Robust Immersive Telepresence and Mobile Telemanipulation: NimbRo wins ANA Avatar XPRIZE Finals

Max Schwarz^{*}, Christian Lenz^{*}, Raphael Memmesheimer, Bastian Pätzold, Andre Rochow, Michael Schreiber, and Sven Behnke

Abstract-Robotic avatar systems promise to bridge distances and reduce the need for travel. We present the updated NimbRo avatar system, winner of the \$5M grand prize at the international ANA Avatar XPRIZE competition, which required participants to build intuitive and immersive robotic telepresence systems that could be operated by briefly trained operators. We describe key improvements for the finals, compared to the system used in the semifinals: To operate without a power- and communications tether, we integrated a battery and a robust redundant wireless communication system. Video and audio data are compressed using low-latency HEVC and Opus codecs. We propose a new locomotion control device with tunable resistance force. To increase flexibility, the robot's upper-body height can be adjusted by the operator. We describe essential monitoring and robustness tools which enabled the success at the competition. Finally, we analyze our performance at the competition finals and discuss lessons learned.

I. INTRODUCTION

Traveling large distances costs money and time; and most forms of travel impact the environment. Reducing the need to travel is thus beneficial for many reasons. While voice calls and video conferencing help, they cannot replace inperson meetings entirely due to lack of immersion and social interaction. Furthermore, many remote tasks require mobility, physical touch, grasping and handling of objects, or even more complex manipulation skills. These requirements cannot be addressed by VR-based conferencing systems that focus on meetings in a virtual space. In contrast, avatar systems allow full immersion into a remote space while also *embodying* the operator in a robotic system, giving them the ability to navigate and physically interact with both the remote environment and persons therein.

The ANA Avatar XPRIZE competition¹ challenged the robotics community to advance the state of the art in avatar systems. Promising a record \$10M prize purse, the competition required teams to build *intuitive* and *robust* robotic avatar systems that allow a human operator to be present in a remote space. The tasks to be solved included social interaction and communication, but also locomotion and complex manipulation. Critically, the systems were to be used and evaluated by operator and recipient judges. In contrast to previous teleoperation competitions, such as the DARPA Robotics Challenge [1] operators could be trained only for a short time how to use the developed avatar systems.

*Equal contribution.



Fig. 1. NimbRo avatar system at the ANA Avatar XPRIZE competition Finals. Stills from the end of our winning final run with the robot holding the correctly retrieved stone (magenta). Top left: Operator judge controlling the avatar. Bottom left: VR view (cropped). Right: Avatar robot in the arena.

In this paper, we present and discuss the updates and extensions of the NimbRo avatar system (Fig. 1) that we made for our highly successful participation in the ANA Avatar XPRIZE Finals in November 2022, where our team won the grand prize². The finals posed new requirements and tasks that resulted in various system extensions and improvements, compared to our earlier system used in the semifinals [2]. The contributions of this paper include:

- 1) hardware integration for tetherless and battery-powered operation of the avatar robot and mobility of the operator station,
- a redundant network stack for robust wireless communication,
- 3) monitoring tools for efficient support crew operations,
- 4) auto-recovery mechanisms for failure tolerance on multiple levels, and
- 5) a thorough analysis of the competition results and lessons learned from our participation at the finals.

II. RELATED WORK

Teleoperated robotic systems are widespread. For example, the DARPA Robotics Challenge [1] resulted in an array of legged, wheeled, and tracked teleoperated robots. We focus our discussion on avatar systems, that is, systems which not only allow teleoperation, but also telepresence immersion in the remote environment and interaction with human recipients, as in the XPRIZE competition.

All authors are with the Autonomous Intelligent Systems group of University of Bonn, Germany; schwarz@ais.uni-bonn.de

https://www.xprize.org/prizes/avatar

²https://www.youtube.com/watch?v=EmESa201q4c



Fig. 2. NimbRo avatar system consisting of the operator station (left) and the avatar robot (right).

Some participants are commercial entities and have no scientific publications, such as Pollen Robotics, who came in second in the finals. Their avatar robot is a heavily modified version of their product *Reachy* with stronger arms (7 DoF, 3.5 kg payload) and a communication head displaying the operator's face. In contrast to our operator station, Pollen uses VR controllers with vibration actuators and a 1 DoF elbow exoskeleton for manipulation. Torque and haptic feedback is thus limited to elbow torque and controller vibration.

Luo *et al.* [3] describe the approach by team Northeastern, who achieved third place at the finals. Their avatar system features an interesting glove and gripper system with hydrostatic actuation, which gives fine-grained force feedback. A wave variable approach handles varying communication latency. The decision to forgo a VR head-mounted display in favor of 2D screens is a simple way to avoid motion sickness caused by network latency, but limits immersiveness in comparison to our system.

Marques *et al.* [4] came in fourth with the AVATRINA system. Mechanically, it is similar to our system with two Franka Emika Panda arms, dexterous hands and VR teleoperation. In contrast to ours, the neck can only rotate and not translate, which makes looking around occlusions impossible and reduces depth perception and immersion.

