Published in final edited form as:

Robot Comput Integr Manuf. 2017 April ; 44: 144–155. doi:10.1016/j.rcim.2016.08.001.

Implementing Speed and Separation Monitoring in Collaborative Robot Workcells

Jeremy A. Marvel and Rick Norcross

Abstract

We provide an overview and guidance for the Speed and Separation Monitoring methodology as presented in the International Organization of Standardization's technical specification 15066 on collaborative robot safety. Such functionality is provided by external, intelligent observer systems integrated into a robotic workcell. The SSM minimum protective distance function equation is discussed in detail, with consideration for the input values, implementation specifications, and performance expectations. We provide analytical analyses and test results of the current equation, discuss considerations for implementing SSM in human-occupied environments, and provide directions for technological advancements toward standardization.

Index Terms

Robot safety; speed and separation monitoring

I. Introduction

Occupational health and safety organizations rely on national and international standards to provide guidance for maintaining safe working environments. For robot workcells, these standards require the separation of man and machine by means of physical barriers with safeguarded entryways. As a consequence, the robot workcell is rigid in design and construction. Changing the layout or use of a robot workcell is expensive and time consuming. Modern manufacturing processes stress agility and flexibility, for which these classical robot workcells are ill-suited. The future of robotics in manufacturing is expected to demonstrate a shift away from the segregation of man and machine, and witness an environment in which operators and robots collaborate directly [1].

The International Organization for Standardization (ISO) Technical Specification (TS) 15066 [2] addresses this need for safety with industrial collaborative robotics. Four different collaborative scenarios are laid out. The first specifies the need and required performance for a safety-rated¹, monitored stop (i.e., a software-monitored “pause” function that prevents the robot from moving without requiring an emergency stop conforming to International Electrical Commission 60204-1 [3]). The second outlines the behaviors expected for hand-guiding a robot's motions via an analog button cell attached to the robot. The third specifies the minimum protective distance between a robot and an operator in the collaborative

¹Safety-ratings for machine components and functions denote compliance with established national and international safety standards.

workspace, below which a safety-rated, controlled stop is issued. The fourth limits the momentum of a robot such that contact with an operator will not result pain or injury.

This report focuses on the third collaborative scenario: “speed and separation monitoring” (SSM). From ISO/TS 15066, the minimum protective distance, S , at time t_0 is:

$$S(t_0) \geq \left(\int_{\tau=t_0}^{\tau=t_0+T_R+T_S} v_H(\tau) d\tau \right) + \left(\int_{\tau=t_0}^{\tau=t_0+T_R} v_R(\tau) d\tau \right) + \left(\int_{\tau=t_0+T_R}^{\tau=t_0+T_R+T_S} v_S(\tau) d\tau \right) + (C + Z_S + Z_R),$$

1

Where v_H is the “directed speed” of the operator (i.e., the rate of travel of the operator toward the robot), v_R is the directed speed of the robot in the direction of the operator, and v_S is the directed speed of the robot in the course of stopping. T_R is the time for the robot system to respond to the operator's presence, while T_S is the time to bring the robot to a safe, controlled stop. Note that T_S is a function of the robot's speed and load. The remaining terms capture measurement uncertainty, where C is an intrusion distance safety margin based on the expected human reach, Z_R is the robot position uncertainty, and Z_S is the operator position (sensor) uncertainty. For all practical safety purposes, the time, t_0 , is considered the *current time*.

Operator safety is maintained via an external observer system integrated into a robotic workcell. The observer system monitors the regions surrounding the robot, and issues commands to slow or stop the robot as humans (“operators”) approach. If the distance between the robot and a detected operator (i.e., the separation distance) is less than the value of S at time t_0 , the safety system initiates a safety-rated, controlled stop. The robot may then resume moving once the separation distance is greater than S .

Equation 1 is derived from a the equation specified in ISO 13855 [4], which specifies the minimum protective distance for stationary machine tools:

$$S = vT + C. \quad 2$$

In Equation 2, v is the approach speed of human body parts, and is analogous to the value of v_H . The value T is the system stopping performance, and is the combination of the time between sensing actuation (i.e., T_R) and the response time of the machine (i.e., T_S). For the sake of comparison, early drafts of ISO/TS 15066 computed S using the linear function,

$$S = (v_H T_R + v_H T_S) + (v_R T_R) + (B) + (C + Z_R + Z_S), \quad 3$$

where B is the distance traveled by the robot while braking. We have put parentheses around the different terms to highlight the one-to-one correlation between Equation 1 and Equation 3. However, it is important to note that Equation 3 assumes a linear relationship, and is thus only a rough approximation of the computation of Equation 1. The first term in parentheses describes the distance the human travels in the time necessary to bring the robot to a full stop from its current speed. The second term describes the distance the robot travels before it initiates the braking sequence. The third term describes the distance the robot travels while it is braking (in Equation 3, this is summarized by the single variable, B). And the fourth term describes the possible distance of intrusion into the robot's work volume as a function of operator reach and the uncertainty of sensors and robot kinematics.

The values of v_H , T_S , B , and C are traceable to existing safety standards. The values of v_H and C are given in ISO 13855, while provisions for evaluating T_S and B are given in ISO 10218-1 [5] (see Section II-A for a brief discussion of the scope of ISO 10218-1). Moreover, T_S and B result from traceable measurements that tie directly back to the robot system under test. The details and responsibility of selecting or otherwise calculating the remaining terms are left to the integrator.

In a previous study [6], we presented a mechanism for evaluating the impact of the SSM threshold on safety and task performance. During simulated and physical trials, it was noted that even a relatively simple SSM algorithm may improve the robot system's safety, but will not provide safety in all circumstances. This illustrates the necessity for validating the collision-avoidance algorithms, models, and sensor systems before relying on them for operator safety. In this study, we decompose and assess the performance of the ISO/TS 15066 SSM minimum protective distance metric, and provide guidance for the implementation and integration of SSM into collaborative robot workcells.

In the following sections, we outline the SSM algorithm, and provide possible considerations when implementing SSM in robot workcells. Discussions are based on numerical analyses and physical trials. Section II presents a brief overview of industrial robot safety and SSM in general terms. Section III evaluates the velocity measurements of the human (v_H) and the robot (v_R). Section IV discusses the factors relating to braking: the time necessary to bring the robot to a complete stop (T_S), and the distance covered by the robot while it is braking (B). Section V discusses the measurement of the reaction time of the safety system, T_R . Section VI discusses the position uncertainties of the robot (Z_R) and human (Z_S). Section VII assesses the value of the intrusion distance, C , which is considered the minimum allowable separation distance before which the robot must come to a safe, controlled stop. Section VIII discusses practical considerations of implementing the SSM algorithm on industrial robots.

II. A Background of Robot Safety

Stemming from a series of unfortunate robot-involved accidents, concerted research and standardization efforts to provide guidelines for robot safety did not begin until the mid 1980s. Over the subsequent decades, major research interest in the topic has ebbed and flowed. Generally speaking, the research has focused on the maintenance of a safe

separation distance between active machinery and humans. Providing a full account of this research is beyond the scope of this report, but several prior efforts are worth mentioning here. This section focuses on three principal areas that are directly pertinent to the SSM functionality: industrial robot safety, robot stopping functionality, and separation maintenance.

