2020 17th International Conference on Ubiquitous Robots (UR)
June 22-26, 2020. Kyoto, Japan

Study on weld seam geometry control for connected gas metal arc welding systems*

S. Mann, R. Glebke, I. Kunze, D. Scheurenberg, R. Sharma, U. Reisgen, K. Wehrle, D. Abel

Abstract — **In this study, a robot welding application for the control of the weld seam geometry by means of in situ image acquisition and robot trajectory correction has been investigated. For this approach, the arc centroid position of the process images was correlated with the weld seam flank ratio of the examined fillet weld seam application. The correction was performed along a selected trajectory plane without violating the welding process stability. In addition, the system as a sum of its partial competences with their different synchronization requirements and corresponding bottlenecks was illuminated. Optimization suggestions arose especially in advanced control engineering as well as in-kernel processing and in-network computing. In summary, it can be stated that, on the one hand, the examined relationship between arc centroid and flank ratio is subject to interference and that image processing in particular has potential for optimization.**

I. INTRODUCTION

Welding production remains one of the largest application areas of industrial robotics with strong prospects for collaborating robots in manual welding production. Harsh process conditions, but also an omnipresent lack of skilled workers in the highly empirical field of welding, aggravated by unfavourable demographic changes in many high-wage countries, are among the drivers of robot-supported welding production. However, major challenges arise in holistic automation approaches to guarantee decisive weld seam quality criteria. The robust and time-deterministic combination of sensors, actuators, and welding process technology is critical, but also strongly dependent on the application. The comprehensive use of digital information sources in the industrial context following initiatives such as Industrie 4.0 opens up new potentials for addressing various robot-supported applications and is thus bringing them into widespread use [1]. On the other hand, however, especially in the case of networked, time-critical systems, the question must be answered as to where and how information should be processed in order to avoid latencies which are too high or to overload networks with high raw data streams. The following work describes a robot welding system which is capable of controlling the weld seam geometry by means of in situ image acquisition. The control tasks as well as image processing are successfully executed centralized on a PC platform. Nonetheless, specific bottlenecks can be identified which reduce the system performance significantly and may be outsourced to different computing competencies in the future.

II. WELD SEAM GEOMETRY CONTROL

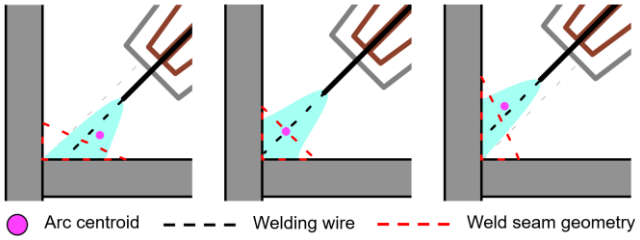

Figure 1: Observing the weld seam geometry by means of the arc

State-of-the-art industrial systems in robot arc welding with closed control loops are, e.g., seam tracking control loops to maintain welding torch movement aligned with the seam. Tolerances of the workpiece geometry, positioning errors, or distortion due to previous weld seams are among the imponderables that can make closed loop seam tracking indispensable. The applied control strategies close their control loops via seam tracking sensors using laser line scanners, or the so-called arc sensor at a corresponding robot interface for the correction of the path planning [2]. Investigations to improve seam tracking via structured light projection have been described in [3] and [4]. However, these control systems are not able to control the actual welding result itself. They lack in particular the application-bound knowledge to map the available sensor data to decisive weld seam quality criteria (application model) as well as a suitable strategy (control strategy) to counteract expected deviations without violating basic welding process stability or other quality criteria. The actual weld seam geometry, as a decisive feature for ensuring the load bearing capacity of the joint, remains therefore untouched in the automated welding system itself. Gao et al. described a system which is observing the weld pool for seam tracking applications, utilizing linear regression and artificial neural networks for tungsten arc welding [5, 6]. Nevertheless, this work lacks integration and validation in a closed control loop setup, especially for GMA welding with its highly dynamic lighting characteristics. An aforementioned, holistic welding system overstrains available

*Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2023 Internet of Production -- 390621612.

S. Mann, R. Sharma, U. Reisgen are with the Welding and Joining Institute (ISF), RWTH Aachen University, Germany (corresponding author: Samuel Mann, phone: +49 241-8097242, email: mann@isf.rwth-aachen.de)

R. Glebke, I. Kunze, K. Wehrle are with the Chair of Communication and Distributed Systems (COMSYS), RWTH Aachen University, Germany

D. Scheurenberg, D. Abel are with the Institute of Automatic Control (IRT), RWTH Aachen University, Germany

robot systems according to demands of flexibility, interfaces, and computing power, which led to the approach to develop a PC-based control system to cover performant image processing and a broad availability of interfaces at one place. This platform remains, on the other hand, hardly real-time capable and will most likely uncover critical communicational and computational demands which come hand in hand with the investigated, holistic automation approach.

