


InsideGNSS

Published by Autonomous Media

GPS | GALILEO | GLONASS | BEIDOU



BOOM!

In-Air Refueling from a Pilotless Drone



ROBUST INTEGRITY | Multi-Sensor Error Characterization

WORKING PAPERS | GNSS Interference Mitigation

WASHINGTON VIEW | Homeland Security Analyzes PNT Threats



L3HARRIS.COM

YOU NEED M-CODE. WE HAVE IT NOW.

At L3Harris, we bring unparalleled innovation to mission-critical solutions across all domains. That includes delivering M-code capabilities for GPS; capabilities like enhanced acquisition, jamming and spoofing immunity, and security features that improve accuracy, authenticity and integrity in contested environments. L3Harris M-code GPS technologies have been independently verified by the U.S. Government and are the first available for integration, test and fielding – all at a competitive price.

We have the M-code solution you need – now.
Discover more at [L3Harris.com/m-code](https://www.l3harris.com/m-code)

USE OF D.O.D. VISUAL INFORMATION DOES NOT
IMPLY OR CONSTITUTE D.O.D. ENDORSEMENT



L3HARRIS™
FAST. FORWARD.

Spirent GSS9000 Series The Ultimate GNSS Simulator

Simulate truly realistic models and trajectories with the industry's first 2 kHz update rate. Be confident your products are ready to succeed in the real world.

VISIT US

spirent.com/pnt

spirentfederal.com

for U.S. Gov/Defense



Spirent™

spirent™
Federal Systems



CONTENTS

JULY/AUGUST 2021 VOLUME 16 NUMBER 4

Published by **Autonomous Media**

ON THE COVER

32 Real-Time Automated Aerial Refueling with Stereo Vision

Overcoming GNSS-Denied Environments

In or Near Combat Areas

James Anderson, Joel Miller, Xiaoyang Wu, Scott Nykl, Clark Taylor and Warren Watkinson

Table of Contents BY THE NUMBERS

EDITORIAL

10 PNT Fulfills Autonomy Promise

ARTICLES

32 Real-Time Automated Aerial Refueling

46 Working Papers: GNSS Interference Mitigation

60 Sensor-Agnostic All-Source Residual Monitoring

DEPARTMENTS

12 360 Degrees

14 Washington View

24 GNSS Solutions

56 State of Play

66 Advertisers Index

66 GNSS Timeline


WORKING PAPERS

46 GNSS Interference Mitigation

Modulations, Measurements and Position Impact

Daniele Borio and Ciro Gioia

This column analyses five interference mitigation techniques, including the Adaptive Notch Filter (ANF) and Pulse Blanking (PB), evaluating their impact on pseudoranges and on the final position and timing solution.



46

60 Evaluation of Sensor-Agnostic All-Source Residual Monitoring for Navigation

Andrew Appleget, Robert C. Leishman and Maj Jonathon Gipson

Unlike two-sensor systems such as GPS-inertial integration, systems of three or more sensors present the problem of ambiguity as to which sensor is adversely affecting the solution. A robust framework is needed to maintain navigation integrity despite the additional sensor modalities.



UNPRECEDENTED PERFORMANCE AT YOUR FINGERTIPS

Introducing the all new Tactical Embedded line.
The best just got smaller.



Tactical-Grade IMU
Heading: 0.05° - 0.1°
Pitch/Roll: 0.015°
GNSS: L1/L2/E1/E5
with RTK/PPK



vectornav.com
+1.512.772.3615

EDITORIAL

10 The Time and the Place

PNT Fulfills Autonomy Promise


COLUMNS



14 Washington View

PNT: Nothing to See?
By Dawn Zoldi

57



With a modular open systems approach

From State of Play: A modular open systems approach (MOSA) means greater availability of suitable PNT replacements.

STATE OF PLAY

56 Resilient PNT for Critical Applications

We need a measured, cost-effective response commensurate with the level of threats and possible consequences.

Logan Scott

DEPARTMENTS

12 360 Degrees

24 GNSS Solutions

How to mitigate a spoofing signal while tracking it for intent analysis?

To detect and mitigate a spoofing attack, it is important to know where a spoofing signal comes from and what it tries to accomplish for its ultimate neutralization.