A. Industrial Robot Safety

The application of safety within different operational environments is typically outlined by technical standards that give specifications, general requirements, and test methodologies for the verification and validation of safety functionality. Although the adherence to standards is voluntary, governing bodies may adopt standards for regulatory purposes. Within the realm of industrial robotics, requirements and provision for maintaining robot safety are given in ISO 10218. Part 1 (ISO 10218-1 [5]) is intended for the manufacturers of robots, and enumerate the safety requirements for hardware and software comprising the robot and its controller. Part 2 (ISO 10218-2 [7]) describes the safety requirements for an industrial robot system, which consists of the industrial robot, the end effectors, and any ancillary devices, equipment, or sensors that support the robot's task. Part 2 also provides guidance for the safety of robots as they pertain to an integrated manufacturing system (as presented in ISO 11161 [8]), and specifies requirements for safeguarding, workcell design, operational modes, maintenance, and so on. In the United States, the national standards for robot safety are provided by the American National Standards Institute (ANSI) and Robotics Industry Association (RIA) in ANSI/RIA R15.06 [9]. In 2012, ANSI/RIA R15.06 was harmonized with the international standards, and effectively contains both parts of ISO 10218.

Traditionally, operator safety is maintained by separating the active machinery from the workforce. Such separations are ensured through physical barriers such as fences and interlocked gates, but other acceptable (as per ISO 10218-2) safeguarding equipment designed to detect the presence of people may be used. Specifically, ISO 10218-2 states that electro-sensitive protective equipment (ESPE, i.e., sensors that generate electrical safety signals) must be in compliance with International Electrotechnical Commission (IEC) standard IEC 61496-1 [10], with particular attention to compliance with IEC/TS 62046 [11] for human presence detection. IEC/TS 62046 specifies the acceptable technologies for human detection as being “active opto-electronic protective devices” (i.e., light curtains, standardized in IEC 61496-2 [12]), “active opto-electronic protective devices responsive to diffuse reflection (i.e., laser detection and ranging devices, standardized in IEC 61496-3 [13]), and pressure-sensitive mats (standardized in ISO 13856-1 [14]). Non-standardized sensors such as passive infrared protective devices (e.g., infrared thermal imaging systems) are permitted for use as part of a larger safety system, but without an associated performance standard they are neither considered reliable enough nor “suitable as the sole means of protection.” Subsequently, such non-standardized sensors are tertiary, and must be accompanied by primary, safety-rated sensing equipment. In every instance of their use, safety sensors are intended to do two things: 1) monitor specific regions for operator presence, and 2) signal the equipment to slow or stop when the operator's presence is detected.

B. Robot Stopping Functions

Stopping any powered, mechanical system relies on any of three possible procedures: 1) removing power to the drives, 2) applying brakes, and 3) actively controlling motors to counter motion. The initiation of a stopping function can be triggered via automatic mechanisms internal to the equipment (e.g., software watchdog systems, electronic monitors, or mechanical limit switches), external safeguards, or manual switches (push buttons, pull cords, or pedal-operated switches).

From IEC 60204-1 [3], stopping equipment effectively falls into two categories. “Controlled stops” reduce the velocity of moving equipment to zero while maintaining electrical control of the actuators. As a result, the motions of the system can be directed until the equipment is fully stopped. In contrast, “uncontrolled stops” remove electrical power from the actuators, effectively removing the ability to control equipment motion while the system stops. Note that the classification of uncontrolled stops makes reference only to the removal of electrical power to the actuators, and does not infer an operational state of any other stopping devices (i.e., mechanical or hydraulic brakes, or antagonist actuators). IEC 60204-1 also specifies three categories of stop functions based on the removal of power, application of brakes, and control of the equipment during the braking function. These three categories are summarized in Table 1. In a stop category 0 (STOP0), the equipment is stopped immediately by removing power to the actuators and applying the brakes. This results in an “uncontrolled stop,” in that the purpose is to stop the equipment as quickly as possible, and does not ensure that the motion of actuators follow a prescribed path. In a stop category 1 (STOP1), the equipment undergoes a controlled stop (i.e., power to the actuators is maintained) until the equipment comes to a complete stop, at which point the power is removed and the brakes applied. And in a stop category 2 (STOP2), the equipment comes to a controlled stop, but the brakes are not applied nor is the power removed.

C. SSM in Literature

Outside of the standards landscape, the research in the field of robot safety has produced a myriad of solutions for collision avoidance and maintaining a safe operational distances between active robot systems and miscellaneous objects inside the robots' work volumes. Most of the literature is focused on identifying and avoiding potential collisions. In [15] Marvel and Bostelman provide a comprehensive review of metrics used to model and evaluate collisions in multiple domains, including robotics and transportation. Many of the metrics identified for articulated and mobile robot platforms were designed specifically for safety with the intention of maintaining a safe separation distance between the hardware and any possible objects or people in the environment. Most safety application attempt to generalize robot safety for SSM applications according to varying metrics including separation distances ([6, 16, 17]), velocity and configuration ([6, 18]), inertia ([19, 20]), or potential severity of impact ([6, 21-24]). Each of these metrics adds to the collective understanding of a robot's impact on its environment, and provides a sound basis for many algorithms that maintain safe operational distances using active collision avoidance. While most of these approaches factor only a singular representative coordinate of the robot (e.g., the tool flange or center of the wheels), some do take into account the entire kinematic chain when enforcing separation. For example, Zanchettin *et al* [25] propose a new metric

specifically designed for maintaining SSM for the entire kinematic chain into a robot's control algorithms.

All SSM applications require some level of workspace monitoring for functionality. While most historical collision avoidance approaches have focused on the control algorithms, recent years have witnessed a boom in research on workspace supervision to accompany the growth in the collaborative robot market. Advances in on-robot sensing (e.g., [26]) and camera-based observer systems (e.g., [27]) have demonstrated considerable promise in alternative sensing technologies for robot safety, and it is likely that such technologies will be critical to ensuring robot safety in unstructured, collaborative environments.

III. Operator and Robot Speed

Validating the speed estimates of the operator and robot is a difficult task. In Equation 3, the speed of the operator, v_H , is assumed to be a worst-case maximum of 1600 mm/s based on the specifications in ISO 13855. ISO 13855 claims a maximum speed of 2000 mm/s for an operator when calculating the minimum protective distance for stationary machinery with the option for using 1600 mm/s in the case that the separation distance is greater than 500 mm. The new application in Equation 1, however, allows for the provision that v_H may be measured directly. If such measurements are not available, v_H is typically assumed to be constant and in the direction of the robot, though more realistic acceleration profiles may be adopted.

As a point of comparison, academic biomechanical research places average normal walking speeds to be between 1250 mm/s and 1510 mm/s [28-30] on level surfaces, average fast walking speeds between 1710 mm/s and 1840 mm/s on level surfaces, and average walking speeds between 390 mm/s and 810 mm/s on incline surfaces [31]. However, significant instantaneous changes in speed and direction are possible given special circumstances, e.g., tripping. Such changes occur without warning and are not accounted for within the assumptions of ISO 13855.

