# Toward Safe Close-Proximity Human-Robot Interaction with Standard Industrial Robots

Przemyslaw A. Lasota, Gregory F. Rossano and Julie A. Shah

*Abstract*— **Allowing humans and robots to interact in close proximity to each other has great potential for increasing the effectiveness of human-robot teams across a large variety of domains. However, as we move toward enabling humans and robots to interact at ever-decreasing distances of separation, effective safety technologies must also be developed. While new, inherently human-safe robot designs have been established, millions of industrial robots are already deployed worldwide, which makes it attractive to develop technologies that can turn these standard industrial robots into human-safe platforms. In this work, we present a real-time safety system capable of allowing safe human-robot interaction at very low distances of separation, without the need for robot hardware modification or replacement. By leveraging known robot joint angle values and accurate measurements of human positioning in the workspace, we can achieve precise robot speed adjustment by utilizing real-time measurements of separation distance. This, in turn, allows for collision prevention in a manner comfortable for the human user. We demonstrate our system achieves latencies below 9.64 ms with 95% probability, 11.10 ms with 99% probability, and 14.08 ms with 99.99% probability, resulting in robust real-time performance.**

## I. Introduction

As the field of robotics continues to advance, an increasing amount of focus is placed on the development of technologies that permit more tightly coupled human-robot interaction (HRI). Enabling humans and robots to work together in close proximity to each other would not only allow for more efficient human-robot collaboration in fields where humans and robots already coexist, but also for the introduction of robots into many previously human-only domains. However, safety will always be the primary concern in any application of HRI, and as various HRI technologies are researched, it is of the highest importance that methods guaranteeing human safety during human-robot interaction are developed in parallel.

One can classify safety into two categories: The first, and most obvious, is physical safety. To maintain physical safety, all unwanted human-robot contact must be prevented, and if contact is required by the task at hand or is inevitable for another reason, the forces exerted by the robot on the human must fall below limits that could cause discomfort or injury.

The second – and often overlooked – category is psychological safety. In the context of human-robot interaction, this means ensuring that human-robot interaction does not cause excessive stress and discomfort for extended periods of time. Take, for example, a hypothetical robotic system capable of moving a sharp end effector at very high speeds within centimeters of a human operator's arm. While the system might be able to prevent unwanted injury via contact, a human working with such a system is likely to be in a state of constant stress and discomfort, which can have very negative long-term health effects [1].

It is therefore critical that methods ensuring both physical and psychological safety are developed and designed to meet international standards. Close-proximity interaction between humans and robots is still a fairly new and developing interaction paradigm, and, as such, formal definitions of safety within this context are still under development. Toward the goal of establishing these definitions, the International Organization for Standardization (ISO) developed the ISO 10218 international standard, entitled "Robots and robotic devices – Safety requirements for industrial robots," which was most recently updated in 2011 [2]. A technical specification (ISO TS 15066), entitled "Robots and robotic devices – Safety requirements for industrial robots – Collaborative operation," which provides information and guidance on how to achieve the safety standards described in ISO 10218 specifically for collaborative robots, is still under development [3].

While specific safety standards are not yet fully defined, the National Institute of Standards and Technology (NIST) has provided guidance on the two main areas of focus of the ISO TS 15066: speed and separation monitoring and power and force limiting. In terms of the former, the guidance intends to enable collaborative robots to track people within a workspace and adjust speed according to the distance of separation between the human and robot. In the second area of focus, the aim is to enable robots to moderate applied forces to ensure that they remain below biomechanical limits. Additionally, the project is intended to develop performance measures to test how well a robot conforms to the required standards [4].

This guidance provided by NIST allowed us to pursue the development of an early implementation of the forthcoming standards. In this work, we describe a low-latency, real-time safety system capable of turning a standard industrial robot into a human-safe platform. The system ensures the physical safety and comfort of

Przemyslaw A. Lasota (email: plasota@mit.edu) and Julie A. Shah are with the Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Gregory F. Rossano is with ABB, Windsor, Connecticut.

the user without the need for specialized actuators or any other modification of the robot's hardware, and is capable of supporting precise stopping thresholds that allow for human-robot interaction at separation distances as low as 6 cm.

