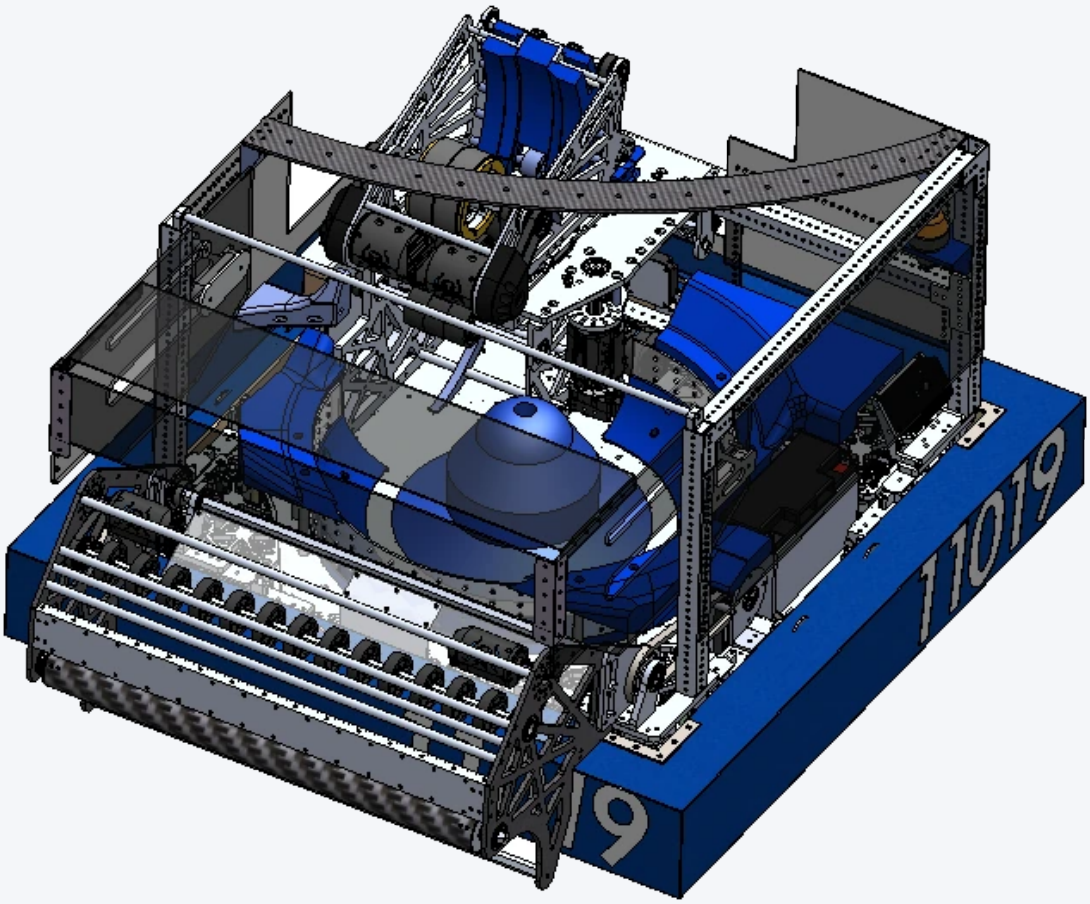
CELESTIAL PIVOT



TECHNICAL MANUAL

FRC Team 11019 Xplore
2026 REBUILT / English edition

Mechanical architecture, moving-shot control, timestamp-aligned localization and autonomous operation.

Contents

03 / Foreword

Development process and the 2026 robot.

04 / Game analysis

Match phases, scoring and the shooting-focused strategy.

05-11 / Mechanical design

Robot layout; design principles; drivetrain; hopper; intake; indexer and rotator; shooter and turret.

12-15 / Moving-shot control

VISTAeroSpin-5D; vacuum trajectories and vision identification; aerodynamic correction and polynomial fitting; velocity compensation and the low-apex comparison.

16-18 / Localization

Capture-time fusion; dual-camera selection; observation quality; CANCoder steering reference; Pigeon 2 and high-rate odometry.

19 / Autonomous operation

Choreo trajectories, closed-loop following, route mirroring and parallel mechanism scheduling.

20 / Document notes

Source, notation and limitations.

Edition note

This English editorial edition follows the supplied binder page by page. Wording and layout have been adapted for clarity. Reported design values are retained; they are not independently verified measurements. Original source files have not been changed.

Foreword



From game analysis to an integrated robot

The technical binder records the game analysis, decisions, results and technical design that guided the 2026 REBUILT robot. The season began with match analysis to establish priorities and robot requirements.

Subsystem development

Those requirements informed subsystem prototypes, development and fabrication. Mechanical systems and software were then integrated for competition operation.

Celestial Pivot

Team 11019 Xplore presents its 2026 robot: Celestial Pivot. The platform combines swerve mobility, a low-mounted intake, expandable storage, rotary feeding and an independently aiming turret.