A significant consideration of the current SSM algorithm is that the actual direction of travel is not factored into the equation. Instead, the ISO/TS 15066 states that the *directed speeds* of the robot and human are used, which does not actually factor travel direction. For instance, a robot moving directly toward a human at 0.1 m/s has the same directed speed as a robot moving at 1.0 m/s but approaches a human moving parallel to it at 0.1 m/s. Note, also, that speed is a measurement of a vector's magnitude, and is therefore always greater than or equal to 0. From a safety perspective, it should be assumed that the operator's motions will follow a worst-case scenario in that the operator will begin moving in the direction of the robot at any time. Therefore, it is possible that even a stationary operator outside of the path of the robot can trigger a safety-rated controlled stop even if the robot is otherwise in no danger of colliding. As the speed of the robot increases, the minimum protective distance increases proportionally, and the resulting bounding circle may envelop the operator before, at, or even after the point of closest approach, as illustrated in Fig. 1. In this example, the robot is actually moving away from the stationary operator, yet a protective stop will nevertheless be issued due to the increasing value of S .

Reporting the robot's speed, v_R , is similarly problematic, as ISO/TS 15066 does not specify the mechanisms by which the time-varying values are measured or conveyed. Even though the Specification states to use the directed speed, such measures are not always available in real time. As such, estimations may be used regarding the robot's and human's respective travel paths. Or it may simply be easier to use the speed of the robot irrespective of its travel direction (e.g., if the robot's trajectory is not known *a priori*). Such speed-reporting options include using the rated maximum speed of the robot, using the commanded velocity, relying on the robot to report its speed at a given time step (e.g., as a percentage of the rated path velocity [32]), or calculating an instantaneous speed based on time-stamped position reports from the robot² or external tracking systems (both of which are subject to reporting delays depending on algorithms used, variable access rights, and communication channel capabilities). Ultimately, there is no agreed-upon guidance regarding the computation of the robot's velocity, and gaining external access to planned system dynamics (i.e., trajectories and speeds) is not common. Thus the onus of choosing which method to utilize is left to the integrator.

The issue may be further complicated by the lack of support for tracking parts attached to or held by the robot. Most industrial robots report the pose of the robot only in terms of either joint configurations or in terms of the Cartesian pose of the attached tooling. Moreover, the robot controller has no knowledge of tools or workpieces held by the robot unless a task-critical tooling offset is provided; still, only that specific point is tracked. As a result, it is possible for a robot's "elbow" to strike something in its environment given that only the tool flange is being tracked. The moment arm of held pieces (e.g., if joint 5 of a 6-axis robot arm moves, the poses of joint 6 and the held workpiece move in Cartesian space as shown in Fig. 2) must be included in the observer system's considerations for SSM. Similarly, not every industrial robot system will report the Cartesian coordinates of all joints or links. In such cases, the robot may be actively moving even though the speed at the tool center point (TCP) is zero (e.g., in a null-space move). Again, the observer system may be required to either directly track the robot's kinematics or infer them from reported joint configurations. Otherwise, providing pose information of the centers of all joints would provide the minimum set of information necessary to determine if any part of the robot, itself, is in danger of colliding with a human or object in the workzone.

Depending on which method is chosen, the errors in reporting v_R may take different forms and magnitudes. The illustration in Fig. 3 shows that even reliable feedback from the robot might be subject to measurement noise. Here, the observer system calculates the instantaneous velocity using the distances between the robot's time-stamped positions, and then smoothed using a low-pass filter. This trial consists of the robot's TCP moving along the Y axis in Cartesian world coordinate space, and then returning to its original position. In the case of the illustrated velocity, the uncertainty in the velocity is a function of the damping of the motion controller, the length of sampling for the low-pass filter, and the data

²It has been observed in some robot systems that the mechanisms by which the robot tracks velocities for control differ from the mechanisms by which velocities are reported externally. Independent (i.e., external) verification of velocities may quantify discrepancies in reported versus actual values.

transmission delay and sampling resolution. The commanded velocity of 1000 mm/s is also shown, demonstrating the discrepancy between the actual and commanded velocities

Externally tracking robots or operators is prone to additional fitting noise (i.e., measurement uncertainty incurred by forcing position estimates based on assumptions of the nature of the measured data) based on the sensors and algorithms utilized, and may be no more reliable. For example, Fig. 4 illustrates the reported path of an operator entering and exiting the workspace at the same location using a laser-based human tracking system. The human's position is estimated based on the centroid of detected objects within the working environment, which varies based on the motions of the legs, shifting of clothing, and sensor noise and missing data. The track measurement errors shown in Fig. 5 are the differences between the individual readings and a fitted spline through the data.

The measurements from external tracking systems may also impose errors given the magnitude of measurement uncertainty. Depending on the rate and size of position sampling window on the position measurement uncertainty (see Section VI), and on the methods used to smooth and accommodate measurement uncertainty, position noise of even a few millimeters may adversely affect the velocity estimate of hundreds of millimeters per second. As an illustrated example, consider a top-down tracking of a 6 degree of freedom (DOF) robot, Fig. 6-A, where the positions of only the robot's base (R_{base}) and TCP (R_{TCP}) are monitored. Assuming a steady speed profile with no acceleration, the base joint is rotated at a constant rate, ω_{base} , while the other joints remain at their current values. The resulting motion consists of the TCP orbiting around the base at a fixed distance, $l = \|R_{TCP} - R_{base}\|$ (Fig. 6-B). Assume that the position of the base and TCP are evaluated at a constant rate, r , such that, given the steady rotational speed, the change in angle is constant at θ .

Based only on this information, two simple algorithmic approximations of the speed of the TCP can be calculated. The first is an instantaneous approximation in which the speed is derived based on discrete pose samples taken over time. Such an approximation is analogous to an outside observer system tracking the TCP, and can be written as

$$v_{TCP} = \left(\frac{r}{N} \sum_{i=0}^N x_i \right) + \sigma, \quad 4$$

where N is the number of samples gathered, $x_i = \sqrt{2l^2(1 - \cos\theta_\Delta)}$, $1 \leq i \leq N$, and σ is the measurement uncertainty. The measurement uncertainty is assumed to have a positive bias with a known distribution that is not necessarily normal. Such a distribution may shift in terms of form and bias as a function of time, intrinsic and extrinsic operational parameters, and combination of sensors used (e.g., in terms of sensor fusion for measuring pose under varying conditions). In cases where the mean distribution for all samples is not known, vendor-provided expected uncertainty models may be substituted.

The second algorithm is an analytical approximation of the TCP speed applying feed-forward control laws known *a priori*, and is based only on the change in angle and the Euclidean distance between the TCP and the robot base. Given the assumption of a constant

ω_{base} , this method is analogous to the ground truth for the robot system, and can be computed as

$$v_{TCP} = \theta_{\Delta} l r. \quad 5$$

Depending on the values of the parameters previously discussed, the differences between the instantaneous and analytical approximations may become more or less pronounced.