## II. Related Work

The task of maintaining safety during human-robot collaboration is multidisciplinary in nature, and thus has been approached in a variety of ways. In terms of the psychological aspect of safety, it has been shown that providing physical safety through collision avoidance is not sufficient to maintain human comfort [5]. Furthermore, it has been shown that several parameters, including separation distance, end effector speed, and advance notice of robot motion, have a significant effect on the mental strain of human operators, even if there is no contact between the human and robot workers. The same research also indicated that having grown accustomed to working with robots does not necessarily diminish these effects [6].

These results illuminate the importance of not only maintaining physical safety, but also ensuring that robot motions are comfortable for the humans interacting with the robot. Prior work that has taken this point into consideration evaluated parameters including the human's field of vision, posture, and kinematics when planning safe and comfortable robot paths [7], [8].

With regard to the physical aspect of safety, work has been done toward minimizing the negative effects of human-robot collision, as well as preventing collision from occurring altogether. Work on collision reaction control strategies has shown that switching to torque control with gravity compensation upon impact can greatly reduce the force exerted on the human in the event of a collision [9]. New types of actuators with variable impedance have been developed, and show great potential in allowing for intrinsically safer robots by reducing joint stiffness when the robot is moving quickly [10].

In the realm of collision prevention, innovation in 3D sensor fusion and the use of dynamic safety zones appears to be a promising method [11]. The ability to predict human actions also has the potential to prevent collision. On the human motion level, recent work has shown that human actions can be predicted from early stages of movement [12]. On the task level, prior work has indicated that observing changes to the entropy rate of a Markov Chain, produced from a task description encoded as a Markov Decision Process, could be utilized to encode the uncertainty of the robot about what action the human will perform next [13]. In other work, the encoding of discrete sets of human and robot actions allowed for the incorporation of task-specific rules and preferences, which could then be utilized to predict likely sequences of human actions [14].

Research and innovation in these various fields has led to the development of new, inherently human-safe robots, such as the RethinkRobotics Baxter, which features force sensors at each joint and Series Elastic Actuators that minimize the force of impact [15]; or ABB's Dual Arm Concept Robot, which has built-in power and speed limitation, as well as software-based collision detection [16]. Besides the creation of brand-new robot designs, work in the field of robot safety has also led to the development of add-on technologies, such as ABB's SafeMove, which, through the addition of external sensing and a software module, provides programmable, complex safe zones by monitoring robot speed and position [17].

While many of these works have yielded very promising results for maintaining human safety in HRI, a great majority of them focus on technologies that can be applied to new robot designs, rather than to existing robotic platforms. While utilizing new robot designs like the Baxter and ABB's Dual Arm Concept Robot and developing more human-safe robots with the technologies mentioned above can ensure human safety in HRI, purchasing new robots or retrofitting existing robots with new hardware components can be cost-prohibitive or physically impossible. With an estimated 1.2 to 1.5 million industrial robots already in use worldwide [18], there is great incentive to design a solution that can turn these robots into human-safe platforms without the need for hardware modification.

In the work mentioned above that does not explicitly require new actuators or arrays of internal robot sensors, safety systems are often designed such that the robot completely avoids a large region where the human is located, or uses approximations of human and robot locations that are too coarse or uncertain to allow for the robot and human to interact in close proximity to one another. As a great deal of industrial work is still performed by humans – even in fields where robots have been successfully integrated, such as the automotive and aerospace industries – many industrial applications stand to benefit from the introduction of robotic assistants that aid human workers. However, this assistance will require close-proximity HRI, which makes the development of a safety system capable of operating effectively at small separation distances attractive.

The goal of this work was, therefore, to build upon prior work in the field in order to overcome the abovementioned drawbacks and create a robot safety system capable of turning current, standard industrial robots into human-safe platforms for close-proximity HRI, without the need for robot hardware modification.

