1-22
Celonis Labs
Beyond

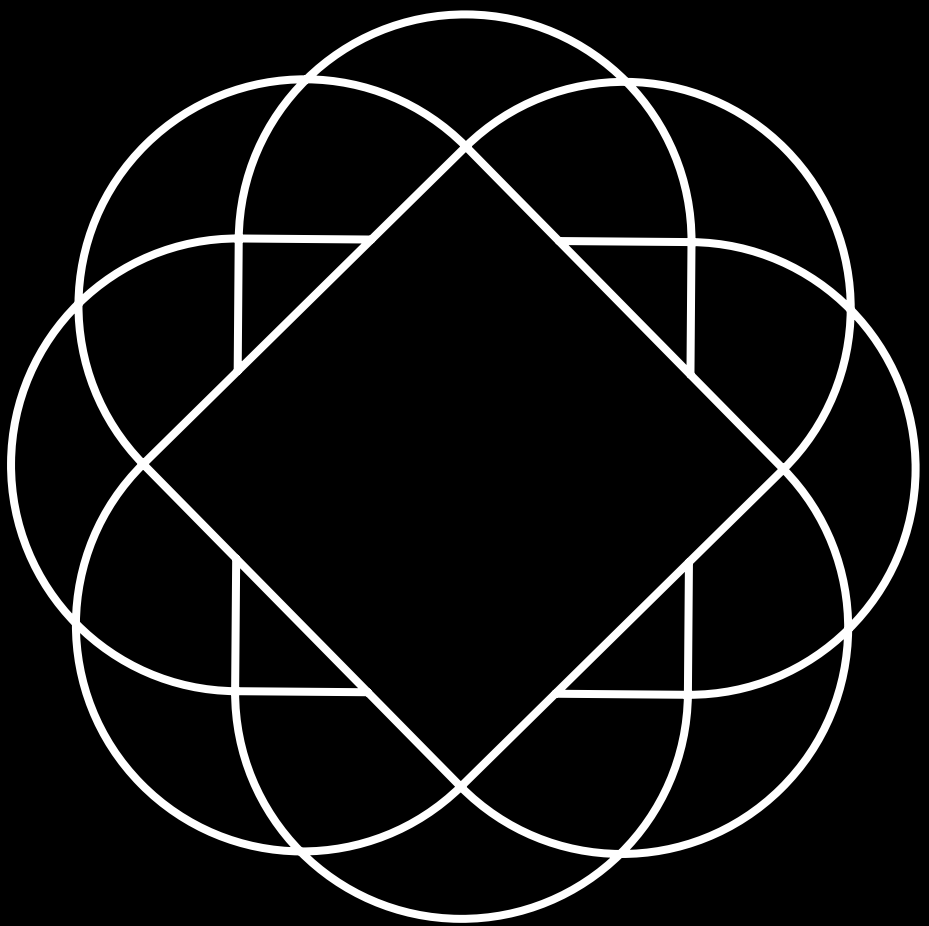

JOURNAL

1–22

JOURNAL

“When you’re smart, people need you.”

Growing up as a nerdy kid who loved math, computers and solving technical challenges, those words from Val Kilmer’s Chris Knight in “Real Genius” spoke to me like few others had. For the first time, I saw silver screen heroes who were thinkers, makers and scientists. Not only was it okay to be smart, it was cool and more important, it was essential. The world needs smart people to tackle today’s tough challenges and think about the problems of tomorrow, which are often caused by the solutions of today.

Celonis Labs is that collection of smart people. Exactly who they are and what they’ve been up to has been a closely guarded secret, even within Celonis. Well, not any more. World, here we come.

This team of engineers, researchers and technologists, led by Eugenio Cassiano, Celonis’s very own mad scientist, is ready to share a small taste of what they’ve been working on in the inaugural issue of Beyond. This journal is a collection of candid, unfettered thoughts, analysis and prototype descriptions from the brightest minds studying process mining, execution management and the future of business.

Beyond, volume 1, issue 1, is divided into following four chapters →

Each chapter examines a current project that Celonis Labs researchers are not only pondering, but for which they have built a working prototype. That’s what makes Celonis Labs one of the most unique think tanks ever assembled to study how the world should work. We not only talk the talk, but walk the walk. From artificial intelligence to augmented and virtual reality, we’re building the technology that will power the way work happens for decades to come.

So let go of your preconceived notions about how business processes work today, and let Celonis Labs walk you through how things will work in the future.

And if you’re wondering why we’re doing this and why now, I’ll leave you with another line from Knight, “It’s a moral imperative.” The world is facing significant challenges and making the processes that run work more sustainable, we are doing our part to build a better future for all.

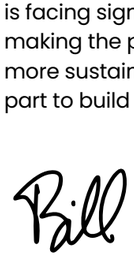

Bill Detwiler
Beyond Editor

1 Business Collaboration Networks **6-7**

2 Zero-trust Sharing **8-10**

3 Real-world events and business domain translation **14-17**

4 Context Awareness in Execution Management **18-21**

Beyond
A publication from Celonis Labs

Editor
Bill Detwiler,
Editor, Technical Content & Ecosystem, Celonis

Celonis Labs
Eugenio Cassiano,
SVP of Celonis Labs

Authors
Andreas Krause,
Head of Innovation Architecture, Celonis Labs
Maximilian Schrupp,
Technology Researcher, Celonis Labs
Patrick Giang,
Technology Researcher, Celonis Labs

Guest Authors
Dr. Michael Rosemann,
Director of the Center for Future Enterprise and Professor for Innovation Systems at the Business School at Queensland University of Technology
Dr. Sarah Güsken,
Academic Research and Innovation Coordinator
Stephan Rossbauer,
Product Manager, Celonis

Jonas Weich,
Product Manager

Design & Production
Jan Henkel,
Art Director, Celonis
Christian Flaccus,
CDO, Celonis

United States - Headquarters
One World Trade Center
87th Floor New York, NY 10007
Phone number: +1 212 847 0692

Germany - Headquarters
Theresienstr. 6
80333 Munich
Germany
Phone number: +49 89 4161596 70

Twitter @Celonis
Linkedin
linkedin.com/company/celonis
Facebook facebook.com/Celonis
Xing xing.com/company/celonis

About Celonis
Celonis is the global leader in execution management. The Celonis Execution Management System provides companies a modern way to run their business processes entirely on data and intelligence. We pioneered the process mining category 10 years ago when we first developed the ability to automatically X-ray processes and find inefficiencies.

Welcome to the first edition of Beyond

The innovation journal by Celonis Labs, which gives you a glimpse on our current research and the concepts we're working on.