Van Erp *et al.* [5] showcase the system of team i-Botics, who came in fifth. Based on Halodi Eve, their robot is considerably more humanoid in shape than the other top five teams, although it operated with a gantry for safety reasons up until the last competition run. During the competition, the system suffered from network connectivity problems. The robot's neck is only 1 DoF, seriously limiting camera motion.

Park *et al.* [6] describe the system of team SNU (8th place). Their avatar robot is fully humanoid in shape, although for finals it was sitting on a holonomic base for fast and safe locomotion. Interestingly, the team decided not to attempt realistic animation of the operator's face, but displayed operator emotions through basic line drawings of the mouth area, which limits identification of the robot with the human

controlling it. The SNU operator station features a unique linear actuator system which seems to enlarge the workspace.

Aside from the XPRIZE competition, there are other notable avatar systems targeting different applications like enabling people with disabilities to work remotely [7], underwater telepresence [8], and space [9]. Takeuchi et al. [7] presented an avatar robot system with focus on enabling people with disabilities to execute physical work remotely. The robot is controlled by mouse and gaze input. In contrast to our proposed system, which aims at transmitting the movements of an operator naturally, their system follows a less immersive approach by controlling motion through a GUI. Lii et al. [9] describe a teleoperated robot controlled by astronauts through a tablet GUI and a 1-DoF haptic joystick. The system can execute assembly and maintenance tasks but has, due to the domain, no human interaction capabilities. TELESAR VI [10] has a long history in the field of telexistence, starting in 1980. The system is designed for remote manipulation and gesturing. Unlike ours, the robot is operated in a stationary seated posture and has two controllable legs. All ten fingers are equipped with multiple sensors for enhanced tactile feedback. However, the force feedback is limited to the fingers.

III. NIMBRO AVATAR SYSTEM

The NimbRo avatar system consists of a robotic operator station and an anthropomorphic avatar robot (see Fig. 2). The human operator sits on a chair and is strapped into two 7 DoF compliant arm exoskeletons (Franka Emika Panda arms) at their palms. 6 DoF force/torque sensors (Nordbo NRS-6050-D80) at the arm wrists are used to provide a weightless feeling for the operator and force/torque feedback to their hands. Finger movements are captured using SenseGlove DK1 hand exoskeletons, which provide force and haptic feedback to the finger tips. The operator's feet are resting on a custom-built 3D pedal device (see Section III-F), which allows omnidirectional control of the robot's base. For visual and auditive immersion, the operator wears a VR head-

mounted display (Valve Index), which is equipped with additional cameras to capture gaze direction, eye opening, and mouth expressions.

The avatar robot is equipped with a holonomic mecanumwheeled base for indoor locomotion. Its spine is a linear actuator that can be used to adapt to different manipulation and communication heights. For bimanual manipulation, the robot features two arms (Franka Emika Panda) in approximately humanoid configuration. Force/torque sensors (On-Robot HEX-E) are mounted on the wrists for force feedback. We chose two different robotic hands (Schunk SVH and SIH), each featuring different capabilities. Please refer to Lenz and Behnke [11] for a detailed description of our arm force feedback telemanipulation controller, including active joint limit avoidance using model-based predictions and an oscillation detection and suppression module.

The robot head is mounted on a robotic arm (UFactory xArm-6), which mirrors the operator's 6D head movement. Together with the head-mounted wide-angle stereo camera pair ($2 \times Basler a2A3840-45ucBAS$) this enables a highly immersive visual 3D experience for the operator. Latencies are mitigated by spherical rendering [12]. The head further carries a screen showing a live animation of the operator's face—a direct video display is not possible, since the operator is wearing the HMD. See [13], [14] for details on the facial animation.

The first integrated NimbRo avatar system (as of February 2021, prior to the ANA Avatar XPRIZE Semifinals) is described in detail in Schwarz *et al.* [2].

The ANA Avatar XPRIZE finals [15] brought new tasks and requirements which necessitated changes to our design. On the one hand, it was clear that the robot had to operate without a tether for communications and power, which necessitated integration of wireless communication and battery power. The robot had to navigate through narrow passages, necessitating to reduce its width. Much emphasis was placed by XPRIZE on haptic perception and we extended our system accordingly. On the other hand, the competition format changed: Whereas in the semifinals, individual scenarios could be attempted multiple times, with manual intervention allowed in-between, the finals called for a continuous mission through ten tasks, with no possibility to skip tasks or to restart the system in case of failure. Additionally, the participants were down-selected over the successive competition days (see Section IV). This meant that considerable focus had to be placed on making the system robust.

We will now detail the changes and improvements to the system since early 2021 up to the finals in November 2022.