Assuming a Gaussian measurement uncertainty (noise) profile with mean $\mu = 0$ and standard deviation σ , the values given in Table 2 indicate potential calculated values for the speed of the base (v_{base}) and TCP (v_{TCP}). The range update frequency, r , is given in a range of 1Hz to 100Hz based on empirical results using industrial robot systems in our facilities, on which actual update frequencies range between 8Hz and 125Hz. Both the base rotation speed and the sampling rate have a direct effect on the differences between the instantaneous (eq. 4) and analytical (eq. 5) approximations. Similarly, while the analytical approximations of the base and TCP speeds are unaffected by measurement uncertainty, the instantaneous speed approximations accumulate considerable measurement errors as a function of uncertainty. The effects of these uncertainties may be limited by increasing the number of samples per speed calculation, but risks biasing measurements toward an artificial homeostasis and any changes in velocity take longer to detect. Further accommodations may be made by means of the robot position uncertainty value, Z_R , as discussed in Section VI. Note, however, that the numbers in Table 2 are based on simulations with perfect information, and that the differences between instantaneous and analytical velocity measurements for physical trials are expected to be larger.

Improving the reliability of the velocity measurements is a matter of establishing conventions. For instance, regarding the assumed value of the operator velocity, v_H , even though the 1600 mm/s is an accurate assessment of normal walking speeds, it may be more prudent to instead use the worst-case assumption of 2000 mm/s from ISO 13855 to accommodate instantaneous rapid motions and the uncertainties associated with operator detection. Although this worst-case assumption is not necessarily accurate at all times, it is clearly safer to assume than the special case average currently chosen, particularly given that it is not clear that the special circumstances involved with selecting 1600 mm/s over 2000 mm/s for stationary machinery is applicable with robot systems. ISO 13855 also explicitly states that the approach of an operator towards the machine will be at a walking speed, and that a faster rate may be used if a risk assessment deems it necessary.

Similarly, with regard to the evaluation of robot velocity, it is advised that one adopts a common convention for measuring and reporting the velocity. One potential solution is to have the robot report the time-stamped Cartesian space velocity of all joints and the tool. The reported velocities may be independently validated using position accuracy validation per an established standard (e.g., [32]). Given that robots with seven or more degrees of freedom are becoming more common, it is recommended that independent minimum protective distances be maintained for all joints. This will provide safe SSM functionality even in special cases such as null space motions (i.e., when the speed of the TCP is 0 mm/s).

Further, given that operator position must be computed using the safety system, it may be prudent to measure changes in the relative distance between the operator and the robot for the purpose of avoiding unnecessary protected stops of the robot when it is moving away from the operator.

IV. Braking Factors

Braking factors for any mechanical system are inherently non-linear, particularly given the variability in stopping functionality. From Section II, recall that IEC 60204-1 [3] outlines three stopping categories based on the removal of power, application of brakes, and control of the equipment during the braking function. For industrial robots, the different stopping reactions of the robot system are triggered in response to different monitored events (trigger) or as a reaction to operator actions. In some cases, certain triggers may result in different stopping behaviors depending on the operational mode or vendor preferences. For instance, the stopping behaviors for different two collaborative, robot arms in the National Institute of Standards and Technology (NIST) Collaborative Robotics laboratory are provided in Table 3 and Table 4. These robots were manufactured by different vendors, and have different capabilities. The first robot (Table 3) is a 7DOF open-chain manipulator that provides two different operational modes: “teach” (10% reduced velocity mode for initial programming and program verification; an enabling device is required to power the robot and run the program), and “automatic” (100% commanded velocity; the operator must select a program, but an enabling device is not required to run the program). The second robot (Table 4) is a 6DOF open-chain manipulator that does not have a “teach” mode.

As mentioned in Section I, the braking time and distance, T_S and B , respectively, result from traceable measurements per Annex B of ISO 10218-1. Both are reported as functions of the robot's initial speed, v_R , and payload (a percentage of the maximum rated load)—a fact that has significant implications discussed in Section VIII—as illustrated in Fig. 7. Many instantiations effectively treat both as constants, though it is recommended that the values for both T_S and B be periodically reevaluated to accurately accommodate the degradation of brake components. While most manufacturers evaluate brake wear during routine maintenance and inspection cycles, not every integrator or end-user subjects robots to regular maintenance. Simple, effective, and inexpensive test methods for assessing robot joint brake wear would increase the likelihood of routine evaluation, thus greatly improving robot braking reliability.

Moreover, the braking distance and time may also be influenced by many intrinsic and extrinsic factors, and may thus further accentuate the nonlinear nature of braking profiles. Such impacting factors include, but are not limited to:

1. Changes in robot configuration (e.g., the effects of gravity due to the mounting of robots may impact motion dynamics)
2. Robot pose (i.e., robots experience different braking characteristics when fully extended than they do when the TCP is closer to the base)

3. Path trajectory (specifically, not all brakes on a given robot behave identically, so the time to stop the robot may vary depending on which joints are moving and at what speeds)
4. Break wear
5. Task-related considerations (e.g., limitations imposed by tooling or application restrictions) that require changes in braking profiles to prevent accidental release or damage to work objects
6. Shifts in the center of mass of the robot's payload

Consequently, it bears consideration that these braking factors should be reassessed to include intrinsic robot configuration parameters to enable more accurate system modeling.

V. Reaction Times

When a monitoring system reports the position of a human operator, it is vital to have an estimate as to how much time has actually passed since the human was at the reported location. If there is a sufficiently long lag between the times when an event occurs and when the event is reported, the delay presents a significant safety risk regardless of the accuracy and precision of the reported data. For instance, if a human walking at 2000 mm/s is reported half a second late, the separation distance between the robot and the human may be 1000 mm shorter than indicated. This delay directly impacts the measurement uncertainty of the human detection/tracking system and affects the robot's reaction time, T_R . To compensate for this larger uncertainty, the separation buffer may be increased. Alternatively, if the effects on measurement uncertainty can be characterized, the effects on the reaction time may be negated by means of predictive filtering (e.g., in Vasquez, Fraichard, and Laugier [33], the authors produce statistical models to predict human and vehicle motions).

Currently, there is no standardized methodology for measuring the reaction time of a human-detecting sensor. Rather, it is treated as a constant value that the integrator must provide. Given the nonlinearity of the human tracking problem, accurately predicting the reaction time of a human-detection system is difficult. Contact-based limit sensors such as pressure-sensitive mats provide rapid notice that a human is present, but suffer from poor resolution, offer little to no tracking capacity, and are easily bypassed or confused. Non-contact human-detection systems such as camera- and laser-based solutions offer higher resolution and the potential for better tracking capabilities, but performance is affected by factors including lighting, processing time, capture frequency, occlusions, and sensor and scene uncertainty.