## III. Implementation

### A. Hardware

The robot used in the implementation and evaluation of the safety system described in this work is the

ABB IRB-120. This is a standard industrial robot with no built-in safety systems for HRI and noncompliant joints, capable of moving at speeds as high as 6.2 m/s [19]. Without an additional safety layer, this type of robot is required to operate within a safety cage – which, when opened, stops the robot immediately. Consequently, it is not capable of safe HRI in its stock form.

A PhaseSpace motion capture system was utilized to sense the position of the human worker within the workspace. This type of active motion capture system provides accurate human localization that is robust to temporary occlusions. While this type of system might not be a viable tracking solution in factory environments, the developed safety system can be utilized with any sensing system capable of providing accurate localization data. As advancements are made in the field of computer vision and new 3D sensing hardware is introduced, the motion capture system may be replaced by a less-intrusive option, if necessary.

The computer platform used to run the safety system software was a standard Windows 7 machine with a Core i7-3610QM 2.3GHz processor.

### B. *Software*

The software implementation of the safety system consists of several subsystems that exchange information in a coordinated, low-latency fashion: the core program, the motion capture software, the robot software running on the ABB IRB-120's controller, and a virtual workspace.

- Core Program: The core program serves as the main logic of the system. It connects to the other sub-components and relays information between them via TCP sockets. The core program also performs some geometrical calculations used by other sub-components and logs data for system analytics.
- Motion Capture Software: This sub-component is responsible for capturing the most recent position of the human. The software takes in raw data from the motion capture system, transforms it into the correct coordinate frame, and relays it to the core program.
- Robot Software: This portion of the system software resides on the robot's controller. It continuously monitors the robot's configuration and relays this information to the core program. The software is also responsible for adjusting the robot's speed according to separation distance data received from the core program. The code runs as a secondary task on the robot's controller, completely independent of the primary task used to command the robot's motions. Consequently, the safety system can run in the background of any task given to the robot, making it easy to use for virtually any task the robot is programmed to perform with a human co-worker.
- Virtual Workspace: This sub-component is responsible for constructing a virtual representation of the workspace shared by the human and robot, based on information received from the motion capture system and robot controller, using OpenRAVE, a robot simulation environment [20]. The current robot position is translated into the virtual workspace based on the known position of the robot's base, its 3D CAD model, and the joint angles received from the robot's controller. In the particular workspace in which the robot and safety system were tested, the only portion of a human worker the robot was able to reach was the right arm and hand. Consequently, the position of the human in the virtual workspace was approximated by two concentric cylinders: one cylinder for the forearm and one larger-diameter cylinder for the hand. The length of the forearm cylinder was adjusted to the particular user's arm length based on information received from the motion capture system. The diameters of the two cylinders were such that the virtual cylinders completely enclosed the user's arm and hand. A sample workspace configuration and corresponding virtual representation are depicted in Fig. 1. Once the configuration of the human and robot is updated in the virtual environment, the separation distance between them is accurately calculated and relayed to the core program, which in turn relays this information to the robot's controller for speed adjustment.

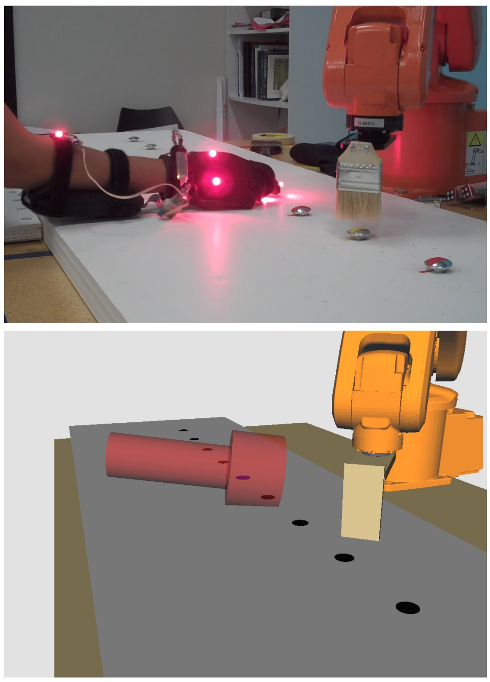

Fig. 1. Real workspace (top) and corresponding virtual representation (bottom)