A. Mobile Operator Station

The finals required a mobile operator station that could be moved into an operator control room and set-up quickly. To this end, we added extendable wheels to it (see Fig. 2). The control computer and a dedicated computer for face animation were integrated into the structure. Additionally, we added a large battery which can supply the operator station for multiple hours. Using this setup, we could prepare

TABLE I Avatar Robot Power Distribution & Average Consumption

Voltage Device	Power	Voltage	Device	Power
Battery Wheels (idle)	16 W 350 W	24V	- xArm motors Hands	50 W 8 W
5V - Base controller	5 W	l	F/T sensors	11 W
$12V \begin{bmatrix} xArm PC \\ Panda PCs \end{bmatrix}$	15 W 145 W	48V ·	- Panda motors	60 W
Total powe	er consumr	tion (avg)	: 660 W	

the operator station for use long before our run and leave everything initialized and switched on during transport to the operator control room.

B. Telemanipulation

Our telemanipulation components, including the arm and hand exoskeletons on the operator side, arms and fivefinger hands on the avatar side, and our force feedback controller have been proven during semifinals [11]. To adapt to the finals requirements, we made three changes. First, we mounted the avatar's arm bases closer together, which reduced the shoulder width and made it easier to maneuver through narrow passages. Second, we equipped the fingertips of the SVH and SIH hands with microswitches and magnet hall sensors, respectively. This allows contact to be measured and then displayed haptically to the operator using the SenseGlove DK1. Finally, we replaced the OnRobot HEX-E force/torque sensors on the operator side with more rigid Nordbo NRS-6050-D80 which offer higher update rate of 1 kHz, resulting in faster response to operator movement.

C. Roughness Sensing

Task 10 of the finals (see Fig. 9) required the ability to remotely feel texture, especially surface roughness. To enable this, we developed an audio-based sensing, detection, and vibrational display system: The left index fingertip is equipped with two microphones: one measuring air vibrations, the other measuring vibrations coupled through the finger. The audio data is transmitted over WiFi (see Section III-D.2) and then analyzed by a CNN, which classifies very short audio segments as rough or smooth. Finally, a vibration actuator on the operator's fingertip displays the classification result, giving the illusion of feeling rough bumps on the object surface with low latency. This low-cost approach is very small and non-invasive on the sensing side. Details of this subsystem are described by Pätzold *et al.* [16].

D. Tetherless Operation

To make the system operable without a tether, wideranging changes to the avatar robot had to be implemented.

1) Power Supply: The robot now carries a RELiON In-Sight 48 V 30 Ah battery, which powers all onboard systems. All power consumers had to be converted to DC power. We identified four voltage rails (see Table I) and installed individual DC-to-DC converters. This required tight integration to fit everything into the robot base (see Fig. 3).





Fig. 4. Network Architecture. The operator station contains a 1 GBit/s ethernet adapter, which is connected to the XPRIZE network (or our own access point during testing). Two separate access points broadcast a WiFi network at 2.4 GHz and 5 GHz, respectively. The avatar control PC is equipped with two PCI-E WiFi adapters, one for each of the networks.

2) Wireless Communication: For the finals, XPRIZE supplied a 2.4 GHz and a 5 GHz WiFi network on the competition arena. To utilize this fully, we incorporated two WiFi adapters on our robot (see Fig. 4). All network adapters (wired on the operator side, wireless on the robot side) are connected via PCI-E to their respective control computers—reducing latency to a minimum and increasing robustness, compared to USB devices. On the robot, we use two Intel AX210 cards, which have excellent driver support. Each WiFi card is connected to two antennas: a main antenna extended up from the spine of the robot (see Fig. 2) and a backup antenna in the base.

To facilitate fine-grained control of routing, i.e. which WiFi band is used for which data stream, each WiFi adapter has its own IP address. Conversely, the operator station uses two IP addresses on its Ethernet interface. Static routes ensure that the correct source address and interface are used for each destination address.

Our network transport is based on nimbro_network³, which transmits data between the otherwise isolated ROS [17] instances. The type and amount of data is statically configured and transmitted over UDP, which avoids any kind of connection handshake. This allows our system to seamlessly begin operation as soon as network connectivity is established and to recover immediately after a network interruption. This capability strongly contributed to our success at the finals (see Section IV).

At semifinals in 2021, our system required around 300 MBit/s downlink bandwidth. At the time, this was fine given a wired 1 GBit/s connection, but now bandwidth had to be reduced for robust WiFi operation. While current WiFi systems can sustain such bandwidths in ideal situations,

TABLE II WIFI BANDWIDTH REQUIREMENTS

Dowr	link fror	n avatar		Uplink to avatar					
Channel	MBit/s	5 GHz 2.4 GHz		Channel	MBit/s	5 GHz 2.4 GHz			
Arm feedback	8.5	\checkmark	×	Arm control	4.9	\checkmark	~		
Transformations	s 4.1	\checkmark	×	Transformations	5 1.4	\checkmark	×		
Main cameras	14.7	\checkmark	×	Operator face	5.7	×	\checkmark		
Hand camera	5.5	×	\checkmark	Audio	0.4	\checkmark	\checkmark		
Diagnostics	0.4	\checkmark	\checkmark						
Audio	0.4	\checkmark	\checkmark						
Total [MBit/s]		28.1	6.3	Total [MBit/s]		6.7	11.0		

they cannot guarantee this in non-line-of-sight or otherwise difficult circumstances. We reduced the required bandwidth significantly by compressing the main video stream $(2 \times 2472 \times 2178 @ 46 Hz)$ with the HEVC video codec instead of transmitting individual JPEG images. We took care not to increase video latency, which is a source of operator disorientation and fatigue. To this end, we perform debayering, white balancing, and color correction in a fused CUDA kernel on the onboard RTX 3070 GPU. Video encoding is then performed using the NVIDIA NVENC library on the GPU as well. On the operator side, compressed video packets are uploaded to the GPU (RTX A6000 48 GB), extracted using NVDEC, and sent to the HMD after spherical rendering [12], which corrects for head movement latencies. The total latency from start of camera exposure to HMD display is under 50 ms. Overall, we achieve considerably lower bandwidth than our old system (see Table II).