Game analysis

Match structure recorded in the binder

Total duration: 2:40. Autonomous: 20 seconds. Teleop: 2:20, including the final 30-second end game. Teleop includes alternating shifts. During those shifts, one HUB is active at a time; the binder estimates approximately 60-70 seconds of scoring opportunity per alliance. The HUB opening is 72 inches high.

Scoring reference

FUEL into an active HUB: 1 point in autonomous or teleop.
TOWER Level 1: 15 points in autonomous; 10 in teleop/end game.
TOWER Level 2: 20 points in teleop/end game.
TOWER Level 3: 30 points in teleop/end game.

Ranking-point targets recorded in the binder

Qualification win: 3 RP.
Energized: 100 FUEL into the HUB, 1 RP.
Supercharged: 360 FUEL into the HUB, 1 RP.
Traversal: at least 50 climbing points, 1 RP.

Strategy trade-off

The team estimates that a Level 3 climb consumes 15-25 seconds, during which shooting could contribute 7-25 points. Against a 30-point climb, that leaves a nominal advantage of 5-25 points before accounting for failed engagement, slipping or timeout risk.

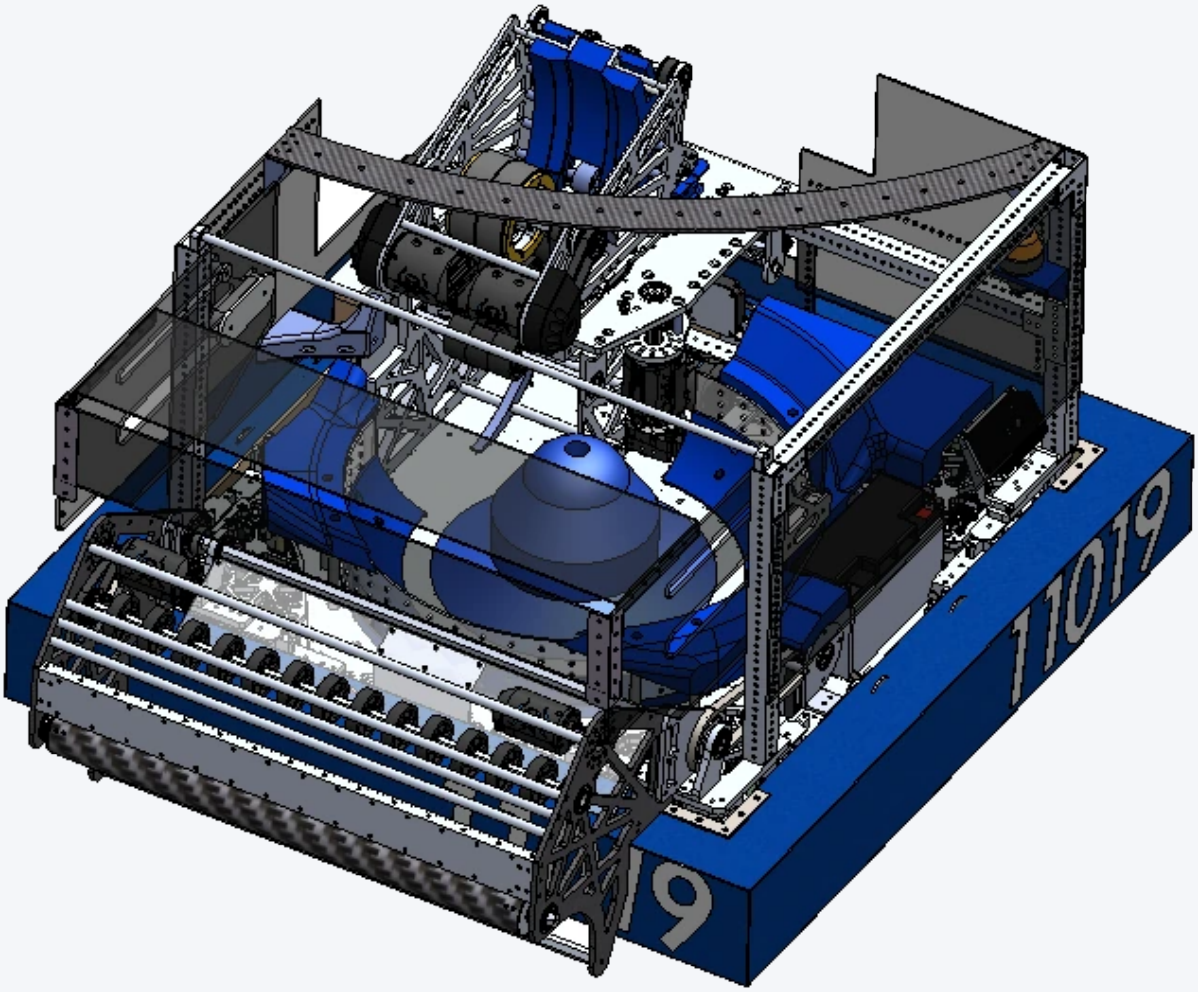
Mechanical implications

The binder estimates a climber adds 5-10 lb and occupies space otherwise available to intake and shooter systems. Winches, ropes and limit devices add integration demands; rearward mass can affect chassis dynamics. The stated strategy therefore prioritizes sustained shooting.