To measure the reaction time, it is necessary to give a controlled signal to the safety system and then measure the latency of the response. For camera-based systems the stimulus may be a sequence of recorded images showing a person entering a protected space, while laser-based systems may require the use of a controlled physical avatar that can be moved into the work zone using repeatable trajectories. Ultimately, the tester should be able to provide both a ground truth and control triggering of the event. Embedded safety systems that perform the acquisition and processing of signal data based on internal or multimodal triggers in a closed manner are difficult to evaluate as these triggers are not always controllable. For such

situations, a second system (e.g., a pressure-sensitive mat) that monitors the boundary of the protected zone will be required to start a timer that is then stopped when the signal from the safety system is received. Multiple repetitions³ are required for statistical significance, as are the reaction time averages, standard deviations, and maximum reported values for assessing the expected and worst-case performances.

For example, in Szabo, Shackleford, Norcross and Marvel [34], the authors introduce a testbed setup in which the various parameters could be derived and the SSM functionality verified. One of the tests presented also measured the reaction time of the robot by setting up a linear motion test in which the robot interrupted a laser beam connected to a safety relay. Stop signals were propagated to the robot's controller via Ethernet links. The process consisted of the following steps:

1. Program the robot to move along a linear path
2. Mount a reflector on the robot
3. Mount the sensor on a stand inside the robot's work volume
4. Connect to the sensor output to the robot's digital inputs used for SSM
5. Build a robot trajectory that passes the sensor
6. Run the robot trajectory and record the robot stop position
7. Repeat step 6 for multiple repetitions at various speeds
8. Determine the robot's positions along the trajectory
9. Evaluate the SSM parameters

The position of the robot after a sensing event is equal to the distance traveled while the robot is responding to the event (T_R) plus the distance traveled while the robot is stopping. Given that the location of the sensing event was constant (albeit unknown), the location of the final, stopped position of the robot, $P_{S,i}$ for a given percentage of maximum velocity, i %, was defined as:

$$P_{S,i} = \left(v_{R,100} T_R + \frac{v_{R,100}^2}{2a} \right) - \left(v_{R,100} T_R (1 - i) \right) - \left(\frac{v_{R,100}^2}{2a} (1 - i)^2 \right), \quad 6$$

where $v_{R,100}$ is the velocity of the robot at 100% maximum speed, and a is the worst-case negative acceleration. The difference between the stopping positions $P_{S,100}$ and $P_{S,i}$ is equivalent to the stopping distance. For multiple velocities, attempts to define the relationships between the stopping positions and speed ratios were quadratic in nature. When solving for this quadratic, the three coefficients reflect functions of the robot's performance parameters:

³The actual number of repetitions will be dependent on the user's accepted tolerances as well as the sensor and sampling uncertainties.

$$c_0 = v_{R,100} T_R + \frac{v_{R,100}^2}{2a} + P_{S,100}, \quad 7$$

$$c_1 = v_{R,100} T_R, \quad 8$$

$$c_2 = \frac{v_{R,100}^2}{2a}. \quad 9$$

Through repeated trials run in the NIST laboratory, the value of T_R for a rail-mounted, 6DOF robot manipulator (Fig. 8) was evaluated to be 0.113 s (+0.019 s at a 0.999 confidence interval).

Despite the inclination to treat the reaction time as a constant, this value should be periodically reassessed to account for wear and calibration degradation of the system and sensors. With traditional industrial robot systems, such wear and degradation is expected to be minimal, and performance should remain deterministic. However, with alternative collaborative robot designs entering the market at an increased frequency, it has been observed that some system and performance degradation (e.g., as a result of non-serviceable sensor drift) may occur. In some cases, robots designed for intrinsically safe power and force limiting (PFL, the limiting of the transfer of pressure and forces to human operators to prevent injury) have inherently nondeterministic behavior given the natural compliance built into the robots' joints.

VI. Measured Uncertainties

The values of Z_R and Z_S relate position uncertainties of the robot and operator, respectively, to the tolerable noise in evaluating the separation distance. Their inclusion is intuitive as large uncertainties in position measurement may allow robot-operator collisions to occur even though the sensor systems believe the two entities are still physically separate. One could argue, however, that the inclusion of the Z_S term is somewhat redundant given the resolution considerations for the workspace-monitoring sensors in the selection of the intrusion distance, C (see Section VII). From a metrological standpoint, however, additional information is needed in the evaluation of these measurement uncertainties.

The expected robot measurement uncertainty may be computed by evaluating the variances of the errors per ISO 9283 [32]. The average variance could be used to satisfy the definition of Z_R , though a worst-case maximum value should be more appropriate for safety purposes. Alternatively, if using an external sensor system to evaluate the position of the robot, measurements of sensor uncertainty would define the value of Z_R . Such applied measurement uncertainties may use vendor-supplied values, when available. Alternatively, these values may be acquired (or a validation of the vendor-supplied values may be

accomplished) by means of standards including the ASTM E2919-14 standard test method for evaluating 6DOF static object pose measurements [35].

It should be noted that there is a significant limitation regarding the evaluation of Z_S : no standard currently exists for human localization or pose estimation measurements. Standards do exist for evaluating the measurement uncertainty of different sensor classes (e.g., active optoelectronic protective devices [12], and active optoelectronic protective devices responsive to diffuse reflection [13]), and it is feasible that such values may be substituted for human position uncertainty until human-positioning standards emerge. However, the measurement uncertainty of the human position is likely larger than the base sensor uncertainty given the variability of a human's surface. A potential solution to reporting Z_S is to use the standard deviation computed during the evaluation of the reaction time and then multiply that value by the speed of the human operator, v_H .

VII. Safety Margins

ISO/TS 15066 defers to ISO 13855 for the valuation of an intrusion distance safety margin for separation, C . ISO 13855 defines necessary separation distances for stationary machinery based on expected approach speeds and biomechanical limitations of human operators. In this sense, stationary machinery is identified as mechanical tools that do not have a set work volume defined by the profile of the machine, e.g., the cabinet of a band saw. The value of C is chosen based on the decision rules outlined in Table 5. Here, d is the detection capability of the presence-identifying sensor in mm, and H is the height of the detection zone. The detection capability, as described by ISO 13855, is a functional limit of sensing that causes the protective equipment to signal the presence of an object. For example, a light curtain's detection capability is measured as the distance between light beams, and a LADAR's detection capability is measured as the disparity of the laser beams at a function of safeguarding distance. ISO 13855 further specifies the lowest value,

$$H=15(d-50)\text{mm}. \quad 10$$

The safety buffer is determined by the selection sensors and by which safety is provided. Limiting the direction of approach, relying on the placement of position-identifying sensors (relative to the floor), or mandating specific safety controlling devices designed for machinery will result in different values for C . For example, if the integrator chooses to ensure safety by limiting the direction of approach to be normal to the machinery and protects the region by multiple laser-based sensors, the value of C must be at least 850 mm. If the operator also requires a two-handed control device, the value of C may be reduced to 250 mm.

As applied to ISO/TS 15066, it is expected that when a protective stop is issued, the robot will come to a halt when the operator is no closer than C mm away. However, given the configuration of the robot system, the applicability of C from ISO 13855 to the robot system is not certain. Namely, because ISO 13855 applies specifically to stationary machinery, it is unclear whether its expectations of operator safety are also applicable to dynamically

reconfigurable platforms such as industrial robots. For example, does the mounting location of a sensor for a robot cell really hold the same level of importance as it does for stationary machinery?

