Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 51-11 (2018) 1403–1408

Asynchronous Communication Platform Concept to Coordinate Large-scale Industrial Processes

Petri Kannisto, David Hästbacka

Tampere University of Technology, P.O. Box 692, FI-33101 Tampere, Finland (e-mail: {petri.kannisto|david.hastbacka}@tut.fi)

Abstract: To save the environment and costs, there is a clear need to optimise and coordinate large-scale industrial processes, but the task is challenging. Complex plants may run multiple unit processes that are operated with separate control systems, although the processes are interdependent. Often, optimisation requires appropriate scheduling for both continuous processes and batch processes. These requirements of coordination, asynchronism and distributed operation are difficult to conform to with the current tools or systems. Therefore, to ease optimisation, this study suggests a communication platform concept. The platform enables the use of mathematical models to manage and utilise knowledge about production processes. To promote adaptation to actual business cases and their future evolution, the platform design emphasises asynchronism, loose coupling of optimisation modules and systems as well as distributed operation. This work has received particular motivation from copper and steel production, as they both require materials refinement in multiple separate but mutually dependent steps. Several contributions are presented for the platform: requirements, architecture design and two proofs of concept. The results suggest that the concept would bring considerable business value in industrial production.

Keywords: Information systems, System architectures, System integration, Knowledge-based systems, Industry automation, Optimization problems

1. INTRODUCTION

Improving the efficiency of industrial production promotes various benefits, including reduced emissions and higher profit margins. While emissions originate from various sources, energy production is one of the major pollutants. Industries consume a large share of the produced energy. Among the member countries of the International Energy Agency, manufacturing and other industries covered 30% of the total energy consumption in 2013 (Quadrelli et al., 2016, p. 6). On the other hand, there is also considerable energy saving potential in various industries at least in the European Union (Chan et al., 2015, p. 188). From the financial point of view, global competition drives enterprises towards constant improvement. Efficiency is often a precondition for survival.

Efficiency may be improved with appropriate communication systems, such as the platform architecture introduced in this paper. Information systems show their power, for instance, when complex computation or the utilisation of large information masses is required. These needs often apply to production plants that run multiple unit processes concurrently. There may be a demand to process large information volumes online, but optimal operation may also require a considerable amount of domain knowledge, which tends to evolve over time. This complexity increases the difficulty of system design, which, consequently, emphasises the importance of an appropriate architecture (Clements et al., 2003, p. 1).

The motivation of this work stems from large-scale process plants that run several unit processes, particularly copper production (from sulphide ores) and steel production. The unit processes are run concurrently and operated with dedicated control systems. Their optimisation may even cover supply chains or other activities beyond the production premises. According to Schlesinger et al., each unit process of copper production requires control over chemical reactions between various substances. The unit processes have mutual dependencies but are still separate. (Schlesinger et al., 2011, pp. 1-12) Due to asynchronism, distribution and process complexity, copper production optimisation as a whole is difficult. On the other hand, while steel refinement is a different process, it also consists of separate production phases. Numerous chemical reactions are controlled and multiple process parameters affect the output quality (Elkader et al., 2015). Therefore, both these cases require extensive knowledge as well as timely, distributed operation. The knowledge also improves constantly, and, on the other hand, any changes in production conditions may make some knowledge obsolete.

Concerning the platform concept, this article presents multiple contributions. First, system architecture requirements are introduced. Second, the related design aspects are considered. Finally, two proofs of concept are demonstrated. The utilised research method is *design science research*: Artefacts are created to solve a problem, after which they are evaluated against the original requirements. The focus of this article is *not* on the actual optimisation

2405-8963 © 2018, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.
Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2018.08.325

but on enabling the efficient utilisation of optimisation tools. The results are qualitative, because there is no way to measure the contribution of an ICT concept during development. Obviously, if any approach enables global optimisation instead of local optimisation, better global results are expected.

This article has the following structure. A review to previous research is provided in Section 2, system requirements in Section 3, design in Section 4, proofs of concept in Section 5 and, finally, discussion in Section 6 followed by a conclusion in Section 7.