Scope

These are the supplied binder's game analysis and design estimates, not a current rules verification or a guarantee of scoring performance.

Robot architecture



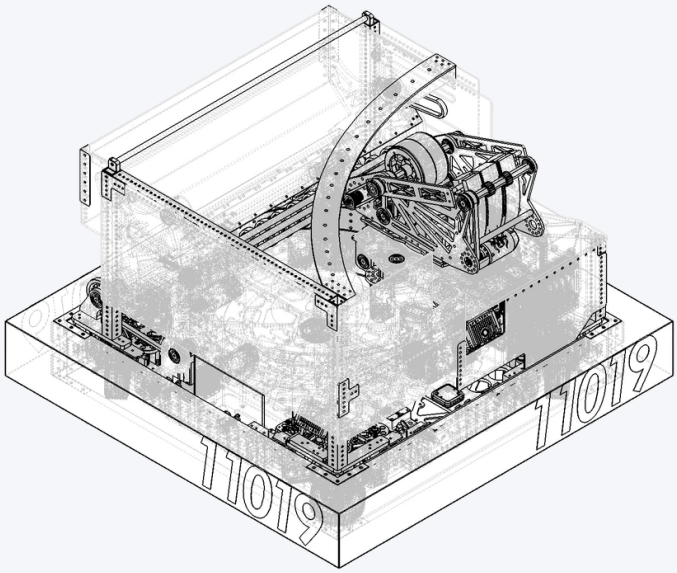
Integrated subsystems

DRIVETRAIN / Four-module swerve drive
INTAKE / Low-mounted dual rollers
HOPPER / Expandable ball storage
INDEXER + ROTATOR / Controlled ball transfer
SHOOTER + TURRET / Flywheel launch and independent aiming

CAD reference

The assembly image is reproduced from the supplied technical binder. It represents documented design geometry; the photograph archive records the assembled robot on the field.

Mechanical design



Design overview

The mechanism architecture addresses collection, storage, feeding and shooting within the REBUILT match cycle. A single rotating turret carries the shooter; the dual-roller intake is positioned at the front and low on the chassis. The turret provides a full-circle aiming range without requiring a matching chassis heading.

Modularity

Subsystems can be removed and serviced independently, supporting rapid inspection and repair between matches.

Materials and mass distribution

Carbon-fiber tubes and machined aluminum components are used where compatible with structural stiffness requirements. The design aims to limit mass and keep the center of gravity low.

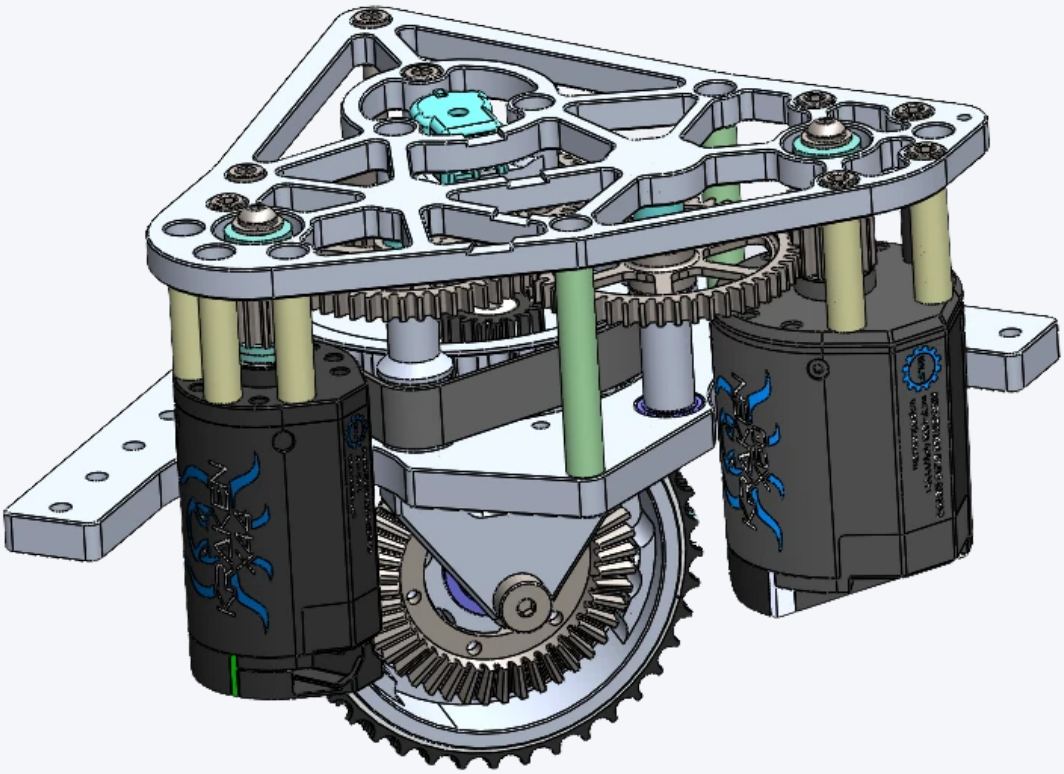
Bearing support and location

Key transmission paths use dual-bearing supports and flange location to establish repeatable mechanical alignment.