In particular, the weld seam geometry in a fillet joint application will be controlled by recording the arc centroid of in situ image acquisition by correcting the robot trajectory along a compensation plane during welding. Figure 1 describes the relationship between the position of the arc center of gravity relative to the welding wire axis and the resulting weld seam geometry. This effect has been mentioned in [7] but not yet examined under the demands of a closed loop control system. Disturbance variables such as the blowing effect or an asymmetrical heat distribution on the workpiece can deflect the arc and thus the arc center of gravity. This in turn, exert influence on the resulting weld seam geometry. The safe guiding of the welding torch in the joint via the previously described control system for seam tracking can therefore not be used on its own to ensure the weld seam geometry.

III. Robot Welding System and Application Control

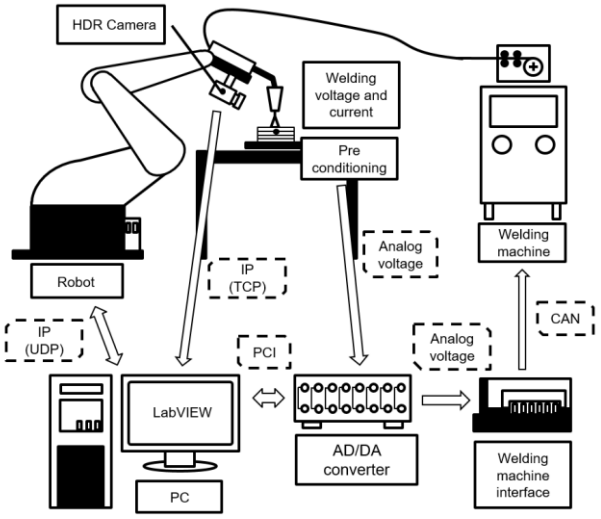

Figure 2: Robot welding system overview

Figure 2 describes the system in use with its components and interfaces. Data processing and control are carried out on a PC, whereby communication with the welding machine (EWM alphaQ 551) takes place via an analogue control interface (RINT). In contrast, the robot (KUKA KR-16-2 KRC4) and the camera (Photonfocus HD1) are IP-based, using UDP for the robot control and TCP for the acquisition of image data.

The extraction of the relative position of the arc centroid from the in situ process recording is done by the image processing strategy described in Chapter IV. The linking of this information with the resulting weld seam geometry and thus the controlled variable is carried out via the so-called application model in Chapter V. The control of the weld seam geometry (control strategy) via the 6-axis industrial robot used is described in more detail in Chapter VI.

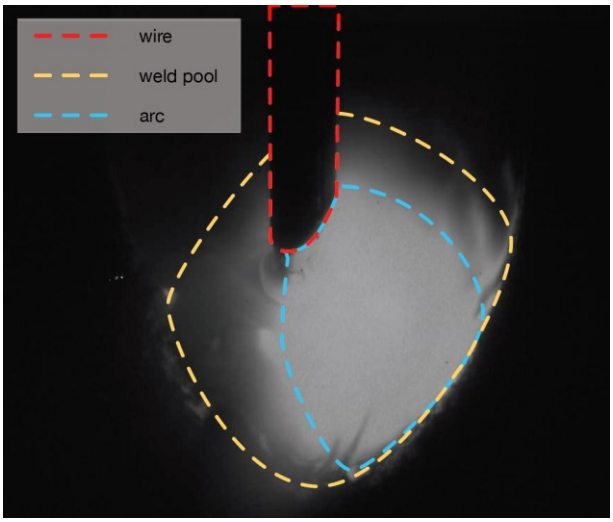

Figure 3: In situ welding process image

Figure 4: Arc centroid determination; intensity depth reduction, threshold, convex hull and centroid determination from left to right

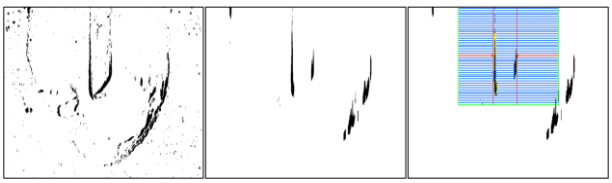

Figure 5: Welding wire position determination; multiplied truncated pyramid filters, asymmetrical median filter and edge detection along search lines from left to right

IV. Image Acquisition and Processing

Figure 3 shows an in situ image of the welding process zone taken by the in situ camera. To reduce the intense radiation of the metal vapor components, a band-pass filter at 830 nm was used for optical image filtering. Clearly visible are the molten pool (yellow border), the wire electrode (red border), and the arc (blue border). The corresponding exposure time in the following investigations was 800 microseconds. The scale of the images used is 12 µm per pixel. The camera position is statically fixed relative to the welding torch, but the wire position can vary due to the cast of the welding wire and wear of the contact tube. To be able to capture the relative movement of the arc centroid x_C to the wire electrode x_W, the latter must therefore also be captured in the image. The relative centroid position is finally calculated by $x_{c-w} = x_C - x_W$.