Celonis Labs is the innovation facility designed for studying and launching advanced applications and experimenting with new ideas and emerging technologies with the ultimate goal of advancing the Execution Management category. At Celonis Labs, we are **walking the talk**. We do not only mention new ideas and technologies in this journal, but we are also sharing the outcomes of our concepts that we believe to have an impact in the future of the Celonis Execution Management System (EMS).

Over the last year, we have seen the rapid evolution in technology at Celonis, starting with process mining, transitioning into process intelligence and finally into the Execution Management System, progressing to become a business performance platform. At Celonis Labs, we focus on identifying and accelerating future use cases and approaches. We keep asking ourselves "How can we push the boundaries **beyond** our current knowledge?" to identify potential new business models that can be incubated in Celonis.

We invite customers and partners, our ecosystem and our academic alliance partners to read through the concepts – Collaborative Execution Management, Zero-Trust Sharing, Real-World Data and Context Awareness – to think about how they apply to their own goals and reach out to us. Our mission is to apply the ideas behind emerging technologies to the problems our customers are trying to solve. We want to inspire everyone with the **art of the possible** that you can ultimately achieve with Celonis.

Let's go!
Eugenio

Eugenio Cassiano
SVP of Celonis Labs

COLLABORATIVE EXECUTION MANAGEMENT

One of the limitations of today’s Execution Management in practice is that most organizations only work with their own internal data.

This fact isn’t surprising given that the software systems providing the process data are also purely internal.

While improving execution within an organization already unlocks immense benefits, we can go much further: most businesses do not operate stand-alone, but are part of structures such as supply chains that mutually benefit multiple entities.

In this chapter, we show how Execution Management unlocks massive potential for collaboration between organizations in the future.

We highlight two possible approaches to collaboration in Execution Management, starting with a thorough way of sharing business object information across a network of organizations.

We also present a different approach that allows collaboration without requiring specific, sensitive information to be shared.

The approaches are not mutually exclusive, and we see both extremely relevant even in the near future.

Editor’s note
Friction is good when you’re an F1 driver, but bad when it’s hurting your business processes. Imagine an urgently needed part is late and you have to call your supplier over and over, too much friction. In Chapters 1 and 2, Andreas Krause, Maximilian Schrupp and Stephan Rossbauer explore how the Celonis EMS can reduce intercompany friction like this by standardizing business objects, increasing intercompany collaboration through a new Process Sharing Network and doing it all security with zero-trust sharing.

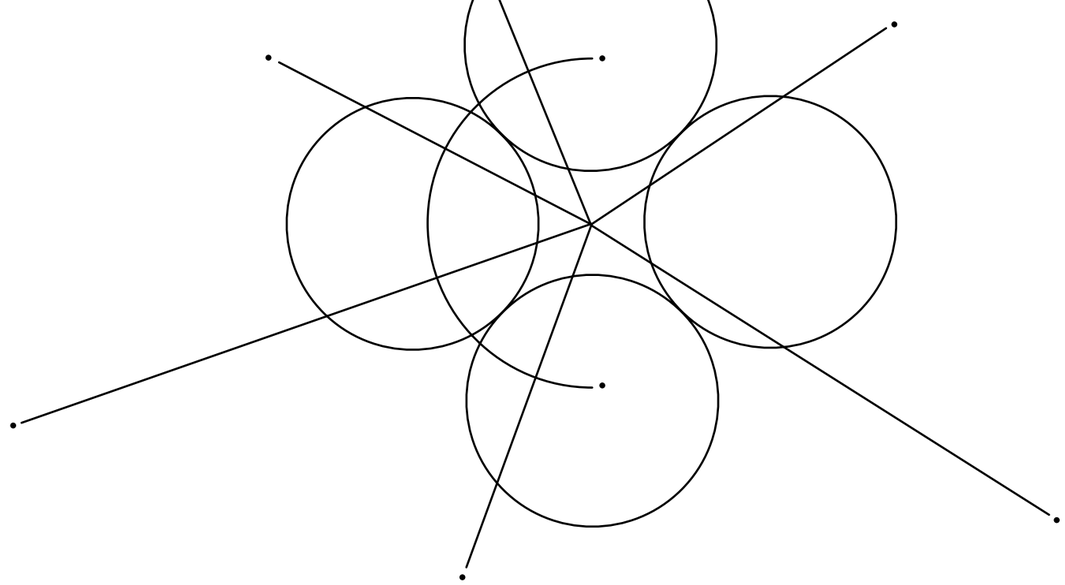

1 Business Collaboration Networks

Companies spend a lot of time sharing and acting on information in their daily operations–think of standard processes like Purchase-to-Pay or Order-to-Cash–where many documents have to be exchanged for each single case.

Andreas Krause
Head of Innovation Architecture at Celonis Labs

During this interaction, the information is stored across companies, in multiple separate systems. Naturally, this can lead to a disconnect, where the involved parties are operating on different truths. To improve the collaboration between the companies, we have to avoid these disconnects.

Business users are facing two challenges through such a disconnect: first, the missing transparency of the status of a business object and the resulting inability to act according to changes or updates; and second, the inconsistency of the information that is attached to different business objects. Both are execution gaps that we believe can be improved in Execution Management by linking important documents better in inter-company processes.

Today, if a company wants to close these execution gaps, integration specialists would have to build a ***specific automation**** for every identified gap, with the logic on how to react and a potential human-in-the-loop only adding more complexity. The necessary, highly detailed planning involves high costs and would still result in inflexible outcomes.

× Specific automation
An existing approach for specific system-to-system integrations is EDI, or Electronic data interchange. EDI is implemented per system and often results in skyrocketing costs due to the necessary coordination and also maintenance efforts.

We devise a different approach with one that is highly standardized at its core, so that a network can begin to develop, where every joining company can integrate with others to close execution gaps at the click of a button. Companies join the network in EMS, which is already built to connect all internal systems across a process. Ultimately, this approach allows a holistic process view of all involved parties and systems.

To also allow more specific use cases that would require non-standard data, we enable companies to leverage optional extensions on top of the standardized core, providing the desired functionality to all interested network participants.

This solution targets operational and strategic staff of our core processes. More interesting, we believe in a solution where digital connections and automations can be built without integration specialists and engineers.

The key principle here is reusability and standardization focused on business objects. An object that is configured once should be reusable with every similar occurrence, e.g. those of suppliers. Further we want to standardize such objects for scenarios where a cross-company match is highly likely and can be automated (or only has to be configured once by each participant).

Not all attributes of an object in the data owner's system can be shared and some are even business secrets. We want to give the data owners absolute control here on what information exactly they make available in the network.

We believe that the Celonis EMS is well positioned to facilitate and simplify improvements to many execution gaps in the field by establishing this Business Collaboration Network and acting as the intelligence layer and single source of truth on top of inter-company processes. Similar to the current structure of the EMS, we imagine a flexible platform with a unified taxonomy and business apps to generate value for specific use cases.