Source inconsistency

The source overview also mentions a rear single-hook climber, while page 4 states a shooting-focused strategy. This edition does not infer whether that climber was fitted to the final competition configuration.

Swerve drivetrain



Drive architecture

The team-built drivetrain uses four-inch MK5R swerve modules. Each module provides independent wheel drive and steering.

Drive train

Motor: Kraken X60
Drive reduction: 6.75:1
Specified maximum speed: 4.58 m/s

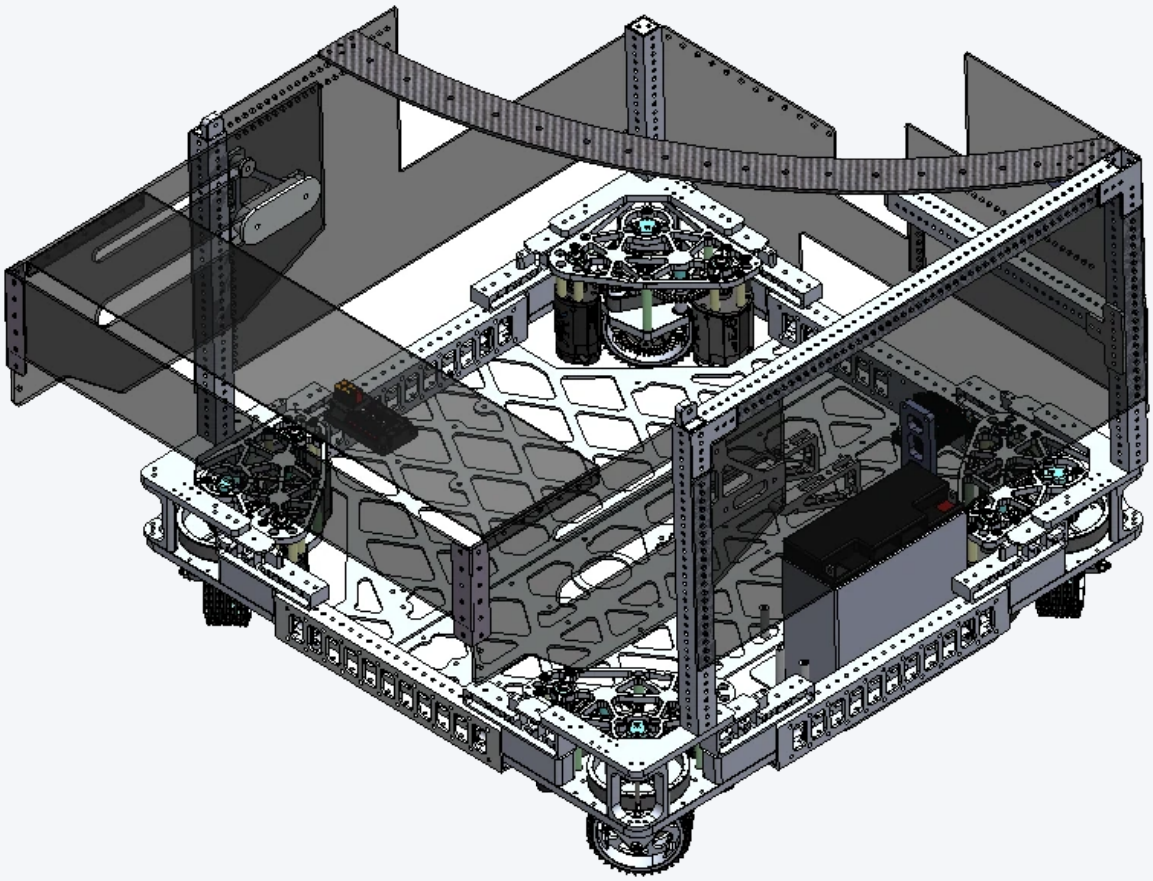
Steering

Motor: Kraken X44
Steering reduction: 24:1
Absolute steering reference: CANCoder

Specification note

The speed is the value stated in the technical binder, not an independently measured top-speed result. Absolute steering feedback and estimator integration are described on pages 17-18.

Expandable hopper



Storage and extension

A sliding mechanism provides 260 mm of extension travel, increasing the available ball-storage volume.

Construction

Smooth polycarbonate inner walls reduce ball friction and help avoid jams. The source identifies 3080 and 5050 aluminum supports mounted directly to the chassis.

Ball retention

Netting at the top limits balls escaping from the hopper during operation. The hopper interfaces with the low-mounted intake and downstream feeding mechanism.