One may argue that the considerations given for C from ISO 13855 are not applicable to the SSM equation as defined in ISO/TS 15066. Specifically, because of the dynamic and reconfigurable nature of industrial robots, the direction of approach is ambiguous, and distinctions between “normal” and “parallel/angled” are inconsistent. Further, given the varied nature of sensor types considered for supervising large robot cells, next-generation safety sensors such as ceiling-mounted cameras may not fit concisely within the bounds of the mounting limitations for stationary machinery. Finally, per the scope of ISO 13855, the controlling devices most frequently used for industrial robots (i.e., two-handed teach pendants) are not applicable for consideration when determining the value of C .

Depending on the application, it may therefore be prudent to evaluate the value of C based on a task-based risk assessment that includes factors such as tooling, robot reach, sensor resolution and update rate, and potential regions of operator approach. One possible solution is to assume the worst-case approach angle from ISO 13855 and then add an extra buffer based on the precision of the monitoring sensor system, for example:

$$c = 850 + d \text{ mm.} \quad 11$$

VIII. Pragmatic Considerations

When implementing SSM, there remain a few matters of practicality regarding the expected technological capabilities of the robot system. Details pertaining to sensor and system requirements, error testing and handling, and timing restrictions may adversely affect the safety and performance of SSM-enabled robot systems.

A. Technology Requirements

It is worth noting that the current standards landscape provide no accounting for the technological requirements for sensing or reporting capabilities of an observer system. What requirements, for instance, should be expected for human sensing precision, uncertainty, or delay? For safety-rated sensors and devices, ISO 13849-1 [36] outlines performance specifications provide “performance levels” based on statistical probabilities of failure. Further, there are few considerations for the restrictions of individual robot systems. Specifically, due to technological limitations, a given robot system may be disqualified for collaborative environment utilization. Example disqualifying conditions include the robot being incapable of reporting velocities in real time, a lack of support for reporting velocities of all joints, unreliable delays in stopping times, or the inability to receive and process external signals to come to a controlled stop.

As an illustrative example, a robot testbed was constructed at NIST that evaluated a SSM algorithm implementation based on tracked LADAR positions of the robot and a computer-

controlled mobile obstacle (Fig. 9). The robot was under-slung on an elevated rail to give it a significantly larger work volume (Fig. 8), which was safeguarded using safety-rated laser sensors (per the requirements of IEC 61496-3 [13]) and an interlocked gate system. The robot configuration is illustrative of those in modern manufacturing, but the robot itself was built using design specifications predating the considerations of ISO/TS 15066. As such, it has no native SSM interface.

The custom configuration experiences substantial robot velocity uncertainties stemming from short sampling intervals (refer to the discussion in Section III regarding this effect), sensor noise, and object correlation errors. A delay of about 220 ms was measured between when the robot was at a specific location and when it reported it was at that same position. Further, it was observed that, depending on the method used to stop the robot safely, the time necessary to bring the robot velocity to 1 mm/s varied considerably. From these observations, one might consider the robot to be incompatible with SSM-based human-robot collaborative safety, with significant uncertainties being unavoidable given the scale and construction of the workcell. However, it could also be demonstrated that, accommodating the system's uncertainties, safe behavior congruent with the SSM methodology is still possible, but requires additional considerations not explicitly stated in the SSM equation.

Fidelity of distance updates and tightness of control loop also drive the robot's performance as a function of S . For instance, assuming a tight control loop with regard to updates from the operator locator, the separation distance and commanded robot velocity are directly influenced by the rate at which updates are presented to the system. As an example, let the robot have a constant acceleration rate of 10.0 m/s^2 . Assume that a human operator maintains a constant separation distance of 1.6 m between himself and the closest point of potential contact, the robot's TCP (i.e., the operator moves at the same speed and in the same direction as the robot's trajectory). As the robot's velocity increases linearly, the minimum protective distance increases proportionately, as seen in Fig. 10. Results for both the current SSM equation (Eq. 1) and the original SSM equation (Eq. 3) are provided in Fig. 10 to illustrate differences in the equations' impacts on v_H and S .

As the update rate from the operator locator decreases, the robot's velocity fluctuates more when the minimum protective distance is violated, the average velocity of the robot decreases, and the amount of time the robot spends recovering increases. At steady state with a 1000 Hz update rate, the robot has an average velocity of 1.6 m/s, whereas the average robot velocities for update rates of 100 Hz and 30 Hz are closer to 1.5 m/s. It is therefore preferable to have as high an update frequency as possible to optimize performance. Given the capabilities of even aged computer technology, rapid responses to a minimum update rate of 100 Hz is not unreasonable, particularly when one considers that the separation distance of an operator and a robot, each moving at 1.0 m/s directly toward one another, can close by over 0.06 m between updates presented at 33 Hz.

B. Verification and Validation

Any given safety system must be periodically evaluated to ensure continued functionality. A critical aspect of this is the need for a mechanism for independently and safely verifying that the observer system is performing per the specifications of the SSM equation. Performing

this verification in a repeatable manner is difficult. This is largely due to the variability of measuring humans and the uncertainty of sensor systems, and determining if and when the robot actually initiated a protective stop.

Qualitatively, one potential verification method would be to simply identify that the robot came to a protected stop before a human surrogate (ideal a biomimetic artifact) came into contact with the robot. Such verification would provide a pass/fail report, and would require an external measurement system with known accuracy and update frequency time synchronized with the robot's observer system.

Quantitatively, there is no methodology currently available that captures the measurement or performance uncertainty of the robot system as a whole. In Section VI, we attempted to address the system's measurement uncertainty in terms of the position uncertainties of the robot (Z_R) and the human operator (Z_S). However, neither of these values account for the variability of the performances of the robot's motors or brakes, the sampling or sensor uncertainties of the monitoring systems as a function of environmental factors, or the locations of payloads due to grasp stability or gripper technology. As such, an integrator may be required to make certain assumptions regarding the rigidity, dependability, and predictability of the robot system.

It is worth noting that the SSM equation and assumptions have not been validated for all possible robot configurations to be safe for human collaboration. As an output of the hazard mitigation process following a risk assessment, the intrusion distance, C , may require additional adjustments to ensure safety. For example, recall from Section VII the discussion of C , and that the current definition may not be sufficient to capture the complexity of reconfigurable systems such as industrial robots. An alternative calculation of C , separate from that prescribed in ISO 13855, may be required such that the value is commensurate with the robot's configuration and payload. Similarly, because ISO 13855 was intended for stationary machine tools, it does not take into consideration any potential hazards posed by the tooling or payload, themselves. Additional buffers may be added to accommodate the maximum possible deflection from the robot's chassis any such tooling may have.

IX. Discussion

The inclusion and discussion of the SSM algorithm in ISO/TS 15066 is an important component of ensuring operator safety around robots. Considerable time and effort have been put into describing how to compute the minimum protective distance threshold. Further, because the applicability of many terms is not clearly defined, there remains some question regarding how to actually go about choosing or finding certain values. As such, although the computation of the minimum protective distance is laid out in detail, the actual implementation of SSM is still largely left to the discretion of the integrator.