Table II also shows that manipulation control and audio are routed over both WiFi networks redundantly. Both data types are very sensitive to interruptions and packet drops, which lead to uncontrolled movement and loss of information, respectively. The redundant transmission minimizes this risk. We also experimented with redundant configuration for our camera stream, but refrained from activating this in the finals due to concerns about 2.4 GHz bandwidth in an arena with many devices of spectators operating in this band.

Audio transmission over WiFi is in itself problematic. Our semifinal solution used the Opus codec with a very low buffer size of 64 samples (at 48 kHz audio), which gives a high packet rate—challenging for WiFi networks. The packet rate can be reduced by increasing the buffer size, but this increases latency and easily leads to echo effects, where the operator can hear their own voice as transmitted by the avatar and then captured by the avatar's microphones. To mitigate this, we integrate an echo cancellation system based on NVIDIA Maxine⁴, which allows us to increase the buffer size to 512, reducing the audio packet rate to roughly 90 Hz.

3) Wireless E-Stop: We integrated an HRI Wireless Emergency Stop to provide a reliable tetherless safety system. Activating the E-Stop depowers the wheels. The avatar stops quickly due to friction, and the base can be pushed around manually. In addition, both Panda arms, the xArm, and both hands hold their current joint positions. The Panda arms can be moved freely using the teach button on their wrist.

³https://github.com/AIS-Bonn/nimbro_network



Fig. 5. Self-centering rudder design with individually tunable springs for intuitive omnidirectional locomotion control. Fig. 2 shows the device during operation.

E. Height Adaption

In contrast to semifinals, where manipulation on table height was required, finals specified a wider height range. To increase flexibility, we integrated a linear joint in the spine of the robot (see Fig. 2). The operator can control the height with a bidirectional Danfoss KEP foot pedal. The pedal contains springs to provide resistance and uses hall effect sensors for angle sensing. The current height and a sideview rendering of the robot is shown to the operator during movement of the actuator to assist the height adjustment.

F. Locomotion Control

For locomotion of the holonomic robot platform, we propose a new feet-controlled 3D-printed rudder design (see Fig. 5). We identified two major problems with the old device: First, the rudder lacked resistance and self-centering, which made it difficult to control. Second, the rudder yielded only position estimates when both feet were placed on the rudder surface. This became especially problematic when the feet were lifted off the rudder and tilted or rotated slightly, necessitating an often unintuitive, lengthy re-initialization step for operators.

To address these issues, we introduced resistance and selfcentering by employing a spring mechanism. The mechanical base of the rudder is built around a ball-bearing joint and a rotational thrust-bearing joint. By employing two joints, we have control over the pitch, roll and yaw axes by using different springs. Springs with different tension allow individual control of the resistance per axis. For absolute position estimates, we attach an HTC Vive tracker to the rudder which receives signals from the VR tracking system that are then translated to movement commands. We place a separator on the rudder's surface to ease blind foot placement.

We equipped the avatar robot with addressable LED strips (90 RGB LEDs) under the base, which indicate the driving direction by illuminating the corresponding side of the robot. The LEDs also show the battery level during charging.

G. Monitoring

Monitoring is an essential part of robust robotics. It allows engineers to analyze problems and find their causes quickly. In our scenario, it was especially important to make sure the system is healthy before starting the run, since from then on, manual intervention was not permitted. During the run, the role of monitoring switches to a safety perspective, allowing the support crew to abort the run in case of danger to the human operator, the robot, or the environment. To be