Ground intake



Dual-roller arrangement

Two 35 mm carbon-tube rollers collect FUEL near floor level. The rollers are placed to support continuous intake across the collection opening.

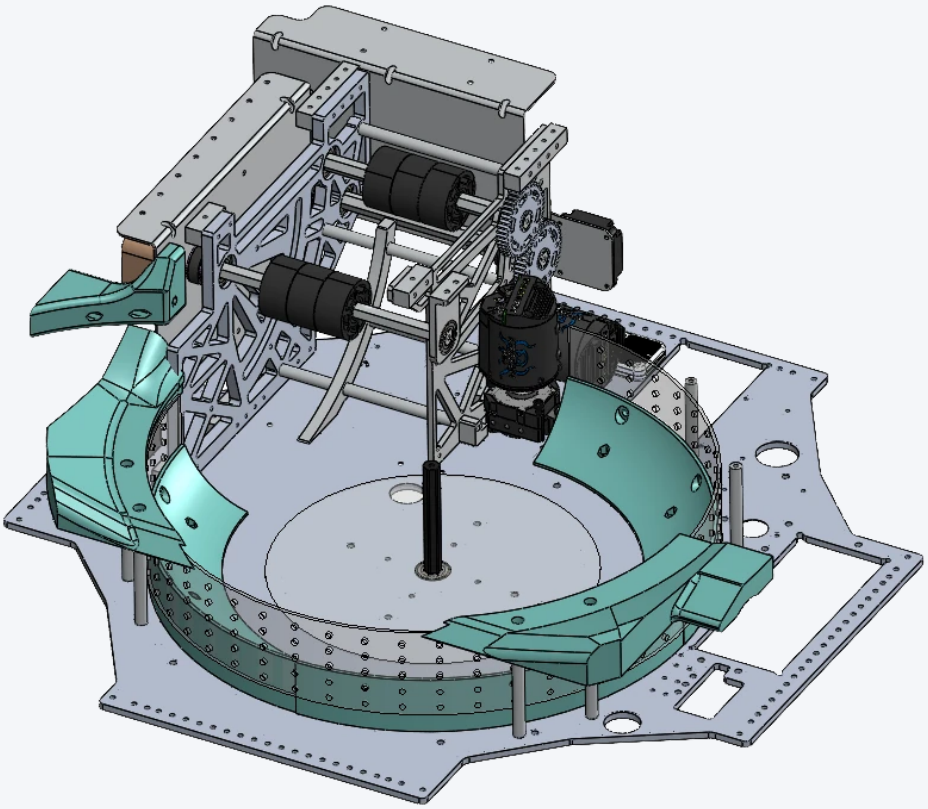
Drive specification

2 x Kraken X44 motors
1:3 gear ratio, as stated in the source
35 mm roller diameter

Integration

The intake is mounted low at the front of the robot and aligned with the hopper entrance. Its collection commands can run alongside chassis motion during autonomous paths.

Indexer and rotator



Ball transfer

Two Kraken X60 motors drive the indexer and rotator. The source specifies gearing of 1:10 and 1:1.5 for the two mechanisms.

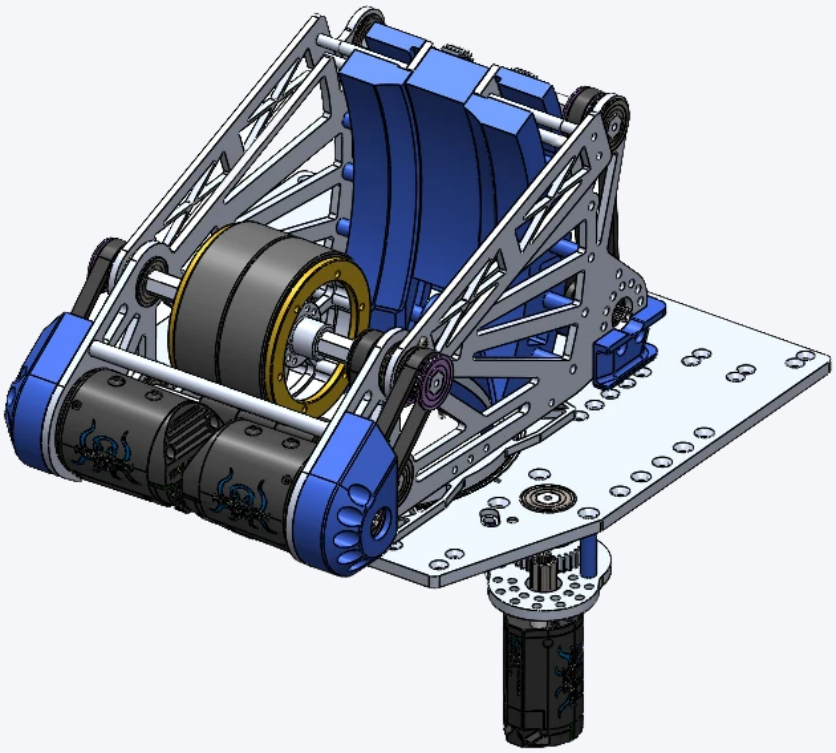
Construction

The rotator combines ABS / PLA printed parts with internal aluminum reinforcement. Printed guides assembled around the perimeter keep FUEL away from dead zones where balls could become trapped.