The issues and considerations raised in this report have all been discussed with ISO Technical Committee 184, Subcommittee 2, Working Group 3, which is responsible for the development of ISO/TS 15066. The pursuant discussions have been fruitful, and it is expected that further dialogues will continue to improve the SSM requirements. This report

is not intended to draw scrutiny to the SSM methodology, but to provide insight into considerations when implementing SSM per the technical specification.

For ease of implementation and verification, one may be left asking, why mandate that the integrator follow a specific equation for the minimum protective distance at all? It may be more succinctly—and just as accurately—stated that an active robot with SSM functionality should be no closer than C mm from a human (where C is defined by a new method separate from ISO 13855), and then provide a test method to validate that such a requirement is followed. The effect is arguably the same (i.e., the robot stops before making contact with the human per the integrator's own recognizance), but the integrator is provided with the tools necessary to evaluate the efficacy of his implementation. The alternative requires more stringent accounting for the measurement uncertainties discussed in this report. This necessitates more effort on the part of the technical specification and of the integrator, but provides the traceability otherwise lacking.

A. Enabling Collaborative Robots in Manufacturing

Currently, robots marketed as being “collaborative” for industrial applications are based on the assumption that such robots will be working independently in the manufacturing environment but in close proximity to human operators. Such robots are designed with PFL in mind, and are typically lighter, slower, and have lower lift capacities and shorter reach than their traditional, non-collaborative counterparts. Integration of such robots must be intentional, with the application task being suitable for robots of their scale.

Converting non-collaborative robots to act collaboratively excludes the PFL functionality, and thus requires SSM integration to maintain safe separation distances between humans and robots. Traditional sensing technologies as described in Section II-A provide a solid foundation for monitoring fenceless workcells, but only in structured environments. Advances in safety-rated camera or on-robot sensing technologies will be critical for enabling SSM in unstructured industrial environments. Human-specific identification and localization, in particular, will be required if the robots are expected to make physical contact with workpieces with pose and identity uncertainty.

B. Open and Ongoing Research

Provisions for ensuring robot safety are far from absolute, and robot-related injuries and fatalities still occur despite decades of advances in sensing and control. Currently, the SSM methodology, while clearly derived from established safety protocols, still lacks verification and validation methodologies to ensure functionality and reliability. Removing the fences from robot workcells is possible given the current state-of-the-art, but additional care must be taken, and safeguards must be carefully selected and integrated into the robot's controller to reduce risk. Again, advancing new monitoring and sensing technologies to the state of being safety-rated is essential, but is likely to be an area of active research for years to come. Moreover, turnkey solutions for safe, collaborative robots will require significant advances in built-in safety systems, application-level learning, and situational awareness to accommodate unstructured manufacturing environments to differentiate humans from the parts and products with which the robots are working.

References

1. Computing Community Consortium. A roadmap for US robotics: From Internet to robotics. 2013. <http://www.us-robotics.us/reports/CCCReport.pdf>
2. International Organization for Standardization. ISO/TS 15066:2016 — Robots and robotic devices — Collaborative robots. 2016
3. International Electrotechnical Commission. IEC 60204-1 — Safety of machinery — Electrical equipment of machines — Part 1: General requirements. 2009
4. International Organization for Standardization. ISO 13855 — Safety of machinery — Positioning of safeguards with respect to the approach speeds of parts of the human body. 2010
5. International Organization for Standardization. ISO 10218-1 — Robots and robotic devices — Safety requirements — Part 1: Robots. 2011
6. Marvel JA. Performance metrics of speed and separation monitoring in shared workspaces. IEEE Transactions on Automation Science and Engineering. 2013; 10(2):405–414.
7. International Organization for Standardization. ISO 10218-2 — Robots and robotic devices — Safety requirements — Part 2: Industrial robot systems and integration. 2011
8. International Organization for Standardization. ISO 11161 — Safety of machinery — Integrated manufacturing systems — Basic requirements. 2007
9. American National Standards Institute; Robotics Industries Association. ANSI/RIA R15.06 - Industrial robots and robot systems - Safety requirements. 2012
10. International Electrotechnical Commission. IEC 61496-1 — Safety of machinery — Electro-sensitive protective equipment — Part 1: General requirements and tests. 2012
11. International Electrotechnical Commission. IEC/TS 62046 – Safety of machinery – Application of protective equipment to detect the presence of persons. 2008
12. International Electrotechnical Commission. IEC 61496-2 — Safety of machinery — Electro-sensitive protective equipment — Part 2: Particular requirements for equipment using active opto-electronic protective devices (AOPDs). 2013
13. International Electrotechnical Commission. IEC 61496-3 — Safety of machinery — Electro-sensitive protective equipment — Part 3: Particular requirements for equipment for active opto-electronic protective devices responsive to diffuse reflection (AOPDDR). 2008
14. International Organization for Standardization. ISO 13856-1 Safety of machinery — Pressure-sensitive protective devices – Part 1: General principles for design and testing of pressure-sensitive mats and pressure-sensitive floors. 2013
15. Marvel JA, Bostelman R. A Cross-Domain Survey of Metrics for Modeling and Evaluating Collisions. International Journal of Advanced Robotic Systems. 2014; 11(142):1–15.
16. Chen M, Zalzal AMS. A genetic approach to motion planning of redundant mobile manipulator systems considering safety and configuration. Journal of Robotic Systems. 1997; 14(7):529–544.
17. Trautman, P.; Krause, A. Unfreezing the robot: Navigation in dense, interacting crowds. IEEE/RSJ International Conference on Intelligent Robots and Systems; 2010; Taipei, Taiwan. p. 797-803.
18. Lacevic, B.; Rocco, P. Kinetostatic danger field - a novel safety assessment for human-robot interaction. IEEE/RSJ International Conference on Intelligent Robots and Systems; 2010; Taipei, Taiwan. p. 2169-2174.
19. Kuli D, Croft EA. Safe planning for human-robot interaction. Journal of Robotic Systems. 2004; 22(7):383–396.
20. Zinn, MR. A new actuation approach for human friendly robotic manipulation. Stanford University; 2005.
21. Nokata M, Ikuta K, Ishii H. Safety-optimizing method of human-care robot design and control. IEEE International Conference on Robotics and Automation. 2002:1991–1996.
22. Bicchi, A.; Peshkin, MA.; Colgate, JE. Springer Handbook of Robotics. Springer; Berlin Heidelberg; 2008. Safety for Physical Human-Robot Interaction; p. 1335-1348.
23. Haddadin S, Albu-Schäffer A, Hirzinger G. Requirements for Safe Robots: Measurements, Analysis and New Insights. International Journal of Robotics Research. 2009; 28(11-12):1507–1527.

