in two different ways: First, there is the option to get sensor data directly form the machine. A Y-splitter is added where the sensor enters the machine control, and the sensor data is also given to a small computer belonging to the Nexeed transparency kit. Second, dedicated sensors are available (light sensors, weight sensors, etc.) that also send data to the Nexeed transparency kit.

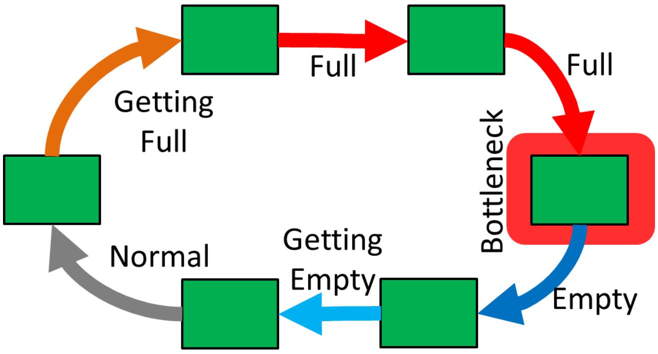

Figure 339: Bosch Nexeed Heat Map (Image Roser)

This bridges the often-cumbersome interface between the machine controller and the MES (manufacturing execution system). This allows, for example, bottleneck detection similar to my bottleneck walk and also sort of a "heat map" on where the material starts to jam up. The sensors kept track of the inventory between the stations. We observed this system also on a diesel engine component line in Feuerbach. As many Bosch lines, this line was highly automated and well organized in a very clean plant.

42.2 Summary

And that is it! The entire five days with a van full of nerds compressed into 6,000 words and quite a few blog posts. Overall, our impression was that while there is definitely quite a bit of boasting and trying to impress, many of the technologies actually work. My personal preference was the ABB Stotz Kontakt plant, since we had a strong feeling that they do Industry 4.0 not to impress the customer, managers, or anybody else, but because they truly want to see what helps them with the performance of the plant. But all the other plants had very interesting points and worked hard to follow their own path.

Many thanks to all the companies and plants that hosted us and shared their ways with us! It was a great week with seven other nerds, and we are already thinking about repeating this exercise next year (maybe in France, or maybe in Silicon Valley?). We will see, and I surely will keep you posted! Now, **go out, make sure you use computer only where they are actually useful rather than only impressive, and organize your industry!**

43 The Inner Workings of Amazon Fulfillment Centers – Part 1

Christoph Roser, October 22, 2019, Original at https://www.allaboutlean.com/amazon-fulfillment-1/

Figure 340: Amazon Logo (Image Amazon for editorial use)

Recently I had the chance to visit two Amazon Fulfillment Centers to take an in-depth look at their inner workings. While many articles about Amazon go over the basics, I will give you a deep dive into the workings of their fulfillment centers. Due to the amount of information, I divided the content across a series of posts. In this first post I will go through their general layout as well as their Kiva robotics system. Please note that most of the images and all of the videos are courtesy of Amazon.

43.1 Locations

Figure 341: Amazon SAT2 (Image Roser)

Amazon has over 175 fulfillment centers all over the world. The first center I visited was the modern SAT2 in San Marcos, Texas. The SAT2 stands for this being the second location close to the San Antonio Airport (SAT). This is a rather modern center established in 2016 using all the fancy Kiva robots (or "Amazon Robotics" as it is called now). It has an area of around 80,000 square meters, with around 2 million shelve slots containing 15-16 million individual goodies.

Figure 342: Amazon FRA3 (Image Roser)

The second one was the FRA3 in Bad Hersfeld, Germany. FRA3 means it is the third location near Frankfurt Airport (FRA). Their numbering is currently up to FRA54. FRA3 is the largest fulfillment center of the 13 (and probably more soon) centers in Germany. It is also the only one that delivers apparel (clothing, shoes, etc.) from within Germany. Established in 2009, it is one of the oldest in Germany, with surprisingly little robotics (the oldest one is the nearby and much smaller FRA1). The entire warehouse contains only one robot for stacking larger boxes for cross-shipment to other warehouses. Quite a contrast to SAT2! It is also crazy that a ten-year-old warehouse is already old.

FRA3 and FRA1 handle what Amazon calls "**sortable**" items (i.e., items that can be sorted and packed easily). Other centers handle "**non-sortable items**" … think washing machines and pianos. There are also **sortation centers** where customer orders arrive in bulk and are repacked for individual customers. **Receive centers** receive large quantities of goods from suppliers and distribute them across the fulfillment center network. **Delivery stations** prepare orders for the last mile, although the majority of orders comes from fulfillment centers directly. Finally, "**specialty**" locations handle items that does not fit the above categories, and also provide support during peak seasons.

43.2 Amazon Robotics (Formerly Kiva Systems)

Figure 343: Amazon Kiva Close up (Image Amazon with permission)

One of the most amazing aspects of Amazon fulfillment is their Amazon Robotics (formerly Kiva Systems) approach to moving material to the stowers and pickers. Hence, before going into more detail on the inner workings, I would like to look at this quite fancy robotic technology. It is estimated that the Kiva robots reduce the fulfillment cost by 40%.

43.2.1 Origin

For the more modern automated storage, Amazon uses Kiva robots (now Amazon Robotics). Amazon bought Kiva Systems in 2012. While Kiva had many different customers before, including The Gap, Walgreens, Staples, Office Depot, and more, Amazon did not renew these contracts, and Kivas are now exclusively used by Amazon. Currently around 26 of the 175 fulfillment centers around the world use Kivas, but both numbers are steadily increasing.

43.2.2 How They Work

A Kiva is a small orange robot about 30cm high and weighing around 130kg. It can lift storage shelves (called **pods**) of up to 750 pounds (340kg). These small Kiva robots carry these specialized shelves to the stowers and pickers for them to add or remove items to/from the pods. Across the US, Amazon has 120,000 KIVA robots. The image below is from MDT2 in Cecil County, Maryland, USA. However, it does not do justice to the enormous space filled with pods in a fulfillment center. I estimated a single floor in one tower has 10,000 bins, easily adding up to over 100,000 bins in a fulfillment center. I will talk more about this later in this series.

Figure 344: Amazon Kiva Storage Locations (Image Maryland GovPics under the CC-BY 2.0 license)

43.2.3 How They Organize Themselves

The picture below is a screenshot from an Amazon press video, showing you the pods of a single floor. You can see an amazing number of pods. Most fulfillment centers have multiple floors like this. You can also see that they keep "roads" empty for access, but that pods are also in larger blocks of 4×6 or similar. Hence not every pod can be accessed right away.

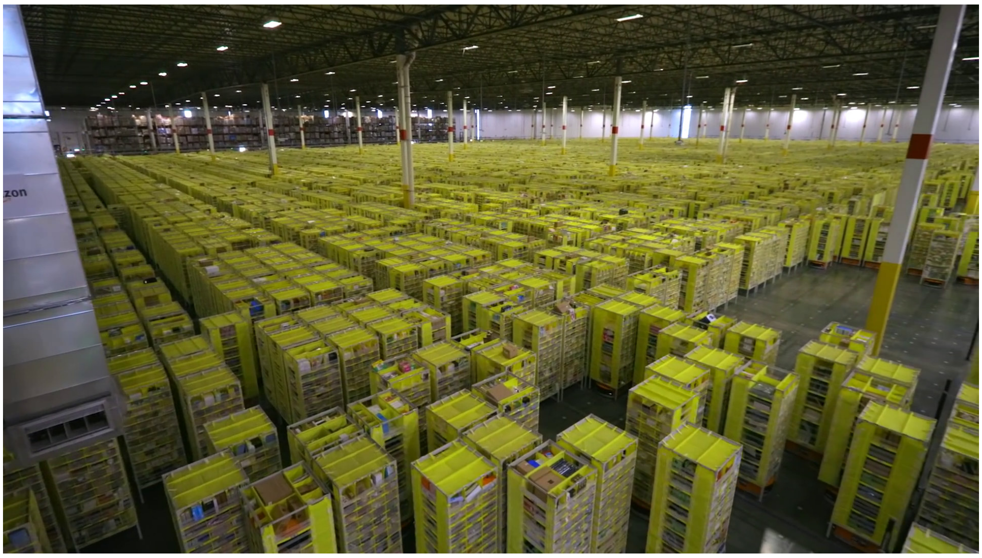

Figure 345: Amazon Kiva Floor Overview (Image Amazon with permission)

Kivas know their location primarily from 2D data codes that are attached to the floor at regular intervals. When a Kiva drives across, it scans for labels to know where it is. Additional sensors help to avoid collisions. Behind it is the **Amazon Fulfillment Technologies (AFT)** software, which I will also talk about in a later post.

The Kiva area and the worker area is clearly separated by a fence. If a maintenance worker has to enter the Kiva area, he wears a signaling device called a **Robotic Tec Vest** that makes all Kivas within 5 meters around him stop in their tracks to avoid a collision, injury, or damage to the goods.

43.2.4 Next-Generation Kivas

Amazon is currently working on the next generation of Kiva robots, nicknamed "**Hercules,**" often abbreviated to **H-Drive**. The development started in 2015 with the goal to make them shorter (i.e., more space for goods above) and stronger. And, or course, smarter.

While the first generation of Kivas were about 30 cm tall, the next generation will be only around 20 cm. While the first generation could lift 750 pounds (340kg), the next generation can lift 1,250 pounds (566kg). Note that there are also heavy-duty Kivas for 3,000-pound pallets (1360kg). Being both smaller and stronger reduces the space needed to store and handle the inventory. This aligns with the goal of Amazon to open more warehouses closer to urban centers where the property prices are higher but where a even faster delivery is possible.

The new design also has significantly less parts (Amazon claims 50% fewer components), making them both easier to maintain and cheaper. They are assembled at the Amazon Kiva headquarters in North Reading, Massachusetts, to be used all over the world.

43.2.5 Airborne?

Figure 346: Drone by Amazon (Image Amazon with permission)

Amazon is also experimenting with airborne delivery using fully electric drones. They are currently able to fly up to 15 miles carrying up to 5 pounds and deliver within 30 minutes. The drone is designed for vertical take off-and landing like a helicopter, but switches to horizontal flight like a plane while en route. Different sensors (hopefully) detect obstacles like chimneys and power lines.

However, you can't expect your parcels to be dropped in your backyard by a drone (yet). There are still many problems to solve, including safety and regulatory airspace management issues for autonomous robotic out-of-sight flight. You will also notice that the video below is showing the drone on a wide open field.

In my next post in this series I will start with the layout and the inbound value stream. Until then, stay tuned, and **go out and organize your industry!**

44 The Inner Workings of Amazon Fulfillment Centers – Part 2

Christoph Roser, October 29, 2019, Original at https://www.allaboutlean.com/amazon-fulfillment-2/

Figure 347: Amazon Manual Storage (Image Álvaro Ibáñez under the CC-BY 2.0 license)

This is the second post in my series on the inner workings of Amazon Fulfillment Centers. In this post I will look at the typical layout of the fulfillment centers and start with the inbound value stream. After all, while we all are looking forward to getting stuff from Amazon, Amazon first has to get the stuff from somewhere else.

Please note that most of the images and all of the videos are courtesy of Amazon.

44.1 Layout

Amazon Fulfillment Centers all have a somewhat similar structure. Below is an overview of the FRA3 layout, which is common for fulfillment centers without robotics.

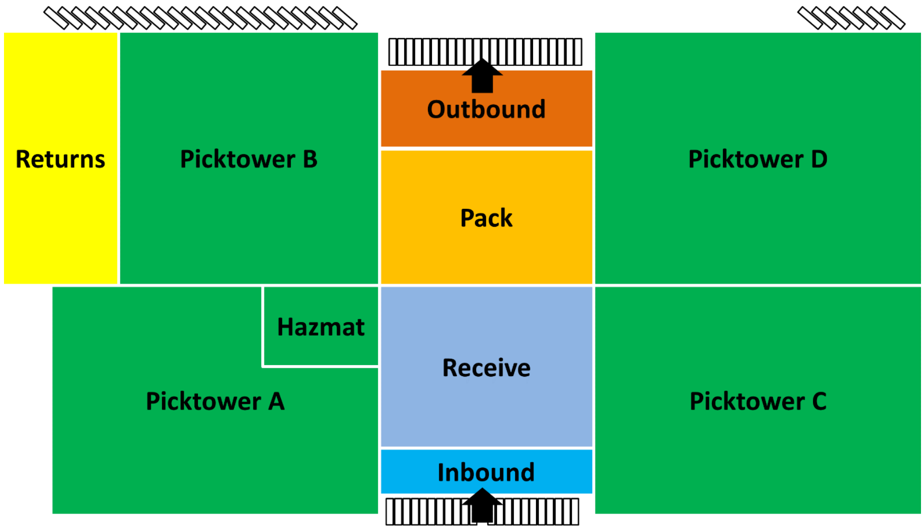

Figure 348: Amazon Bad Hersfeld Layout (Image Roser)

The storage towers where all the goods are stored are placed around a center with the material flow. Some fulfillment centers still have a manual picking process, and the items are stored in shelves, often on multiple levels. Larger items are in compartmentalized cardboard boxes. Items in high demand like new computer games may also be stored on an entire pallet. The image below is from Leipzig, Germany.

Figure 349: Inside an Amazon Pick Tower (Image Amazon with permission)

Naturally there are the docks for the inbound trucks, which are unloaded in the inbound area. In the receive area, the inbound goods are unpacked and distributed into the different towers. The inbound and outbound docks are usually on the ground floor, usually with many more docks for outbound than for inbound. The receive and pack areas usually occupy multiple floors. In SAT2, for example, there were two floors for receiving and two floors for packing. The upper floors also have the cafeteria and very few offices.

The goods you send back for whatever reasons are handled in the returns area. Some fulfillment centers also have a hazmat area. There is surprisingly little separate office space, and most functions like HR merely have a couple of desks somewhere on the warehouse floor. Both locations I visited were climatized. There was a lot of visual management to help the workers in their jobs.

Figure 350: Amazon FRA3 Entry Bridge (Image Roser)

On a side note, a lot of fulfillment centers seem to have a tower outside the building and a bridge for people to enter on the second floor. I believe that the ground floor space is the most valuable one to Amazon for inbound and outbound shipments, and hence the entrance for people is on the second floor (as is the cafeteria). Additionally, this keeps foot traffic and vehicle traffic separated for safety reasons.

44.2 Inbound Value Stream

44.2.1 Unloading and Unboxing

The inbound process starts with the trucks arriving at the inbound docks. These trucks may be from suppliers, or also from other fulfillment centers. Below are the docks at FRA3. Since trucks don't drive on Sundays in Germany, Monday morning is the least busy time for the inbound area. A traffic light system combined with a movable stop sign (see below) aims to prevent a truck from leaving before it is safe and the doors are properly closed.

Figure 351: Amazon Inbound docks Bad Hersfeld (Image Roser)

The trucks are unloaded and the cardboard boxes are moved to the receive area. The boxes are scanned and opened and the individual goods are put into black boxes (called totes) or on silver carts. The packaging is disposed of and the boxes and carts are distributed to the pick towers.

44.2.2 Manual Stowing

Figure 352: Amazon Fulfillment Stowing (Image Amazon with permission)

Here we have to distinguish between normal shelves and manual stowing, and robotized pods. For manual stowing stowers walk though the shelves and place the items on the shelves. There is no assigned location for a product, the stowers simply place it wherever they find space. Hence the items are stored quite randomly. This is called chaotic storage. To keep track of the locations they take the item, scan the shelf, scan the item, and place it on the shelf. If there is no space on the shelf, then they ask the computer for a new location. Below is a closeup of such a shelf for books showing level D for locations 250 and 255. The numbers represent 10cm intervals along the shelves, hence the shelves below are around 50cm wide and around 25m from the start of the shelf. As you can see, there are multiple different products on the same shelf.

Figure 353: Amazon Shelve Storage Detail (Image Álvaro Ibáñez under the CC-BY 2.0 license)

44.2.3 Robotic Stowing

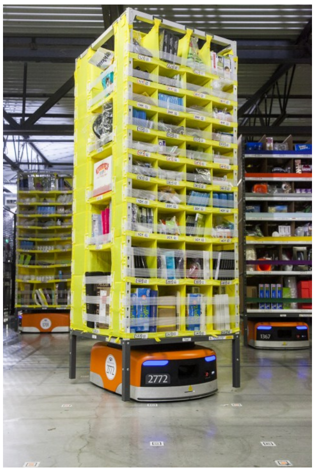

Figure 354: Amazon Kiva with Pod (Image Amazon with permission)

In robotic stowing, it is still a human who puts the item onto the shelf. (Although Amazon is researching robots for stowing items. Maybe in 5 years?) The difference is that the stower no longer walks to the shelf, but a small orange robot lifts up the entire shelf and brings it to the stower. The technical term used at Amazon for these shelves are "pods," and they have tens of thousands of these pods in a warehouse.