/otto/sysmon/state		/anna/sysmon/state 🗸 🗸 🗸		Otto				
CPU	Usage 38.94%	Battery	88% 1:59 690.50W	Movement	On	Off		
HDD	Usage 40% (527G free)	CPU	Usage 16.44%		-	- 44		
USB		Temperature	CPU: 73° PCH: 66° SSD: 43°	Send Cmds	On	Off		
Network		HDD	Usage 27% (642G free)					
Index cam	52.0 Hz (delay 0.07s)	USB		Head Control	On	Off		
Mouth cam	56.4 Hz (delay 0.08s)	Ping		ficua control	011	011		
Eye Left	25.7 Hz (delay 0.10s)	Network		Right Hand	On			
Eye Right	26.0 Hz (delay 0.11s)	Basler Left	46.1 Hz (delay 0.10s)	Left Hand	On	Off		
Operator Cam	29.1 Hz (delay 0.09s)	Basler Right	46.8 Hz (delay 0.09s)					
Arm Left TF	Delay: 0.01s	Brio Front	19.4 Hz (delay 0.12s)	Force / Torque				
Arm Left Comm		Brio Rear	20.0 Hz (delay 0.13s)	Otto	On			
Arm Right TF	Delay: 0.01s	Hand Cam	15.0 Hz (delay 0.11s)	Anna Parallanda	0.5			
Arm Right Comm		Hand Left		Anna Feedback	On	Oli		
Glove Left	96.0 Hz (delay 0.06s)	Hand Right	49.1 Hz (delay 0.05s)	Anna Limits	On			
Glove Right	96.1 Hz (delay 0.06s)	Magnet		Atlas				
FT left	921.7 Hz (delay 0.06s)	SVH Contact	193.8 Hz (delay 0.04s)	Acias				
FT right	933.8 Hz (delay 0.06s)	Head	Delay: 0.02s	Drive	On			
Rudder	Ready	Arm Left	Delay: 0.03s	Spine	On	Off		
Pedal	47.6 Hz (delay 0.08s)	Arm Right	Delay: 0.02s					
Eye Tracking	50.1 Hz (delay 0.11s)	FT left	483.4 Hz (delay 0.04s)	Recording				
VR Calibration	Trackers/Arms not working	FT right	483.4 Hz (delay 0.04s)	Record	On			
Audio	Running	Wheels Delay: 0.04s		Rus				
Jamulus Otto	Registered on server	Spine 0.90m (37%)		Kuli				
Jamulus Recording	Recording	Audio	Running	Y Offset		Off		
HDMI	57.8 Hz (delay 0.07s)	Face display						
Bagfile	Recording	E-Stop						
		Paofilo	Percording					

Fig. 6. System Monitoring GUI. Left: Operator Station status. Each line corresponds to a system check. The red check indicates an issue with the VR trackers mounted on the exoskeleton—caused by a support crew member occluding the line-of-sight. Center: Avatar robot status. Right: Control buttons that enable/disable individual system components.



Fig. 7. Network status. Left: Overview with individual network flows. Green boxes indicate hosts in the network. Right: Individual data groups can be configured to use 5 GHz and/or 2.4 GHz streams. The packet/s rates on the right indicate packet drops due to WiFi congestion.

able to monitor the highly complex avatar system with one glance, we developed an integrated GUI. Because it contains a multitude of video streams and complex plots, the standard ROS GUI, rqt, was not suitable, as it is not optimized for high-bandwidth display. Instead, we developed a GUI based on imgui⁵, an immediate-mode GUI toolkit with OpenGL bindings. This allows us to decode and display the video streams directly on the GPU. The GUI follows the rqt paradigm with windows that are arranged via drag & drop.

The most important monitoring display is shown in Fig. 6. Both operator station and avatar robot run a sysmon node, which performs several checks with 1 Hz. These checks range from "Is hardware device X connected?" over "Does component Y produce data?" to "Is the operator station properly calibrated?". The intention is simple: If all checks are successful, the support crew can start the run with confidence. Indeed, our policy was that every time an undetected error or misconfiguration led to a sub-optimal test run, a specific check for this condition was added. Overall, checks are similar to unit tests in software engineering, but monitor the live system in hardware *and* software.

The network subsystem is monitored and configured through two GUI components (see Fig. 7). An overview visualization shows bandwidths and parameters of each individual network connection. For the two WiFi connections, signal strengths are also visualized as green bars. A small control box allows switching data groups between the two

```
<sup>5</sup>https://github.com/ocornut/imgui
```



Fig. 8. Camera streams. Left: Raw wide-angle camera stream (left eye) from the robot. Right: Eye cameras, mouth camera, and reconstructed animated face of the operator.

WiFi connections. This ability allows quick trouble-shooting by giving instant feedback on bandwidths and packet drops.

Finally, a section of the GUI with camera streams (see Fig. 8) together with headsets providing audio feedback give situational awareness to the support crew.

H. System Robustness

Ensuring support crew situational awareness and the connectionless network system are features that make the system more robust. However, there are many problems that can occur during a run, where manual intervention is not possible without aborting the trial. For this reason, we added autorecovery mechanisms on multiple layers.

First, the Franka Emika Panda arms have independent safety systems which detect unsafe situations and either perform a soft-stop (braking with motor power) or hardstop (engaging hardware brakes and switching off motor power). Since the operator can trigger both, e.g. by hitting an object with high speed, it is desirable to recover from these conditions. To this end, we modified the Panda firmware to be able to trigger recovery from an autonomous observer, which restarts the arms as long as the manual E-Stop is not triggered. During restart of the arm, the operator is shown a 3D model of the arm to indicate that the arm is restarting and they should wait until the process is finished. The arm pose is then softly faded to the current operator pose and operation can continue (see [11] for more details).