66 Advertisers Index

66 GNSS Timeline

Calendar of Events

MIDS[®] - Mobile Interference Detection System

Detect **Interference & Jamming signals** with our GNSS signal recorder and our GNSSA-LHG (Low Horizon Gain Antenna) as well as **Spoofing signals** with our GNSSA-DCP (RHCP/LHCP Dual Circularly Polarized Antenna)
Full GNSS L-Band coverage
Receive up to two GNSS bands (e.g. L1/E1/G1/B1 and/or L5/E5/G3/B2 signals)

GNSS Resilience
The Future begins here!



Technology licensed by Fraunhofer IIS
<https://teleorbit.eu>





AHEAD OF WHAT'S POSSIBLE™

IMU Solutions for the Harshest Conditions

Analog Devices delivers IMU sensor solutions that maintain high precision even while operating in the harshest conditions. Featuring robust accuracy, full calibration, and plug and play functionality, ADI's IMUs provide the level of navigation and stabilization you need to deploy complex, high performance designs.



Small Size for
Flexibility



Unparalleled
Speed to Market



Robust, Reliable,
Repeatable Performance

Find your IMU at analog.com/IMU

Editorial Advisory Council

VIDAL ASHKENAZI

Nottingham Scientific Ltd., Nottingham, United Kingdom

JOHN BETZ

MITRE Corporation, Bedford, Massachusetts, USA

PASCAL CAMPAGNE

France Développement Conseil, Vincennes, France

MARIO CAPORALE

Italian Institute of Navigation, Rome, Italy

MARCO FALCONE

European Space Agency, Noordwijk, The Netherlands

SERGIO GRECO

Thales Alenia Space, Rome, Italy

JEAN-LUC ISSLER

CNES, Toulouse, France

CHANGDON KEE

Seoul National University, Seoul, Korea

MIKHAIL KRASILSHCHIKOV

Moscow Aviation Institute, Moscow, Russia

SANG JEONG LEE

Chungnam National University, Daejeon, Korea

MARCO LISI

ESA, Belgium

JULES MCNEFF

Overlook Systems Technologies, Inc., Vienna, Virginia, USA

PRATAP MISRA

Tufts University, Medford, Massachusetts, USA

BRAD PARKINSON

Stanford University, Palo Alto, California, USA

TONY PRATT

Professor and Consultant, United Kingdom

SERGEY G. REVNIVYKH

ISS Reshetnev, Zheleznogorsk, Russian Federation

MARTIN RIPPLE

Frequentis AG, Australia

CHRIS RIZOS

University of New South Wales, Sydney, Australia

TOM STANSELL

Stansell Consulting, Rancho Palos Verdes, California, USA

JACK TAYLOR

The Boeing Company, Colorado Springs, Colorado USA

JÖRN TJADEN

European Space Agency, Noordwijk, The Netherlands

A.J. VAN DIERENDONCK

AJ Systems, Los Altos, California, USA

FRANTISEK VEJRAZKA

Czech Technical University, Prague, Czech Republic

PHIL WARD

Navward Consulting, Garland, Texas, USA

CHRISTOPHER K. WILSON

Vehicle data and technology consultant, California, USA

LINYUAN XIA

Sun Yat-Sen University, Guangzhou, China

AKIO YASUDA

Tokyo University of Marine Science and Technology, Tokyo, Japan

InsideGNSS

GPS | GALILEO | GLONASS | BEIDOU

ENGINEERING SOLUTIONS FROM THE GLOBAL NAVIGATION SATELLITE SYSTEM COMMUNITY

July/August 2021 Volume 16/Number 4

Published by Autonomous Media

EDITORIAL

Editor-in-Chief **Alan Cameron** alan@insidegnss.com

Editor Emeritus **Glen Gibbons** glen@insidegnss.com

Creative Director **Christine Waring**

Contributing Editor for "Working Papers"
Günter Hein Günter.Hein@unibw-muenchen.de

Contributing Editors for "GNSS Solutions"
Sam Pullen spullen@stanford.edu and **Di Qiu**