### C. Implementation Discussion

The decision to use this type of implementation scheme – a virtual environment constructed according to the position of the human and robot in the workspace – as opposed to, for example, a purely vision-based approach, was made in order to fully leverage the known robot configuration, rather than approximate it via another method. Since the position of the human is also accurately known, the virtual environment implementation scheme allows for very accurate separation distance measurements in real-time, resulting in precise robot speed control. The robot's speed was adjusted according to a function of the form:

$$\alpha(d) = \begin{cases} 1-\beta\left(\mathrm{d}-\mathrm{d_{stop}}\right)^{\gamma} & \{d \mid d_{stop} \le d \le d_{slow}\} \\ 0 & \{d \mid d > d_{slow}\} \\ 1 & \{d \mid d < d_{stop}\} \end{cases}$$

In the above form, $\alpha$ represents the percentage reduction in the robot's speed expressed as a decimal, d is the current separation distance between the human and robot, $d_{stop}$ is the distance at which the robot should be stopped, $d_{slow}$ is the distance at which the robot's deceleration begins, and $\beta$ and $\gamma$ are tuning parameters that define the behavior of the speed reduction function; for example, how quickly the speed should drop off and whether the bulk of the reduction should occur near $d_{stop}$ or $d_{slow}$.

The ability to precisely control the speed of the robot as a continuous function of separation distance allows the safety system to be effective at very low separation distances. At moderate robot speeds and with the proper choice of parameters in the speed reduction function, the safety system is effective with the tested hardware at $d_{stop}$ values as low as 6 cm. This means that the human and robot can safely perform tasks in very close proximity to one another, which is not possible with other safety systems that incorporate coarse, discretized workspace occupancy approximations or discrete safety zones.

An additional benefit of precise robot speed control based on separation distance is that the system can be tuned such that the deceleration of the robot occurs at a rate comfortable for the human. By properly tuning the parameters, we can have the robot ease to a stop gently and smoothly, as opposed to stopping abruptly or using coarse "slow" and "stop" zones that cause sudden changes in robot speed. Fig. 2 depicts three possible modes of speed reduction based on the tuning of $\beta$ and $\gamma$ in the speed reduction equation, with $d_{slow}$=15 cm and $d_{stop}$=6 cm. The green dashed line in the figure represents the strategy of reducing speed slowly when the slow-down threshold is passed, then quickly reducing speed to zero when the stop threshold approaches. The red dotted line represents the opposite strategy: rapidly decreasing speed once the reduction threshold is passed, then gradually easing to a stop. The blue solid line represents a balanced approach between these two modes.

Which mode is appropriate is left for the end user to determine, as their choice might depend on the robot's speed, the tool it is holding (which can include considerations for its inertia, sharpness, and other attributes), the potential for pinch points during assembly, or other task-dependent parameters. The freedom to finely tune the deceleration behavior of the robot allows the user to adjust it so that the interaction is comfortable and stress-free, even at small distances of separation.

Another key benefit of this type of implementation scheme is that it does not require any robot hardware modification. While other systems require special actuators or the retrofitting of robots with force and torque sensors, this implementation can be utilized with standard, unmodified industrial robots. This makes our safety system easy and cheap to implement for organizations that already own industrial robots, as they are able to turn their previously dangerous industrial robots into human-safe platforms capable of close-proximity HRI at a very low cost.

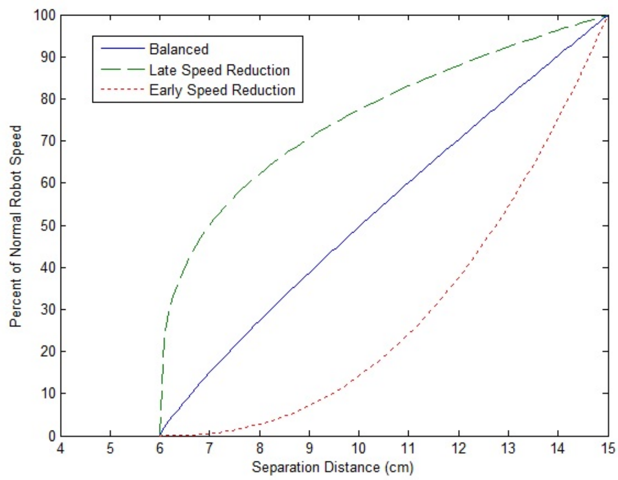


Fig. 2. Three possible modes of speed reduction as a function of separation distance