Secondly, many hard- and software problems can be solved by simply restarting the affected processes [18]. As a simple example, restarting a device driver ROS node will recover after a transient disconnection of the device, without the need to make the driver node itself robust against such events. We stringently use the respawn feature of the ROS launch system to ensure that all nodes are automatically restarted whenever they exit. Watchdogs are integrated that force nodes which do not produce output to exit.

Finally, as a last line of defense, the main control PC is equipped with an external watchdog device. Our software running on the control PC regularly resets this watchdog. Should the system hang completely (which happened once during testing), the watchdog device will force a reset of the computer. Consequently, the software is configured to autostart again, automatically resuming operations. The complete boot-and-recovery process takes less than one minute.

TABLE III Results of the ANA Avatar XPRIZE Finals

			Time		
Rank	Team	Total	Task	Judged	[mm:ss]
1	NimbRo	15.0	10	5.0	05:50
2	Pollen Robotics	15.0	10	5.0	10:50
3	Team Northeastern [19]	14.5	10	4.5	21:09
4	AVATRINA [4]	14.5	10	4.5	24:47
5	i-Botics [5]	14.0	9	5.0	25:00
6	Team UNIST	13.5	9	4.5	25:00
7	Inbiodroid	13.0	8	5.0	25:00
8	Team SNU [6]	12.5	8	4.5	25:00
9	AlterEgo [20]	12.5	8	4.5	25:00
10	Dragon Tree Labs	11.0	7	4.0	25:00
11	Avatar-Hubo [21]	9.5	6	3.5	25:00
12	Last Mile	9.0	5	4.0	25:00

IV. EVALUATION

The ANA Avatar XPRIZE Finals took place in November 2022 in Long Beach, CA, USA. After several downselections over three years, 17 teams from 10 countries competed in the finals for a prize purse of \$8 million. The developed avatar systems were evaluated by untrained operator and recipient judges in a series of ten tasks over one qualification and two competition days. Only the top 16 teams and ties (qualification day) and top 12 teams (first competition day) advanced to the next day. Both judges were selected from the international expert panel and unknown to the teams until 60 min before the competition run. Teams had 45 min to train and familiarize the operator judge with their system. The operator judge controlled the avatar robot located in the arena from the operator control room, out of range of direct visual or auditory feedback. Information could only be exchanged between both locations via the avatar system. The operator had up to 25 min to complete all ten tasks (see Fig. 9). During the competition run, teams were not allowed to interact with the system or the judges.

A. Analysis of Competition Scores

Each competition run was scored based on task performance and judge experience. Table III shows the final scores of the top 12 teams. One point was awarded for each task completed. In addition, up to five points were given based on the experience of both judges. Each judge scored up to one point per criterion if they felt the operator was present at the remote site and if they could clearly see and hear each other through the system. The final point was given by the operator judge if the avatar system was easy and comfortable to use. The maximum score of both competition days counted as the final result. Ties were broken by completion time.

Four systems completed all ten tasks. Pollen Robotics' and our system were the only ones to solve all tasks on both days. Most of the systems received 4.5 or the maximum 5 points for the judge experience. Our team NimbRo won the competition with a perfect score of 15 points and the fastest completion time of 5:50 min—almost twice as fast as the runner-up Pollen Robotics, who also received a perfect score with a time of 10:50 min.



Fig. 9. Tasks of the ANA Avatar XPRIZE finals. ■ T1: Short locomotion (approx. 10 m) to the mission control desk. ■ T2: The operator introduces themselves to the mission commander. ■ T3: The operator receives mission details and confirms the tasks. ■ T4: Activate the power switch. ■ T5: Approx. 40 m of locomotion. ■ T6: Select a canister by weight (approx. 1.2 kg). ■ T7: Place the canister in the designated slot. ■ T8: Navigate around obstacles. ■ T9: Grasp and use the power drill to unscrew the hex bolt. ■ T10: Select a rough-textured stone based on touch and retrieve it.



Fig. 10. Per-task execution time for the top six competition runs solving all ten tasks in the ANA Avatar XPRIZE finals.

B. Task Completion Times

We extracted the per task completion times for both competition days from the official video feed⁶ for the six runs completing all ten tasks (see Figure 10 and Table IV). Both of our competition runs were faster than any other successful run. As our operator judge on Day 1 solved all tasks in 8:15 min, giving us a comfortable lead, our operator judge on Day 2 was instructed to take more risks. In addition, we greatly increased our avatar's maximum base velocity for Day 2, resulting in much faster execution times for all tasks involving larger locomotion (Tasks 1, 4, 5, and 8). We encountered a minor network issue during Task 9 on Day 1 (see Section IV-C) which explains our longer execution time of 1:56 min, compared to 1:04 min on Day 2. All remaining tasks (2, 3, 6, 7, and 10) were solved within the same time (\pm 4 sec.) on both days, showing the robustness of our system.