A. *Image Processing - Arc Centroid Position*

The consecutive image processing steps are depicted in Figure 4: The brightness intensity is first reduced to four values, with the highest intensity of the image being assigned to the arc. A convex envelope is then formed to equalize the uneven contour and thus increase the stability of the centroid position.

B. *Image Processing - Wire Electrode Position*

Figure 5 describes the corresponding process. The edges of the wire electrode are first highlighted by multiplying two directional 7x7 truncated pyramid filters. The edges from left to right and once the edges from right to left are emphasized. In the next step, an asymmetric median filter (Δx=3, Δy=50) along the vertical axis was applied to reduce noise in the image that cannot be attributed to the wire electrode running vertically in the image. In the last step, the position of the wire contour can be averaged by contrast edge detectors along a number of search lines (blue) in the image. The optical distortion has been calibrated via test patterns and can be considered negligible, especially orthogonally towards the wire axis.

V. Application Model

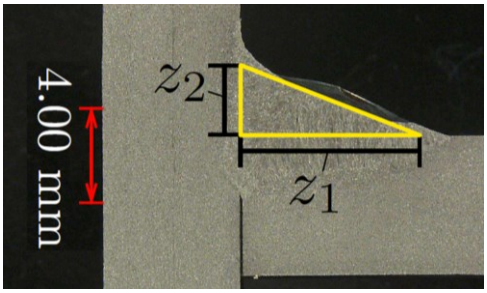

Figure 6: Weld seam leg definition in a fillet weld seam cross section

The application model maps the sensor-observable variable (here: the position of the arc centroid relative to the wire electrode) to the application variable (here: the flank ratio $r = \frac{z_1}{z_2}$ of the fillet weld seam) (see Figure 6). In order to determine this relationship, three series of measurements were carried out in which the welding torch was guided at an angle to the joint path in order to produce different flank geometries. For each series of measurements, 19 cross-sectional cuts were made, and the corresponding flank ratios were compared to the measured centroid-wire offset (see Figure 7).

Table 1: Welding boundary conditions:

Boundary Condition	Value
Filler wire	EN ISO 14341-A: G 3Si1; 1.2mm
Shielding Gas	Ar: 98%, CO2: 2%
Sheet metal	S235 JR; 6mm
Welding position	PB

Table 2: Welding parameters

Parameter	Unit	Value
Wire feeding speed	m/min	7.5
Welding speed	m/min	0.1
Welding current	A	234
Welding voltage	V	25.9

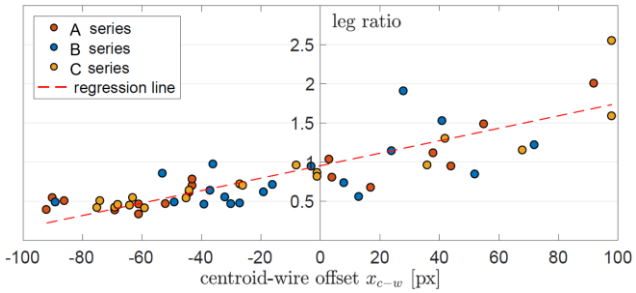

Figure 7: Weld seam leg ratio over centroid-wire offset

All welding tests were performed according to the parameters and boundary conditions shown in Table 1 and Table 2.

First, Figure 7 shows an approximately linear relationship between the centroid wire offset and the leg ratio, which can be represented by a linear regression. The straight line intersects the leg ratio 1 at a centroid wire offset of 0, which corresponds to a symmetrical weld seam and was to be expected. On the other hand, the large number of measuring points is at a leg ratio <1, which can be explained by the asymmetrical sheet arrangement and the associated change in blow effect and heat distribution as well as the welding position. However, the overall movement bandwidth (±100pxl) of the centroid wire offset must be considered low as it describes a total movement of ±1.2mm given the pixel ratio in Chapter IV. In contrast, undefined asymmetries in the weld seam according to [8] are already outside the highest permissible evaluation group with a leg difference $h > 1.5\ mm + 0.15\ a$.