While there are definitely hurdles to overcome, the opportunity of creating a Process Sharing Network is tempting and the implementation of a first version is a candidate for the next Celonis Labs prototype.

2

Zero-trust Sharing

Companies are hesitating to share their sensitive, internal data, even if there are incentives to do so, as concerns of exploitation outweigh potential benefits.

Maximilian Schrupp
Technology Researcher
at Celonis Labs

This is not different in Execution Management, where insightful process models and even case-specific information are available. Many companies see no way to collaborate on improving their shared execution gaps.

Let's identify the underlying challenge here: missing trust. Companies that want to collaborate either have to build trust and start from there, or find a way to realize a solution for an execution gap without having to rely on trust.

Here we talk of our vision of a zero-trust solution for sharing process information in Execution Management, where there is either no sensitive data involved or the data remains inaccessible to the respective other side.

Encryption is a way to protect data, but traditional encryption means that you either have full access to the data or no access at all. Still, there is an area of research that allows calculations or transformations of encrypted data without having access to the plain text, called ***Homomorphic Encryption***×. While this is fundamentally interesting to Computer Science, and would potentially solve this challenge as well, it's still way too slow today.

Our approach here is simple reduction: boiling the information down to the essentials, depending on the use case. In a very basic setup, these can be key numbers for comparison, but also more complex forms are possible, especially when it comes to process models.

We focus on the overall model of a process here. There seems to be a fundamental potential for collaborative applications based on sharing these process models. But again, to do so, we must find a way to avoid exploitation by the insights another party gains.

Currently, information about the nature of processes is shared across companies mostly in a theoretical form. There is lots of work done in the field of business process modeling and re-engineering, but this is abstract and the created models are incomplete and outdated when compared to the results we achieve with the Celonis EMS.

A promising way to reduce the information of process data models we have in the Celonis EMS is statistics. Under the assumption that the general workflow of a process, i.e. which events are part of it, is not considered a secret, we want to create a statistical model of the process that can be shared.

The idea is to reduce the sensible details to a model of possible events and the probabilities of which event comes next. These probabilities are calculated flexibly based on the events that happened before. In simple words, for each event that happens in a specific case, a dedicated decision engine provides probabilities for the next step. The information of which events can follow which is preserved, thus allowing a granular picture of the process.

The engines are trained with the real cases of a company, and there are ***technologies available to do so today***×. This training is not a one-shot setup step, but rather an ongoing, continuous action, always providing an up-to-date and complete model of all process variants known to the company. After the training, the model only contains a set of abstract engines for each activity and no case-specific information is included.

This statistical model will essentially work like a simulator: when presented with enough different, made-up cases, the distribution of resulting, simulated process variants will be close to those found in real cases.

Sharing such a model seems realistic to companies: it's powerful enough to allow applications based on comparison between similar processes, and might even allow the creation of a cross-company process model by combining multiple single models. And most importantly, this model doesn't contain any sensitive case-specific information. While this example still has limits, especially when the general process flow is considered sensitive as well, we see a huge potential for new applications here.

Building a prototype and exploring zero-trust sharing applications based on these models is a candidate for an upcoming Celonis Labs prototype. Stay tuned for more insights on this exciting topic!

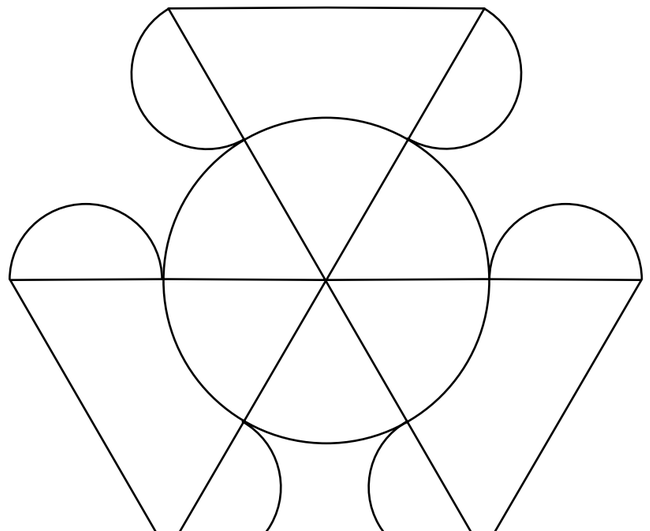

× Homomorphic Encryption is an encryption method to allow data to stay encrypted while calculations or transformations are being performed. The plaintext of the data remains hidden and doesn't need to be decrypted for processing. A third party can provide services for data that is never revealed. After processing, the data owner can decrypt the results. In practice today, it works best with pure integer numbers. There is still a long way to go, as the same actions that take milliseconds today in unencrypted form take days or weeks in Homomorphic Encryption.

× Some of the **technologies** we are researching to use here are Markov Chains, Long short-term memory or Autoregressive Models.

Stephan Rossbauer
Product Manager
at Celonis, on process benchmarking

COMMENT

The Celonis EMS today is very good at giving Celonis customers objective, measured data on the performance of their processes.

Following these insights, the important questions then are: Is the performance strong or poor? How big is the opportunity to improve the performance of the process? What does world class performance look like and how big is the gap? Is my process the equivalent of Usain Bolt running the 100m dash in 9.58s and every additional hour of training invested will only produce marginal improvements or is my process performing at collegiate level and even just one hour of strength training per week will increase the performance by a big margin?

The answers to these questions are hidden in the collective data already processed by the Celonis EMS and I believe everyone can benefit from such a collective intelligence. The idea is to share aggregate process KPIs (e.g. Days Payable Outstanding) and in return participants will receive information in which percentile their performance sits (e.g. your DPO is 28 days, which is in the 80th percentile of all participants).

This information would then allow them to put their process performance into context and evaluate improvement opportunities more accurately and efficiently.
This approach uses statistics to abstract the information and ensure desensitization so that the data shared is no longer connected to any individual participant, by following a simple rule: each datapoint in the shared data must be created by the aggregation of at least 10 individual data points.