### D. Latency Improvement

A key requirement for the robustness of this type of safety system is low latency. This means that the amount of time necessary to perform a complete cycle of the safety system, from sensing to robot speed adjustment, must be very low. Consequently, several measures were taken to improve the latency of the described system.

First, to decrease the amount of time needed to perform minimum distance calculations, the CAD model of the robot was reduced in quality. An increase of 0.75 mm in the maximum deviation of the model reduced the number of polygons in the CAD model quite drastically (~50% reduction in file size). This led to significantly faster separation distance calculation at a very low penalty to model accuracy.

The second key improvement in latency resulted from network optimization. The various subsystems described in Section 3.2 reside on different physical machines, and so they must communicate with each other via an internal LAN. In order to ensure that all pertinent data is delivered successfully and in the correct sequence, Transmission Control Protocol (TCP) sockets were chosen as a mode of connection and transmission. While User Datagram Protocol (UDP) sockets can provide faster communication, this protocol does not guarantee successful delivery or correct sequence and is not supported by the RAPID programming language used to program ABB robots [21].

Due to the design of the system, very small packets of data are continuously sent over the TCP sockets, and each successive transmission must finish before the next one begins. This led to slow transmission speeds due to Nagle's Algorithm, a network optimization algorithm designed to prevent network congestion by chunking small packets together and sending them all at once [22]. Instead of sending a small packet immediately upon generation, Nagle's Algorithm instructs the program to wait for more data to send, reducing bandwidth at a cost to latency. As the continuous transmission of small packets is required by the safety system, Nagle's Algorithm was disabled for all socket connections, leading to significant improvements in latency.

## IV. System Latency Evaluation

While our safety system is capable of robot speed adjustment in real-time, there are no hard guarantees for latency (i.e., it is not "hard real-time"). Consequently, it is desirable to evaluate the latency based on collected performance data to ensure that latency, on average, is at a sufficiently low level to yield consistent performance. To allow for such evaluation, the safety system was utilized during an array of human-subject experiments involving the ABB IRB-120 robot.

Safety system latencies were recorded for a total of 174 experiment runs, resulting in approximately 1.8 million latency measurements over the course of 3 hours. The average latency was 6.13 ms, with a maximum latency of 389.6 ms. While the maximum latency was substantially higher than the average, large deviations from the average happened very rarely. Based on statistical analysis of the collected data, assuming the distribution is normal with a mean of 6.13 ms and a standard deviation of 2.14 ms, latencies are expected to be below 9.64 ms with 95% probability, below 11.10 ms with 99% probability, and below 14.08 ms with 99.99% probability.

In order to visualize these results, a histogram of latencies was constructed. Fig. 3 depicts the overall histogram of latencies, indicating a very large peak around the average latency and only sporadic high latency jumps. The source of these jumps is not known, but they may be caused by infrequent network connection drops. Fig. 4 shows a zoomed-in view of the histogram near the average latency.

While no hard guarantees on latency are made by the current implementation of the system, based on these results, we can say with high confidence that the average latency is very low and the sporadic jumps in latency do not occur often enough to significantly degrade system performance. In fact, with the robot used in this implementation, there is an inherent delay of 300-500 ms from the time the robot is ordered to reduce speed until it begins to execute the order [21], making the approximately 6 ms latency of the safety system an insignificant contribution to overall system latency. Since the frequency of high latencies is low and the sporadic jumps do not jeopardize the performance of the system, the system is classified as soft real-time according to the definition of the IEEE Technical Committee on Real-Time Systems [23].