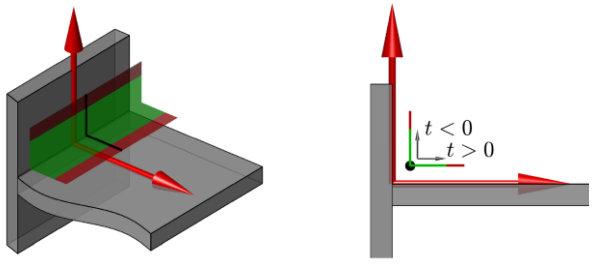

Figure 8: Trajectory surface with correction parameter t

VI. Control Strategy

The robot was controlled via the KUKA.PLC mxAutomation interface by sending linear motion commands to the buffer of the robot controller. In order to be able to guarantee an almost constant welding speed, the movement commands are ground in the robot internal controller. This also means, however, that at least one motion command must be present in advance in the buffer to enable a corresponding grinding. In the standard configuration of the KUKA robot, even two motion commands must be present in advance. This fact alone results in a dead time of at least two motion commands.

The flank ratio correction of the system was accomplished by the motion correction of the robot. For this purpose, the movement horizontal and vertical to the weld joint was

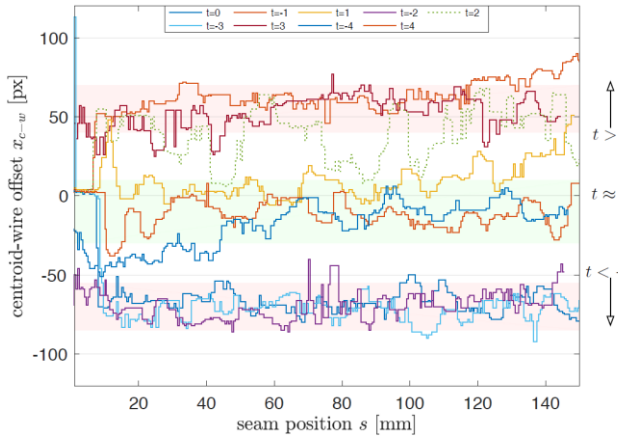

Figure 9: Step response for different step heights

reduced to an auxiliary variable t, which lies on a previously defined trajectory plane (see Figure 8).

To check the ability to adjust the flank ratio, step responses from t=-4 to t=4 along the trajectory plane were performed and averaged over 5 images per median (see Figure 9). On the one hand, the results show a strong noise, which can be attributed to process irregularities and the associated measurement uncertainties of the image processing, which are strongly pronounced even with the applied median filtering over 5 images. In addition, the step tests tend to jump to maxima and minima at approx. ±75pxl at deflections greater than 1 and smaller than -1, which could also be observed in [7].

Different deflections in these areas marked in red are therefore difficult to differentiate. The weld seam geometry control is thus reduced to the control of a symmetrical flank ratio.

VII. System Synchronization

As already shown in Figure 9, the detection of x_{c-w} is overlaid with a high degree of noise due to the high process dynamics of gas metal arc welding. This difficulty can be counteracted with more complex image processing methods or with appropriate smoothing of any interference over several images. However, the latter reduces the transient resolution of x_{c-w}, while more complex image processing methods require higher computing times. Initially, the used camera limits the acquisition rate to 54fps, since the radiation intensity at the used exposure time of 800µs theoretically allows image frequencies of 1250fps. Nonetheless, the duration of the image processing turned out to be much more decisive. With up to 150ms, the duration of the applied processing algorithms was significantly higher than the acquisition rate at 25fps or 40ms respectively, but was also, depending on the image, inconstant and sometimes only 30ms long.

Welding robot path planning was carried out online, and allowed a minimum length of 2 mm of the path increments to ensure continuous speed with an update time of 200 ms. This limited the absolute motion speed to 0.6m/min, which is still acceptable for arc welding. However, the torch was corrected transversely to the joint, resulting in a lower welding speed than 0.6m/min in favor of lateral movement. In the practical tests it was also found that the control of the robot via a non-real-time PC in an application with deterministic requirements, especially via the UDP protocol that cannot register packet losses, is faulty and can lead to malfunctions. If the buffer cannot be filled with movement commands in time, the welding movement will stall.