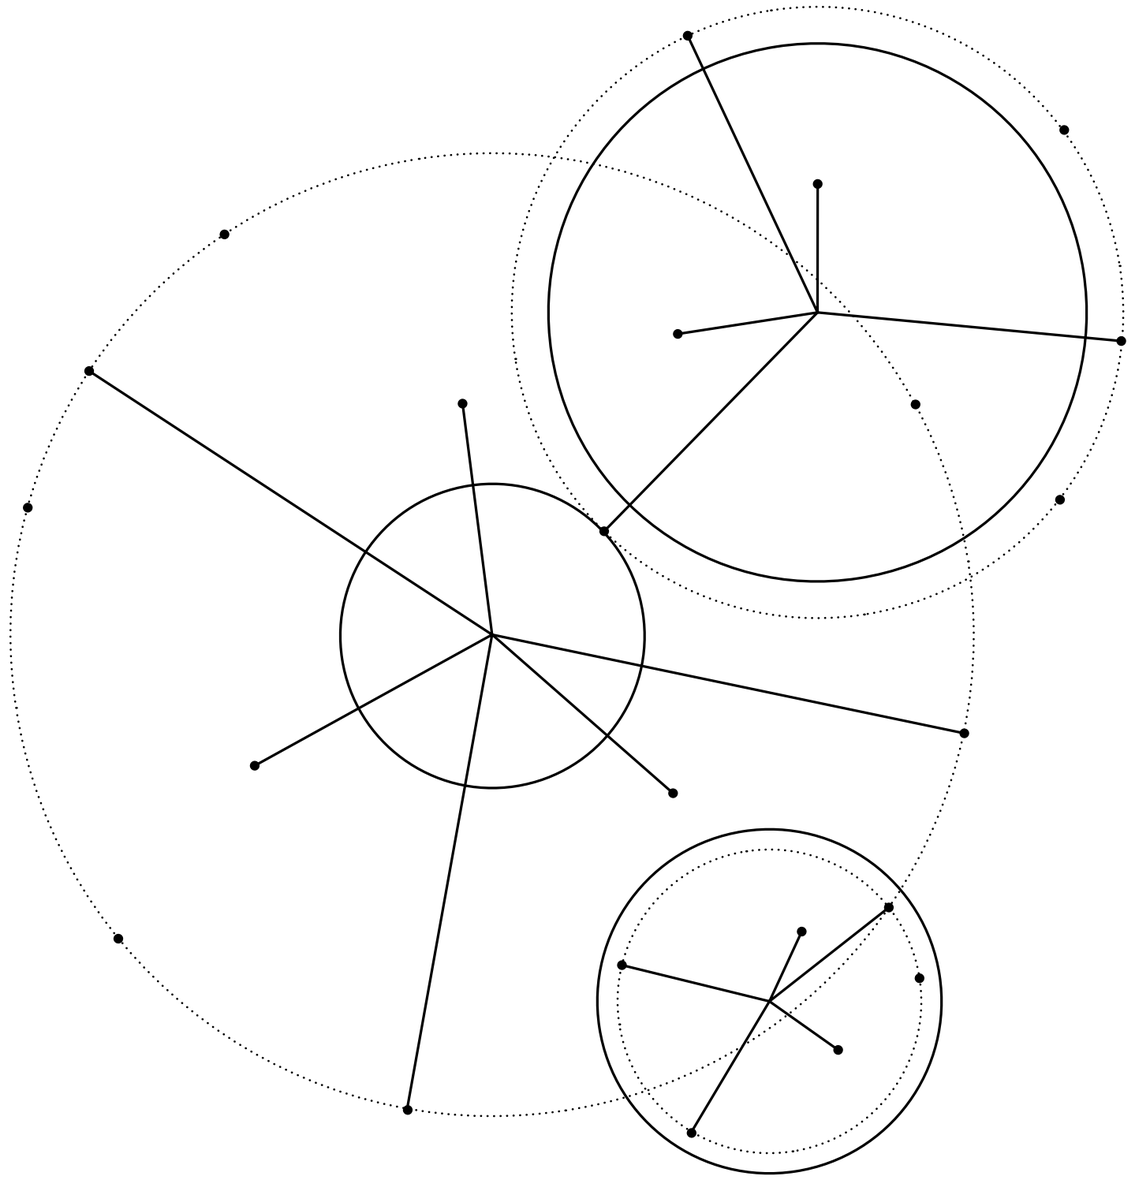

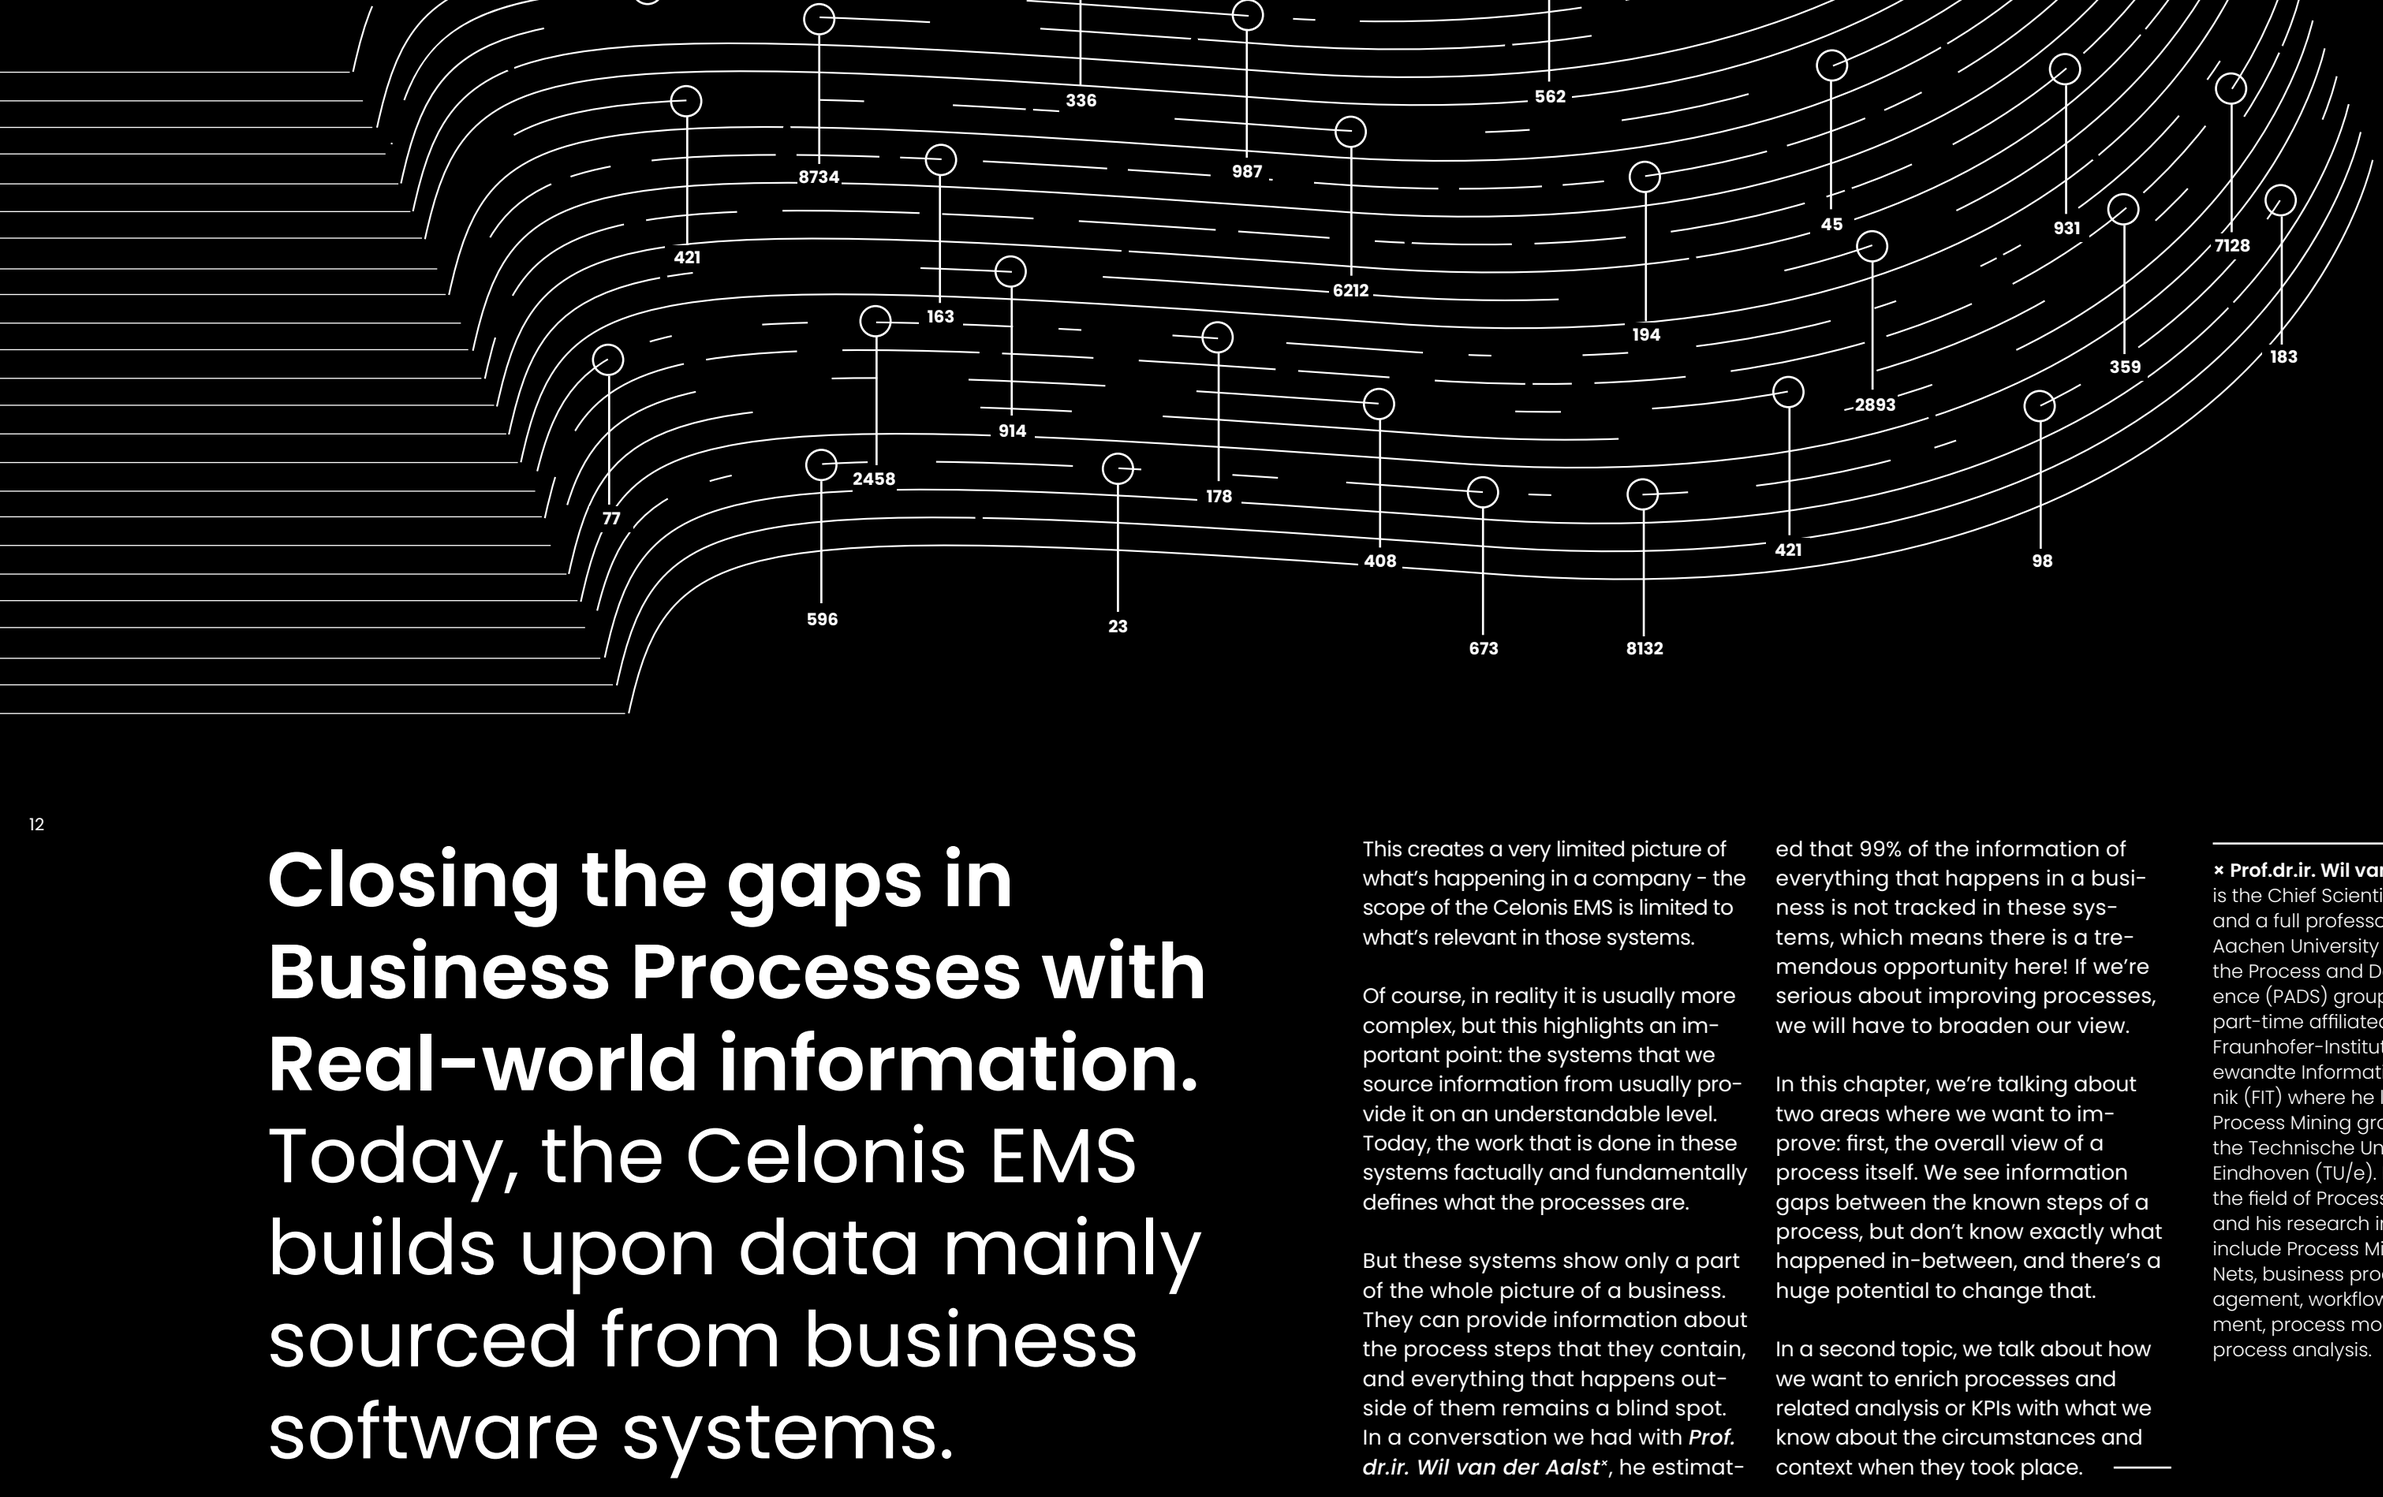

Closing the gaps in Business Processes with Real-world information.

Today, the Celonis EMS builds upon data mainly sourced from business software systems.

This creates a very limited picture of what's happening in a company - the scope of the Celonis EMS is limited to what's relevant in those systems.

Of course, in reality it is usually more complex, but this highlights an important point: the systems that we source information from usually provide it on an understandable level. Today, the work that is done in these systems factually and fundamentally defines what the processes are.

But these systems show only a part of the whole picture of a business. They can provide information about the process steps that they contain, and everything that happens outside of them remains a blind spot. In a conversation we had with ***Prof. dr.ir. Wil van der Aalst****, he estimated that 99% of the information of everything that happens in a business is not tracked in these systems, which means there is a tremendous opportunity here! If we're serious about improving processes, we will have to broaden our view.

In this chapter, we're talking about two areas where we want to improve: first, the overall view of a process itself. We see information gaps between the known steps of a process, but don't know exactly what happened in-between, and there's a huge potential to change that.

In a second topic, we talk about how we want to enrich processes and related analysis or KPIs with what we know about the circumstances and context when they took place.