Contributing Editor for "Washington View"
Dawn K. Zoldi

Contributing Editor for "Brussels View"
Peter Gutierrez peter@insidegnss.com

Advisory Editor **Abe Peck** abe@insideunmannedsystems.com

Contributing Editor for "GNSS & the Law"
Ingo Baumann ingo.baumann@bho-legal.com

Technical Editor **Fiona Walter**
Circulation Director **Jan Edwards-Pullen**

ADVERTISING SALES AND BUSINESS DEVELOPMENT

Publisher **Richard Fischer** richard@insidegnss.com

Mobile: 609-240-1590

Office: 732-741-1964 richard@insidegnss.com

Ad Services **Gina McGuinness** gina@insidegnss.com

Mobile: 732-456-4911

**PUBLISHED BY INSIDE GNSS MEDIA & RESEARCH,
a wholly owned subsidiary of Autonomous Media, LLC.**

157 Broad Street, Suite 307, Red Bank, New Jersey 07701 USA

Telephone: 732-741-1964



Follow us on Twitter @insideGNSS

Copyright 2021 Inside GNSS Media & Research LLC. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical (including by Internet, photocopy, recording, or information storage and retrieval), without written permission. Authorization is granted to photocopy items, with attribution, for internal/educational or personal non-commercial use. For all other uses, contact Richard Fischer.

INSIDE GNSS magazine (ISSN 1559-503X) (Online version ISSN 2329-2970) is published bi-monthly by Autonomous Media, LLC, 157 Broad Street, Suite 307, Red Bank, NJ 07701. Print subscriptions are free to qualified USA subscribers. Periodical Postage Pending paid at Red Bank, NJ and additional mailing offices. POSTMASTER send address changes to INSIDE GNSS, PO BOX 92288, Long Beach, CA 90809. *Inside GNSS* is a registered trademark of Autonomous Media. *INSIDE GNSS* does not verify any claims or other information in any of the advertisements or technical articles contained in the publication and cannot take responsibility for any losses or other damages incurred by readers in reliance on such content.

Subscribe Online

FREE one-year subscriptions to the print and/or digital versions are available to qualified readers who work in GNSS-related companies, organizations, research institutes, government agencies, and the military services.

You may also change your address, renew, or unsubscribe online:

WWW.INSIDEGNSS.COM/SUBSCRIPTIONSERVICES

Assured PNT Solutions

Products for Mounted, Dis-mounted, and Airborne applications,
and solutions for Legacy Equipment retrofit.

RSR Transcoder™: no more worries about A-PNT compliance.

Add SAASM, M-Code, CSAC, and INS capability to any legacy GPS Receiver.

The RSR GNSS Transcoder™ allows retrofitting of ANY legacy GPS receiver to next-generation A-PNT capability, simply by replacing the GPS antenna. Gluelessly.

The Transcoder can take any NMEA baseband PNT/PVT signal from any positioning source and convert this to a GPS L1 RF antenna signal in real-time. This allows retrofitting of any GPS system with Mil-GPS assured capability as well as CSAC atomic clock holdover capability, INS positioning, as well as detecting- and mitigating jamming and spoofing events with its internal GPS-FireWall capability. Making your existing equipment compliant with A-PNT requirements has never been easier.



Time and Frequency:

Cesium, Rubidium, and OCXO references with
Concurrent GNSS, SAASM, or M-Code disciplining.

Rugged, Battle-Proven Timing- and Frequency-References with deep sub microsecond/day holdover.

=> Jackson Labs Tech. is shipping the industries' first integrated SAASM/OCXO/CSAC module in quantity today, with a demonstrated upgrade path to M-code.

=> The Ruggedized Low Noise Rubidium GNSS reference can track up to three concurrent GNSS systems, and provides sub 500ns holdover capability/24Hrs with up to -114dBc/Hz @ 1Hz noise.

Select from over 38 different product types to fit your application. Our references are highly customizable, cover a large application variety, and are very cost effective.



The SAASM HD CSAC GPSDO and the FireFly-IIA SAASM have been granted Security Approvals by the Global Positioning System Directorate.

PNT Fulfills Autonomy Promise for Military



ALAN CAMERON
EDITOR IN CHIEF

In the beginning, GPS was envisioned as a military program. A sign on the wall in the Joint Program Office where the system came into being during the early 1970s read:

“The mission of this Program Office is to

- drop 5 bombs in the same hole
- and build a cheap set that navigates
- and don’t you forget it!”

The first goal was motivated by a desire to reduce or eliminate the civilian casualties and collateral infrastructure damage inherent in armed conflict.