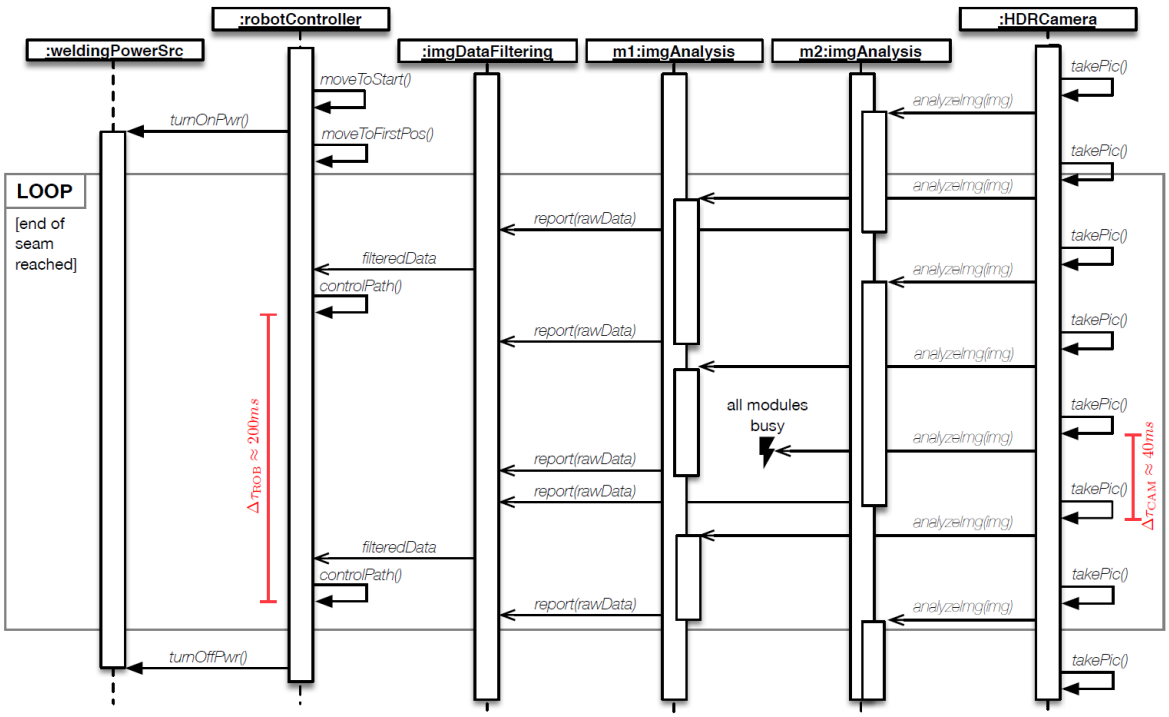

Figure 10: System synchronization sequence

In summary, the following limiting factors can be identified which must be taken into account in a controlled system:

- The determination of the centroid wire offset is subject to interference, so that a high frame rate must be used to smooth the interference.
- The image processing duration is variable with 30-150ms and mainly longer than the image acquisition time.
- Motion commands can be sent to the robot every 200ms.
- The motion control of the robot requires two motion commands in advance to maintain an approximately constant motion speed.

To cope with these influences, a synchronization sequence as shown in Figure 10 was chosen.

Starting from the right, the process images (*HDRCamera*) are captured at a frame rate of 25fps or 40ms and thus represent a compromise between acceptable image processing duration in the following process steps and a necessary amount of data for median averaging. The available images are then processed from a buffer in two parallel image processing modules (*imgAnalysis*). The resulting intervals of the calculated centroid wire offset are then no longer constant. This circumstance is being counteracted by averaging over the variable amount of data in the data filter module (*imgDataFiltering*) which is then fed to the robot controller every 200ms. If all image processing modules are occupied, the image is rejected. Since the maximum image processing time is 150ms below the robot control interval of 200ms, it can be ensured that at least one processed image is available for the next welding torch correction.

VIII. Closed Loop Control Results

As the application model described in Figure 7 consists mainly of a linear relation, a simple PI controller seemed suitable for a first approach. The following tests were carried out after step tests for empirical controller parameterization with the controller structure shown in Figure 11. A controller setting of

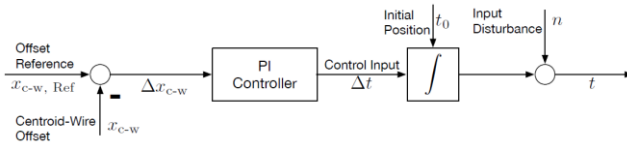

Figure 11: Fundamental controller structure

KP=0.01 and KI=0.01 resulted in oscillation, as seen in Figure 12, whereas KP=0.001 and KI=0.01 converged with a slight overshooting (Figure 13). For the evaluation of the controller, a disturbance variable of +4t and -4t was initiated after each 150mm of weld length and controlled to $x_{c-w} = 0$. The detected centroid wire offset, the position of the robot in the compensation plane, and the welding result of two experiments are shown in Figures 14 to 16 for a short-circuiting process of mild steel and a shielding gas composition of 8% CO2 and 92% Argon. The absolute welding torch position has been obtained directly from the robot controller.