2. RELATED WORK

Operator assistance systems are a tool to help employees perform better by utilising existing production-related knowledge. To manage such knowledge, some authors have researched information system architectures: Kannisto et al. (2014) and Kannisto et al. (2017) have considered machinery operation, whereas Ullrich et al. (2015) had manufacturing as the domain. Assistance systems may also utilise algorithms to learn from actual employee workflows, as researched by Bleser et al. (2015).

Various authors have researched system architecture aspects in production planning or scheduling. Shen et al. (2007) have suggested a cross-enterprise manufacturing scheduling architecture where agents and web services would be utilised for communication. Leitão and Restivo have considered the distributed application of scheduling knowledge. In their approach, coordinating scheduling is normally used; local scheduling is only a resort when centralised scheduling fails. (Leitão and Restivo, 2008) Li (2010) has considered production scheduling and its implementation with service components. Scheduling requires production tracking, which can be implemented with RFID tags in piece goods manufacturing (Zhong et al., 2013; Guo et al., 2015).

As far as is known, the previous research does not address the combination of system architectures, large-scale industrial plants and knowledge application in production coordination. According to Engell and Harjunkoski (2012), there have been no proper means to integrate at least advanced control and production scheduling. To integrate optimisation systems with production systems, ISO 15746-2 (2017) specifies information exchange requirements, workflows and service definitions. However, although the standard saves some design effort, it still does not provide an actual solution. Therefore, this work aims at filling a clear research gap.

3. EFFICIENCY CHALLENGES IN LARGE-SCALE PROCESSES

Fig. 1 provides an overview of the complexity of copper refinement (from sulphide ores) and steel refinement. Both the refinement processes cover multiple phases. Some of the phases are continuous processes while the others are batch processes. Either way, each phase must be controlled so that the entire chain may work appropriately, and it is clear that the output and performance of one phase may largely reflect to the others. Various refinement phases also produce impurities that have to be collected and possibly further processed (e.g., the capture of sulphur oxide as documented by Schlesinger et al. (2011)). Overall, the two refinement processes consist of concurrently operated unit processes with mutual dependencies, and the related optimisation is difficult, requiring a considerable amount of knowledge to be managed.

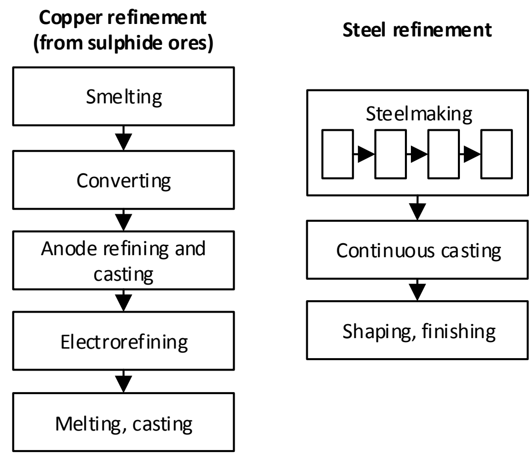

Fig. 1. Copper and steel refinement processes simplified. Adapted from (Schlesinger et al., 2011, pp. 1-12) and (Overview of the steelmaking process, 2013).

During the work, multiple requirements have appeared from copper refinement. A major problem is that it is difficult to coordinate unit process operation; however, due to unit process dependencies, it is clear that coordination would be beneficial. To coordinate operator work, there must be an appropriate mechanism, such as an advisory system. An advisory system should also assist unit process operation in detailed process-specific matters, as each single unit process is complex as such. Effective advisory also requires application of mathematical models; only then, it is possible to consider process dynamics properly. Furthermore, unit process scheduling is essential due to various timing requirements. Finally, there is also a requirement of reacting to various events in an asynchronous manner; for instance, finishing a particular batch process may be a triggering precondition for other operations.

Respectively, steel refinement has its specific requirements. Various parameters affect the output and its quality. Quality control requires process-related knowledge, and mathematical models should be applied to imitate process dynamics. Coordination is also required to operate the overall refinement efficiently. Finally, quality control requires production monitoring which enables the observation of process outputs, thus providing essential data.