The shorter tasks 1-3 (locomotion and communication with the recipient judge) and Task 7 (placing the canister into the designated slot) were consistently solved with similar execution times across the top six competition runs. AVATRINA's system had a much slower drive compared to the top three teams, as evidenced by longer execution times for the locomotion tasks. They also had problems during the manipulation in Task 6, which resulted in a software restart, costing 2:10 min. Pollen Robotics' longer locomotion time (Task 5) on Day 2 was similarly due to a reset of the operator control. Larger differences in individual task execution times are due to subsystem failures or sub-optimal

TABLE IV TASK COMPLETION TIMES IN ANA AVAIAD XPRIZE FINALS

TASK COMPLETION TIMES IN ANA AVAIAK AI KIZE TINALS											
Team / Run	T1	T2	Т3	T4	Т5	T6	T7	T8	Т9	T10	Total
Avatrina D1	0:28	0:23	2:03	1:45	3:10	6:17	0:19	2:24	3:10	4:48	24:47
Northeastern D2	0:16	0:19	1:47	0:52	1:14	1:05	0:15	1:00	4:54	9:27	21:09
Pollen Rob. D1	0:10	0:09	1:39	0:40	1:15	0:53	0:14	0:50	5:06	2:24	13:20
Pollen Rob. D2	0:15	0:09	1:43	0:49	2:02	1:15	0:18	0:51	1:28	1:59	10:50
NimbRo D1	0:18	0:10	1:35	0:52	1:00	0:22	0:06	0:50	1:56	1:06	8:15
NimbRo D2	0:08	0:09	1:31	0:23	0:32	0:26	0:09	0:26	1:04	1:02	5:50
NimbRo D2-D1	-0:10	-0:01	-0:04	-0:29	-0:28	+0:04	+0:03	-0:24	-0:52	-0:04	-2:25

We show times of the top six competition run in minutes. D1/D2: Day 1 / Day 2.

grasp poses in case of the drill (Task 9): Both Pollen Robotics on Day 1 and Team Northeastern lost the first drill, requiring grasping the second drill. Team Northeastern then struggled to reach the stones in the box for Task 10, maybe due to their kinematics with the wrist above the hand. The left arm shut down completely due to collision with the top bar and could not recover. However, the operator managed to retrieve the correct stone with the right arm after several attempts.

Despite the small sample size, this time analysis shows the robustness and reliability of our system. Comparing our completion times with the rest of the competition suggests that our system is easier and faster to use and provides sufficient intuitive situational awareness to the operator.

C. System Failures & Recovery

During all three competition runs, the arm recovery system was put to the test. On qualification day, we had switched off most traffic on the 2.4 GHz band since it had proven unstable in the team garages due to channel congestion. During the run, there were certain intervals with higher packet jitter on the 5 GHz band. Since redundancy on the arm commands



Fig. 11. Network latency events during finals. Delays over 100 ms in the arm command channel (red), temporarily put the arms in pause mode. From Day 1 on, the arm commands were transmitted redundantly. Additional logs from Day 1 show video decoding errors resulting from packet loss (green).

was unavailable, the arm controllers on the robot disabled themselves when not receiving commands (see Fig. 11). Thankfully, the packet jitter was only transient, so that autorestarting the controllers was possible. Additionally, during the final task, the left wrist hit the boundary of the stone box with enough force to disable the arm. Pressure on the arm prevented immediate auto-recovery but the operator resolved the situation by lowering the torso, reducing the pressure and allowing the arm to restart.

After experiencing the packet jitter on qualification day, we enabled redundancy for the arm commands, which greatly reduced this error class from seven instances in qualification to two on Day 1 and only one on Day 2. This demonstrates the utility of the redundant communication system.

Arm shutdowns due to excessive force happened once on Day 1 again in the stone box, and on Day 2 while putting a canister back on the table. In both cases, auto-recovery immediately succeeded and the run could continue without problems.

D. Lessons Learned

1) Robustness: Despite the extensive experience available in our group, this is the most complex system we ever built. Good monitoring, failure tolerance, and auto-recovery were extremely important for success. In contrast to other teams, technical problems did not cause major delays or failures.

2) Frequent testing under competition conditions: Having a high test frequency allowed us to identify and address several issues that were annoying or uncomfortable for operators, e.g. resulting in the improved rudder device. Frequent tests also helped the support crew to establish routine in efficiently training the operators.

3) 1:1 correspondence is best: The connection between operator and avatar needs to be as close to identity as possible. Avoiding any scaling, offsetting or 3D processing helps operators to quickly immerse into the system. In particular, correct hand-eye transformations let operators identify the avatar's hands as their own.

V. CONCLUSION

We presented the extended and updated NimbRo avatar system, which won the ANA Avatar XPRIZE finals. Key improvements, compared to the semifinals system [2], such as a new base design, a linear actuator to adjust the torso height, haptic perception, monitoring tools, failure tolerance, and robust wireless communication enabled this success.