The results show a rather slow convergence and overshooting in Figure 14 after the deflection of 4t. The presentation of the measured centroid wire offset in Figure 15 also describes a

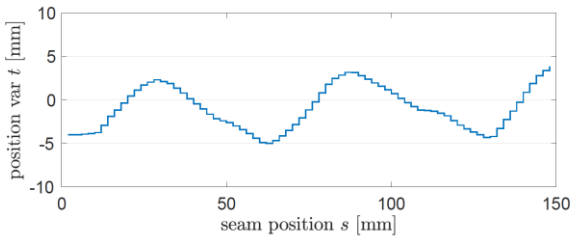

Figure 12: Closed loop control experiment KP = 0.001, KI = 0.001

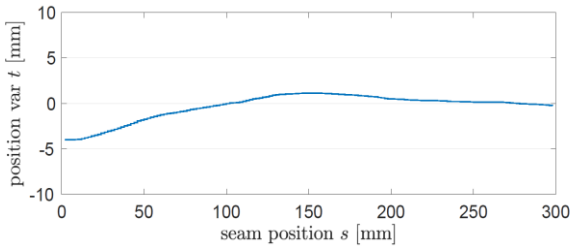

Figure 13: Closed loop control experiment KP = 0.001, KI = 0.01

very noisy measurement and the non-linearity of the controlled system described in Figure 9. A comparable system performance could be observed in all experiments. This is even more critical due to the highly dynamic lighting characteristics of the short-circuiting arc. The reference x_{c-w} can be examined at around 25 px (dashed red line) rather than at 0 px due to the welding position.

The same tests were conducted as well for a more stable spray transfer process as seen in Figures 17 to 19 under the same boundary conditions. The reference x_{c-w} can be examined, again at around 25 px, pointing towards the welding position, rather than noise. Besides one overshooting behavior while compensating -4t, the overall performance can be described as slightly better.

This study has been proofing the overall concept but there remains room for further investigations and optimizations as described in the following chapter.

IX. Outlook – Advanced Control Methods

The presented work shows that a basic control of the motion correction of the robot with a simple PI controller is possible. However, the performance of the process control can be improved through the expansion of the existing PI controller to a PID controller. The additional derivative improves the damping and therefore flattens the offset trajectory. The parameterization of the controller, to achieve best disturbance rejection, can be carried out by the Ziegler–Nichols tuning method. This heuristic tuning method approximates the controlled system as a first order transfer function with dead time and is based on the step response of the system [9]. Therefore, no analytical model of the process is needed.

Since the controller and its parameterization explicitly take the dead time and disturbances of the process into account, an improved performance, especially concerning convergence and overshooting of the welding torch position, is expected.