Anyway, these pods are carried to the stower stations, which look very similar to the picking stations further below. The stower has the boxes from the inbound area and scans the different items. Similar to the manual stowers, they can place it on any free slot on the pod in a chaotic storage. The stower adds the item, while a motion sensor tracks the location in the pod. After inserting the item into the pod and clicking confirm, the stower gets the next location or pod for the next item. If you watch this, it goes very smoothly. Below is a brief video of Kivas in action. They navigate using QR codes on the floor.

The Video by Amazon is available on AllAboutLean.com as "Amazon Kivas Video" at https://www.allaboutlean.com/wp-content/uploads/2019/10/Amazon-Kivas.mp4

If there is not enough space at the pod location, the stower must not put it on another spot on the same pod, but instead gets a new location or pod by the computer. Problem solvers will look at why the slot in the pod is already full. Similarly, if the stower notices damage, the item also goes into an "amnesty bin" to be checked out later.

In the next two posts I will look at the outbound value stream (i.e., how you get your parcels). The outbound process is probably the most significant value stream in Amazon. This will include the picking (manual and robotics), packing, SLAM (Scan, Label, Apply, Manifest), and outbound sorting and loading on top of a number of additional possible steps like seasonal Christmas picking, shipments between fulfillment centers, the handling of packages with multiple items, a re-picking that some centers seem to have, gift wrapping, and more. Until then stay tuned, look at your inbound value streams, and **organize your industry!**

45 The Inner Workings of Amazon Fulfillment Centers – Part 3

Christoph Roser, November 05, 2019, Original at https://www.allaboutlean.com/amazon-fulfillment-3/

Figure 355: Amazon Warehouse Floor (Image Amazon with permission)

This is the third post on my series on the inner workings of the Amazon Fulfillment Center. In this post I start with the highly interesting process of the outbound value stream (i.e., how the goodies go from storage to your door). Since this is the core process, the next post will continue this outbound value stream.

Please note that most of the images and all of the videos are courtesy of Amazon.

45.1 Outbound Value Stream

The outbound value stream is the largest part of the work. It all starts with the picking process, where the desired goods are picked up in the picktowers. Modern fulfillment centers use robots to aid the picking process, while older centers and centers with goods that are difficult to handle (bottles, clothing, etc.) use manual picking.

45.1.1 Manual Picking

Let's start with manual picking. As you probably imagine, someone is walking along the shelves picking up stuff. They push a cart with yellow boxes and have a handheld scanner. The scanner tells them the shelf location for the next pick. Depending on the product, this may be anywhere from 30cm intervals as shown below, or with a new shelf area starting every 3 meters. Since this is a random storage system, the shelf area may include other products, and the same product may also be on other shelves. The picker scans the shelf barcode followed by the item. If everything is okay, the picker moves on to the next location.

Figure 356: Amazon manual picking (Image Amazon with permission)

Pickers work on multiple orders at the same time. If a parcel has multiple items (a multi-pack), multiple pickers may work on the same parcel. Hence the boxes of the pickers usually contain items for more than one order.

If there is any problem (i.e., the item is not found on the shelf), the picker gets a new item from the scanner and nothing is picked. Even if the picker sees the item in an adjacent location, he must not take it. Specialized problem solvers will have a look at this shelf to figure out what went wrong. If there is a wrong item or damaged item, an un-scannable item, or stuff that is still part of a larger box, it goes into an "amnesty bin" to be sorted out later by problem solvers.

The items in higher demand are stored closer to the packing stations so that the pickers have to walk less. Oversized items, hazardous goods, and goods with different storage needs (e.g., clothes on a hanger) are stored in different areas with shelves or racks that suit the product. Heavier goods weighing more than 15kg are handled using vacuum lifts or with two people. They also try to have one picker handling only one area of the pick tower, to reduce time-consuming moves between floors (called zone-picking).

45.1.2 Robotic Picking

Figure 357: Amazon Robotic Pick Station (Image Amazon with permission)

Similar to robotic stowing, the robotic picking does not use a robot to grab items off the shelf. This is currently still too unreliable and slow. However, Amazon works to also automate the picking, and is thinking about using more standardized packing to make the job easier for robots. But as for now, taking items off the shelf is still done by humans. The same small Kiva robots that carry the pods for the stowing also carry the pods for the picking process.

The worker at the picking station sees on a monitor the next item including a picture of the item, how many of it, and which slot in the yellow pod the item is. The worker simply grabs the item out of the pod, scans it, and puts it into one of multiple totes as determined by a computer. Similar to a pick-by-light, the placement is indicated by a light above the tote. The image above shows a hand scanner, but the SAT2 I visited had a scanner over every box. Hence similar to a supermarket checkout, the worker simply held the item under the scanner before dropping it into the tote. A future planned change is to include spotlights above the worker directed to the pod, showing which slot in the pod has the desired item. This may make the system even faster. Below is a short video for the picking from pods.

The Video by Amazon is available on AllAboutLean.com as "Amazon Pick Sequence Video" at https://www.allaboutlean.com/wp-content/uploads/2019/10/Amazon-Pick-Sequence.mp4

Similarly to the stowing, if anything is out of order, the worker is not instructed to fix it, but instead to give a signal to problem solvers who will sort it out. To help with this process, a camera takes a photo of the side of the pod that was picked after the picking. An artificial intelligence algorithm tries to find out if the photo is consistent with the expected pod content. If the algorithm senses a difference (a "delta"), it may be sent to a human problem solver. This is not a cross-checking of the picker, but only a quality-control process.

Every now and then, the yellow (or sometimes black) totes are pushed onto a conveyor belt for the next stations, which may be packing if there will be only one item in a box (single item), or another rearranging process for multi-packs having more than one product in an box.

45.1.3 Seasonal: Christmas Picking

During high season (especially Christmas), the normal number of picking and packing stations do not suffice. Additional stations that are normally stored away are set up for more pickers to work. Often, these additional stations are not connected to the conveyor network, and a separate worker picks up the totes with a cart and brings them over to the conveyor system.

45.1.4 Additional Stages: Inter-Center Shipment

Figure 358: Amazon Robot (Image Amazon with permission)

Sometimes an ordered product is only available in a fulfillment center farther away. In this case it is often too expensive to ship a small parcel directly to the customer over long distance. The cheaper option may be to ship a large box containing many such items from one center to another and pack it in a center closer to the customer. Such shipments between fulfillment centers are often done using large black boxes that are stacked on a pallet. This pallet is deconstructed in the target fulfillment center. The only robot used in FRA3 stacked such black boxes onto a pallet (and it is named "Chuckle Berry"). The robot in the image is a similar robot for handling pallets in the USA.

45.1.5 Additional Stages: Multi-Pack

Between picking and packing, there may be additional stages possible. One of these is the multi-pack. If a customer gets only one item in a parcel, the yellow totes go directly to the packing stage. These are either small orders or Prime customers that get faster delivery with more parcels. In any case, if the customer gets more than one item per box, the picked items need to be rearranged. This is done in a separate station.

The yellow totes arrive from picking. Another worker at SAT2 takes items out of these totes, scans them, and puts them into a slot in an adjacent shelf as indicated by lights. Each of these slots is one parcel with multiple items. When one order is completed in a slot on the shelf, a light on the other side indicates for another worker to take these items out and put them into another box, a gray tray. This tray goes then to packing.

At FRA3 the system was slightly different, where the items were picked out of the tote into a slot of a shelf. When the shelf is full, the entire shelf is then wheeled to a packing station. This example of the movable shelf below is from MAD4 in San Fernando de Henares, Spain. Since FRA3 also handles fashion items like clothing and shoes, it is more likely to have multi-packs, with an average of 2.18 items per parcel. People seem to order different sizes to see what fits, and return the rest. As a result, FRA3 also has a higher-than-average return of items, but 95% of it was suitable for resale.

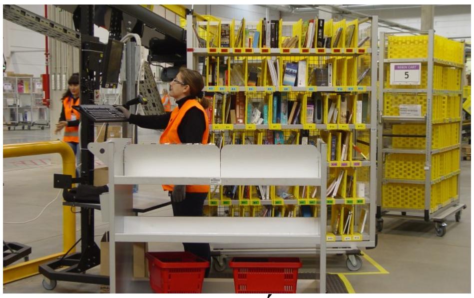

Figure 359: Amazon Multi Pack Shelf (Image Álvaro Ibáñez under the CC-BY 2.0 license)

45.1.6 Additional Stages: Re-Pick

It seems in some fulfillment centers they make the totes from picking rather full, and then re-pick them a second time into red boxes for individual orders. This is shown below.

The Video by Amazon is available on AllAboutLean.com as "Amazon Re-Picking Video" at https://www.allaboutlean.com/wp-content/uploads/2019/10/Amazon-Re-Picking.mp4

45.1.7 Additional Stages: Giftwrap

For many products the customer can select an optional gift wrap. The products then go to a gift-wrapping station.

Figure 360: Amazon Gift Wrapping (Image Amazon with permission)

In the next post I will continue with the outbound process: pack, SLAM, and outbound. Until then stay tuned, look at your material flow, and **organize your industry!**

46 The Inner Workings of Amazon Fulfillment Centers – Part 4

Christoph Roser, November 12, 2019, Original at https://www.allaboutlean.com/amazon-fulfillment-4/

Figure 361: Amazon Warehouse Floor (Image Amazon with permission)

This fourth post in the series on the inner workings of the Amazon Fulfillment Center will continue to look at the outbound material flow, including pack, SLAM, and the loading of the trucks. In the next and last posts, I will also look at the software behind it as well as some surrounding processes.

Please note that most of the images and all of the videos are courtesy of Amazon.

46.1.1 Pack

Figure 362: Amazon Packing (Image Amazon with permission)

Finally, the item goes to the pack station. A fulfillment center has easily 100 pack stations or more, although not all of them are active all the time. During Christmas more stations are added temporarily.

At the pack station, the item is taken out of the yellow tote and scanned. A monitor displays which size carton or envelope to use. A dispenser provides packing tape in the right length. The employee builds the box and tapes the bottom, inserts the item, and tapes it closed. At the very end, a white barcode is added. This barcode links the parcel to all the necessary information. The actual address label is added later during the SLAM process. Below is a brief video of the packing process.

The Video by Amazon is available on AllAboutLean.com as "Amazon Packing Video" at https://www.allaboutlean.com/wp-content/uploads/2019/10/Amazon-Packing.mp4

The tote is conveyed back to be refilled at picking. The parcel goes on a conveyor belt to the SLAM process. Often, the packing process is supplied with new packing material from the other side. Hence it is usually two rows of packers with one supply path in between. The image below is the view from the supplier of cardboard.

Figure 363: Amazon Packing lane (Image Amazon with permission)

46.1.2 SLAM

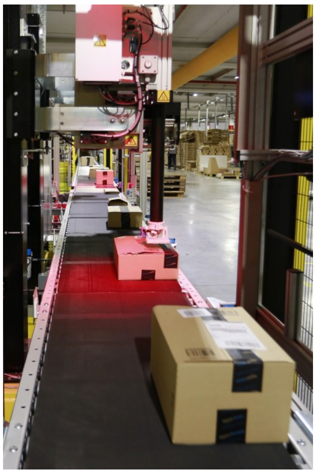

Figure 364: SLAM process at Amazon (Image Amazon with permission)

SLAM has nothing to do with wildly slamming your parcel around (that is done during delivery), but stands for Scan, Label, Apply, Manifest. It is also a final quality check that weighs the parcel and compares the weight with the expected weight of the items and packaging. Again, if there is a discrepancy, the parcel is checked again manually.

During packing, a barcode was attached to the parcel. This links the parcel with the related information, but this is only machine-readable. During SLAM, a machine scans this code, prints a proper shipping label that humans can also read, and attaches this label to the parcel. The image on the right shows the SLAM machine, and the image below shows a parcel with the packing barcode in the upper part and the SLAM label below.

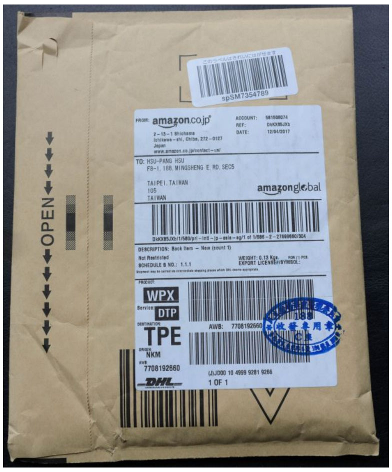

Figure 365: Amazon Parcel Taiwan (Image 玄史生 in public domain)

The machine looks like it is pressing on the parcel. However, it is moving only slightly above the parcel, and the adhesive label is attached using air pressure instead. A fulfillment center may have somewhere around 10 SLAM machines.

46.1.3 Outbound

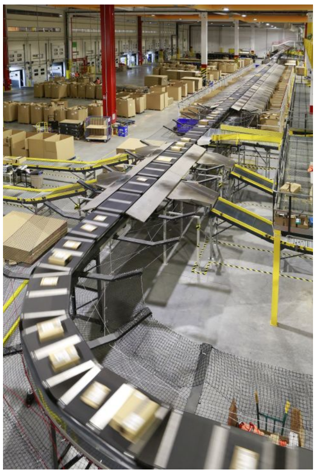

Figure 366: Amazon Fulfillment Conveyor Belt (Image Amazon with permission)

Now the parcel is ready for shipping. A conveyor belt with speeds of up to 30km/h brings the items to a sorter. There seem to be different sorting systems. The first picture here on the right from MAD4 seems to have a cross belt (a small conveyor belt going sideways on top of the big conveyor belt). There is one parcel per segment. When the parcel passes the correct slot the cross belt activates and moves the parcel to the right or to the left into the correct chute. I have seen similar belts also at the Haneda Chronogate logistic terminal.

Figure 367: Amazon Fulfillment Conveyor Belt (Image Amazon with permission)

The two locations I visited had a slightly different belt as shown here on the left. Instead of a small cross belt, they used sliders to push the parcel off the belt into a chute.

Figure 368: Amazon Cork Screw Slide (Image Amazon with permission)

Finally, one photo I saw seems to have indicated a manual sorting, although I am not sure about the details of this photo.

Depending on the floor of the sorting, the parcels may also slide down a cork screw slide to the level with the outbound loading docks. Like pretty much all conveyor belts I have seen, items are dropped occasionally. Every now and then a parcel falls off. Hence, someone checks the belt occasionally and inserts dropped parcels back into the process.

46.1.4 Outbound Loading Docks

The parcels are now sorted and ready to be loaded into the delivery trucks. The image below shows the outbound area, with the conveyor belts extending into the trucks. The cork screw slides on the right come from the sorting area. FRA3 was under reconstruction and temporarily loaded the trucks manually until the new conveyors were installed.

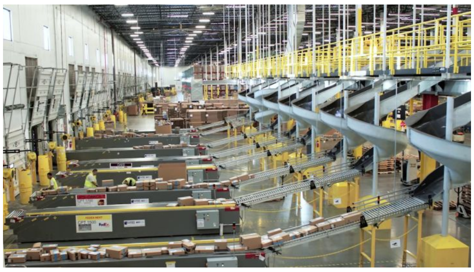

Figure 369: Amazon Outbound (Image Amazon with permission)

Depending on the type of truck, the loading can feel like a 3D version of Tetris, trying to fit as many packages into the truck as possible. FRA3 ships, on average, 900,000 to 1 million items per week, with a peak demand of 300,000 items in 24 hours during Christmas.

Figure 370: Amazon Tetris Truck (Image Amazon with permission)

46.1.5 Shipping

Figure 371: Amazon Delivery Van (Image Amazon with permission)

Parcels are shipped with all types of logistic companies like FedEx, DHL, Hermes, UPS, DPD, Royal Mail, US Postal Service, and whatever is available in that country. More and more, Amazon is also using their own delivery trucks for the last mile. This is of course a threat for the other delivery companies, who make good business with Amazon on the risk that this business can evaporate very quickly. A small fleet of aircraft also handle longer distance transport between centers. Lately during delivery I was able to see a live map with the location of the delivery van, and how many more stops there are before it arrives at my location. This made accepting parcels a lot easier for me.