There are also general ICT requirements that apply to both cases. First, production plants are often distributed, and some facilities may even be physically located in another plant; thus, connectivity over the Internet is required. Second, there is considerable variety among industrial back-end systems and information sources, such as measurement data providers. Thus, connectivity mechanisms must be adaptable to arbitrary information interfaces. Third, to facilitate system maintenance and updates and to improve scalability, the integration techniques should enable loose coupling between system components (as designed strict dependencies increase the related difficulty). Any updates in one module – such as applying fresh process knowledge – should not needlessly require updates in other modules. Fourth, generic or standard data structures should be favoured for three main reasons: some specification effort is saved if previous work is reused; generic structures contribute to interoperability between system components; and the solution becomes more applicable in various industrial domains.

The list below summarises the requirements arisen from the cases.

- Enable both unit process optimisation and plant-wide coordinating optimisation
- Utilise knowledge both at the unit process level and coordinating level
- Utilise mathematical models to capture knowledge and imitate process dynamics
- Enable optimisation in a reactive, asynchronous way
- Enable connectivity over the Internet
- Enable integration with any information system
- Favour standard information structures in communication
- Utilise technologies that enable loose coupling (to facilitate updates and to improve scalability)

The requirements are applicable to the various domains of industrial production. Due to the heterogeneity of production facilities, the resulting design should be generic and conceptual, as a high degree of adaptability is necessary.

4. SYSTEM DESIGN TO ENABLE COORDINATING OPTIMISATION

4.1 Interoperability Layer Concept

Due to concurrent operation and distribution of ICT units, interoperability matters are in the core of the design. Back-end information systems provide the data necessary for operation and optimisation. Due to coordination, even optimisation modules may have to communicate mutually. There must also be functionality to utilise mathematical models in optimisation tasks. User interfaces show information to process operators.

The layout of ICT systems varies between production facilities. To keep the concept consistent, a high abstraction level is necessary. Fig. 2 illustrates the concept of an *interoperability layer* that enables communication between arbitrarily organised systems and modules. Although the concept is simple, the interoperability layer requires careful design; heterogeneity is inevitable, and so are the communication needs of each module. However, as communication-related complexity is kept – as much as possible – at a dedicated level in the architecture design, further module development becomes easier. This choice follows the *separation of concerns* principle of software design.

4.2 Communication Stack and Technologies

Fig. 3 illustrates the required details as the interoperability layer concept is developed to a platform design. The platform requirements are numerous and so are platform components.

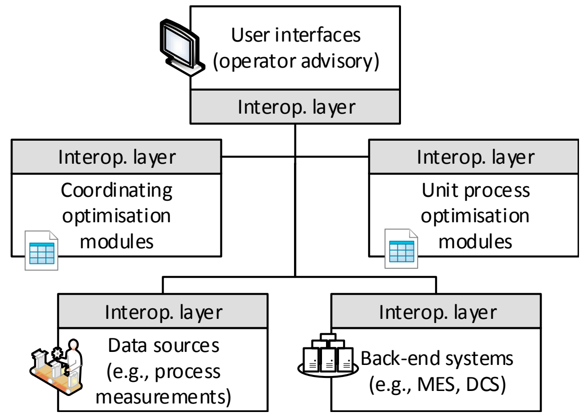

Fig. 2. An interoperability layer enables the interaction of each system module. (MES = Manufacturing Execution System, DCS = Distributed Control System)

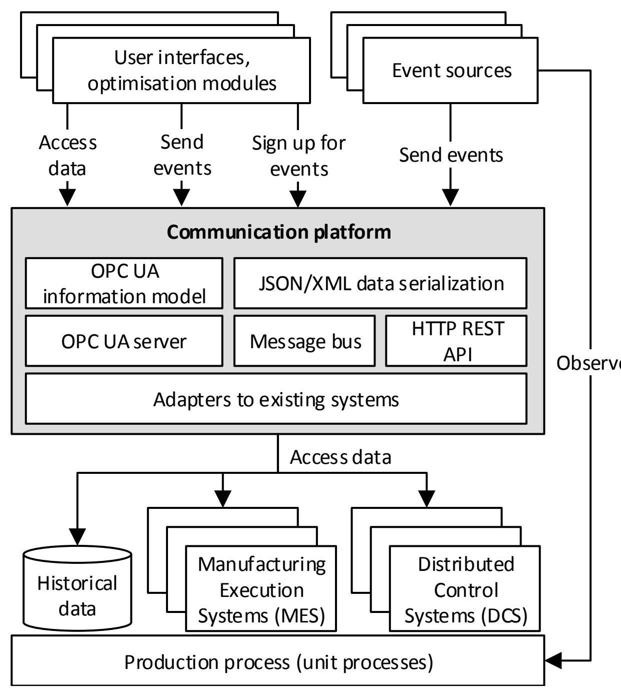

Fig. 3. The technologies of the communication platform and its position related to other systems.