It didn’t take long for strategists to see beyond precise targeting to further ambitious uses for precise navigation. Prime among these was robotic resupply of the front lines. Any armed force is every bit as much a logistical enterprise as a fighting one.

Resupply of ships at sea; refueling combat aircraft in the air; reprovisioning advanced units in remote, difficult or conflicted terrain.

Now, nearly 50 years later, those visions have emerged as realities, thanks to the

receiver aircraft, piloted, maneuvers its rigidly mounted fuel probe into the tanker’s trailing drogue.

In-flight refueling requires sustained, minimal separation between paired aircraft, as little as 20 feet at airspeeds of, well, the Navy doesn’t like to say, but let’s presume in excess of 200 miles per hour, at least. Very little room for error.

The delicate feat was achieved using differential GPS, but—of course—GPS is one of the first things to be denied or challenged in combat environments. Therefore different sensor packages must be readied. Our cover story describes one of these, a stereo camera vision-based package.

It’s still PNT. And that’s what we’re about.

MORE ROBUSTNESS TO THE FORE. A second story in this issue, also originating from the Air Force Institute of Technology, explores integrity problems that will unavoidably arise when three or more sensors are employed on a navigation platform, as will surely be the case when drones or robotic vehicles are dispatched into conflicted environments.


GPS/GNSS integrity issues have been studied and designed against for years. That’s why we have receiver autonomous integrity monitoring (RAIM).

Integrity for GPS-inertial integrations is also well understood. But when three or more positioning sensors are employed, as is already done in many autonomous platforms, and there is a disagreement among them—how do you know which to believe?

MEANWHILE. There’s one more GNSS-driven autonomous military navigation suite I would have liked to stuff into this issue, but I don’t have the full info yet. Believe me, I’m chasing it. If anyone knows a good source, please contact me.

The U.S. Naval Air Warfare Center used a Blue Water UAV prototype from Skyways to develop resupply efforts for submarines and other ships over long distances, using small UAVs.

Demonstrations of long-range ship-to-ship and shore-to-ship cargo transport will soon get underway at NAS Patuxent River in Maryland, after customization of the UAV for the requirements of military sealift operations. The Blue Water UAV will take part in additional experiments in the Atlantic Ocean with the Navy’s fleet in 2021.

Exciting times. 

GPS IS ONE OF THE FIRST THINGS TO BE DENIED OR CHALLENGED IN COMBAT ENVIRONMENTS.

autonomous potential inherent in GPS/GNSS. Prototypes have been tested and some may come online as soon as this year.

Refueling aircraft are considered force multipliers because they expand the combat radius of attack, allow patrol aircraft to remain airborne longer and enable aircraft to carry heavier payloads. Planes are also most vulnerable when they land to refuel themselves; replay the Battle of Midway.

But aerial refueling is a very delicate undertaking indeed. And to do it without a pilot onboard the tanker?

Challenging, to say the least.

On June 4, the Boeing MQ-25 T1 test asset transferred fuel to a U.S. Navy F/A-18 Super Hornet: the first time in history that an unmanned aircraft has refueled another aircraft. The magazine cover depicts this.

The MQ-25 stably follows a GNSS-based flight path while extending its fuel drogue, or boom. The

Visit us at
**AUVSI
XPONENTIAL**
BOOTH I956

MOTION SENSING SPACE

IN ANY

NEW 9DoF IMU
DMU41



PRECISE, TOUGH MOTION SENSORS AND
SYSTEMS FOR LOW-DRIFT PERFORMANCE
IN THE HARSHTEST ENVIRONMENTS

Precision MEMS
technology

Ultra-compact,
easily integrated

Inertial sensors for all
performance requirements



siliconsensing.com

SILICON
SENSING®

360 DEGREES

News from the
world of GNSS



GPS III launch in June.
Photo courtesy of SpaceX.

Los Angeles Air Force Base, California

M-Code Comes Closer: As Soon as Latest Launched Satellite Turns On

The U.S. Space Force brought broadcast of the modernized encrypted M-Code signal one step closer to global availability for authorized military users with the launch of the fifth GPS III satellite on June 17. In contrast to most GPS launches, the rocket blast-off was actually moved forward from its originally scheduled date in July. No reason was given for this break in protocol.