46.1.6 A Note on Terminology

In SAT2 they had different colored boxes with different names for different stages in the process. For this they went quite a bit through synonyms for boxes. Let me give you the sequence again: The KIVA shelves are called "**pods.**" These are picked into yellow "**totes.**" These totes go directly to pack, unless it is a multi-item package. In this case these totes are re-picked into gray "**trays.**" Single pack totes and multi-pack trays go to packing, where they finally go into a cardboard "**box**" (or smaller items into envelopes). Problems and damaged goods go into an "**amnesty bin,**" which are handled by problem solvers. All these different names probably help with communication rather than just calling all of them a box. In FRA3, however, they used the same box for the entire path from pick to pack, plus large black boxes for shipments between fulfillment centers. All centers also had shelves on wheels. Larger items are also moved by forklifts.

This fourth post in this series on Amazon Fulfillment centers closes the outbound process. However, there is more. In my next post I will talk a bit about the software behind this process that makes all this possible. I also will show you some surrounding processes. Until then, stay tuned, and **go out and organize your industry!**

47 The Inner Workings of Amazon Fulfillment Centers – Part 5

Christoph Roser, November 19, 2019, Original at https://www.allaboutlean.com/amazon-fulfillment-5/

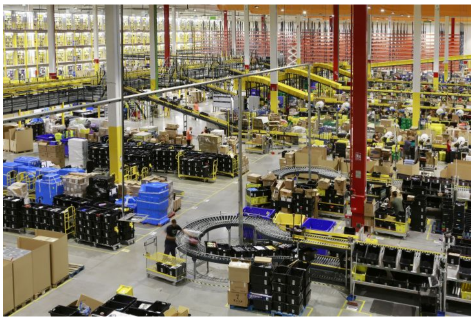

Figure 372: Amazon Warehouse Floor (Image Amazon with permission)

This is the fifth post on my series on the inner workings of the Amazon Fulfillment Center. Here I will look at the software that runs behind all the processes and makes this performance possible. Other companies would probably plaster the label "Industry 4.0" all over this, but at Amazon they just do it.

Please note that most of the images and all of the videos are courtesy of Amazon.

47.1 The Heart of Fulfillment: Amazon Fulfillment Technologies (AFT)

Figure 373: Amazon Order Form (Image Amazon with permission)

Above I talked a lot about the physical aspect of the Amazon Fulfillment Center. While this is nicely done, the real power of Amazon fulfillment is invisible in the software, called **Amazon Fulfillment Technologies (AFT**). They claim it is the largest fulfillment execution engine in the world, and is their connection between the physical and virtual world. Due to the huge number of widely different products, this is a highly challenging and complex task. In a single day, Amazon handles around 37 million orders.

Most other companies would plaster the label "Industry 4.0" all over this, but Amazon does not really mention Industry 4.0. They just do it. I like that. Nevertheless it includes a large number of advanced machine learning technologies for forecasting, text comprehension, image recognition, translation, speech recognition, and a lot more.

47.1.1 Example: Bin Images

Figure 374: Amazon Bin Image (Image Amazon with permission)

One example are the images from the pods used for machine learning. The tape holding the items in the pod is transparent to make identification easier. After picking, photos are taken of the pod shelves (called bins). One of these you can see here on the right. Machine learning is used to answer questions like

- How many objects are in the bin?
- Is there a certain product in the bin?
- How many of a certain product are in the bin?
- etc.

A data set with half a million bin images is [available via Github]() for machine learning. I looked at a few images, and for quite a few of them I would have trouble myself to answer correctly.

47.1.2 Aurora Database

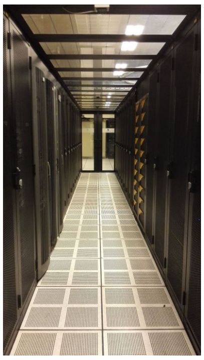

Figure 375: (Generic) Server Room (Image Pixabay in public domain)

While initially based on Oracle, the Amazon fulfillment database eventually became just too big to manage. In May 2017 they switched to their own **Amazon Aurora** system, although they turned off their last Oracle database only in October 2019. Aurora is a cloud-based MySQL- and PostgreSQL-compatible database (both are open source database management systems). Amazon claims that it is five times faster than a normal MySQL database and three times faster than a PostgreSQL database. Amazon also claims it costs them [as little 1/10th of the previous Oracle database](), both in licensing (60% reduction) and administration (70% reduction). The maximum size per database seems to be 64TB, although Amazon migrated around 75 petabyte of data from Oracle to Aurora (including also other databases like Alexa, Prime Video, Amazon Music, and more).

47.1.3 Process

A computer decides what to pick, when to pick, activates the Kivas if available, and organizes the shipment. And this computer system is thoroughly refined. Based on a customer's order, it makes an initial plan but refines this plan at every step based on minute details such as the likely route of the truck, the available space on a truck, and many more aspects that I can only imagine. Hence the fastest and/or cheapest path to the customer is frequently updated based on new information.

This system includes actual artificial intelligence. While AI is talked about a lot, it is still rarely seen in industrial applications of logistics and production management. One of the features is also machine vision to analyze the pictures of the Kiva pods or other products.

47.1.4 Performance

Figure 376: Computer Performance (Image Murrstock with permission)

In FRA3, the average time between an order and the corresponding parcel leaving the docks is an astonishing 2 hours and 45 minutes – and FRA3 is one of the slower fulfillment centers due to the handling of clothes and the manual picking. Their goal is to improve this to 2:30 hours. The current record of Amazon between the order of the product by the customer and the ringing of the doorbell of said customer is … hold on to your seat … 3 minutes!

I need more time if I want to get something from my basement! One of my past jobs considered it as Just in Time to get a part from the warehouse across the street within 3 days. And Amazon delivered a parcel in 3 minutes! This is quite a statement on the power of their physical side of the fulfillment, but even more of the software side. Granted, the customer lived right across from the fulfillment center, and also granted this is the best time out of who-knows-how-many gazillion parcels delivered, but I find this highly impressive.

In some locations, Amazon even delivers regularly within 2 hours. This Amazon Prime Now delivers a wide range of goods within 2 hours for Amazon Prime customers in selected locations. Many larger cities in the USA are covered, as well as some cities in the rest of the world. Only a few years ago, a delivery time from ordering to ringing the doorbell within a few hours was unbelievable for end customers, unless you pretty much paid a taxi to bring the stuff to you. Now Amazon does it for free!

In a previous article on Industry 4.0 – What Works, What Doesn't, I talked about where Industry 4.0 makes sense. One of the areas was logistics. Industry 4.0 can be very expensive. Logistics is a segment where there is an economy of scale. If you have programmed one Kiva, then you have programmed them all. If you have the software back-end for one fulfillment center, then (with minor adjustments) you have it for all. Hence it makes sense for Amazon to pour significant resources into their Amazon Fulfillment Technologies software, since the benefit can be shared across their hundreds of fulfillment centers. Overall, an amazing piece of work that I believe will continue to be improved on.

While still a lot is unknown to people outside of Amazon, I consider this Amazon Fulfillment Technologies (AFT) exceptional. Getting so much data together, keeping track of it, and updating it in real time is quite a challenge. But Amazon has the resources to push this software (they are hiring a lot), and the implementation has company-wide benefits. Overall pretty cool. Now, **go out, and organize your industry!**

48 The Inner Workings of Amazon Fulfillment Centers – Part 6

Christoph Roser, November 26, 2019, Original at https://www.allaboutlean.com/amazon-fulfillment-6/

Figure 377: Amazon Employee with Shelf (Image Amazon with permission)

This is the sixth and last post on my series on the inner workings of the Amazon Fulfillment Center. Here I will look at some supporting processes as well as the all-important employee satisfaction. I will look at the process of taking inventory, security, their interesting office locations on the warehouse floor, Amazon Go stores, and Employee Satisfaction.

Please note that most of the images and all of the videos are courtesy of Amazon.

48.1 Taking Inventory

Figure 378: Amazon Shelves (Image Amazon with permission)

Like in any warehouse, there are occasionally differences between the data and the reality. Sometimes extra items appear (maybe a packing error or an over-delivery by the supplier). More commonly, however, items disappear. This could be lost, damaged, under-delivered, sent to the customer in error, or even stolen.

Amazon has standardized processes on how to fix such issues. First of all, if something pops up, there are problem solvers who investigate the issue. If a shelf is so full that the person stowing can't add the item, it is checked. If a item should be on the shelf but the picker can't find it, it is also checked. This is a permanent ongoing process.

There is also the process of inventory taking, where people check the content of a shelf and compare it with the data (i.e., what is there and what should be there). In FRA 3 they claimed (anecdotally) that they find 1 error in a few dozen meters of shelf – which to me sounds too good to be true. But anyway, they have a perpetual inventory taking that is common for retail and similar inventories.

48.2 Security

As is expected for a warehouse with lots of desirable goods, these warehouses need security. Going into the building requires a chip card through a security gate. Within the building, going into the actual warehouse is no problem, but going out has also some security. In FRA3 in Germany, these were simple RFID scanners similar to a shopping mall, and they do not do body scans or check pockets of their workers. In SAT2 in Texas, it was closer to an airport security with body scans and x-rays of luggage. This is aimed to reduce theft, although it is unable to completely prevent it.

Figure 379: My new laptop… (Image Roser)

I order through Amazon frequently, and recently I had the first case of what I suspect to be theft. I ordered a nice new laptop, but when the box arrived it contained only chopped straw for the bedding of small pets. The box was still sealed with the original packaging tape and did not appear to have been tampered with. The two bags of bedding weighted 2kg, which was approximately the weight of the laptop, hence the double-checking during the SLAM step did not notice anything. This may be either a very weird packaging error or some employee figured out how to get a free laptop. Due to the price tag, the service agent from Amazon needed to escalate my issue to its supervisor, but they sent me a replacement without much hassle (cheers to the Amazon Support Desk!). They did, however, start an internal investigation to figure out what happened to the laptop, and my sense of justice hopes that they catch the perpetrator.

48.3 Offices

One thing I liked was how Amazon integrated the offices with the shop floor. In all locations I visited, and on some photos like MAD4 below, most offices were not separate rooms, but simply tables and computers put in a corner of the warehouse. No dividers, no walls, no doors, completely accessible.

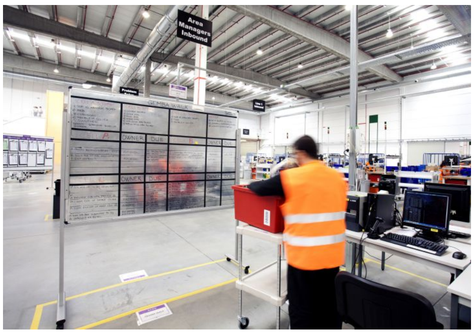

Figure 380: Amazon Warehouse Office (Image Amazon with permission)

For example, HR at Amazon Fulfillment often has a desk at the main entrance. Every employee coming in or going out walks past the HR desk, without any separating doors or walls. If the employee needs or wants to interact with HR, then they are right there and can (hopefully) help the employee.

Similarly, problem solvers, quality people, team corners, and improvement teams are right where they should be – on the warehouse floor (or in lean terms, the genba). If you look closely at the image above, you will see a gemba walk board that also has a list of problems for different pick towers, a responsible, and a due date … and it is all hand written! I quite like that approach.

Like all open offices, and especially offices on a shop or warehouse floor, there is the issue of noise. This is a downside of such open offices, but maybe they work with noise-canceling headphones, which become more common in open offices.

48.4 Employee Satisfaction

Figure 381: Worker at Amazon (Image Amazon with permission)

Finally, I would like to talk about employee satisfaction. Amazon got some critical press lately, including a video on Last Week Tonight by John Oliver, lamenting the difficult conditions in the warehouses, pressure to perform, and difficulties of accessing toilets.

Personally, however, I did not notice such problems. Granted, working in a warehouse is tough and repetitive work, and the pay could always be better. However, I don't think conditions are worse at Amazon than than at let's say Wal-Mart or other warehouses. Observing the people working during my tour, I did not notice any glum or worn-out people. My tour guides were quite cheerful and appeared happy with their jobs at both locations. The working pace also seems to be acceptable. All centers I visited were air-conditioned, and the guide in SAT2 noted that the Wal-Mart warehouse nearby is not. But then, I have never worked at Amazon, and my observations are only casual. At both locations they claimed that they have toilets within the warehouse that can be accessed without a security check.

Figure 382: Amazon Picking from Pods (Image Amazon with permission)

A brief glance at company review sites like Glassdoor or Kununu showed reasonably good reviews with 3.8 and 3.66 stars out of 5 respectively. While this does not put Amazon in the top 100, it is in my view a good rating. Of course it is a mix of warehouse workers and white-collar workers are often more satisfied than blue collar workers, but this is also true for other companies. For comparison, at the time of this writing, Wal-Mart had 3.2 stars on Glassdoor, FedEx 3.7, UPS 3.4, United States Postal Service 3.0, Union Pacific 2.2, and DHL 3.7, to pick just a few logistic companies. Overall, I believe the negative press on the working conditions is exaggerated.

48.5 Amazon Go

Figure 383: Amazon Go Seattle (Image Sikander Iqbal under the CC-BY-SA 4.0 license)

While talking about fancy new things, I would also briefly like to mention Amazon Go. The largest online retailer Amazon goes brick-and-mortar. Amazon Go is a chain of convenience stores with 17 locations in the USA. The first one opened 2017, although only for Amazon employees as a beta-test. The general public can enter such stores since 2018 if they have the Amazon Go app installed.

This store is – sort of – the retail version of a fulfillment warehouse, and stuffed with all kinds of sensors, computers, and artificial intelligence. The goal is to automate a retail purchase as much as possible. The customer merely picks up whatever he wants and takes it with him. In a normal store this would be outrigger theft, but at Amazon Go, cameras monitor the customers, shelves measure their weight, RFID chips detect items, etc. If a customer takes an item off the shelf, the store detects this and adds the item to the Amazon Go app shopping cart. If the customer puts it back, the item is removed from the shopping cart. When the customer leaves the store, the items in the shopping cart are paid for automatically through the app.

The concept is quite revolutionary, and can cut down retail labor cost even more, as well as losses due to theft. According to some news, Amazon plans to open thousands of stores within the next few years.

48.6 Amazon Fulfillment Tours

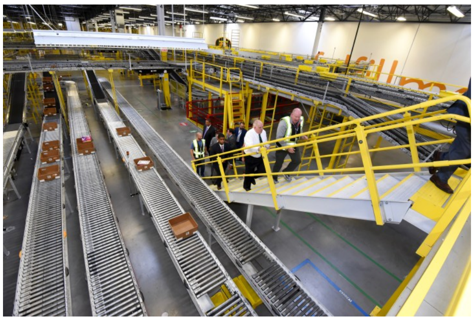

Figure 384: Amazon Tour (Image Joe Andrucyk under the CC-BY 2.0 license)

If you would like to also see an Amazon Fulfillment Center, many of them offer tours. On their website you can book tours in different centers. Similar tours are also available in other countries (e.g. Germany). I thoroughly enjoyed the tours and the insights I got in a well-organized operation. Now, **go out, get your goodies to the customer, and organize your industry!**

49 One Up One Down – Approach to Manage Manual Production Lines

Christoph Roser, December 03, 2019, Original at https://www.allaboutlean.com/one-up-one-down/

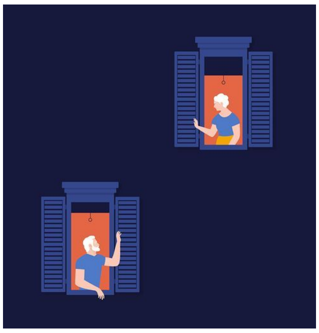

Figure 385: Up Down Window Neighbor (Image Moremar with permission)

Production lines have fluctuations. Sometimes production takes longer, sometimes shorter, than the average. This makes the line balancing tricky. Besides using a simple buffer between workstations, it is also possible to adjust capacity. Other approaches I have written about include the rabbit chase and the bucket brigade. Here I present a variation of the bucket brigade called "**One Up One Down**."

49.1 The Bucket Brigade

In a bucket brigade, people move along the line with a part until they meet the next worker (or the end of the line) and hand over the part. Afterward they move back until they meet the previous worker (or the beginning of the line) to get a new part to work on. This is visualized in the animation below.

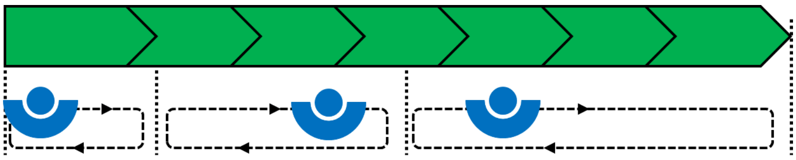

Figure 386: Animated Bucket Brigade Loops. The original image can be found at https://www.allaboutlean.com/bucket-brigade-1/ (Image Roser)

The challenge is that all workers need to be trained at all stations. Especially for longer or more complex lines, this can be tricky and time consuming. Newly hired employees may also need more training time before they can be put on the line.