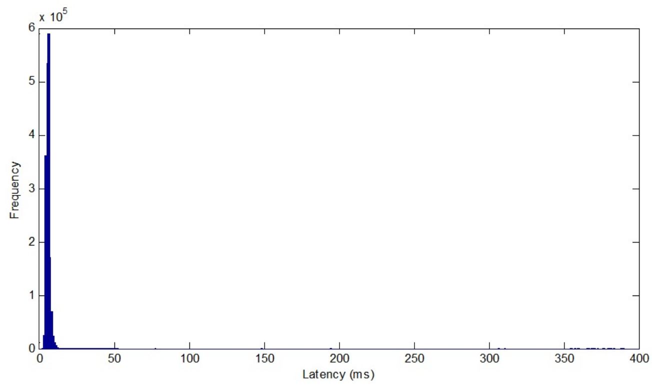


Fig. 3. Histogram of system latencies. Note the very low number of high latencies on the right side of the graph.

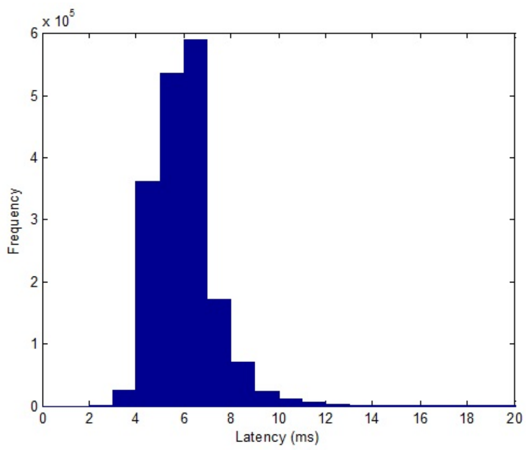


Fig. 4. Portion of histogram of system latencies shown in Fig. 3 near the average latency

## V. Conclusions and Future Work

In this work, we described a novel implementation of a real-time safety system capable of turning a standard industrial robot into a human-safe platform. We showed that this implementation does not require any robot hardware modifications, such as special actuators or internal force and torque sensors, making our safety system inexpensive and easy to implement in domains where robots are already present. By leveraging known robot joint angles and utilizing accurate human localization, we showed that we can construct a virtual representation of the workspace that allows for the calculation of accurate separation distance data in real-time. We then described how this information can be used to precisely control robot speed, allowing for safe HRI at distances of separation as low as 6 cm, as well as for robot deceleration comfortable for the human worker. Finally, we demonstrated the benefit of deploying various latency improvement strategies, which resulted in system latencies falling below 9.64 ms with 95% probability, below 11.10 ms with 99% probability, and below 14.08 ms with 99.99% probability.

While the latency is low compared to the inherent delay between speed adjustment commands and the execution of those commands by the robot used in this work, the lack of formal guarantees on system latency means that this is not a "hard real-time" system. In the future, we plan to incorporate a middleware solution, such as OROCOS RTT [24], to allow for hard real-time operation. The addition of middleware will also add a level of platform independence, making it easier to apply the system to a variety of robots and 3D sensors.

The inherent delay between speed adjustment commands and the execution of those commands by the robot utilized in this work can be expected with other industrial robots. As such, there is a substantial delay that cannot be removed through latency improvement within the safety system. We must overcome this limitation in order to improve system performance and allow for even closer HRI or higher robot speeds.

Consequently, another future direction of our work is to augment the safety system through the prediction of future locations of the human and robot. If robot trajectories are known and we can accurately model where a human might move to in the next 300-500 ms using current motion or previously learned motion models, we can attempt to look a few hundred milliseconds "into the future" and adjust robot speed based on this information. Such an approach would help to overcome the limitation imposed by the inherent speed adjustment delays, but would be highly dependent on the accuracy of the future location prediction.

## Acknowledgments

This work is supported by the ABB Group. We especially thank Thomas Fuhlbrigge, Gregory Rossano, and Carlos Martinez for their continuing support.

## References

[1] B. S. McEwen and E. Stellar, "Stress and the individual: mechanisms leading to disease," *Archives of Internal Medicine*, vol. 153, no. 18, pp. 2093–2101, 1993.

[2] "ISO 10218-1:2011 Robots and robotic devices – Safety requirements for industrial robots – Part 1: Robots." International Organization for Standardization, `http://www.iso.org`, 2011.