*** Prof.dr.ir. Wil van der Aalst** is the Chief Scientist of Celonis and a full professor at RWTH Aachen University leading the Process and Data Science (PADS) group. He is also part-time affiliated with the Fraunhofer-Institut für Angewandte Informationstechnik (FIT) where he leads FIT's Process Mining group and the Technische Universiteit Eindhoven (TU/e). He coined the field of Process Mining and his research interests include Process Mining, Petri Nets, business process management, workflow management, process modeling, and process analysis.

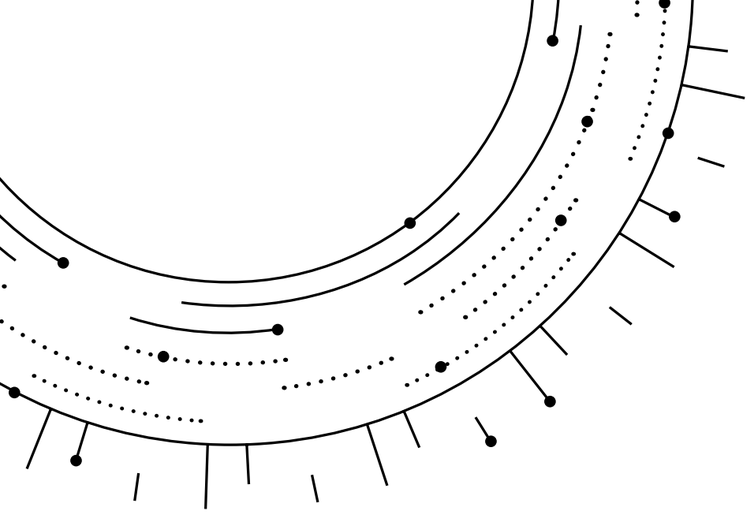

3

Real-world events and business domain translation

Closing blind spots of information in Execution Management is a challenging task.

Maximilian Schrupp
Technology Researcher at Celonis Labs

To broaden our view and reduce the 99% blind spots, we want to introduce events from real-world data sources to the Celonis EMS. At Celonis Labs, we believe that there are many things happening in the real world that are either not digitized at all or can't be integrated into a process view because the data is in a totally different domain than the one a business user is working in.

Let's have a look at an example: in manufacturing companies, to fulfill an order, often the items are produced on-demand (i.e. they are not held in stock). An order process is started, and shortly after, also a production process. Now a disconnect happens: on the shop floor, the manufacturing starts. The business user working on the order has to wait until all items are produced before taking the next steps. But what if something happens during the production that affects the underlying order? This should be reflected immediately, in a meaningful way, so we can act quickly. But how can we achieve that?

In general, we want to provide information to business users that they don't have today, to ultimately allow actions in the Celonis EMS that will improve their work. There are many different users working on all the different processes out there, and we have to take into account that the information we provide needs to be really helpful and meaningful in the context of each different process. We won't help a user working in order management with the information of individual manufacturing steps or issues, but rather with higher-level events when a certain order is actually affected.

We need to watch all relevant sources of information and automatically recognize and translate events into the business domain of the process we want to enrich. This **business domain translation** is a complex task: we start out with the target, a desired event we would like to introduce to a certain process. Next, we need to identify and connect data sources that can provide the data containing the necessary information. And finally, the events need to be automatically recognized in real-time from the different source ***data streams**** and integrated into the Celonis EMS. This is a new way of introducing process events, as we are tailoring exactly the events that are helpful for a certain domain.

We want to introduce real-world data sources in the same step as the event recognition and domain translation. This allows us to not only provide domain-tailored events to the business users, but also events containing information that was never integrated in a process before.