49.2 One Up One Down

The One Up One Down concept gives each worker a "home station," but the worker is also qualified in the adjacent stations. Depending on the workload, the worker can do the work at the prior workstation or at the subsequent workstation. This is similar to the bucket brigade, but the worker is qualified only in a small range of workstations.

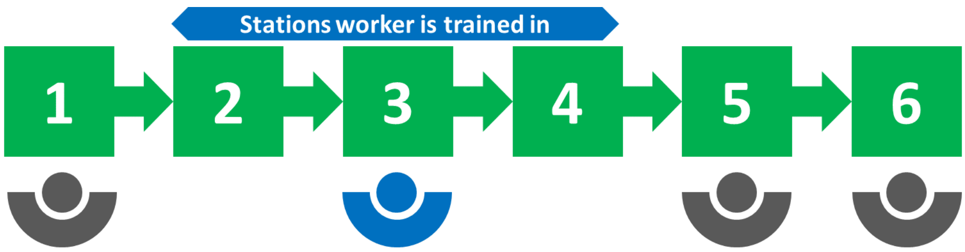

Figure 387: UpDown Trained Stations (Image Roser)

Of course, if the next process is already being worked on by another worker, the current worker cannot join but instead moves back to see if he can work on one of the preceding processes. If these are also already occupied, the worker has to wait until the preceding process releases a part for him to work on. It may be possible, however, that the worker, instead of waiting, can assist the other worker. This can be especially true for larger products.

With multiple workers trained in adjacent workstations, it is possible to balance out some cycle-time fluctuations by having the worker move temporarily to an adjacent workstation. Please note that this requires that the workers are qualified for overlapping stations.

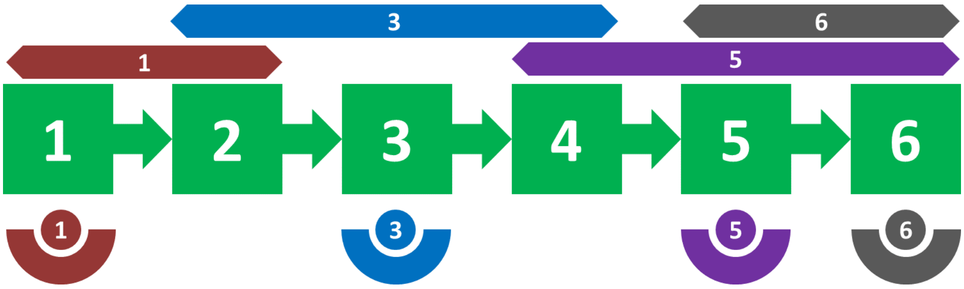

Figure 388: UpDown Trained Stations Multiple (Image Roser)

Please note that an overlap is highly recommended. If for some reason the qualifications of the workers do not overlap (like worker #6 below), you will lose the flexibility to handle cycle-time fluctuations. In the example below, worker #6 always has to do two stations. If he is faster than his colleagues, he has to wait. If he is slower, the others have to wait for him. Alternatively, you need a material buffer to decouple these fluctuations. It is doable, but you lose the benefit of the flexibility.

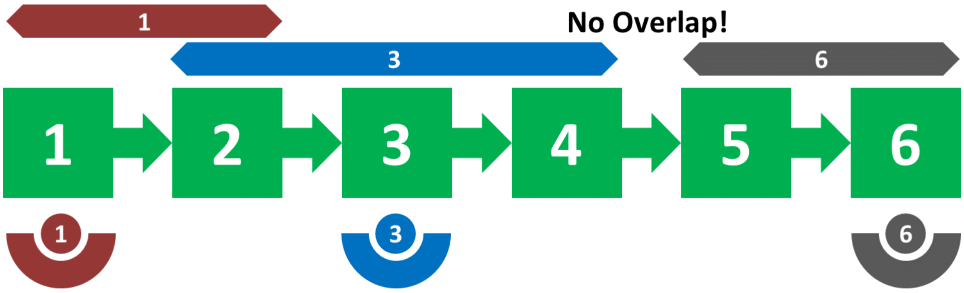

Figure 389: UpDown Trained Stations No Overlap (Image Roser)

Similarly, if all stations are manned, then you also get no benefit from an overlap in qualification. Since all stations are occupied, a worker can never move a station up or down if there is a waiting time. On the plus side, this system will produce the most parts, since every station is working (almost) all the time. Although, here, too, workers may be able to assist other workers if they have time available.

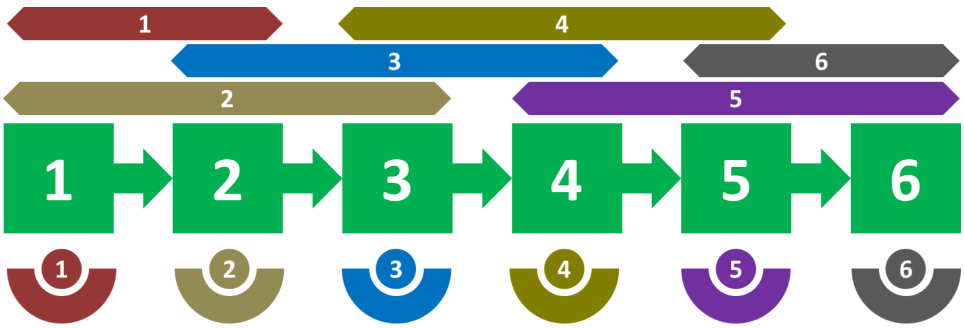

Figure 390: UpDown Trained Stations All (Image Roser)

In any case, for the One Up One Down approach, you will need a few empty workstations (sometimes called “holes”) for the workers to move. The more empty stations you have while still having overlapping qualifications, the better you will be able to adjust to cycle-time variations. At the same time, more empty stations means a lower output since the work content is distributed among fewer workers.

49.3 Defined Range Up and Down

You can also train workers in more than just the adjacent station. In the example below, workers #3 and #5 are qualified in four stations. This will give you even more flexibility to adapt to cycle-time fluctuations, but will also require you to spend more time on training.

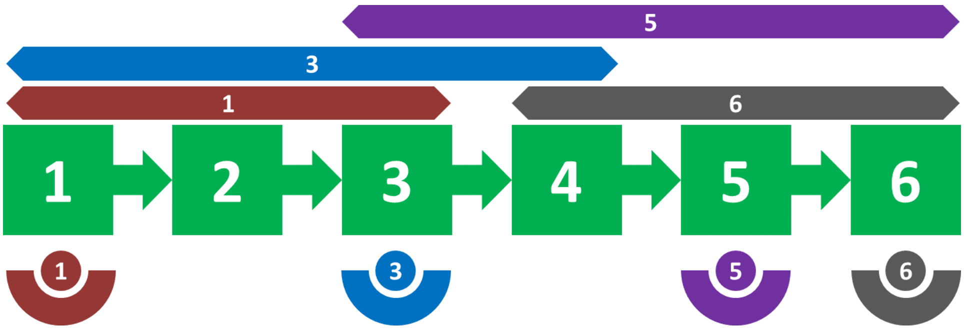

Figure 391: UpDown Trained Stations Wide Range (Image Roser)

If you continue this qualification, you eventually will end up with a normal bucket brigade. All workers are qualified in all stations, and can move back and forth freely as required by the cycle-time fluctuations. For more details see my two posts on the The Lean Bucket Brigade Part 1 and Part 2.

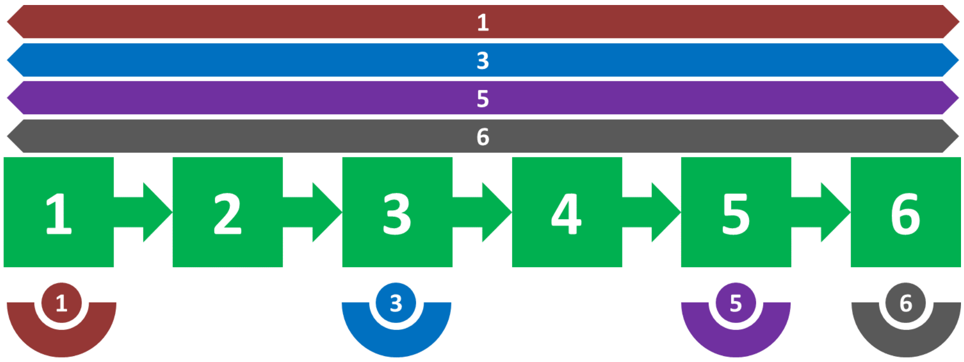

Figure 392: UpDown Trained Stations Bucket Brigade (Image Roser)

49.4 Caveat: Make Sure Cycle Is Completed!

There is one thing you need to pay attention to: **Each worker must complete a full cycle before moving to another station!** The handover of parts or workstations must be only at the end of a cycle for this part or workstation. If you hand over the part halfway, you will have a significant risk of something being forgotten. The first worker can think that the next one will do a sub-task, while the next worker can think the previous one has done it already. In the end you may have parts that miss a sub-step due to a botched handover. Workers often mistakenly believe that this will not happen, but believe me, sooner or later you will have a part missing some work steps. It probably will be easiest to avoid any halfway handovers, and allow handovers only for completed processes with a part or workstation.

49.5 Where to Put the Gaps

The approach is often pretty straightforward, but sub-par assignments may still lose flexibility. As mentioned above, you need some holes to allow workers to fluctuate between stations, and an overlap of the qualifications of adjacent workers. If possible, try to distribute these overlaps across the production line. If the overlaps are clustered together or if the workers are clustered together, you again lose some flexibility. In the example below, stations #3 and #4 can be worked on by three workers each, whereas the other stations have less qualified workers. At the same time, workers #3 and #4 are often sandwiched between workers #2 and #5, and may not be able to move to adjacent stations very often. In the example below, consider reassigning different overlaps to distribute them more evenly.

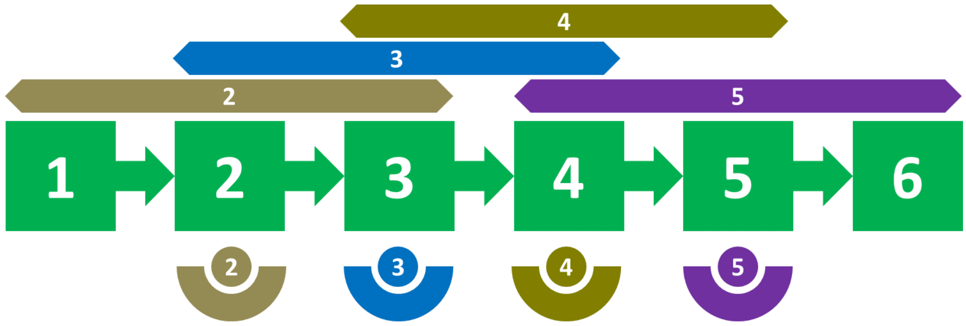

Figure 393: UpDown Trained Stations Locations (Image Roser)

This location of the gabs or holes needs to be considered when training workers. It is of no use if the workers are trained in a way that ends up with no worker being qualified for a station, or the qualifications for one worker being far apart, or the gaps being ill-distributed as in the example above.

49.6 On Timing, Customer Takt, Station Takt, and Worker Takt

It also helps if the workstations all have a similar workload (i.e., a similar cycle time). This is called line balancing, and I have written a whole series of posts on line balancing. This way you can also run the system with different numbers of workers. If you need a lot of parts, man the line completely. If you need less, run the line with less workers. At a bare-bone minimum, you can run the line with one person (although in this case this one has to be qualified for all stations). While due to the improved flexibility you do not need a perfectly balanced line, try to avoid larger disparities between cycle times unless you have either calculated or tested it beforehand.

If you plan to run the line with all stations manned, the customer takt must not be faster than the takt of each process. If you plan to run the line with less workers, the customer takt must not be faster than the average work content of each worker (plus a bit for walking).

Assume the above example of a line with 6 stations and a work content of 1 minute per station. If all stations are manned, then each station (and each worker) has a takt of 1 minute, and your line will produce one part every minute (For simplicity sake I ignore OEE losses here). Changing the number of workers usually does not change the work content, hence the takt time (i.e. the output of the line) changes. If you have only four workers manning the line, each station still has 1 minute of work content on average. With four people doing six minutes worth of work, your takt time just increased to 6/4 = 1.5 minutes. For three workers your takt time changes to 6/3 = 2 minutes. At a bare bone minimum of one worker, your takt changes to 6/1 = 6 minutes.

So, this is it on the One Up One Down approach, a variation of the bucket brigade. I hope this post was interesting for you, and you can use this in your production system. Now, go out, balance your line, manage your fluctuations, increase your flexibility, and **organize your industry!**

P.S.: This blog post is based on a suggestion by, and with input from, Richard Rahn from Leonardo Group Americas. Many thanks, Richard!

50 More Reasons for Working Less

Christoph Roser, December 10, 2019, Original at https://www.allaboutlean.com/more-reasons-for-working-less/

Figure 394: Working Woman looking at Clock (Image Pra Chid with permission)

You probably hate long drudging hours working in the office and feeling really worn out at the end of the day. Maybe you're sitting in an office working right now (and of course your reading my blog means you are working), knowing that you will be worn out at the end of the day. This post is especially for you, looking at work hours and productivity. The good news is: **Less is more, but don't overdo it!**

50.1 Industrial Revolution

Figure 395: Typical early Industrial Revolution employee … (Image Lewis Hine and Roser in public domain)

During the early years of the Industrial Revolution, the goal of the factory owner was to have his employees work as long as possible. Due to the lack of electric light, this usually meant from sunrise to sunset, sometimes even reaching sixteen hours during summer. The factory owners justified this by reasoning that if the worker was working, he had no time to get drunk and make other trouble – although I suspect that in the back of their minds was also the idea that longer work for the same wage meant more profit. With the onset of electric light around 1880, there were attempts for even longer working hours.

Unsurprisingly, workers did NOT like to work sixteen hours per day, and in 1830 unions in England discussed an eight-hour day … meaning they cut the maximum work time in half. But it was not quickly adopted, and only in 1847 did the Factory Act decree a ten-hour workday … but this was not really enforced. Only from 1867 onward was a ten-hour workday common, and only from 1900 onward was it a standard. The British Navy even switched to an eight-hour workday in 1891 for many of its employees.

Most people expected that a reduction in work hours increases productivity per hour, due to the workers being better rested. What they absolutely did NOT expect was that not only did the hourly productivity go up, the **daily productivity** went up too! Frequently workers were able to produce more in eight hours than they did in fourteen to sixteen hours before. Nowadays you would call this a win-win situation, where the factory owner got more stuff produced and the worker got more free time for leisure and other private activities. Their work-life balance increased significantly.

50.2 Frederick Taylor in the USA

Figure 396: Frederick Winslow Taylor (Image unknown author in public domain)

Similarly, Frederick Winslow Taylor (the father of Scientific Management and inventor of HSS Steel) experimented with work hours in the USA. As a consultant he often advocated a reduction of work hours to improve productivity. In a plant where women were inspecting ball bearings he gradually reduced daily work time from 10.5 hours to 10 to 9.5 to 9 and finally to 8.5 hours. He found that the **daily output increased by 33%, and quality got better too!**

50.3 Five-Day Work Week

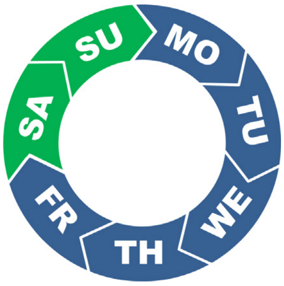

Figure 397: 5 Day Work Week (Image Roser)

Throughout history, not only the daily work time but also the number of weekly workdays have been reduced. Around 1850, many people worked six or even seven days per week. Eventually companies and countries started reducing the work days. Ford reduced the work week to five days in 1926. The Fair Labor Act in the US (1938) reduced the work week to forty hours, resulting in a two-day weekend. In Germany, unions successfully promoted a forty-hour five-day work week. Their campaign had a poster with a young kid stating "On Saturday Daddy belongs to me!"

As a result, consumption increased. People had more time to spend money on stuff.

50.4 Microsoft Japan Four-Day Work Week

Microsoft

Figure 398: Microsoft Logo (Image Microsoft for editorial use)

There are also modern-day examples. Just recently Microsoft Japan experimented with a four-day work week. Now, Japan has some very brutal working hours (see my post on The Dark Side of Japanese Working Society). In August 2019, Microsoft Japan tried out a radical four-day work week without reducing pay. They gave all employees special paid leave for Friday. Effectively, they reduced the work time by 20% (actually 25.4% since there were five Fridays in August).