[3] "ISO/DTS 15066 Robots and robotic devices – Safety requirements for industrial robots – Collaborative operation." International Organization for Standardization, `http://www.iso.org`, 2014.

[4] "Safety of human-robot systems in fixed workcell environments." The National Institute of Standards and Technology (NIST), `http://www.nist.gov`, 2014.

[5] E. Meisner, V. Isler, and J. Trinkle, "Controller design for human-robot interaction," *Autonomous Robots*, vol. 24, no. 2, pp. 123–134, 2008.

[6] T. Arai, R. Kato, and M. Fujita, "Assessment of operator stress induced by robot collaboration in assembly," *CIRP Annals-Manufacturing Technology*, vol. 59, no. 1, pp. 5–8, 2010.

[7] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon, "A human aware mobile robot motion planner," *Robotics, IEEE Transactions on*, vol. 23, no. 5, pp. 874–883, 2007.

[8] E. A. Sisbot, L. F. Marin-Urias, X. Broquere, D. Sidobre, and R. Alami, "Synthesizing robot motions adapted to human presence," *International Journal of Social Robotics*, vol. 2, no. 3, pp. 329–343, 2010.

[9] S. Haddadin, A. Albu-Schaffer, A. De Luca, and G. Hirzinger, "Collision detection and reaction: A contribution to safe physical human-robot interaction," in *Proc. IEEE/RSJ International Conference Intelligent Robots and Systems (IROS)*, pp. 3356–3363, IEEE, 2008.

[10] G. Tonietti, R. Schiavi, and A. Bicchi, "Design and control of a variable stiffness actuator for safe and fast physical human/robot interaction," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, pp. 526–531, IEEE, 2005.

[11] P. Rybski, P. Anderson-Sprecher, D. Huber, C. Niessl, and R. Simmons, "Sensor fusion for human safety in industrial workcells," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3612–3619, IEEE, 2012.

[12] J. Mainprice and D. Berenson, "A human-robot collaborative manipulation planning framework that reasons on early prediction of human motion," in *Proc. IEEE International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2013.

[13] S. Nikolaidis, P. Lasota, G. Rossano, C. Martinez, T. Fuhlbrigge, and J. Shah, "Human-robot collaboration in manufacturing: Quantitative evaluation of predictable, convergent joint action," in *Proc. International Symposium on Robotics (ISR)*, pp. 1–6, IEEE, 2013.

[14] R. Alami, R. Chatila, A. Clodic, S. Fleury, M. Herrb, V. Montreuil, and E. A. Sisbot, "Towards human-aware cognitive robots," in *Proc. The Fifth International Cognitive Robotics Workshop (The AAAI-06 Workshop on Cognitive Robotics)*, 2006.

[15] "Baxter safety and compliance overview." Rethink Robotics, `http://www.rethinkrobotics.com`, 2014.

[16] "Dual-arm concept robot: Safety." ABB Group, `http://www.abb.us`, 2012.

[17] S. Kock, J. Bredahl, P. J. Eriksson, M. Myhr, and K. Behnisch, "Taming the robot: better safety without higher fences," *Doc no. 9AKK105152A2830* *`http://www.abb.com`*, 2006.

[18] "World robotics industrial robots 2013 (executive summary)," tech. rep., International Federation of Robotics (IFR), 2013.

[19] N. D. Keijser, "ABB's smallest robot - IRB 120 for flexible & compact production," Product presentation IRB 120 Revision E, `http://www.abb.com`.

[20] R. Diankov and J. Kuffner, "Openrave: A planning architecture for autonomous robotics," *Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-08-34*, p. 79, 2008.

[21] "RAPID reference manual." ABB Group, `http://www.abb.com`, 2014.

[22] J. Nagle, "Congestion control in IP/TCP internetworks," 1984.

[23] "Terminology and notation." IEEE Technical Committee on Real-Time Systems `http://www.tcrts.org/`, 2014.

[24] "The Orocos real-time toolkit." The Orocos Project, `http://www.orocos.org/rtt`, 2014.