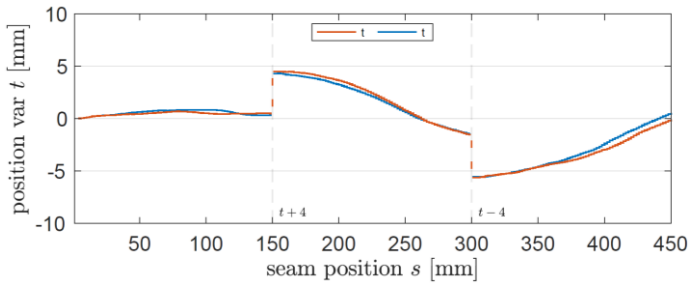

Figure 14: Absolute welding torch position – short circuiting arc

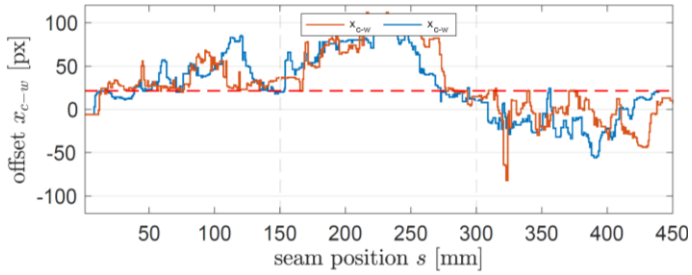

Figure 15: Centroid wire offset – short circuiting arc

Figure 16: Weld seam image – short circuiting arc

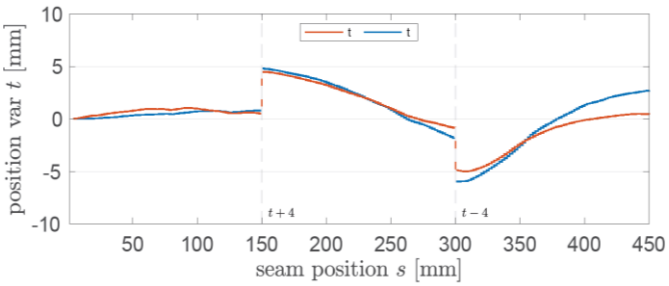

Figure 17: Absolute welding torch position – spray arc

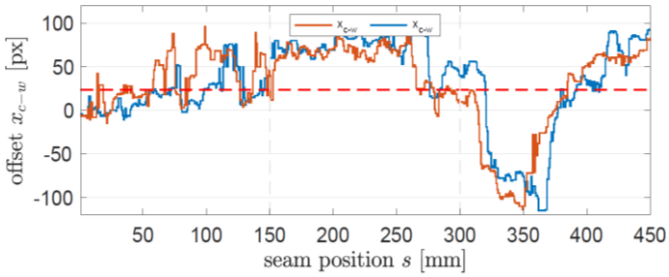

Figure 18: Centroid wire offset – spray arc

Figure 19: Weld seam image – spray arc

A. *Model Predictive Control*

A more advanced control method which would significantly improve the performance of the system, is the model predictive control (MPC). It is based on an iterative, finite horizon optimization of a cost function [10], whereas the cost function objective in this case, is to minimize the position offset while keeping the input changes small. The optimization calculates the optimal input values for the complete horizon, though only the first input is implemented. In the next time step, the optimization process is repeated. An MPC also takes constraints, disturbances, and dead time into account and can also handle nonlinear models.

However, this method requires a sufficiently precise model of the process and the current dynamic state of the process. The application model and relationship between the centroid wire offset and the leg ratio, therefore, need further refinement in modeling to describe the dynamics of the process more precisely. In addition, a model of the dynamics of the motion control of the robot has to include any delays in the execution of the movement commands of the welding torch.

X. Outlook – Image Analysis Latency Reduction

While our initial results suggest that image analysis is able to improve the control performance, we note that with values between 30ms and 150ms, the duration and variation of image analysis times make it difficult to achieve even better control performance in our setting. Lower delays and variance (jitter) would enable us to reduce the update intervals with which we send new commands to the robot and thus potentially improve the precision of our welding process.

There are multiple ways of achieving lower and more stable image processing latencies, based on the observation that the camera pictures are sent over a network link to the controlling PC, which then analyzes the pictures with application software. This places the burden of dealing with both the network protocols and the image analysis on the PC's CPU. Reducing the amount of work carried out by the CPU may thus increase processing speeds. LabVIEW already allows outsourcing computer vision-related tasks to graphics processing units (GPUs), and utilizing this method would enable us to open up more image processing modules to cope with the high frame rate of the camera. However, in general-purpose operating systems such as Windows and Linux, it would still be up to the operating system kernel (and hence, the CPU) to extract the image data from the network interface card (NIC) and direct it towards the GPUs [11]. Recent developments in operating systems and networking research try mitigating this limiting factor and can be combined with our approach to varying degrees and with varying effort.

A. *In-Kernel Processing and Short Paths*

Exhibiting only a minimal basic networking-related API to applications, operating system kernels abstract the handling of protocols such as TCP/IP and provide isolation between multiple applications accessing the network. However, this comes at the cost of having to copy received data (or data to transmit) multiple times before the destination application (or NIC) is reached, which incurs additional latency [12].

There are techniques to offload small-scale programs operating on packets directly into the Linux kernel [13], which would avoid the copies. However, these so-called (e)BPF programs are restricted in their expressiveness [14] and therefore most likely not suited for our complex processing pipelines. A more viable approach is to apply a specific kernel bypassing technique such as netmap [12], which allows

accessing NIC buffers without having to go through the kernel first. In this scenario, we can keep our original control- and image processing applications with only minimal modifications. However, circumventing the kernel means that we need to re-implement the TCP/IP protocols ourselves or resort to existing variants such as mTCP [15], either way adding additional complexity to our implementation.

Netmap only covers the aspect of NIC-CPU interaction. Outsourcing the computer vision-specific tasks to a GPU is thus still CPU-bound. Recently, however, there have been efforts to establish a similar shortcut between the NIC and GPUs [11,16] which allow us to directly access camera pictures in the NIC buffers from within the GPU. However, as in netmap, this comes at the cost of re-implementing the network stack, and we consider the question whether such stacks can be successfully (and efficiently) incorporated into GPU programs a matter of further research. However, we believe that if we are able to pass the camera images directly to the GPU, this could yield considerable improvements in both latency and jitter and offer the possibility to tighten the control loop such that we can issue commands at a higher pace.

B. In-Network Computing

An even more radical approach than kernel bypassing or GPU offloading is to directly perform computations on the NICs or even inside the network connecting the camera with the PC. This "in-network-computing" approach has only recently come to attention following the introduction of programmable network hardware (both NICs and switches/routers). Similar to (e)BPF programs, the programs on such hardware target reliably high data rates and thus are limited in their capabilities (e.g., no native floating point arithmetic and limited memory) but offer the potential to reduce the amount of image data that needs to be passed to and processed by the PC by performing (pre-)processing steps while the data is still in transit. A recent paper has already shown that a basic edge detection mechanism is feasible on networking hardware [17], although at the expense of sacrificing some of the available bandwidth. It would be interesting to investigate whether such preprocessing steps could also be applied in our setting.