The result was again surprising. **They determined an overall 40% productivity increase despite reducing work time by 25%!** (Although it is not quite clear how "productivity" was measured.) On hard data they measured 58.7% fewer pages printed, and 23.1% less power consumption. Meetings also became much shorter. Unsurprisingly, 92% of the employees said that they liked the four-day work week, and truly enjoyed the extra day off. Unfortunately for the employees, Microsoft went back to a normal five-day work week at the end of the trial. It is to be seen (hoped) that they may repeat this trial or even implement it permanently.

Many other places experimented with a four-day workweek, as for example parts of Utah, Hawaii, Gambia, Romania, New Zealand, and the UK. Many employees enjoyed the extra day off, but others felt more pressure to complete the work in less time.

50.5 Statistical Data

Figure 399: Idle Worker (Image Milkos with permission)

There seems also to be a statistical relation when comparing the work hours, GDP, and happiness of different countries. Below is a plot of the GDP per person (purchasing power parity) in 2018 (latest data) against the annual work hours (2016, latest data). There seems to be a clear trend that richer countries work less. The full data set is available here as an Excel file.

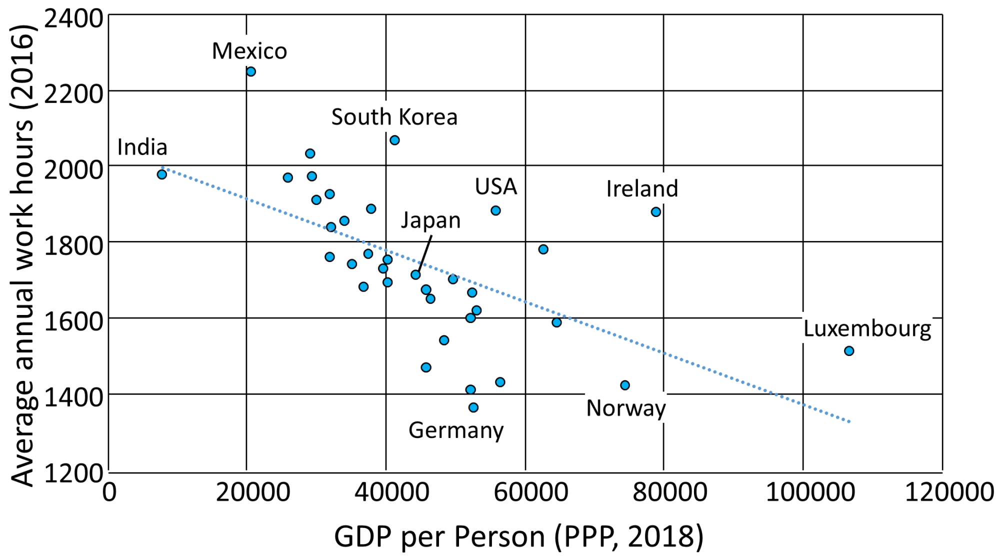

Figure 400: GDP and Work Hours (Image Roser)

Similarly, plotting the happiness of these countries (as measured by the World Happiness Index) against GDP also shows a clear relation. There is also a strong relation between work hours and happiness.

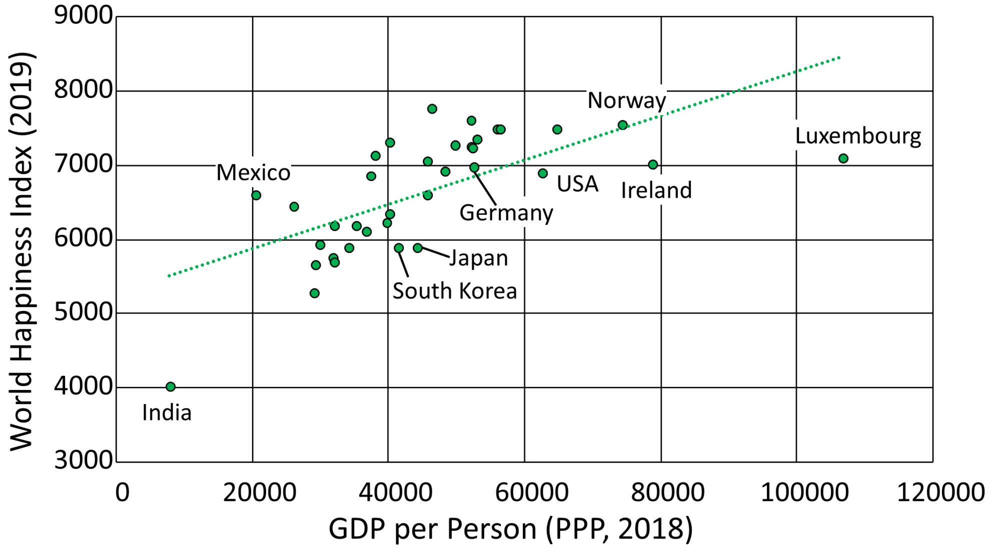

Figure 401: GDP and Happiness (Image Roser)

50.6 Correlation Is Not Causation

Now, looking at this data, the extrapolation is clear: If nobody works, we would be the richest and happiest country in the world … although while I feel that while this linear extrapolation is attractive, it is probably not true. I doubt that we would be wealthy if nobody worked. Hence, the trend of reducing work hours while producing even more probably would not go on indefinitely. There are also examples where a reduction in work hours did not work out well, as for example the thirty-five-hour work week in France.

It may also apply only to a lesser degree for some situations. If you work at an assembly line that has a certain takt, then the output is indeed very closely correlated to the working time. Yet, measuring productivity in office jobs is tricky, and hard numbers are difficult to come by. I personally believe that a good mix of work and leisure is important, but I find it really hard to give concrete advice on working hours. It probably differs from industry to industry, from task to task, and even from person to person. But I am convinced that more working hours is not always better. In any case, despite the lack of solid advice, I hope this blog post was interesting to you. Now **go out, take time to relax, and organize your industry!**

P.S.: Many thanks to Rapinder Sawhney from the University of Tennessee for his great presentation at the ELEC 2019 in Milan that gave me the inspiration for this post!

51 When and How to Use Extra Kanban

Christoph Roser, December 17, 2019, Original at https://www.allaboutlean.com/extra-kanban/

Figure 402: Dog with bandage and first aid kit (Image smrm1977 with permission)

Kanban (and pull systems in general) are a beautiful way to manage production. While the number of kanban cards should be verified periodically, for certain situations, however, it may be sensible to have additional kanban prepared. These extra kanban are added to the system on short notice to alleviate symptoms of other problems. However, please do not see this as a new and cool additional feature to your kanban system. Instead, if you have to use extra kanban, then something is going wrong. Such extra kanban are only an emergency fix and do not solve the underlying problem. Still, in some cases you do need an emergency fix. Let me explain:

51.1 Sequence of Pull Systems

Often, kanban (or CONWIP) loops are in sequence to establish pull across the entire value stream. The image below is a simple example with three subsequent kanban loops. The supermarket acts as a buffer against fluctuations between the three processes.

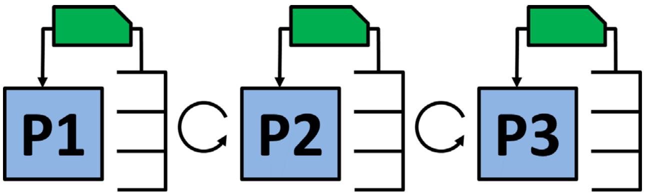

Figure 403: Three Kanban Loops in a Row (Image Roser)

51.2 The Original Problem: Unusual Fluctuations

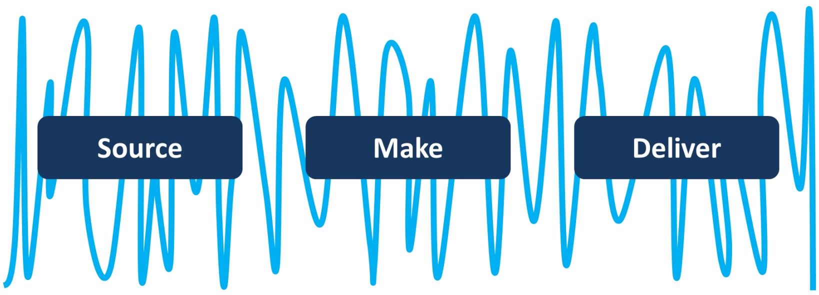

Figure 404: Source Make Deliver Fluctuations (Image Roser)

Ideally, supermarkets are set up to buffer fluctuations in supply and demand. Sometimes the production process may be faster and the supermarket fills up a bit. Sometimes it is slower and the supermarket empties a bit. Similarly, if the customer orders more or less than usual, the supermarket becomes a bit emptier or fuller. There should be a product in the supermarket whenever the subsequent process or customer is requesting one. If this often does not happen, then you either have too few kanban or your supplying process is too slow and can't keep up.

However, sometimes you have rarer and larger fluctuations. These could be **unexpected**. For example, a machine suddenly breaks down and it will take days to get the spare parts. These could also be **expected**, in which case they can be planned for. For example, you know your ski boots sell a lot in fall but very little in spring (seasonality). Or you have a planned annual maintenance that shuts down a machine for a week. In any case, you have rare and large fluctuations. It would make no sense to set up the pull system to cover these few and far-between but extraordinary fluctuations, just to have tons of material sitting around for the rest of the time.

51.3 The Quick Fix: Extra Kanban

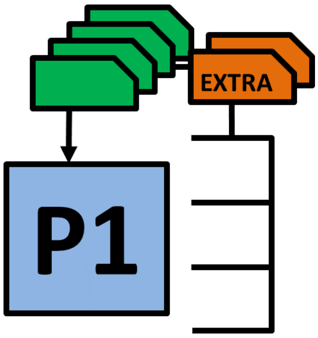

Figure 405: Extra Kanban Extra (Image Roser)

This brings me to my quick fix: Extra kanban! You keep a stack of extra kanban to insert into the system temporarily, and then remove these cards again once the fluctuation has passed. You do build up inventory, but only temporarily for the duration of the extraordinary event, and remove kanban again after the disruption has passed.

51.4 When It Is Useful…

However, these extra kanban can help you only in some cases, not all. The goal of these extra kanban is to decouple these fluctuations with extra material. However, this is not always possible.

- It depends on whether the fluctuations temporarily **reduce capacity or demand** (e.g., breakdowns, maintenance) or if they **increase capacity or demand** (e.g., seasonality).
- It also depends on the **location of the bottleneck**. It makes no sense to increase inventory if the bottleneck will never be able to catch up anyway.
- Finally, it depends on whether you **know the disruption beforehand** and can react before it happens (e.g., seasonality, planned maintenance) or if you **know it only after it hits you** (e.g., breakdowns).

Overall, extra kanban can only help if the build-up of an extra capacity buffer (i.e. the extra kanban) are helpful. Let me give you a few examples. The example below shows three pull loops in sequence. The process in the last loop has a breakdown that will take longer to fix. The bottleneck is somewhere before this disrupted process. In this case it is viable to temporarily increase the number of kanban in the pull loop before the disruption. The bottleneck can keep on working and build up inventory. After the disruption ends, the subsequent processes can catch up with the previous processes and reduce the inventory again. This reduces the overall loss of capacity due to the disruption.

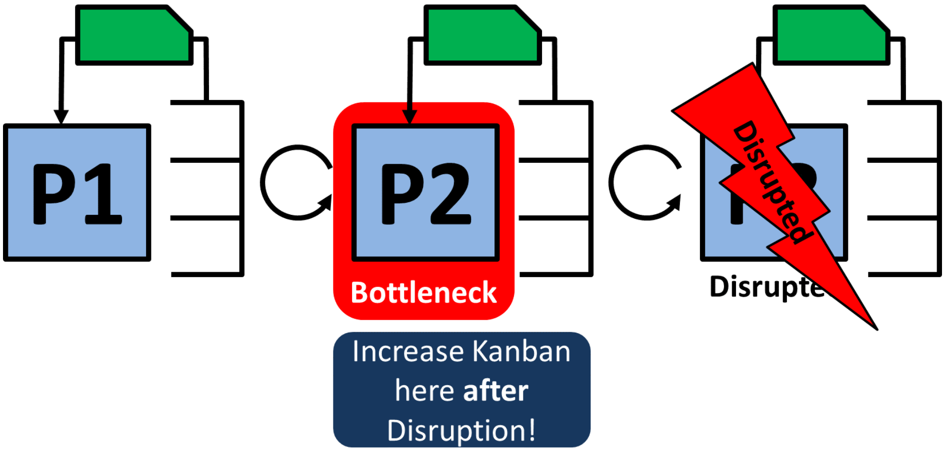

Figure 406: Extra Kanban Disrupt after BN (Image Roser)

It is different, however, if the disruption happens in a process in front of the bottleneck. Here we would need to know the disruption beforehand to build up inventory in the disrupted loop **before** the disruption. During the disruption, the bottleneck keeps on working using the inventory built up in preparation for the disruption. Naturally this works only for disruptions that are known beforehand, like planned maintenance.

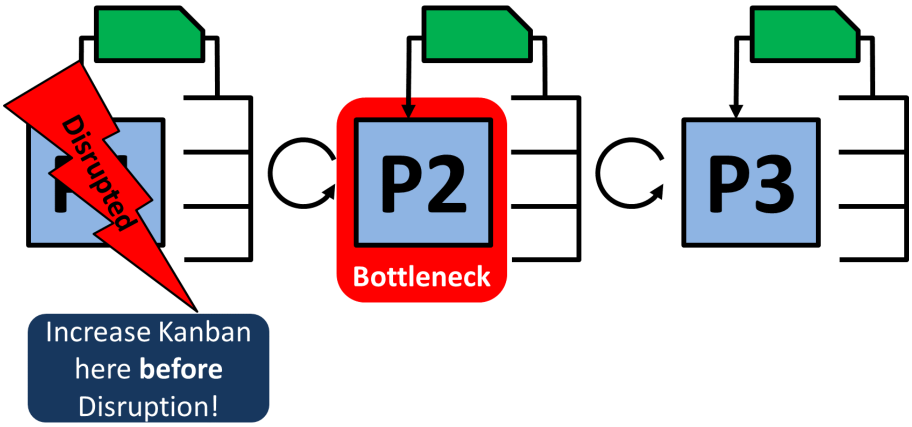

Figure 407: Extra Kanban Disrupt before BN (Image Roser)

It is also possible that a disruption does not decrease capacity but increases capacity or demand. The most common example is seasonality, where the customer temporarily orders more parts that the pull system is set up for. Usually this is resolved through a seasonal adjustment of the pull systems and the provided production capacity. It is also possible to build up an inventory beforehand using extra kanban if the capacity is not enough to satisfy peak demand. Here, too, we would need to know this beforehand to prepare.

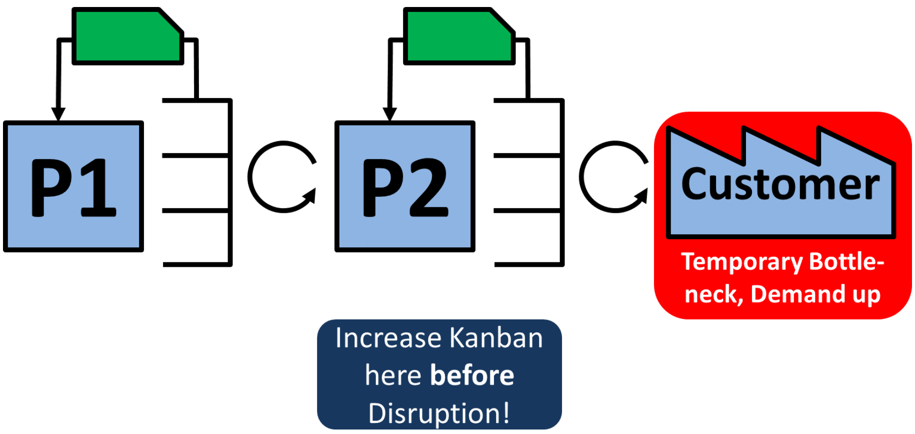

Figure 408: Extra Kanban Customer Up (Image Roser)

In some cases disruptions cannot be remedied with extra kanban. This could be a disruption before the bottleneck, as shown farther above, where we do not know the disruption beforehand. In this case we cannot make additional parts for the bottleneck after the fact.

Below is another example, where the disruption happens at the bottleneck. Having more parts before the bottleneck won't help, since the bottleneck can't use more parts than provided by P1 anyway. At best we can stock up on finished goods to keep on supplying the customer, but only if we know the disruption beforehand.