System role

The assembly transfers stored balls toward the shooter. Feeding is scheduled with intake, turret, hood and flywheel commands as part of the robot control sequence.

Shooter and turret



Flywheel drive

Two Kraken X60 motors drive metal-core, high-friction 50A polyurethane flywheels through 1:1 gearing.

Hood adjustment

The hood has a 25-degree adjustment range. Its backplate is printed in PLA-CF. Hood and flywheel setpoints support the required scoring or passing trajectory.

Turret actuation

The turret has a 360-degree rotation range. One Kraken X44 drives it through a 126-tooth belt pulley.

Control interface

The ballistic model determines the required launch state. The control system converts that demand into turret position, hood position and flywheel setpoints; the mechanism remains responsible for reproducing those setpoints.

VISTAeroSpin-5D

Design rationale

The team compared its approach with distance-to-hood/flywheel lookup interpolation and preset time-of-flight compensation. The binder identifies limited sample density and nonlinear changes between calibration points as constraints, particularly when chassis velocity, drag, spin and mechanism losses vary.

Earlier iteration

An earlier approach repeated the flight-time calculation 20 times to converge the motion-compensation solution. The team reports that this did not address the underlying limitations of lookup interpolation sufficiently.

Physical inputs

The selected system starts from target position, robot state and a physical trajectory, rather than using a table of previously successful shots as the primary model or forcing the model to match such a table.

Name and architecture

VISTA: Visual Spatiotemporal Tracking and Adaptation.
AeroSpin: aerodynamic drag and Magnus-effect compensation.
5D: fifth-order polynomial real-time controller.

Five stages

1. High-speed visual detection.
2. Camera-calibrated trajectory reconstruction.
3. Aerodynamic physical simulation.
4. Fifth-order model approximation.
5. Real-time shooting outputs on roboRIO.

Operating scope

The design objective is continuous prediction across distance and robot velocity. This document does not establish validated operating bounds, raw calibration data or a measured accuracy guarantee.

Trajectory and vision

Vacuum initial solution

For horizontal distance dx, target height difference dz and a selected apex height hmax measured from the release point, a vacuum trajectory provides an initial launch state. The source repeats dx for the height difference; dz is used here to distinguish the vertical quantity.

Launch-state relationships

vz = sqrt(2 g hmax)
t = [vz + sqrt(vz^2 - 2 g dz)] / g
vx = dx / t
v0 = sqrt(vx^2 + vz^2)
launch angle = atan2(vz, vx)
These relationships select the descending target crossing and require a physically reachable apex.

Visual identification

A high-speed perception module based on a deep neural network, with manual refinement, detects the FUEL ball frame by frame and associates detections over time. Camera calibration converts image coordinates into spatial trajectory observations.

What is identified

The trajectories describe initial speed, launch angle, apex height, descent shape and abnormal collisions. They support physical identification rather than forming a table of successful shot settings.

Figure note

The website includes an English-label trajectory illustration. Its generatively translated plotting marks are not a substitute for the original video, calibration or numerical samples.

Aerodynamic model

Beyond the vacuum solution

Drag reduces ball speed during flight. Flywheel and hood contact introduce spin; the Magnus effect can alter lift or lateral motion. Mechanical energy loss and angle-mapping errors also affect the required launch state.

Velocity corrections

vx,required = vx,vacuum + delta_vx(dx, hmax)
vz,required = vz,vacuum + delta_vz(dx, hmax)
The correction terms come from offline numerical simulation rather than manual lookup.

Offline computation

The binder reports approximately 534 million underlying physical simulation calculations across combinations of distance, apex height, launch speed, launch angle, drag and spin parameters. This figure is a simulation count, not a count of physical shots.

Fifth-order approximation

The high-fidelity integration model is approximated using fifth-order polynomials with horizontal distance dx and apex hmax as inputs:
delta_v = sum c_{ij} dx^{i} $hmax^{j}$, for i + j <= 5
This reduces repeated integration inside the roboRIO control loop.

Model selection

Higher-order terms can represent coupled changes in distance, apex height and velocity correction. The team compared third- and fifth-order fit residuals and selected fifth order to balance expressive capacity and generalization, rather than increasing the order indefinitely.

Evidence limitation

Raw fitting samples, coefficients and an independently validated error envelope were not supplied with this website. English-label fit figures are translated illustrations; their plotted marks must not be used for numerical analysis.

Moving-shot compensation

Field-relative velocity

Flight time is solved from the current physical trajectory, not a fixed TOF table. The required ball velocity is first expressed in field coordinates. Current chassis translation is then subtracted:
v_launch/robot = v_required/field - v_chassis/field
The result is converted to executable turret, hood and flywheel setpoints. Compensation varies with speed, distance and trajectory state.