A message bus enables loosely coupled, asynchronous communication as well as event-based functionality. Instead of connecting modules point-to-point, loose coupling is maximised as the message bus may deliver data between peers in an arbitrary manner. This enables easy reconfiguration or rearrangements (as expected if a system instance has a long lifespan). Asynchronism is enabled with message queues managed by the message bus with two means. First, after submitting a message, its sender does not have to care about its further processing. Second, each recipient is free to process incoming messages in any way or at any time without delaying other nodes. For the message bus protocol, a good candidate is AMQP (Advanced Message Queueing Protocol). It provides a centralised approach; with high loads, the bus may become a bottleneck, but it can be scaled up as needed. A significant advantage of centralised messaging is its easier management.

Synchronous communication is also required in some use cases. A use case for it may be, for instance, a user request for production performance reports or another case requiring request-response messaging. HTTP REST API provides a synchronous data access mode. REST (Representative State Transfer) is an architectural style that refers to modelling interfaces as resources rather than operations. In some cases, the synchronous API may also wrap asynchronous communication via the message bus for more straightforward access.

OPC UA provides a means to expose industrial process measurements or other equipment-related information in a standard manner. Its standard information model, end-to-end security as well as wide tool support facilitate integration, and its client-server approach is a common messaging pattern.

As data delivery formats are chosen, commonly agreed, digitally structured, text-based formats should be favoured. One candidate is B2MML (2013), which covers ISA-95 for manufacturing management (e.g., scheduling or resource allocation) and ISA-88 for batch process details (e.g., recipes or batch process monitoring). As measurement data is delivered, a format such as Observations and Measurements (2013) is suitable. The related data formats are generic enough to be suitable for arbitrary cases, but additional specification is then required for utilisation. Although multiple physical formats may be used (including OPC UA, JSON or XML), a similar logical structure should be utilised for each. Then, any mappings between the formats become straightforward, and some design effort is saved.

The concept enables arbitrary adapters to any back-end system. The vendors of each physical system may favour their proprietary data formats. The data access purposes may cover, for instance, historical data, online process data exposed by DCS systems or manufacturing operations related information from MES systems. These issues are to be covered in future research.

5. IMPLEMENTING ASYNCHRONOUS COMMUNICATION

5.1 Coordinating Scheduling

The first demonstrated use case focuses on coordinating production scheduling. In the case, a unit process receives a schedule from a coordinating optimisation module. The case emphasises events and asynchronism: when a unit process finishes a batch, the coordinating optimiser receives a notification and reacts by sending a refreshed schedule.

The related practical example is the coordination of Peirce-Smith Converters (PSC) in a copper refinery (in Fig. 1, they perform the "converting" phase). There are multiple PSCs that are run concurrently, and the coordinating optimiser module calculates approximate batch schedules for them. No actual production systems are included in this implementation. Instead, communication is demonstrated with simple applications that imitate real systems.

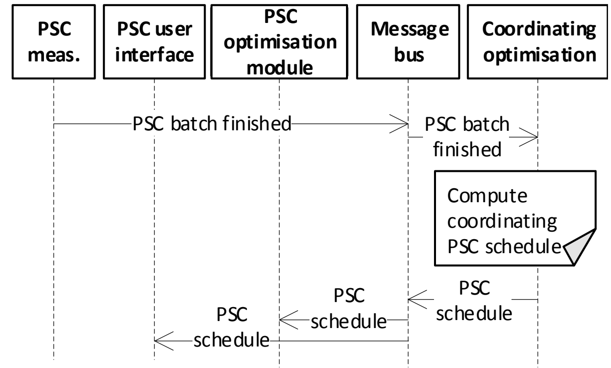

Fig. 4. Module interaction in the coordinating optimisation demonstration.