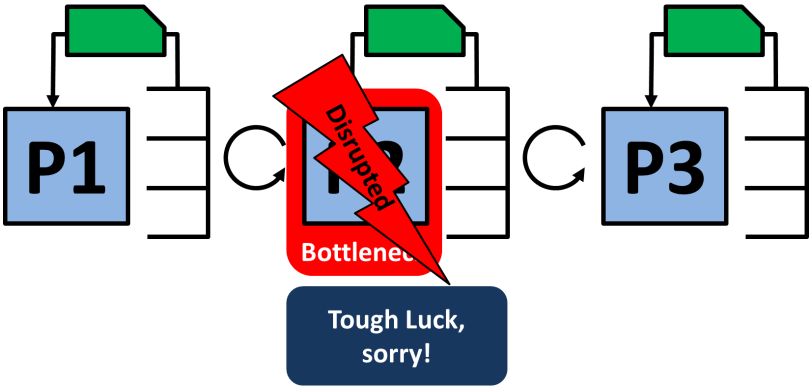

Figure 409: Extra Kanban Disrupt at BN (Image Roser)

Overall you have to think it through for your system to see if additional kanban can help or not. In the worst case, you create additional inventory without any benefit for the entire production system.

51.5 How to Do It

It helps to have these extra kanban clearly marked as temporary. Besides some additional label, they could also be of a different color for easier identification.

Another question is which products to prepare. If the disruption is after the bottleneck, it is sensible to build up inventory (i.e., add extra kanban) for high-runner products. In this case you are more confident that you actually need the material. Low runners can then be prioritized using the available capacity. If the disruption is before the bottleneck, however, you may have to add extra kanban for all items produced.

You should also think about how many extra kanban you should add. It could be enough to cover the disruption, but you may be constrained by the available storage space. Often it can make sense to cover only part of the disruption to avoid excess storage and handling fees.

The additional kanban cards are added in the list of kanban cards that have to be produced. Ideally they are not a big block of the same part type, but mixed in suitable lot sizes. When you want to remove cards, you take away only cards that are not attached to an item. If such a card is attached to an item in the supermarket, then you have to wait until the item is taken out of the supermarket. For low runners you may also swap an extra kanban card with a normal one taken out of the production queue before production.

51.6 Why It Is Only a Quick Fix

Please note that this is not a fancy new method for pull, but a quick fix of the symptoms of another problem. Much better would be, of course, to prevent the disruption in the first place. Better than having extra kanbans is to have no larger breakdowns, or to improve and shorten maintenance time. Hence, using extra kanban is usually a sign of problems rather than excellence. Now, **go out, reduce fluctuations so you don't need extra kanban cards, and organize your industry!**

P.S.: This post was inspired by a discussion with Karl-Ludwig Blocher.

52 115 Years after the Birth of Joseph Juran

Christoph Roser, December 24, 2019, Original at https://www.allaboutlean.com/115-years-after-joseph-juran/

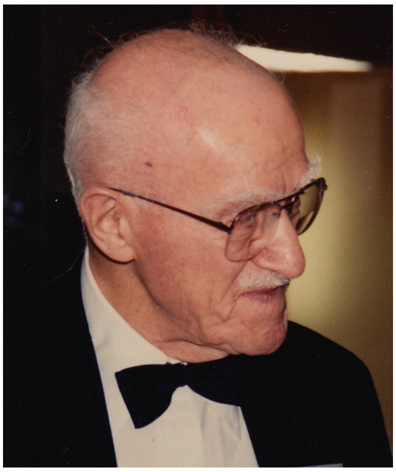

Figure 410: Joseph Juran (Image FELDSPATH under the CC-BY-SA 4.0 license)

On Christmas Eve 115 years ago, Joseph Moses Juran (December 24, 1904 – February 28, 2008) was born. He was a highly respected and very influential quality guru. His work not only helped the United States, but also changed Japan, possibly even more than that of his better-known colleague Edwards Deming. Time to look back on his life's impact on the world.

52.1 Early Life

Figure 411: Brăila, Romania around 1900 (Image unknown author in public domain)

Joseph Juran was born in the city of Brăila, Romania, in 1904 as one of six children of to Jakob and Gitel Juran. When he was three, his family moved to Gura Humora. In 1912 the family immigrated to Minneapolis, Minnesota, hence he and the rest of his Jewish family escaped the horror of the Holocaust.

Figure 412: Juran family around 1910. Joseph is next to his mother. (Image unknown author in public domain)

What follows is a muster story of a successful immigrant starting out in poverty but becoming one of the most respected authorities in his field. The family lived in a shack, and Joseph pitched in by selling newspapers … where he was also cheated by one of his customers. Later on, he worked numerous jobs, including as a grocer assistant, a (highly under-qualified) bookkeeper, a wrapper of packages, a janitor, a printer, an editor, a shoe salesman, and more.

As a student he was an avid reader and also very good at mathematics. The 1920, the same year his mother died, he graduated high school. He then obtained a bachelor's degree in engineering from the University of Minnesota in 1924.

52.2 Working at Western Electric's Hawthorne Works

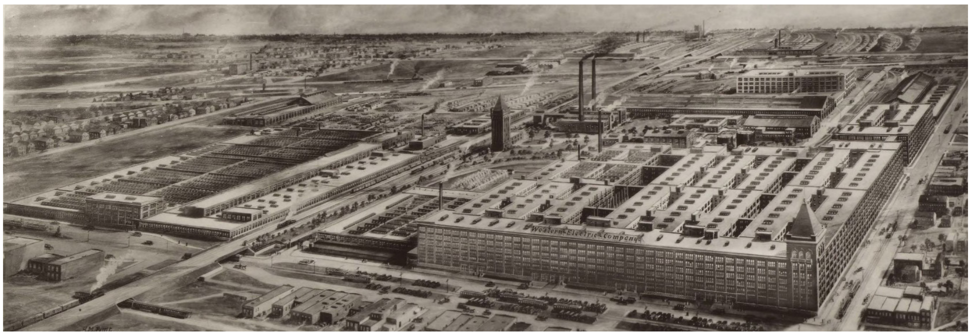

Figure 413: Hawthorne Works, Western Electric, 1925 (Image Western Electric Company in public domain)

In 1924, Joseph left his dad's home and started working at the Western Electric (later AT&T) Hawthorne Works, where he was assigned the inspection branch (nowadays known as quality control). Later he became a "troubleshooter" who investigated complaints.

Juran climbed up the hierarchy, eventually becoming head of the works inspection control division. However, in retrospect he considered himself a pretty bad manager. He also learned more and more about quality. By today's standards, Hawthorne's quality was not impressive, but by the standards of 1920, it was a leader in the field of quality. Juran was involved (but not instrumental) in the development of statistical quality control, nowadays often known by its abbreviation, SPC. He did, however, develop the training courses and trained many employees on SPC, including top management. As a side note, Juran also used his statistical skills to win quite a bit of money at roulette.

A later colleague, W. Edwards Deming, also started at Hawthorne in 1924, but they did not meet there. Well-known Walter Andrew Shewhart had joined the company only slightly earlier in 1918, and developed statistical control charts as well as gave the inspiration for Plan Do Check Act (PDCA). Shewhart influenced both Juran and Deming.

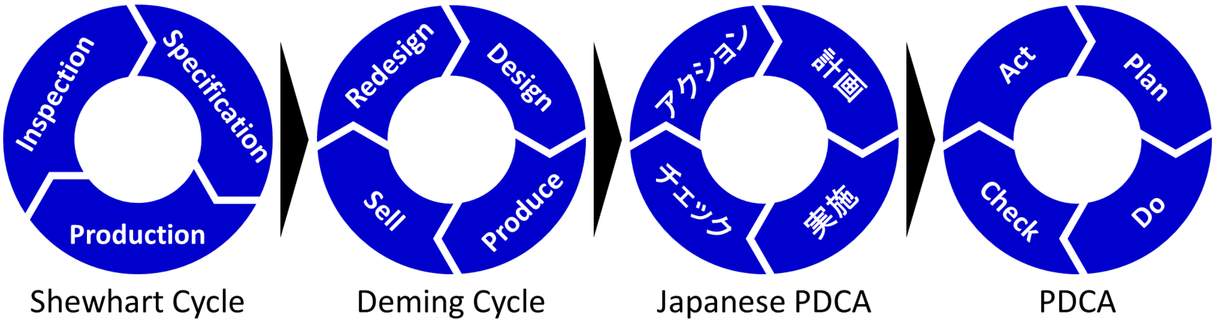

Figure 414: PDCA History (Image Roser)

The Hawthorne plant is also famous for the Hawthorne experiments, (purportedly) showing that merely observing workers improves quality and productivity, although the data is based only on five female workers, two of whom were exchanged during the study. Subsequent research sheds doubt on these results. Neither Juran nor Deming were involved in these experiments.

While at Hawthorne, Juran married the love of his life, Sadie Shapiro, in 1927. Soon he had children: Robert Arnold (1928), Sylvia Louise (1930), and Charles Edward (1931). The family survived the Depression of 1930 well. While his salary was cut from $380 to $285, it was still a good salary, especially with prices going down similarly. Nevertheless, as back-up insurance, Juran enrolled in law school in 1931, then graduated in 1935 and was admitted to the Illinois bar in 1936 (although he never practiced). In 1937, still with AT&T, he moved to New York, although his career stalled (by his own account in retrospect, due to his own stupidity for starting a fight with his boss for no reason whatsoever).

52.3 Writing, War, Academia, and Consulting…

However, in New York he got in contact with many industrial societies like the American Society of Mechanical Engineering (ASME), the Society for the Advancement of Management (SAM), and the American Management Association (AMA), and the National Industrial Confidence Board (NCIB). He soon enjoyed writing papers for the journals of these societies.

During World War II he was requested, in 1941, by one of his former employees to join the statistics division. Juran worked in the Lend-Lease Administration and Foreign Economic Administration. Toward the end of the war, Juran decided to go freelance, leaving AT&T in 1945 and becoming an independent consultant, initially as part of a consultancy but later on as a quite-successful freelance consultant. He consulted for various companies and industries, including Gillette razors, watch makers, carpet companies, automotive components, latex, optical instruments, food, Otis elevators, Xerox copiers, the US Navy, and many more. In 1951 he published his first edition of the now famous Quality Control Handbook.

After the war, he taught for a few years at New York University as an adjunct professor. He also became a good friend of another giant in consulting, Peter Drucker (1909 – 2005, a well-known management consultant).

52.4 Juran in Japan

In 1953 Juran followed an invitation to Japan to lecture about quality. This was the start of frequent trips to Japan. He visited numerous companies and also gave many well-received lectures on quality control. He was supported by the Japanese Union of Scientists and Engineers (JUSE), especially its managing director, Ken-Ichi Koyanagi. He expanded his lectures from focusing on quality to teaching about management in general. In Japan he is revered for his teaching, and many Japanese believe that the teaching from Juran was more helpful and better to apply than the teaching of Deming.

Overall, Juran visited Japan at least ten times, all but one following an invitation from Japan. His last visit was in 1990. Afterward, at age 86, his family told him to no longer travel internationally due to his age.

52.5 Friction with Deming

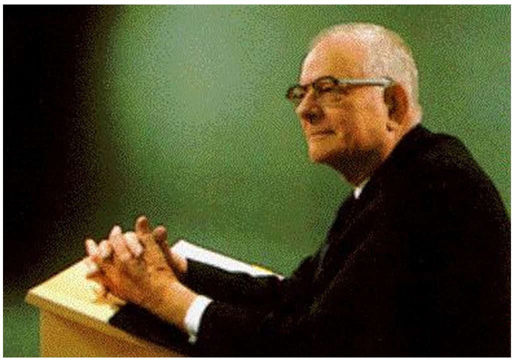

Figure 415: W. Edwards Deming (Image FDA in public domain)

During Juran's time in Japan, another even-better-known person was also working frequently in Japan: W. Edwards Deming. Their biographies have many similarities: they worked at Hawthorne, were influenced by Shewhart, and worked on quality, although Deming looked more at the mathematical side.

However, they never saw each other at Hawthorne and first met during the 1940s in Washington when Juran was the chair of a meeting that Deming attended. Initially they liked each other.

Both were invited to Japan to teach quality. However, Deming was little known outside Japan until a 1980 TV documentary *If Japan Can, Why Can't We?* turned him into an instant celebrity in the US too.

However, it seems that they frequently disagreed on topics and over time grew apart. Deming believed all quality issues can be solved with statistics, whereas Juran believed that you need more than statistics to improve quality. Juran was shocked when Deming talked badly about the US's Malcolm Baldrige National Quality Award simply because it was not named after Deming like the Deming Prize in Japan. Juran eventually believed that Deming vastly overestimated his contribution to Japanese quality, thinking that he (Deming) was the only reason Japan succeeded in quality. Juran, on the other hand, thought that he had a larger influence on quality in Japan.

52.6 His Legacy

Juran himself started the Juran Institute in 1979, which still teaches quality. He semi-retired at the respectable age of 90. Nevertheless, he wrote his (highly readable and recommended by me) autobiography Architect of Quality at the age of 92. He died on February 28, 2008, at age 103 due to a stroke.

Juran is considered one of the great minds in quality control. Some called him "the man who taught quality to the Japanese." He is often mentioned in Japan along with Deming. Many of his books were translated into Japanese by JUSE, as were his lectures and papers.

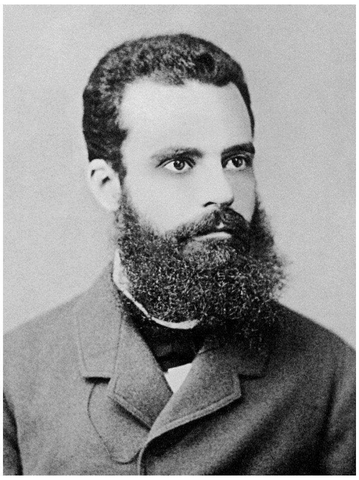

Figure 416: Vilfredo Pareto (Image unknown author in public domain)

Juran did not invent the Pareto principle – that was the Italian Vilfredo Pareto (1848 – 1923) – but Juran rediscovered it and applied to to quality. He moved quality away from mere statistics, but included the human aspect in quality, promoting the training of managers in quality aspects. In his view, bad quality did not originate with the workers but was a flaw in management.

His approach to improve quality is known as the **Juran Trilogy**, consisting of **quality planning** (product or process design that enables quality); **quality control** (measuring and controlling quality aspects; SPC would be found here); and **quality improvement** (take actions to improve quality).

Overall, Juran had a profound impact on quality, and his methods and teachings are still crucial for modern-day quality. Now, **go out, improve your quality, and organize your industry!**

52.7 Sources

- Juran, Joseph M. Architect of Quality. McGraw-Hill, 2003
- Defeo, Joseph, and Joseph M. Juran. Juran's Quality Handbook: The Complete Guide to Performance Excellence. 5th edition. New York: McGraw-Hill Education, 1999.

53 Image Credits

Most of the images in this book were created by me, Christoph Roser. Many other images are used under the Creative Commons license, which I support by making almost all of my images available under this free license and by uploading many of my images to the Wikimedia Commons. The Creative Commons license requires giving credit to the source, if possible with links, and a link to the original license text. Here are the links to the Creative Commons licenses used in this book.

- CC-BY 2.0 at https://creativecommons.org/licenses/by/2.0/
- CC-BY 3.0 at https://creativecommons.org/licenses/by/3.0/
- CC-BY 4.0 at https://creativecommons.org/licenses/by/4.0/
- CC-BY 3.0 Brazil at https://creativecommons.org/licenses/by/3.0/br/
- CC-BY 4.0 at https://creativecommons.org/licenses/by/4.0/
- CC-BY-SA 2.0 at https://creativecommons.org/licenses/by-sa/2.0/
- CC-BY-SA 2.5 at https://creativecommons.org/licenses/by-sa/2.5/
- CC-BY-SA 3.0 at https://creativecommons.org/licenses/by-sa/3.0/
- CC-BY-SA 3.0 Netherlands at https://creativecommons.org/licenses/by-sa/3.0/nl/
- CC-BY-SA 4.0 at https://creativecommons.org/licenses/by-sa/4.0/

A complete list of creative commons licenses can be found at https://creativecommons.org/licenses/. Here is the list of images from other authors used in this book, including the author (if known), the license, and the source. Many thanks to all the other image authors for providing their images under a free license!