The acquisition of real-world data is not a new challenge, with technologies like (industrial) IoT and digital sensors of all kinds, including cameras, as well as techniques like Computer Vision and Machine Learning, we have a large toolset readily available. The data we receive from sensors are raw values without any specific meaning. We want to quickly boil this data down to essential information, depending on which business domain we target. We need to ***combine all sources**** we need to provide the necessary set of information to an event recognition engine. This event recognition engine is a service where all the pieces of data are observed and evaluated in real time to constantly decide if our desired event just occurred - and once so, the event is integrated in the respective process in EMS.

In our manufacturing example, we target an order-to-cash process as well as a manufacturing process. They need information on a different level, and while individual production steps are integrated as events in the manufacturing process, only events such as Production of order started, Production of order finished or Production of order delayed will enrich the order-to-cash process. We are starting out with a simple prototype: using cameras as sensors, we passively monitor what happens in an exemplary factory, a demonstrative production line we set up using educational robot arms. Again, the goal is to recognize events translated in the domain of our target processes. To achieve this for the order process, we also have to know which item we detected in the factory belongs to which order.

The opportunity we see here is tempting. We can extend the cross-system functionality of today's EMS to a new level: by connecting all available (or necessary) data sources, similar to what happens today with business systems, event recognition can become even

× Our excitement around real-world events is shared by Dr. Michael Rosemann, who emphasizes the great importance in the sensemaking and business meaning of what he calls signals. He particularly notes that while some of these signals might be weak, i.e indicate but not confirm an event, amplifying them by combining different data sources is crucial to reach a high validity.

Dr. Michael Rosemann is the Director of the Center for Future Enterprise and Professor for Innovation Systems at the Business School at Queensland University of Technology.

× Data streaming
is the continuous flow of data from a source system to a target system. This is the opposite of batch processing, where a set of data points are collected from the source system and provided as a bundle to the target.

more powerful. Seemingly unrelated sources can be combined. Another aspect is the real-time nature of these events, here we align with the general direction the EMS is developing towards. By design, the recognized events are available shortly after creation and allow use cases where a reaction in a timely manner makes a big difference. In our factory example, alerting a customer about a late order delivery directly after something goes wrong in the respective production allows all involved parties to react as early as possible. Partial deliveries or reprioritizing the production can suddenly become structured, well-communicated and comprehensible.

There are countless interesting data sources to explore. Imagine what happens when we combine shipping information a carrier provides with the data of cargo ship or plane trackers, will it be possible to identify misdeliveries early and introduce this in the EMS? We can talk about many examples here, but we rather want to prove that this concept can really work, so our prototype is already being implemented. While our showcase factory is just an example to prove the key aspects of this idea and to make the principles behind it more tangible, some of our customers already look in similar directions, and we're looking forward to validating this concept together. We see huge opportunities such as improvements in time-to-action or allowing powerful automations that were not possible before. Stay tuned for the outcomes of this project!

COMMENT

Dr. Sarah Güsken
Academic Research and Innovation Coordinator

Jonas Weich
Product Manager

AI4Pro A research project in Academic Alliance

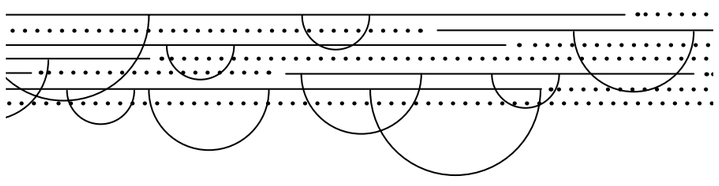

Within the Celonis Academic Alliance, we are currently realizing the visionary long-term project ***AI4Pro**** together with our partners. A diverse team is working on improving **spatially distributed manufacturing processes** by using the means of process mining. These processes consist of the daily procedures that happen during manufacturing. Initially, there was no digital system in place that somehow represented these procedures. To explore what happens during manufacturing from a process perspective and to also find execution gaps, we need to recognize the underlying events and integrate them in the EMS first - similar to what was explained already in this topic.

For us, the first step to create the events is the same: find and connect data sources. We combine multiple sensors via ***sensor fusion**** to a single data source and use the resulting mobile sensor devices to track each single item during manufacturing. As a result, we receive highly detailed data such as location, temperature or even magnetic measurements that we can assign to the correct item right from the beginning. In process mining terms, this item would be the so-called case.

The data streams are consumed by different systems. Most importantly, in an event recognition engine: depending on the manufacturing step, the necessary data streams and different timeframes are evaluated in an AI model to recognize the right production steps to build a view of the process. Here we also receive the recognition certainty along with each event, allowing us to filter out obvious misinformation very early.

In the next step, the recognized events are streamed to the EMS in real time, together with the raw data. Pioneering the EMS-based exploration of raw data was an important learning in the project, as it still contains important information that is not reflected on the event level. This enables the reproduction and validation of the recognition engine's operation and therefore strongly supports improving this key component.

Integrating the events, the recognition certainty and also the raw data sources allows us to create a powerful digital twin of both the production process and the recognition engine in EMS. This brings transparency to a new level, for example knowing exactly where each product component is in the manufacturing process. Even the recognition certainty history of each event is available, enabling novel analysis and filtering methods based on certainty for the first time, e.g. declaring reported events as invalid once better information becomes available. Besides that, several more innovative applications are made possible in the EMS, allowing to explore and improve a manufacturing process in its real-world context.

AI4Pro has already generated most interesting learnings and insights, and the exciting project still has over a year to go!

× AI4Pro
is the short name for the project "Application of Artificial Intelligence methods for Process Optimization and Prediction via Process Mining based on Cyber Physical Systems in Manufacturing Environments". This project was started in 2020 together with the Fraunhofer IIS Institute, Maxsyma, AST-X, PASS and Rauschert in 2020 and is funded by the Bavarian State Ministry for Economy, Regional Development and Energy.

× Sensor fusion
is a way of combining more than one sensor to achieve results that are more meaningful than each sensor individually. The sensors measure the same target, but provide different views, e.g. different kinds of data, different locations or perspectives. The challenge is to combine measurements taken at exactly the same time across sensors.

Editor's note
If we collect enough data, can we predict the future? According to Einstein's theory of relativity, maybe not. But that hasn't stopped Maximilian Schrupp, Dr. Sarah Güsken and Jonas Weich from examining how we can use the vast amounts of data collected from digital sensors and IoT devices that permeate business systems to improve process efficiency and perhaps predict the winner of the next World Cup. (I made that last part up, but I'm keeping my fingers crossed.)

4

Context Awareness in Execution Management

In the previous topic, we discussed how the EMS allows us to build process models and how we could introduce new types of events, sourced from the real-world.

Patrick Giang
Technology Researcher
at Celonis Labs

Here, we talk about a similar opportunity of introducing new information to our view on processes: the process context.

Currently a process is mainly a series of events that happened in a certain order, but there are always external conditions influencing a process as well. These circumstances range from the time of day when the process happened up to major global crises that affect almost everything.