Fig. 4 illustrates module interaction in the example. As the coordinating optimisation logic receives the information that a PSC batch has finished, it refreshes the schedule that coordinates PSCs. It submits the schedule to both a PSC optimiser (for more detailed scheduling) and a PSC user interface (to be displayed to operators). All the message traffic is transferred via an asynchronous message bus; that is, there are no direct links between the modules. In a more extensive case, the coordinating optimiser would also compute the schedules of other unit processes, but this example is limited to PSCs.

The indirect messaging mechanism has been implemented with a *topic-based* approach. That is, at design time, specific messaging topics are agreed on. For instance, the topic for coordinating PSC schedules may be "coordination.schedules.psc". The producer of those schedules (i.e., the coordinating optimiser) submits PSC schedules to that topic, and any consumer modules sign up for the topic. The topic-based message traffic is coordinated by the centralised message bus; thus, there are no direct dependencies between message producers and consumers.

The demonstration utilises open, platform-independent technologies. The information about batch finishing is structured utilising the Observations and Measurements schemata. The coordinating PSC schedule is structured as a B2MML production schedule. Finally, the message bus is RabbitMQ, which utilises AMQP protocol. All AMQP traffic is encrypted, and user access control is also applied – thus, the modules may communicate securely over the Internet.

5.2 Unit Process Scheduling with Calculation Module

In contrast to coordinating scheduling, the detailed scheduling example has more focus on a single unit process. A detailed schedule provides comprehensive instructions to process operators. In the example, there is a Flash Smelt Furnace (FSF) that performs the smelting unit process of copper refinement (see Fig. 1), providing material input for PSC batches. Fig. 5 illustrates the case. First, a measurements module produces an FSF batch composition estimate, which is then delivered to a PSC optimisation module; PSC optimisation utilises the composition information to estimate the PSC processing needs of the batch. Using the composition estimate, the PSC optimiser mod-

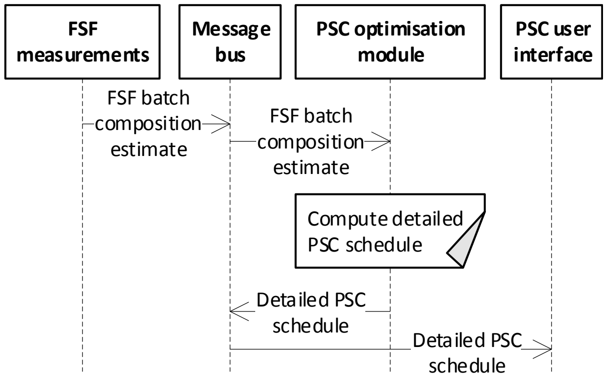

Fig. 5. Module interaction in the unit process scheduling demonstration.

ule calculates a fresh PSC schedule. Finally, the message bus delivers the schedule to the PSC operator interface. Similar to the first scenario, this demonstration does not integrate actual production systems either but runs simple applications that imitate real systems.

The technology set of the previous demonstration is extended with a calculation module. B2MML holds scheduling information, Observations and Measurements schemata provide the structures for measurements, RabbitMQ delivers messages and messaging is based on topics. However, an additional Matlab-based calculation module has been integrated to generate the schedules. Although Matlab may not be the best choice for production environments, it demonstrates the integration of mathematical models. Matlab also enables the utilisation of extensive simulations.

6. DISCUSSION

The identified concept requirements stem from two practical cases in process industry. To truly evaluate if the concept covers the varying needs of process industry, more cases should be studied. Still, it is noteworthy that, due to the heterogeneity of organisations and facilities, some adaptation and customisation is required for each physical plant anyway. In addition, especially the copper case has a high degree of distribution, asynchronism and process complexity. Consequently, its requirements likely provide an extensive coverage in platform functionality and features.

The concept meets its requirements well. The envisioned platform enables production coordination, knowledge modelling in optimisation modules, distribution, synchronous and asynchronous communication as well as loose coupling. The concept is also generally adaptable to other production plants.