REFERENCES

- E. Krotkov, D. Hackett, L. Jackel, M. Perschbacher, J. Pippine, J. Strauss, G. Pratt, and C. Orlowski, "The DARPA robotics challenge finals: Results and perspectives," *The DARPA Robotics Challenge Finals: Humanoid Robots To The Rescue*, pp. 1–26, 2018.
- [2] M. Schwarz, C. Lenz, A. Rochow, M. Schreiber, and S. Behnke, "NimbRo Avatar: Interactive immersive telepresence with forcefeedback telemanipulation," in *International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 5312–5319.
- [3] R. Luo, C. Wang, C. Keil, D. Nguyen, H. Mayne, S. Alt, E. Schwarm, M. Evelyn, T. Padır, and J. P. Whitney, "Team Northeastern's approach to ANA XPRIZE Avatar final testing: A holistic approach to telepresence and lessons learned," *arXiv:2303.04932*, 2023.
- [4] J. M. Marques, N Patrick, Y. Zhu, N. Malhotra, and K. Hauser, "Commodity telepresence with the AvaTRINA nursebot in the ANA Avatar XPRIZE semifinals," in *RSS Workshop: Perspectives on the ANA Avatar XPRIZE Competition*, 2022.
- [5] J. B. Van Erp, C. Sallaberry, C. Brekelmans, D. Dresscher, F. Ter Haar, G. Englebienne, J. Van Bruggen, J. De Greeff, L. F. S. Pereira, A. Toet, *et al.*, "What comes after telepresence? Embodiment, social presence and transporting one's functional and social self," in *Int. Conference on Systems, Man, and Cybernetics (SMC)*, 2022.
- [6] B. Park, J. Jung, J. Sim, S Kim, J Ahn, D. Lim, D. Kim, M. Kim, S. Park, E Sung, *et al.*, "Team SNU's avatar system for teleoperation using humanoid robot: ANA Avatar XPRIZE competition," in *RSS Workshop: Perspectives on the ANA Avatar XPRIZE Comp.*, 2022.
- [7] K. Takeuchi, Y. Yamazaki, and K. Yoshifuji, "Avatar work: Telework for disabled people unable to go outside by using avatar robots," in *Companion of Int. Conference on Human-Robot Interaction*, 2020.
- [8] O. Khatib et al., "Ocean one: A robotic avatar for oceanic discovery," IEEE Robotics & Automation Magazine, 2016.
- [9] N. Y. Lii, D. Leidner, P. Birkenkampf, B. Pleintinger, R. Bayer, and T. Krueger, "Toward scalable intuitive telecommand of robots for space deployment with METERON SUPVIS Justin," in *Symp. on Adv. Space Tech. for Robotics and Automation (ASTRA)*, 2017.
- [10] S. Tachi, Y. Inoue, and F. Kato, "TELESAR VI: Telexistence surrogate anthropomorphic robot VI," *International Journal of Humanoid Robotics*, vol. 17, no. 05, p. 2050019, 2020.
- [11] C. Lenz and S. Behnke, "Bimanual telemanipulation with force and haptic feedback through an anthropomorphic avatar system," *Robotics and Autonomous Systems*, vol. 161, p. 104 338, 2023.
- [12] M. Schwarz and S. Behnke, "Low-latency immersive 6D televisualization with spherical rendering," in *International Conference on Humanoid Robots (Humanoids)*, 2021.
- [13] A. Rochow, M. Schwarz, M. Schreiber, and S. Behnke, "VR facial animation for immersive telepresence avatars," in *International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [14] A. Rochow, M. Schwarz, and S. Behnke, "Attention-based VR facial animation with visual mouth camera guidance for immersive telepresence avatars," in *International Conference on Intelligent Robots* and Systems (IROS), 2023.
- [15] XPRIZE Inc., "ANA Avatar XPRIZE rules and regulations," 2022.
- [16] B. Pätzold, A. Rochow, M. Schreiber, R. Memmesheimer, C. Lenz, M. Schwarz, and S. Behnke, "Audio-based roughness sensing and tactile feedback for haptic perception in telepresence," in *Int. Conference on Systems, Man, and Cybernetics (SMC)*, 2023.
- [17] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "ROS: An open-source robot operating system," in *International Conference on Robotics and Automation (ICRA) Workshop on Open Source Robotics*, 2009.
- [18] G. Candea, S. Kawamoto, Y. Fujiki, G. Friedman, and A. Fox, "Microreboot - a technique for cheap recovery," in *Symposium on Operating Systems Design and Implementation (OSDI)*, 2004.
- [19] R. Luo, C. Wang, E. Schwarm, C. Keil, E. Mendoza, P. Kaveti, S. Alt, H. Singh, T. Padir, and J. P. Whitney, "Towards robot avatars: Systems and methods for teleinteraction at Avatar XPRIZE semifinals," in *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2022.
- [20] G. Lentini, A. Settimi, D. Caporale, M. Garabini, G. Grioli, L. Pallottino, M. G. Catalano, and A. Bicchi, "Alter-Ego: A mobile robot with a functionally anthropomorphic upper body designed for physical interaction," *IEEE Robotics & Automation Magazine*, 2019.
- [21] J. C. Vaz, A. Dave, N. Kassai, N. Kosanovic, and P. Y. Oh, "Immersive auditory-visual real-time avatar system of ANA Avatar XPRIZE finalist Avatar-Hubo," in *IEEE International Conference on* Advanced Robotics and Its Social Impacts (ARSO), 2022.