Low-apex comparison reported in the binder

The source compares an apex parameter of 2.2 m with 4.2288 m in a vacuum model. Conditions: target height difference 1.2288 m, ball mass 0.27 kg, horizontal range 1-5 m.

Reported theoretical results

Initial kinetic energy reduction: 30.8-47.2%.
Flight time: approximately 1.71 s to 1.12 s.
Flight-time reduction: approximately 34.8%.
Near-range instantaneous flywheel replenishment demand: up to approximately 47% lower in the theoretical comparison.

Limits of the comparison

The energy calculation excludes aerodynamic drag, mechanism losses and field disturbances. It is not a measured energy-use or hit-rate result. The numeric values are reproduced as reported, rather than presented as a new independent validation.

Website demonstration

The simplified interactive example uses release height 0.65 m, target height 1.84 m, apex 2.2 m and g = 9.81 m/s^2. It sums ascent and descent times and sets lateral drift to speed multiplied by flight time. These presets differ from the binder comparison and do not run the complete physical model.

Capture-time vision fusion



From observation to current pose

A camera observation may describe a robot position tens to hundreds of milliseconds before its arrival at roboRIO. Its timestamp identifies the image-capture time. The estimator locates that time in odometry history, applies the vision correction there, then replays wheel and gyro increments to the present.

Complementary measurements

Vision corrects the historical field position. Odometry propagates the corrected state forward. This avoids directly treating a delayed observation as a measurement of the current position.

Two independent Limelight inputs

Cameras mounted at different positions observe AprilTags from redundant viewpoints. Each independently supplies raw timestamped NetworkTables pose samples.

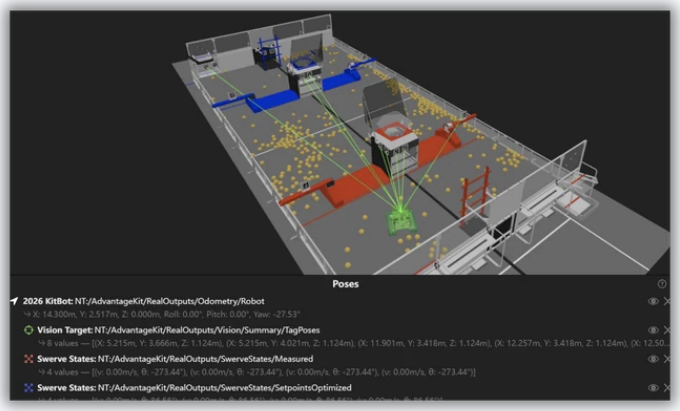
MegaTag selection

MegaTag 1 and MegaTag 2 may originate from the same camera frame. The system prefers MT2, falling back to MT1 only when MT2 is unavailable for that camera. A frame is not fused twice.

Avoiding duplicated gyro information

MT2 uses Pigeon 2 heading to help resolve single-tag ambiguity. The estimator gives useful weight to its translation but does not feed MT2 visual heading back into the estimator, avoiding reuse of the same gyro information.

Observation quality



Adaptive uncertainty

Vision standard deviation is adjusted for mean tag image area, chassis linear velocity, chassis angular velocity and per-camera calibration. Smaller tag area or faster motion increases uncertainty. Each camera has its own calibration coefficient.

Rejection conditions

Reject observations with no tags, high single-tag ambiguity, abnormal pose height, a location outside the field or excessive rotation during MT2 processing. Recorded thresholds: ambiguity no greater than 0.8; height deviation no greater than 0.75 m; MT2 gyro angular speed no greater than 360 degrees/s.

Logging, simulation and replay

The vision implementation uses an AdvantageKit IO boundary. The physical robot uses Limelight IO; simulation uses PhotonVision cameras. AdvantageScope replays the logs. Per-camera raw inputs, tags, candidate poses, accepted/rejected outcomes, rejection reasons and innovations are recorded.

CANCoder absolute steering reference

Each swerve module uses a 12-bit CANCoder: 4096 discrete positions per revolution. This establishes the module orientation at startup and combines with the TalonFX steering loop through FusedCANcoder feedback. A consistent absolute zero keeps all four modules in the same mechanical reference frame after a restart.

Odometry and estimation

Pigeon 2 inertial measurements

Pigeon 2 supplies yaw, yaw rate, roll, pitch and three-axis acceleration. The implementation includes these signals in the Phoenix high-rate sampling thread and applies calibrated heading scale factors to yaw and yaw rate.

Timestamped odometry

The reported sampling rate is up to 250 Hz on CAN FD and 100 Hz otherwise. Each sample is stamped using FPGA time minus average CAN latency. Four drive positions, four steering positions and gyro heading enter swerve kinematics on a common timeline.

Between main-loop updates

High-rate samples reconstruct chassis movement between 20 ms main-loop updates rather than relying only on the main loop to represent intervening motion.