As optimisation modules are integrated, the utilisation of the platform causes some overhead in the related design, but the overhead is compensated by the various advantages of the platform. It could be argued that a point-to-point approach would be a faster, more straightforward way to implement integration. However, the problem of point-to-point is the lack of scalability and flexibility in terms of changes; in complex point-to-point scenarios, modifications are expensive and difficult to perform. Industrial information systems have traditionally been independent and isolated from each other (ISO 15746-2, 2017). However, due to business evolution, adaptability is more important in future enterprise information systems than today (Weichhart et al., 2016). These adaptability aspects are contributed to by the platform; as all the messaging is performed in a loosely coupled manner, it is easy to manage systems integration, module dependencies and information routing. Although some production systems may remain the same for a long time, optimisation knowledge will likely evolve anyway.

The specification of information structures utilisation likely requires additional work. A challenge of the platform is that multiple information structure standards are required. Although appropriate standards have been discovered for the proofs of concept, they likely provide no all-round coverage within industrial production. Furthermore, while B2MML provides a good basis for numerous needs, it is a system-level specification and, thus, ignorant of equipment level information such as measurement values. In addition, B2MML structures are somewhat loose. Instead of only instructing practitioners to use a highly adaptable standard, there should also be recommendations about its detailed application. Similarly, although Observations and Measurements seems well suitable for the system, there should also be guidelines that determine which of its fields are utilised and how.

The implemented demonstrations show how a part of the concept may be implemented. However, while functional and relevant, they also have a limited scope. First, they only focus on one case (i.e., copper refinement), and second, their coverage within that case is limited. Still, there is focus on asynchronous communication, which is typically more complex and more difficult to implement than synchronous functionality. Implementing synchronous, high-level service interfaces (e.g., REST) and related clients is straightforward with modern development tools. In addition, communication scenarios are expected to repeat certain patterns, and an extensive number of mutually resembling experiments would bring little conceptual value.

Compared to the traditional monolithic control systems, the proposed generic, asynchronous communication platform would significantly increase the possibilities of production optimisation. It would be possible to build systems that consider both global optimisation and the local requirements or restrictions of unit processes. Large-scale optimisation tasks would be taken to a new level due to improved communication and easier information access.

7. CONCLUSION

This article has presented a platform concept to enable coordinating production optimisation. Asynchronism, distribution and integration of various systems and modules were emphasised in particular. First, the requirements of the platform were considered, then, a platform design was introduced and, finally, two proof-of-concept implementations were demonstrated. Although the results still have a limited scope and coverage, the platform design seems promising. The practical industrial cases (copper and steel refinement) likely provide sufficient challenges to recognise the required features of a generic communication platform that suits for other industrial cases as well. With the proposed platform, it would be possible to utilise production-related knowledge in an unprecedented way, which has notable potential to improve efficiency.

Various issues still require further research to improve the coverage of the platform concept. The data formats to deliver information between modules could be further studied due to the limitations of the current scenarios. The integration of actual production systems could also be considered. Furthermore, related to the cases, the number of the scenarios covered could be extended.

ACKNOWLEDGEMENTS

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 723661. This study reflects only the authors' views, and the Commission is not responsible for any use that may be made of the information contained therein. The authors want to express their sincere gratitude to the funder and the project partners in the COCOP project (Coordinating Optimisation of Complex Industrial Processes, `https://www.cocop-spire.eu/`).

Furthermore, the authors are grateful to the Academy of Finland for their funding (grant 310098).

REFERENCES

B2MML (2013). Business to manufacturing markup language. MESA international. `http://www.mesa.org/en/B2MML.asp`. Accessed: 24th Oct 2017.

Bleser, G., Damen, D., Behera, A., Hendeby, G., Mura, K., Miezal, M., Gee, A., Petersen, N., Maes, G., Domingues, H., Gorecky, D., Almeida, L., Mayol-Cuevas, W., Calway, A., Cohn, A.G., Hogg, D.C., and Stricker, D. (2015). Cognitive learning, monitoring and assistance of industrial workflows using egocentric sensor networks. *PLOS ONE*, 10(6), 1–41. doi:10.1371/journal.pone.0127769.

Chan, Y., Kantamaneni, R., and Allington, M. (2015). Study on energy efficiency and energy saving potential in industry from possible policy mechanisms. Technical report, ICF Consulting Limited, London.