Consider for example the recent incident at the Suez canal, when a ship blocked the traffic at one the most important shipping routes for an extended period of time. Naturally, supply chain processes were heavily impacted by this incident. Another, and one of the most prevalent crises of our recent history, is the COVID pandemic with a global impact resulting in many major problems, including countless business failures. In sales processes, in-person meetings have decreased whereas calls and virtual meetings have increased. It's likely that these kinds of major events will trigger permanent changes in how businesses are run.

Such major events or crises seem to become more frequent, and it's even more important to keep an eye on these when evaluating our processes. But there are also ordinary external factors affecting our processes that we need to include as well.

We want to take this great opportunity to make our processes context-aware. By introducing the process context in EMS, we see two general applications: enriching the process analytics, and allowing execution based on the contextual information.

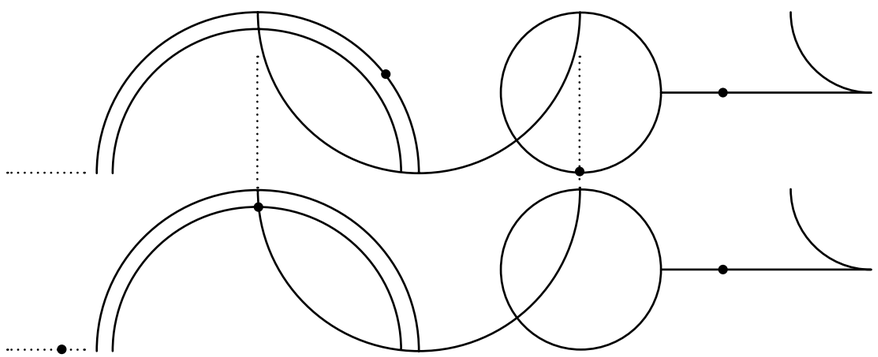

Let's start with the analytics. With the context we are not only able to know which external influences were present, but can also analyze how one or all processes in a business were impacted by a certain process context. This allows answering questions like "What was the impact of the Suez canal blocking on my process?" By comparing cases prior, during and after the context was relevant, we can find out if and how the performance of processes was affected. This can be as simple as creating the same KPIs separately for the periods of the three phases, but also more fine-grained analyses are possible. The learnings from these comparisons are straight-forward and can be understood easily. Ultimately, we hope to make processes more resilient to these external influences.

This is the point where we see another application of the process context in EMS: the ability to react to these external factors. We can use our insights and clear evidence from our past analyses to develop strategies to react to such external conditions. The emergence of certain process contexts can be used to trigger actions and measures to maximize process efficiency and close execution gaps. This is very similar to what is possible in EMS today for process events, and we hope to also achieve improvements in metrics such as time-to-action.

The way to bring the context into the EMS is the introduction of context tags that can be assigned to every case of a process, similar to how hashtags work in social media. Each context tag represents a certain external factor. The assignment can be automated based on matching event timestamps, with the limitation of ***regionally restricted circumstances**** that would also require the introduction of location information in a process itself. But also manual tagging is thinkable, as the tagging happens on a per-case basis and is thus not too cumbersome. In the end, we receive all process cases tagged with the context information that was relevant during the time or at the location where the processes were executed.

In general, the EMS can centrally provide common context tags, such as the Suez canal blocking or the COVID crisis, for automatic tagging. There are sources that curate and pre-aggregate such information that can easily be integrated today. Of course, businesses can also add their own tags, and might even publicly share tags that are non-confidential.

Making processes in EMS context-aware by introducing context tags is a powerful feature we want to explore in an upcoming Celonis Labs prototype. Measuring the impact of external factors suddenly becomes possible. It's another step forward to relieving users from an unstructured part of their work, as some knowledge about external influences is often already "in people's heads", but there is no way of including this information in their operational tools. The digitization of this information allows users to focus on the execution - on gaining insights, becoming creative, crafting strategies and acting.

*** Many affecting factors** are regionally restricted. Examples are local political situations, ecological disasters or huge sports competitions. Even the weather can be an affecting factor in some cases.

Editor's note
Context is everything. Just ask any politician, novelist, comedian or the person who received that last mistaken text you sent. In Chapter 4, Patrick Giang and Dr. Michael Rosemann explain how they are working to make business processes context-aware and why we should.

Context awareness in business processes is a very broad topic and there are countless examples.

Dr. Michael Rosemann
Director of the Center for Future Enterprise and Professor for Innovation Systems at the Business School at Queensland University of Technology.

For instance, a global organization often runs its processes differently depending on the location, retailers vary their processes based on temporal contexts such as winter sales or Black Friday, and procurement departments will re-configure their parameters when the national interest rate changes as keeping stock might become more or less expensive. However, despite the significant role of process context, it has not yet made it into process management and execution solutions.

This shortcoming is amplified by the fact that a short ***time-to-process**** is important to customers. An example is an Australian insurance company that experiences a threefold claim volume in the case of a tropical storm. By monitoring relevant weather forecasts, they automated and accelerated a previously manual process to switch to rapid lodgement and the hiring of additional call center agents. Only such a rapid adoption guarantees a great customer experience.

Understanding which context matters requires correlation with process mining. In a large project with various hospitals we studied the context of patient waiting periods to identify if it was depending on the patient, the doctor or rather on time or location. Knowing the relevant context can improve situations like this by conducting a root-cause analysis and to intervene adequately.

The conceptualisation of context is what we call **context mining**. It is still in its infancy and the resulting context management is also still upcoming. In our research, we differentiate four types of context in context mining, depending how close they are to the immediate process concerns.

The **immediate context** deals with those elements that are direct and well captured input to a process. For example, a process in a hospital will depend on the characteristics of a patient.

The **internal context** represents business-specific decisions and variables. For example, initiatives such as post-COVID priorities will influence the way processes are executed. The **external context** captures the immediate environment of a business. A supplier might be on strike, cost of external capital might increase, or competitors might change their strategies.

Finally, the **environmental context** provides macro-economic changes. The example of the Suez canal blockade mentioned above falls into this category, and so do weather changes, new regulations like national pandemic responses or socio-demographic shifts such as veganism.

Organizations are encouraged to capture important context variables and to correlate these with their relevant processes to allow a quick adaptation to the new circumstances. Adding the context as a conceptual element to business process models will enable automatic monitoring and the triggering of the required changes immediately via event-condition-action rules.

× Time-to-process
is the time it takes from the first event happening in reality to the start of the actual process in software systems. Today, a process starts with a user creating the first input in software that is associated with a certain process case, but this is not necessarily the initial trigger of the process. In the example, the start of the process really is the first time when the tropical storm is known to come.

Celonis Labs


C