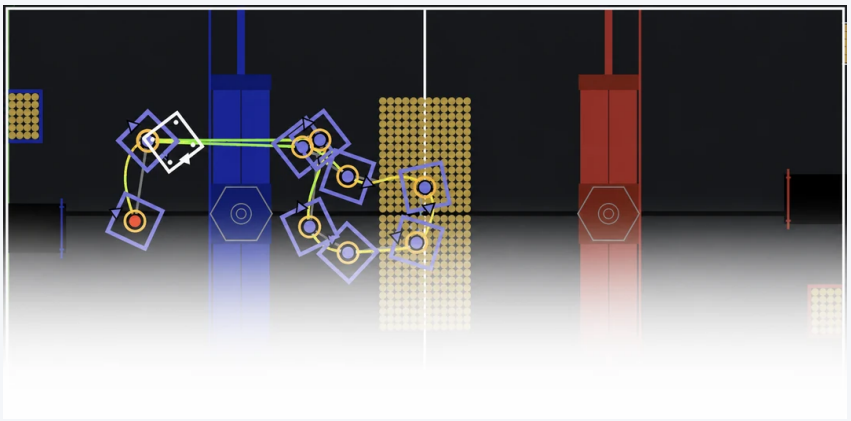
Covariance-weighted fusion

WPILib SwerveDrivePoseEstimator supplies the Kalman-like fusion layer. Wheel travel and gyro rotation provide frequent motion prediction. AprilTag observations provide less frequent corrections in absolute field coordinates.

Independent uncertainty models

The inputs are not equally averaged. State uncertainty is assigned to the motion model, while observation uncertainty is calculated for each vision frame. Large tags and steady movement permit stronger visual correction; distant tags or rapid movement reduce visual weight. Delayed corrections are applied at capture time and propagated forward.

Autonomous coordination



Offline planning and closed-loop following

Choreo trajectories contain field pose, translation velocity, angular velocity, acceleration and module wheel-force demands. Optimization runs before the match. At runtime, the PathPlanner holonomic controller compares the target trajectory with estimated pose and converts position and heading errors into chassis velocity commands.

Initialization and mirroring

Before an autonomous routine, turret, hood and extension home in parallel. Odometry initializes from the first Choreo segment's starting holonomic pose. Consistent transforms mirror field position, heading, lateral velocity and angular velocity across alliance or left/right starting configurations.

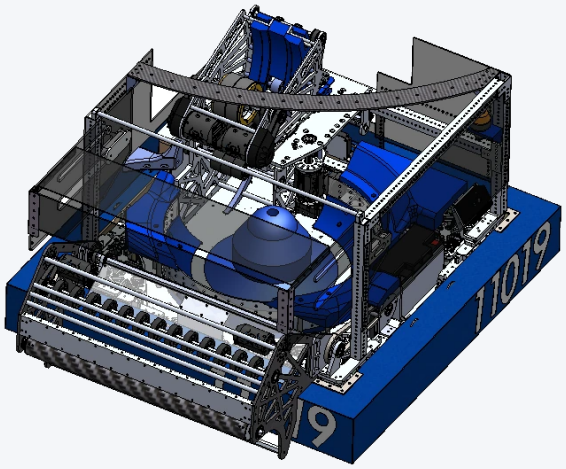
Parallel mechanism scheduling

Path segments define stage boundaries. Intake, Indexer, Extension, Turret, Hood and Flywheel commands run alongside them. Collection mechanisms deploy while driving toward FUEL. Turret tracking and flywheel preparation overlap transit toward the scoring area. Segments form repeated collection, transfer and scoring cycles.

Deterministic transitions and replay

The path command determines stage duration. When it ends, the prior task is stopped and the next stage begins in a defined order. AdvantageKit records active trajectories, target and estimated poses, and command starts, completions and interruptions. AdvantageScope aligns motion and mechanism events to distinguish path design, tracking and timing errors.

Xplore continues.



Source and edition

Based on Team 11019's supplied 2026 Sanya technical binder. This English edition retains the original 20-page topic sequence for website references. It is an editorial adaptation, not a new team-issued engineering certification.

Images

Robot photographs and CAD views come from the supplied source. English-label charts shown on the website were localized using generative image editing and are clearly identified as translated illustrations, not raw numerical evidence. The two material studies are AI-generated and do not depict actual robot parts.

Technical boundaries

Mechanism dimensions, gear ratios, thresholds and performance figures are reported source values. Theoretical vacuum-model comparisons are not measured hardware results. Raw control coefficients, CAD manufacturing validation and physical performance testing are outside this document.

Notation and source issues

The repeated horizontal-distance symbol in the source trajectory discussion is distinguished as dx and dz here. The page 4 shooting strategy and page 6 climber mention are not treated as proof of a final climber configuration. Approximate energy and timing values are preserved as reported.

TEAM 11019 / XPLORE

CELESTIAL PIVOT
2026 FRC REBUILT