- Figure 1 by MichaelMaggs under the CC-BY-SA 2.5 license, available at https://commons.wikimedia.org/wiki/File:Chess_piece_-_White_pawn.JPG
- Figure 2 by Jim.belk in public domain, available at https://commons.wikimedia.org/wiki/File:One-Twist_Trefoil.png
- Figure 3 by Hamilton Public Library in public domain, available at https://www.flickr.com/photos/hpllocalhistory/14072356987
- Figure 4 by Clker-Free-Vector-Images in public domain, available at https://pixabay.com/vectors/checklist-lists-business-form-41335/
- Figure 5 by Wolfgang Zwanzger with permission, available at https://www.allaboutlean.com/muda-mura-muri/overworked-crying-businessman/
- Figure 6 by Netfalls with permission, available at https://www.allaboutlean.com/hoshin-kanri-1/a-serious-man-in-a-black-suit-dark-background/
- Figure 7 by Joakant in public domain, available at https://pixabay.com/photos/divers-scuba-divers-diving-668777/
- Figure 9 by Kyle Gese in public domain, available at https://www.spangdahlem.af.mil/News/Article-Display/Article/293755/checkmate-saber/
- Figure 12 by Jim.belk in public domain, available at https://commons.wikimedia.org/wiki/File:One-Twist_Trefoil.png
- Figure 13 by Alan Light under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:ChessSet.jpg
- Figure 14 by Marietjieopp with permission, available at https://www.allaboutlean.com/hoshin-kanri-2/a-pile-of-wooden-chest-pieces-isolated-on-a-wooden-background/
- Figure 15 by Roser adapted from Cherubino under the CC-BY-SA 4.0 license, available at https://commons.wikimedia.org/wiki/File:Four_Poplars_in_four_seasons.JPG
- Figure 16 by Poznyakov with permission, available at https://www.allaboutlean.com/creative-provocation-and-others/mad-business-people-office-in-high-school-celebrate-successful-p/
- Figure 18 by White Spirit with permission, available at https://www.allaboutlean.com/hoshin-kanri-3/white-chess-set/
- Figure 19 by Artem Alexandrovich with permission, available at https://www.allaboutlean.com/hoshin-kanri-3/bigstock-private-jet/
- Figure 20 by NASA/Pat Rawlings in public domain, available at https://commons.wikimedia.org/wiki/File:Humans_on_Mars.jpg

- Figure 21 by Rido81 with permission, available at https://www.allaboutlean.com/hoshin-kanri-3/portrait-of-successful-group-of-business-people-at-modern-office/
- Figure 23 by Mangostar with permission, available at https://www.allaboutlean.com/hoshin-kanri-3/bigstock-understanding-manager/
- Figure 26 by Dr Jacek Filek in public domain, available at https://commons.wikimedia.org/wiki/File:Chess_for_Three_-_Hexagonal_Board.jpg
- Figure 27 by bramgino with permission, available at https://www.allaboutlean.com/hoshin-kanri-x-matrix/fotolia_sneaky-consultant-x-matrix/
- Figure 30 by OpenClipart-Vectors in public domain, available at https://pixabay.com/vectors/comic-characters-critical-thought-2026313/
- Figure 31 by Glen Scarborough under the CC-BY-SA 2.0 license, available at https://www.flickr.com/photos/photographerglen/5921505892
- Figure 33 by Hike The Monicas under the CC-BY-SA 4.0 license, available at https://commons.wikimedia.org/wiki/File:True_North_Mount_Allen.fw.png
- Figure 35 by banderchenno with permission, available at https://www.allaboutlean.com/kanri-noryoku-program/black-and-white-chess-pieces-fallen-on-chessboard/
- Figure 36 by Gajus with permission, available at https://www.allaboutlean.com/kanri-noryoku-program/business-hierarchy-ranking-and-strategy-concept/
- Figure 37 by ronstik with permission, available at https://www.allaboutlean.com/kanri-noryoku-program/business-leadership-teamwork-power-and-confidence-concept/
- Figure 38 by STILLFX with permission, available at https://www.allaboutlean.com/kanri-noryoku-program/black-chess-pieces-on-board/
- Figure 43 by Thomas Karol in public domain, available at https://www.ellsworth.af.mil/News/Photos/igphoto/2001807165/
- Figure 44 by style-photographs with permission, available at https://www.allaboutlean.com/digital-kanban/african-worker-as-warehouse-specialist-apprentice-controls-the-i/
- Figure 45 by Rawpixel.com with permission, available at https://www.allaboutlean.com/digital-kanban/hands-holding-kaizen-word-in-balloon-letters/
- Figure 46 by studiostoks with permission, available at https://www.allaboutlean.com/lean-and-industry-4-0/robot-fight/
- Figure 47 by MaxPixel in public domain, available at https://www.maxpixel.net/Seemed-Amusement-Park-Roller-Coaster-Theme-Park-3100041
- Figure 57 by Peter Craven under the CC-BY 2.0 license, available at https://commons.wikimedia.org/wiki/File:Estonia_sand_stockpiles_under_conveyor_%286256989788%29.jpg
- Figure 59 by TerriC in public domain, available at https://pixabay.com/photos/yogurt-breakfast-berries-food-3018152/
- Figure 60 by Frank Reynolds in public domain, available at https://commons.wikimedia.org/wiki/File:Uriah_Heep_from_David_Copperfield_art_by_Frank_Reynolds.jpg
- Figure 61 by WavebreakMediaMicro with permission, available at https://www.allaboutlean.com/just-in-sequence-definition/warehouse-worker-checking-the-inventory/
- Figure 66 by Boffy b under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:IBM_PC_5150.jpg
- Figure 67 by MaxPixel in public domain, available at https://www.maxpixel.net/Android-Robot-Cybernetics-Cyborg-Futuristic-3310191
- Figure 68 by Minerva Studio with permission, available at https://www.allaboutlean.com/toyota-lot-size-one/angry-businessman/
- Figure 69 by bisonlux under the CC-BY 2.0 License, available at https://commons.wikimedia.org/wiki/File:Praia_de_Copacabana_-_Rio_de_Janeiro,_Brasil.jpg
- Figure 79 by logoboom with permission, available at https://www.allaboutlean.com/cobacabana/manager-with-playing-cards/
- Figure 80 by Morio under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:1924_Non-Stop_Shuttle_Change_Toyoda_Automatic_Loom,_Type_G_1.jpg
- Figure 89 by Toyota with permission, available at https://www.allaboutlean.com/toyoda-model-g/weft-fork-mechanism/
- Figure 94 by unknown author in public domain, available at https://www.allaboutlean.com/150-years-albert-kahn/1896-nettleton-kahn-trowbridge/
- Figure 95 by Trussed Concrete Steel Company for editorial use, available at https://commons.wikimedia.org/wiki/File:Kahn_System_brand_logo.jpg

- Figure 96 by chevin in public domain, available at https://commons.wikimedia.org/wiki/File:Cromford_1771_mill.jpg
- Figure 97 by unknown author in public domain, available at https://commons.wikimedia.org/wiki/File:Packard_plant_no_10_construction.jpg
- Figure 98 by Andrew Jameson under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:Highland_Park_Ford_plant.jpg
- Figure 99 by Zacharias L. under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:Giengen_Fabrikhalle_Steiff.jpg
- Figure 100 by Detroit Publishing Co. in public domain, available at https://commons.wikimedia.org/wiki/File:River_Rouge_aerial_4a25915r.jpg
- Figure 101 by Alfred T. Palmer in public domain, available at https://commons.wikimedia.org/wiki/File:River_Rouge_tool_and_die8b00276r.jpg
- Figure 102 by Michael Barera under the CC-BY-SA 4.0 license, available at https://www.allaboutlean.com/150-years-albert-kahn/william-l-clements-library/
- Figure 103 by Alexas_Fotos in public domain, available at https://pixabay.com/photos/time-is-money-coins-currency-euro-1601988/
- Figure 104 by Malene Thyssen under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:Running_hare.jpg
- Figure 105 by 3Dman_eu in public domain, available at https://pixabay.com/illustrations/time-money-horizontal-speculate-1019889/
- Figure 106 by Petr Kratochvil in public domain, available at https://www.publicdomainpictures.net/en/view-image.php?image=13018&picture=gold-coin-stack
- Figure 110 by MaxPixel in public domain, available at https://www.maxpixel.net/Clock-Figures-Coins-Funny-Bank-Note-Time-Is-Money-3344117
- Figure 111 by Petr Kratochvil in public domain, available at https://www.publicdomainpictures.net/en/view-image.php?image=13018&picture=gold-coin-stack
- Figure 113 by MaxPixel in public domain, available at https://www.maxpixel.net/University-Professor-Education-Teaching-Coach-1696568
- Figure 115 by Free Clip Art and Roser under the CC-BY-SA 4.0 license, available at https://commons.wikimedia.org/wiki/File:Black_Delivery_Man_Cartoon.svg
- Figure 133 by PathDoc with permission, available at https://www.allaboutlean.com/weak-parallel-fifo/ignorance-and-arrogance-portrait-of-a-young-man-shrugging-shoul/
- Figure 138 by Minerva Studio with permission, available at https://www.allaboutlean.com/on-firefighting/angry-senior-man-pointing-his-finger-at-somebody/
- Figure 139 by Minerva Studio with permission, available at https://www.allaboutlean.com/toyota-lot-size-one/angry-businessman/
- Figure 140 by MarianSigler {bla}, Mediatus in public domain, available at https://commons.wikimedia.org/wiki/File:Zeichen_125_-_Gegenverkehr,_StVO_1992.svg
- Figure 141 by Open Knowledge Foundation under the CC-BY 2.0 license, available at https://commons.wikimedia.org/wiki/File:Johnellis.jpg
- Figure 146 by Bundesanstalt für Arbeit for editorial use, available at https://www.allaboutlean.com/diseconomies-of-scale/arbeitsamt-logo-alt-neu/
- Figure 147 Passat by IFCAR in public domain, available at https://commons.wikimedia.org/wiki/File:2006-07_Volkswagen_Passat_Wagon.jpg and Skoda by OSX in public domain, available at https://commons.wikimedia.org/wiki/File:2009-2010_%C5%A0koda_Octavia_TSI_station_wagon_01.jpg
- Figure 148 by Wim van Rossem, Anefo under the CC-BY-SA 3.0 Netherlands license, available at https://commons.wikimedia.org/wiki/File:Cyril_Northcote_Parkinson_1961.jpg
- Figure 150 by Microsoft with permission, available at https://www.allaboutlean.com/losing-kanban/magnifying-glass/
- Figure 151 by Tomascastelazo modified by Roser under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:Birds_on_the_wire.jpg
- Figure 153 by Traimak with permission, available at https://www.allaboutlean.com/where-lean-went-wrong/team-of-three-successful-stylish-young-lawyers-on-a-black-backgr/
- Figure 154 by Tyler Olson with permission, available at https://www.allaboutlean.com/chalk-circle/male-supervisor-showing-something-to-colleague-at-warehouse/
- Figure 155 by Raymundo Valencia in public domain, available at https://www.fs.usda.gov/detailfull/r5/home/?cid=FSEPRD539719&width=full
- Figure 158 by alphaspirit with permission, available at https://www.allaboutlean.com/supervisor-workload/bigstock-drowning-in-boxes/

- Figure 167 by metropolitan Museum of Art in public domain, available at https://www.metmuseum.org/art/collection/search/552391
- Figure 169 by loubeat under the CC-BY 2.0 license, available at https://www.flickr.com/photos/thekcb/5734618076
- Figure 174 by Kenishirotie with permission, available at https://www.allaboutlean.com/how-to-do-brainstorming/solution-for-problem-for-business-metaphor/
- Figure 175 by Metropolitan Museum of Art in public domain, available at https://www.metmuseum.org/art/collection/search/322903
- Figure 176 3 lug nuts by Bindydad123 under the CC-BY-SA 4.0 license, available at https://commons.wikimedia.org/wiki/File:Proton_Tiara_in_Penang,_Malaysia_(3).jpg, 4 lug nuts by PXHere in public domain, available at https://pxhere.com/en/photo/714078, and 5 lug nuts by PXHere in public domain, available at https://pxhere.com/en/photo/1182792
- Figure 178 by Ford Motor Co. under the CC-BY-SA 2.0 license, available at https://commons.wikimedia.org/wiki/File:Kansas_City_Assembly.png
- Figure 179 by Mirco Vacca with permission, available at https://www.allaboutlean.com/mixed-model-sequencing-2/poznan-poland-december-18-2018-two-angry-lego-men-with-hats/
- Figure 181 by Metropolitan Museum of Art in public domain, available at https://www.metmuseum.org/art/collection/search/552392
- Figure 183 by Mariette M. Adams in public domain, available at https://media.defense.gov/2016/Jan/14/2001503076/-1/-1/0/160106-F-GH858-999.JPG
- Figure 189 by Metropolitan Museum of Art in public domain, available at https://www.metmuseum.org/art/collection/search/545189
- Figure 191 by Vladimir Voronin with permission, available at https://www.allaboutlean.com/production-speed-measurements/fotolia_metronom_xs-2/
- Figure 194 by Open Knowledge Foundation under the CC-BY 2.0 license, available at https://commons.wikimedia.org/wiki/File:Johnellis.jpg
- Figure 195 by Metropolitan Museum of Art in public domain, available at https://www.allaboutlean.com/mixed-model-sequencing-5/beads-frankisch/
- Figure 196 by Andy Hay under the CC-BY 2.0 license, available at https://www.flickr.com/photos/andyhay/10513264765
- Figure 198 by Wayhome Studio with permission, available at https://www.allaboutlean.com/disabled-employees-in-manufacturing-omron-taiyo-in-japan-part-2/content-female-and-male-technicians-in-special-uniform-keep-hand/
- Figure 199 by mipan with permission, available at https://www.allaboutlean.com/timing-of-flow-lines-2/time-spiral/
- Figure 200 by WavebreakMediaMicro with permission, available at https://www.allaboutlean.com/just-in-sequence-definition/warehouse-worker-checking-the-inventory/
- Figure 204 by Metropolitan Museum of Art in public domain, available at https://www.metmuseum.org/art/collection/search/587882
- Figure 208 by Andrii Zastrozhnov with permission, available at https://www.allaboutlean.com/mixed-model-sequencing-6/difficult-business-choice-long-entangled-way-concept/
- Figure 209 by alphaspirit with permission, available at https://www.allaboutlean.com/dealing-with-uncertainty/businessman-choosing-the-right-door/
- Figure 210 by Metropolitan Museum of Art in public domain, available at https://www.allaboutlean.com/mixed-model-sequencing-6/animal-shaped-beads-on-a-string/
- Figure 212 by Vladimir Voronin with permission, available at https://www.allaboutlean.com/production-speed-measurements/fotolia_metronom_xs-2/
- Figure 213 by Mirco Vacca with permission, available at https://www.allaboutlean.com/mixed-model-sequencing-7/poznan-poland-february-13-2019-two-lego-construction-worker/
- Figure 216 3 lug nuts by Bindydad123 under the CC-BY-SA 4.0 license, available at https://commons.wikimedia.org/wiki/File:Proton_Tiara_in_Penang,_Malaysia_(3).jpg, 4 lug nuts by PXHere in public domain, available at https://pxhere.com/en/photo/714078, and 5 lug nuts by PXHere in public domain, available at https://pxhere.com/en/photo/1182792
- Figure 178 by Ford Motor Co. under the CC-BY-SA 2.0 license, available at https://commons.wikimedia.org/wiki/File:Kansas_City_Assembly.png
- Figure 218 by Metropolitan Museum of Art in public domain, available at https://www.metmuseum.org/art/collection/search/454086
- Figure 226 by Metropolitan Museum of Art in public domain, available at https://www.metmuseum.org/art/collection/search/243888