XI. Summary and Conclusion

The article presents a novel approach towards holistic, connected welding systems by introducing an application model and a control strategy for weld seam geometry control. The results show that the proposed closed loop flank ratio control for a fillet weld application could be successfully implemented but leaves room for optimization. A synchronization sequence has been introduced in order to manage different latencies of image acquisition (40ms), image processing (30ms-150ms), and robot control cycles (200ms). Especially the image processing time as well as the control design approach open up opportunities for model predictive control, in-kernel processing and short paths or in-network computing.

XII. Acknowledgement

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2023 Internet of Production -- 390621612.

References

[1] U. Reisgen, S. Mann, K. Middeldorf, R. Sharma, G. Buchholz, and K. Willms, "Connected, digitalized welding production—Industrie 4.0 in gas metal arc welding," Weld World, vol. 63, no. 4, pp. 1121–1131, 2019, doi: 10.1007/s40194-019-00723-2.

[2] Sensoren für das vollmechanische Lichtbogenschweißen, DVS 0927-1:2005-02, 2005.

[3] Rongqiang Du, Yanling Xu, Zhen Hou, Jun Shu, and Shanben Chen, "Strong noise image processing for vision-based seam tracking in robotic gas metal arc welding," Int J Adv Manuf Technol, vol. 101, no. 5, pp. 2135–2149, 2019, doi: 10.1007/s00170-018-3115-2.

[4] de Xu, L. Wang, and M. Tan, "Image Processing and Visual Control Method for Arc Welding Robot," in 2003 International Symposium on Compound Semiconductors: Post-conference proceedings, 2004, pp. 727–732.

[5] X. Gao, D. Ding, T. Bai, and S. Katayama, "Weld-pool image centroid algorithm for seam-tracking vision model in arc-welding process," IET Image Process., vol. 5, no. 5, p. 410, 2011, doi: 10.1049/iet-ipr.2009.0231.

[6] X. Gao, D. Ding, T. Bai, and S. Katayama, "Weld pool image centroid algorithm for seam tracking in arc welding process," in IEEE International Workshop on Imaging Systems and Techniques: IST '09 ; 11 - 12 May 2009, Shenzen, China, Shenzhen, China, 2009, pp. 385–390.

[7] U. Reisgen, K. J. Willms, and M. Purrio, "Online-Schmelzbad-diagnostik zur Schweißnahtqualitätsüberwachung und Vermeidung von Schweißnahtfehlern beim MSG-Schweißen; Roboter 2015, 2015, pp. 30–34.

[8] DIN EN ISO 5817:2014-06, Schweißen-Schmelzschweißverbindungen an Stahl, Nickel, Titan und deren Legierungen (ohne Strahlschweißen) - Bewertungsgruppen von Unregelmäßigkeiten (ISO_5817:2014); Deutsche Fassung EN_ISO_5817:2014.

[9] J. G. Ziegler and N. B. Nichols, "Optimum Settings for Automatic Controllers," Transactions of the ASME, no. 64, pp. 759–768, 1942, doi: 10.1115/1.2899060.

[10] C. E. García, D. M. Prett, and M. Morari, "Model predictive control: Theory and practice—A survey," Automatica, vol. 25, no. 3, pp. 335–348, 1989, doi: 10.1016/0005-1098(89)90002-2.

[11] A. Nguyen, Y. Fujii, Y. Iida, T. Azumi, N. Nishio, and S. Kato, "Reducing Data Copies between GPUs and NICs," in 2014 IEEE ICCPS, Hong Kong, China, Aug. 2014 - Aug. 2014, pp. 37–42.

[12] Luigi Rizzo, "netmap: a novel framework for fast packet I/O," in 2012 USENIX Annual Technical Conference.

[13] S. McCanne and van Jacobson, "The BSD Packet Filter: A New Architecture for User-Level Packet Capture," Conference Proceedings on USENIX Winter 1993, 1993, p. 2.

[14] S. Jouet, R. Cziva, and D. P. Pezaros, "Arbitrary packet matching in OpenFlow," in 2015 IEEE 16th International Conference on High Performance Switching and Routing (HPSR), Budapest, Hungary, Jul. 2015 - Jul. 2015, pp. 1–6.

[15] E. Jeong et al., "mTCP: a Highly Scalable User-level TCP Stack for Multicore Systems," in 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14), 2014, pp. 489–502.

[16] M. Silberstein et al., "GPUnet," ACM Trans. Comput. Syst., vol. 34, no. 3, pp. 1–31, 2016, doi: 10.1145/2963098.

[17] R. Glebke, J. Krude, I. Kunze, J. Rüth, F. Senger, and K. Wehrle, "Towards Executing Computer Vision Functionality on Programmable Network Devices," in Proceedings of the 1st ACM CoNEXT Workshop on Emerging in-Network Computing Paradigms - ENCP '19, Orlando, FL, USA, 2019, pp. 15–20.