24. Althoff D, et al. Safety assessment of robot trajectories for navigation in uncertain and dynamic environments. *Autonomous Robots*. 2012; 32(3):285–302.
25. Zanchettin AM, et al. Safety in human-robot collaborative manufacturing environments: Metrics and control. *IEEE Transactions on Automation Science and Engineering*. 2016; 13(2):882–892.
26. Avanzini GB, et al. Safety control of industrial robots based on a distributed distance sensor. *IEEE Transactions on Control Systems Technology*. 2014; 22(6):2127–2140.
27. Vicentini, F., et al. 41st International Symposium on Robotics. Munich, Germany: 2014. Dynamic safety in collaborative robot workspaces through a network of devices fulfilling functional safety requirements; p. 1-7.
28. Knoblauch RL, Pietrucha MT, Nitzburg M. Field studies of pedestrian walking speed and start-up time. *Transportation Research Record: Journal of the Transportation Research Board*. 1996; 1538:27–38.
29. Young SB. Evaluation of pedestrian walking speeds in airport terminals. *Transportation Research Record: Journal of the Transportation Research Board*. 1999; 1674:20–26.
30. Carey, N. Establishing pedestrian walking speeds. Portland State University; 2005.
31. Fujiyama, T.; Tyler, N. An explicit study on walking speeds of pedestrians on stairs. 10th International Conference on Mobility and Transport for Elderly and Disabled; 2004; Hamamatsu, Japan.
32. International Organization for Standardization. ISO 9283 — Manipulating industrial robots — Performance criteria and related test methods. 1998
33. Vasquez DA, Fraichard T, Laugier C. Growing Hidden Markov Models: An incremental tool for learning and predicting human and vehicle motion. *International Journal of Robotics Research*. 2009; 28(11-12):1486–1506.
34. Szabo, S., et al. NISTIR-7851: A Testbed for Evaluation of Speed and Separation Monitoring in a Human Robot Collaborative Environment. National Institute of Standards and Technology; 2012.
35. ASTM International. ASTM E2919-14 - Standard Test Method for Evaluating the Performance of Systems that Measure Static, Six Degrees of Freedom (6DOF), Pose. ASTM International; 2014.
36. International Organization for Standardization. ISO 13849-1 — Safety of machinery — Safety-related parts of control systems — Part 1: General principles for design. 2006

Fig. 1.

Post-event “collision.” A stationary operator, H , triggers a safety-rated stop after the robot has already passed at time $t4$ because the minimum protective distance—shown by the co-centric bounding circles at each time step—increases due to an increase in the robot velocity.

Fig. 2.

Example scenario in which the rotational speed of a part being held by the robot is greater than that of the robot itself. Here, the Cartesian poses of joints 1 (J1) through 5 (J5, red dot) are stationary, whereas Cartesian poses of the tool flange at the end of joint 6 (J6) and the attached workpiece (green bar) change. The rotational velocity of the tip of the held workpiece (turquoise bar), r mm from a J5 motion is much higher than that of the tool flange for a given angular velocity, ω rad/sec.

Fig. 3. Example derived velocity of an actual industrial robot arm based on position measurements reported by the robot and smoothed using a low-pass filter (LPF).

Fig. 4.
Operator path through the shared workspace using a LADAR-based tracking system from [5]. Lateral position offsets are due to sensor noise and errors in centroid fitting.

Fig. 5.

The track measurement error from [5] is the Euclidean distance between the measured positions and a fitted spline of the operator path.

Fig. 6.

A top down perspective of a 6DOF robot (A) results in a simplified model of the robot where the base and TCP are tracked. Adjusting the base joint results in a circular motion profile with a radius equal to the Euclidean distance between the base and the TCP (B), an angular change, θ , and a TCP motion distance of x .

Fig. 7. Example parameterized plot for stopping from ISO 10218-1, illustrating the influence of initial speed (X axis) and payload (a) on the stopping time (Y axis).

Fig. 8.
Schematic drawing of the robot configuration used in the NIST robot testbed.

Fig. 9. Two retro-reflective targets, one representing a simulated human (right) and the other indicating the position of the robot (left) were tracked for SSM implementation testing.

Fig. 10.

Given identical behaviors and initial parameters, the performance of a robot system may be adversely affected by reduced update rates from the safety system. Here, as the update rate slows, the closer the robot's reactive drop in velocity approaches 0 as a human maintaining a constant separation distance is detected as the minimum protective distance grows with increases in the robot's velocity.

Table 1

Stopping categories as defined in IEC 60204-1.

Stop Category	Controlled (Yes/No)	Power Removed (Yes/No)	Brakes Applied (Yes/No)
STOP0	No	Yes	Yes
STOP1	Yes	Yes	Yes
STOP2	Yes	No	No

Table 2

Example calculation of v_R (TCP and robot base, right yellow columns) using instantaneous and analytical approximations as different variables change values. Changes from the base set of parameters (row 1) are indicated by grey boxes.

N	ω_{base} (deg/s)	ω_{TCP} (deg/s)	l (mm)	σ (mm)	r (Hz)	v_{base} (mm/s, instant)	v_{base} (mm/s, analytic)	v_{TCP} (mm/s, instant)	v_{TCP} (mm/s analytic)
4	20	0	1.5	0	100	0	0	524	524
4	20	0	1.5	0	10	0	0	524	524
4	20	0	1.5	0	1	0	0	521	524
4	50	0	1.5	0	100	0	0	1309	1309
4	20	0	1.0	0	100	0	0	349	349
4	20	0	1.5	1	100	39.6	0	493	524
4	20	0	1.5	3	100	98.205	0	605.411	523.599
50	20	0	1.5	3	100	12.473	0	550.177	523.599

Table 3
Stopping response for a 7DOF collaborative, industrial robot with teach and automatic modes

Trigger	Stop Category Response	
	Teach Mode	Automatic Mode
Safety gate opened	--	STOP1
E-Stop pressed	STOP0	STOP1
Enable withdrawn	STOP0	--
Start key released	STOP2	--
"Drives Off" key pressed	STOP0	STOP0
STOP key pressed	STOP2	STOP2
Operating mode changed	STOP0	STOP0
Encoder error	STOP0	STOP0
Motion enable canceled	STOP2	STOP2
Controller turned off / power failure	STOP0	STOP0

Table 4
Stopping response for a 6DOF collaborative, industrial robot without a teach mode

Trigger	Stop Category Response
Safeguard	STOP2
E-Stop pressed	STOP1
Encoder error	STOP0
Motion enable canceled	STOP0
Controller turned off / power failure	STOP0

Table 5

Decision criteria that determine the value of the intrusion distance, C , from ISO 13855. Choosing to provide safety by limiting the direction of approach, relying on the placement of sensors, or requiring a specific control device determines the safety buffer.

Decision Criteria	Conditions		C
Direction of Approach	Normal	$d < 40$ mm	$\max(8(d - 14), 0)$ mm
		$40 \text{ mm} < d < 70$ mm	850 mm
		Multiple Separate Beams	850 mm
		Single-Height Beam	1200 mm
	Parallel/Angled		$\max((1200 - 0.4H), 850)$ mm
Sensor Mounting	General/Step		$1200 \text{ mm} - 0.4H$
	Floor		1200 mm
Control Device	Two-Handed		250 mm