Clements, P., Bachmann, F., Bass, L., Garlan, D., Ivers, J., Little, R., Nord, R., and Stafford, J. (2003). *Documenting software architectures: views and beyond*. Pearson Ed.

Elkader, M., Fathy, A., Megahed, G., Eissa, M., Shama, S., and El-Sharkawy, A. (2015). Influence of direct-reduced iron percentage in EAF charge mix on EAF operation parameters. Technical report, steel-grips.com.

Engell, S. and Harjunkoski, I. (2012). Optimal operation: Scheduling, advanced control and their integration. *Computers & Chemical Engineering*, 47, 121 – 133. doi:10.1016/j.compchemeng.2012.06.039. FOCAPO 2012.

Guo, Z., Ngai, E., Yang, C., and Liang, X. (2015). An RFID-based intelligent decision support system architecture for production monitoring and scheduling in a distributed manufacturing environment. *Int J Prod Econ*, 159, 16 – 28. doi:10.1016/j.ijpe.2014.09.004.

ISO 15746-2 (2017). ISO 15746: Automation systems and integration – integration of advanced process control and optimization capabilities for manufacturing systems – part 2: Activity models and information exchange. Standard, ISO.

Kannisto, P., Hästbacka, D., and Kuikka, S. (2017). System architecture for mastering machine parameter optimisation. *Computers in Industry*, 85, 39 – 47. doi:10.1016/j.compind.2016.12.006.

Kannisto, P., Hästbacka, D., Palmroth, L., and Kuikka, S. (2014). Distributed knowledge management architecture and rule based reasoning for mobile machine operator performance assessment. In *Proceedings of the 16th International Conference on Enterprise Information Systems*, 440–449. doi:10.5220/0004870004400449.

Leitão, P. and Restivo, F. (2008). A holonic approach to dynamic manufacturing scheduling. *Robotics and Computer-Integrated Manufacturing*, 24(5), 625 – 634. doi:10.1016/j.rcim.2007.09.005.

Li, B. (2010). *Research on the Production Scheduling Management System Based on SOA*, 286–294. Springer Berlin Heidelberg. doi:10.1007/978-3-642-16515-3_36.

Observations and Measurements (2013). Version 2.0. Open geospatial consortium. `http://www.opengeospatial.org/standards/om`. Accessed: 9th Nov 2017.

Overview of the steelmaking process (2013). Worldsteel association. `https://www.worldsteel.org/en/dam/jcr:177c8e5c-e02a-4e08-9dc6-cce7372b41c2/Overview+of+the+Steelmaking+Process_poster.pdf`. Accessed: 9th Nov 2017.

Quadrelli, R., Tonolo, G., Ziebinska, U., and Burghgraeve, S. (2016). Energy efficiency indicators highlights 2016. Technical report, International Energy Agency.

Schlesinger, M.E., King, M.J., Sole, K.C., and Davenport, W.G. (2011). *Extractive Metallurgy of Copper*. Elsevier, Oxford, 5th edition.

Shen, W., Hao, Q., Wang, S., Li, Y., and Ghenniwa, H. (2007). An agent-based service-oriented integration architecture for collaborative intelligent manufacturing. *Robotics and Computer-Integrated Manufacturing*, 23(3), 315 – 325. doi:10.1016/j.rcim.2006.02.009.

Ullrich, C., Aust, M., Kreggenfeld, N., Kahl, D., Prinz, C., and Schwantzer, S. (2015). Assistance- and knowledge-services for smart production. In *Proc. of the 15th Int. Conf. on Knowledge Technologies and Data-driven Business*, i-KNOW '15, 40:1–40:4. ACM, New York, NY, USA. doi:10.1145/2809563.2809574.

Weichhart, G., Molina, A., Chen, D., Whitman, L.E., and Vernadat, F. (2016). Challenges and current developments for sensing, smart and sustainable enterprise systems. *Computers in Industry*, 79, 34 – 46. doi:10.1016/j.compind.2015.07.002.

Zhong, R.Y., Li, Z., Pang, L., Pan, Y., Qu, T., and Huang, G.Q. (2013). RFID-enabled real-time advanced planning and scheduling shell for production decision making. *International Journal of Computer Integrated Manufacturing*, 26(7), 649–662. doi:10.1080/0951192X.2012.749532.