- Figure 237 by Metropolitan Museum of Art in public domain, available at https://www.metmuseum.org/art/collection/search/325321
- Figure 244 by MaxPixel in public domain, available at https://www.maxpixel.net/Girl-Female-Person-Woman-Model-Pearl-Necklace-918787
- Figure 245 by Metropolitan Museum of Art in public domain, available at https://www.metmuseum.org/art/collection/search/570193
- Figure 252 by Metropolitan Museum of Art in public domain, available at https://www.metmuseum.org/art/collection/search/557242
- Figure 256 by Boffy b under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:IBM_PC_5150.jpg
- Figure 257 by RitaE in public domain, available at https://pixabay.com/photos/garbage-waste-container-waste-2729608/
- Figure 258 by Tverdokhlib with permission, available at https://www.allaboutlean.com/cardboard-engineering-preparation/dream-career-adventure-education-pilot-travel-airdrome-ima/
- Figure 260 by Roadell Hickman under the CC-BY-SA 4.0 license, available at https://www.flickr.com/photos/usnavy/5373594914
- Figure 262 by Vera Kratochvil in public domain, available at https://www.publicdomainpictures.net/en/view-image.php?image=9097&picture=empty-cardboard-boxes
- Figure 263 by Maxpixel in public domain, available at https://www.maxpixel.net/Cutter-Knife-Carpet-Knife-Tool-1668382
- Figure 265 by Jan Hagelskamp1 under the CC-BY 4.0 license, available at https://commons.wikimedia.org/wiki/File:Edding_Flipchart_Marker.jpg
- Figure 266 by Roadell Hickman under the CC-BY-SA 4.0 license, available at https://www.flickr.com/photos/usnavy/5373594914
- Figure 267 by Casito under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:WA_archery_target_with_arrows.jpg
- Figure 268 by Vittaya_25 with permission, available at https://www.allaboutlean.com/cardboard-engineering-workshop/construction-worker-team-planning-about-building-plan-with-bluep/
- Figure 269 by geralt in public domain, available at https://pixabay.com/photos/hand-hands-keep-time-clock-226708/
- Figure 270 by Clker-Free-Vector-Images in public domain, available at https://pixabay.com/vectors/checklist-lists-business-form-41335/
- Figure 271 by Ross Dunn under the CC-BY-SA 2.0 license, available at https://live.staticflickr.com/4291/35755157191_58fdf713f0_b.jpg
- Figure 272 by Susan L. Davis in public domain, available at https://www.airforcemedicine.af.mil/News/Photos/igphoto/2000899183/
- Figure 273 by Susan L. Davis in public domain, available at https://www.grandforks.af.mil/News/Article/763661/gf-afb-holds-first-global-cardboard-challenge/
- Figure 274 by CHOReograPH with permission, available at https://www.allaboutlean.com/cardboard-engineering-alternatives/children-in-astronaut-pilot-and-super-hero-costumes-are-laughin/
- Figure 275 by Hermann Ultraschall with permission, available at https://www.herrmannultraschall.com/en/
- Figure 276 by Hermann Ultraschall with permission, available at https://www.herrmannultraschall.com/en/
- Figure 284 by Everblock with permission, available at https://www.everblocksystems.com/
- Figure 285 by R3DT with permission, available at https://www.allaboutlean.com/cardboard-engineering-alternatives/r3dt-screenshot-cropped/
- Figure 287 by Jane Biriukova with permission, available at https://www.allaboutlean.com/6th-birthday/birthday-cake-with-burning-candles-and-age-6-candle-in-the-dark/
- Figure 292 by Feedspot with permission, available at https://www.allaboutlean.com/awards/top-20-lean-manufacturing-blogs_crop/
- Figure 293 by Kanbanize for editorial use, available at https://www.allaboutlean.com/awards/kanbanize-logo/
- Figure 297 by OpenStreetMap contributors under the CC-BY-SA 2.0 license, available at https://www.openstreetmap.org/#map=4/21.84/82.79
- Figure 298 by Berthold Werner in public domain, available at https://www.allaboutlean.com/i4-0-tour-germany-1/rothenburg-ob-der-tauber-sieberstorlinks-und-koboltor-rechts/
- Figure 299 by Joachim Köhler under the CC-BY 3.0 license, available at https://commons.wikimedia.org/wiki/File:Neckarsulm-AudiForum-Innen-SW.JPG
- Figure 300 by ABB for editorial use, available at https://commons.wikimedia.org/wiki/File:ABB-Logo.svg
- Figure 302 by ABB under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:Residual_Current_Device_F202.jpg

- Figure 304 by Ktm250-1150gs under the CC-BY-SA 3.0 license, available at https://commons.wikimedia.org/wiki/File:Stack_Light_Reference_Image.png
- Figure 306 by Anthony O'Neil under the CC-BY-SA 2.0 license, available at https://www.geograph.org.uk/photo/5303775
- Figure 307 by Hoshinim under the CC-BY-SA 4.0 license, available at https://commons.wikimedia.org/wiki/File:Onetech_Asia_Hololens_Demo_idea_real_estate.jpg
- Figure 308 by Trumpf for editorial use, available at https://commons.wikimedia.org/wiki/File:Logo_Trumpf.svg
- Figure 315 by Bosch for editorial use, available at https://commons.wikimedia.org/wiki/File:Bosch_logo.png
- Figure 319 by Bosch with permission, available at https://www.bosch-presse.de/pressportal/de/de/halbleiterfertigung-in-300-millimeter-technologie-174533.html
- Figure 320 by Bosch with permission, available at https://www.bosch-presse.de/pressportal/de/de/halbleiterfertigung-142472.html
- Figure 321 by Kärcher for editorial use, available at https://commons.wikimedia.org/wiki/File:K%C3%A4rcher_Logo_2015.svg and by Siemens for editorial use, available at https://commons.wikimedia.org/wiki/File:Siemens-logo.svg
- Figure 322 by Kärcher with permission, available at https://www.allaboutlean.com/i4-0-tour-germany-1/van-of-nerds-at-karcher/
- Figure 323 by Kärcher with permission, available at https://www.kaercher.com/de/
- Figure 325 by Kärcher with permission, available at https://www.kaercher.com/de/
- Figure 327 by Siemens with permission, available at https://press.siemens.com/global/de?offset=2&division=&search=amberg
- Figure 328 by Pjt56 under the CC-BY-SA 4.0 license, available at https://commons.wikimedia.org/wiki/File:ARENA2036-pjt.jpg
- Figure 334 by Tobias Schwerdt with permission, available at https://www.allaboutlean.com/i4-0-tour-germany-6/hs-karlsruhe-robogistics-schwerdt/
- Figure 335 by R3DT with permission, available at https://www.allaboutlean.com/cardboard-engineering-alternatives/r3dt-screenshot-cropped/
- Figure 338 by Bosch with permission, available at https://www.bosch-presse.de/pressportal/de/de/sotware-macht-die-arbeit-in-produktion-und-logistik-einfacher-186049.html
- Figure 340 by Amazon for editorial use, available at https://www.allaboutlean.com/amazon-fulfillment-1/amazon-logo/
- Figure 343 by Amazon with permission, available at https://press.aboutamazon.com/images-videos
- Figure 344 by Maryland GovPics under the CC-BY 2.0 license, available at https://www.flickr.com/photos/mdgovpics/36490655883/in/album-72157686813447104/
- Figure 345 by Amazon with permission, available at https://press.aboutamazon.com/images-videos
- Figure 346 by Amazon with permission, available at https://amazon-presse.de/
- Figure 347 by Álvaro Ibáñez under the CC-BY 2.0 license, available at https://commons.wikimedia.org/wiki/File:Amazon_Espa%C3%B1a_por_dentro_(San_Fernando_de_Henares).JPG
- Figure 349 by Amazon with permission, available at https://www.allaboutlean.com/amazon-fulfillment-1/amazon-pick-towers-inside/
- Figure 352 by Amazon with permission, available at https://amazon-presse.de/
- Figure 353 by Álvaro Ibáñez under the CC-BY 2.0 license, available at https://commons.wikimedia.org/wiki/File:Amazon_Espa%C3%B1a_por_dentro_(19).jpg
- Figure 354 by Amazon with permission, available at https://www.allaboutlean.com/amazon-fulfillment-1/amazon-kiva-with-pod/
- Figure 355 by Amazon with permission, available at https://amazon-prensa.es/Log-stica/Centros-Log-sticos/Presskit/amazon/es/Imagenes-Prensa/Centro-Logistico/download/es/Imagenes-Prensa/Centro-Logistico/Panorama-interior-centro-logistico-Amazon.es/
- Figure 356 by Amazon with permission, available at https://amazon-presse.de/
- Figure 357 by Amazon with permission, available at https://www.allaboutlean.com/amazon-fulfillment-1/amazon-robotic-pick-station/
- Figure 358 by Amazon with permission, available at https://press.aboutamazon.com/images-videos
- Figure 359 by Álvaro Ibáñez under the CC-BY 2.0 license, available at https://commons.wikimedia.org/wiki/File:Amazon_Espa%C3%B1a_por_dentro_(14).jpg
- Figure 360 by Amazon with permission, available at https://www.allaboutlean.com/amazon-fulfillment-1/amazon-gift-wrapping/

- Figure 361 by Amazon with permission, available at https://amazon-prensa.es/Log-stica/Centros-Log-sticos/Presskit/amazon/es/Imagenes-Prensa/Centro-Logistico/download/es/Imagenes-Prensa/Centro-Logistico/Panorama-interior-centro-logistico-Amazon.es/
- Figure 362 by Amazon with permission, available at https://amazon-presse.de/
- Figure 363 by Amazon with permission, available at https://amazon-prensa.es/Log-stica/Centros-Log-sticos/Presskit/amazon/es/Imagenes-Prensa/Centro-Logistico/download/es/Imagenes-Prensa/Centro-Logistico/Panorama-interior-centro-logistico-Amazon.es/
- Figure 364 by Amazon with permission, available at https://amazon-prensa.es/Log-stica/Centros-Log-sticos/Presskit/amazon/es/Imagenes-Prensa/Centro-Logistico/download/es/Imagenes-Prensa/Centro-Logistico/Panorama-interior-centro-logistico-Amazon.es/
- Figure 365 by 玄史生 in public domain, available at https://commons.wikimedia.org/wiki/File:B5_Size_Package_of_Amazon.co.jp_Sended_to_Taiwan_20170414b.jpg
- Figure 366 by Amazon with permission, available at https://amazon-prensa.es/Log-stica/Centros-Log-sticos/Presskit/amazon/es/Imagenes-Prensa/Centro-Logistico/download/es/Imagenes-Prensa/Centro-Logistico/Panorama-interior-centro-logistico-Amazon.es/
- Figure 367 by Amazon with permission, available at https://www.allaboutlean.com/amazon-fulfillment-1/amazon-conveyor-2/
- Figure 368 by Amazon with permission, available at https://amazon-prensa.es/Log-stica/Centros-Log-sticos/Presskit/amazon/es/Imagenes-Prensa/Centro-Logistico/download/es/Imagenes-Prensa/Centro-Logistico/Panorama-interior-centro-logistico-Amazon.es/
- Figure 369 by Amazon with permission, available at https://press.aboutamazon.com/images-videos
- Figure 370 by Amazon with permission, available at https://press.aboutamazon.com/images-videos
- Figure 371 by Amazon with permission, available at https://amazon-presse.de/
- Figure 372 by Amazon with permission, available at https://amazon-prensa.es/Log-stica/Centros-Log-sticos/Presskit/amazon/es/Imagenes-Prensa/Centro-Logistico/download/es/Imagenes-Prensa/Centro-Logistico/Panorama-interior-centro-logistico-Amazon.es/
- Figure 373 by Amazon with permission, available at https://www.allaboutlean.com/amazon-fulfillment-1/amazon-order-form/
- Figure 374 by Amazon with permission, available at https://www.allaboutlean.com/amazon-fulfillment-5/amazon-bin-image/
- Figure 375 by Pixabay in public domain, available at https://www.pexels.com/photo/computers-data-center-server-room-servers-270518/
- Figure 376 by Murrstock with permission, available at https://www.allaboutlean.com/amazon-fulfillment-5/kpi-key-performance-indicator-increase-optimisation-business-and/
- Figure 377 by Amazon with permission, available at https://amazon-presse.de/
- Figure 378 by Amazon with permission, available at https://amazon-prensa.es/Log-stica/Centros-Log-sticos/Presskit/amazon/es/Imagenes-Prensa/Centro-Logistico/download/es/Imagenes-Prensa/Centro-Logistico/Panorama-interior-centro-logistico-Amazon.es/
- Figure 380 by Amazon with permission, available at https://amazon-prensa.es/Log-stica/Centros-Log-sticos/Presskit/amazon/es/Imagenes-Prensa/Centro-Logistico/download/es/Imagenes-Prensa/Centro-Logistico/Panorama-interior-centro-logistico-Amazon.es/
- Figure 381 by Amazon with permission, available at https://amazon-presse.de/
- Figure 382 by Amazon with permission, available at https://amazon-presse.de/
- Figure 383 by Sikander Iqbal under the CC-BY-SA 4.0 license, available at https://commons.wikimedia.org/wiki/File:Amazon_Go_-_Seattle_(20180804111407).jpg
- Figure 384 by Joe Andrucyk under the CC-BY 2.0 license, available at https://www.flickr.com/photos/mdgovpics/36467011734/in/album-72157686813447104/
- Figure 385 by Moremar with permission, available at https://www.allaboutlean.com/one-up-one-down/up-down-window-neighbor/
- Figure 394 by Pra Chid with permission, available at https://www.allaboutlean.com/more-reasons-for-working-less/professional-business-woman-seeing-wrist-watch-to-check-time-at/
- Figure 395 by Lewis Hine and Roser in public domain, available at https://www.allaboutlean.com/faster-better-cheaper/addie-card-colorized_v03/
- Figure 396 by unknown author in public domain, available at https://commons.wikimedia.org/wiki/File:F._Taylor_1856-1915.jpg
- Figure 398 by Microsoft for editorial use, available at https://commons.wikimedia.org/wiki/File:Microsoft_logo_(2012).svg
- Figure 399 by Milkos with permission, available at https://www.allaboutlean.com/more-reasons-for-working-less/idleness-at-work-top-view-of-carefree-afro-employee-using-smart/

- Figure 402 by smrm1977 with permission, available at https://www.allaboutlean.com/extra-kanban/sick-dog-injured-and-funny-black-puppy-lying-down-bitting-and-r/
- Figure 410 by FELDSPATH under the CC-BY-SA 4.0 license, available at https://commons.wikimedia.org/wiki/File:Joseph_Juran.jpg
- Figure 411 by unknown author in public domain, available at https://commons.wikimedia.org/wiki/File:Br%C4%83ila_19th_century.jpg
- Figure 412 by unknown author in public domain, available at https://www.allaboutlean.com/115-years-after-joseph-juran/juran-family-1910/
- Figure 413 by Western Electric Company in public domain, available at https://commons.wikimedia.org/wiki/File:Hawthorne,_Illinois_Works_of_the_Western_Electric_Company,_1925.jpg
- Figure 415 by FDA in public domain, available at https://commons.wikimedia.org/wiki/File:W._Edwards_Deming.jpg
- Figure 416 by unknown author in public domain, available at https://commons.wikimedia.org/wiki/File:Vilfredo_Pareto_1870s2.jpg

54 Author

Figure 417: Christoph Roser (Image Roser)

Prof. Dr. Christoph Roser is an expert for lean production and a professor for production management at the University of Applied Sciences in Karlsruhe, Germany. He studied automation engineering at the University of Applied Sciences in Ulm, Germany, and completed his Ph.D. in mechanical engineering at the University of Massachusetts, researching flexible design methodologies. Afterward he worked for five years at the Toyota Central Research and Development Laboratories in Nagoya, Japan, studying the Toyota Production System and developing bottleneck detection and buffer allocation methods. Following Toyota, he joined McKinsey & Company in Munich, Germany, specializing in lean manufacturing and driving numerous projects in all segments of industry. Before becoming a professor, he worked for the Robert Bosch GmbH, Germany, first as a lean expert for research and training, then using his expertise as a production logistics manager in the Bosch Thermotechnik Division. In 2013, he was appointed professor for production management at the University of Applied Sciences in Karlsruhe to continue his research and teaching on lean manufacturing.

Throughout his career Dr. Roser has worked on lean projects in almost two hundred different plants, including automotive, machine construction, solar cells, chip manufacturing, gas turbine industry, paper making, logistics, power tools, heating, packaging, food processing, white goods, security technology, finance, and many more. He is an award-winning author of over fifty academic publications. Besides research, teaching, and consulting on lean manufacturing, he is very interested in different approaches to manufacturing organization, both historical and current. He blogs about his experiences and research on AllAboutLean.com. He also published his first book, “Faster, Better, Cheaper,” on the history of manufacturing.

Prof. Dr. Christoph Roser is an expert for lean production; Toyota, McKinsey, and Bosch Alumni, and professor for Production Management at the Karlsruhe University of Applied Sciences. He is interested in everything related to lean manufacturing, bottleneck detection and management, as well as historic developments of manufacturing. His first book is “Faster, Better, Cheaper” on the history of manufacturing.

Having successfully written my award-winning blog, AllAboutLean.com, for over six years now, I decided to make my blog posts available as collections. There will be one book of collected blog posts per year, from 2013 to 2019. This way you can store these blog posts conveniently on your computer should my website ever go offline. This also allows you a more professional citation to an article in a book, rather than just a blog, if you wish to use my works for academic publications.
If you like my writing, please check out also my other books on AllAboutLean.com. In any case, keep on reading my blog ☺.