GPS III Space Vehicle-5 (SV-5), built by Lockheed Martin, brings the number of M-code broadcasting satellites in the GPS constellation to 24, critical for global access to the jam-resistant signal. This number includes earlier GPS IIR-M and GPS IIF satellites as well as five of the Third Generation. GPS III SV-5 will replace one of the early models.

The code will begin broadcasting once the satellite is operational, which should be two weeks after launch, according

to Col. Edward Byrne, senior materiel leader, Medium Earth Orbit Space Systems Division at Space and Missile Systems Center (SMC).

SMC plans to orbit 10 GPS III satellites and then update to a follow-on version called GPS IIIF.

Full operational use of M-code must still await full operational capability of the updating ground control system OCX, scheduled for Q3 of 2023.

“Digital capabilities will roll in over the next year to take advantage of the GPS III capabilities,” added Byrne. “That will allow us to declare IOC [initial operational capability] for the constellation,” he said. “OCX and the user equipment piece do not come online until the third quarter of 2023; that is when we would expect to have our initial operational capability for the GPS enterprise across across all segments: space, ground and user equipment.” 


Ottawa, Canada and Berlin, Germany

Canada, Germany First Allies to Receive M-Code User Equipment

The first M-code enabled receiver cards were delivered to a U.S. ally in February 2021. Canada took possession of an unspecified number of Military Code-capable GPS receiver cards for the purposes of laboratory and field testing. The “loan,” in the terms of an official government release, constitutes the first fulfillment of a U.S. Space Force Space and Missile Systems Center (SMC) three-year multinational Project Arrangement with partnering nations.

The arrangement, established in close coordination with the Department of Defense, Chief Information Officer and the Deputy Under Secretary of the Air Force for International Affairs, became effective in December 2020 when Canada became the first co-signer of the document. France, Germany, the Republic of Korea and the United Kingdom are projected to receive Military GPS User Equipment (MGUE) Increment 1 technology. All partnering nations will conduct laboratory and field tests to evaluate the performance and compatibility of MGUE Increment 1 products with their respective platforms and share their findings and lessons learned. Australia, Italy, the Netherlands, and Sweden have expressed interest and intent in joining the agreement later this year.

In late June, BAE Systems, Inc. received the first contract from the Space and Missile Systems Center’s Space Production Corps to deliver M-Code MGUE to Germany.

The German order focuses on the Miniature PLGR Engine–M-Code (MPETM-M), the smallest, highest-performance M-Code GPS receiver for ground applications available today. 

See Additional News Stories

at www.insidegnss.com/news

- Beyond GPS: Air Force Rethinks Position, Navigation and Timing
- Army Funds Research & Development Small-Sat Payload for GPS-Denied Nav, Guidance & Control
- Galileo EU Defence (GEODE), The Biggest Galileo Application Ever Launched
- Anti-Jamming Kits for Dismounted Soldiers ... *and more.*



Need to create custom GNSS test scenarios? Simulate anywhere with SatGen

SatGen simulation software generates bespoke GNSS test scenarios at the time, date and location of your choice:

- Compatible with all LabSat simulators
- Multi-Constellation
- Multi-Frequency
- Static or high-dynamic scenarios (including space simulations)
- Real-time simulations (only available with LabSat Real-Time simulator)





This column breaks down the 286-page DHS report written by the RAND Corporation, “Analyzing a More Resilient National Positioning, Navigation, and Timing Capability.” The document seems to both complement (no pun) and contradict the Department of Transportation’s (DOT) earlier Complementary PNT Report. It is unclear how Congress will react to it, or how it may ultimately impact federal budgets and programs.

PNT: Nothing to See?

A Department of Homeland (DHS)-chartered May 2021 report concludes that PNT threat and resilience concerns are not as dire as some have made them out to be, and that funds for backup could be spent elsewhere. Why this runs counter to other recent government reports is not clear, nor is the fallout from this divergence of Congressionally mandated views. The Department of Transportation has distanced itself a bit from this report by the RAND Corporation—and even its issuer, the DHS, seems to have done so.

DAWN M.K. ZOLDI (COLONEL USAF, RET.)



Dawn M.K. Zoldi
(Colonel, USAF, Retired)
is a licensed attorney
and a 25-year Air
Force veteran. She

is an internationally recognized expert on advanced technology law and policy, a recipient of the Woman to Watch in UAS (Leadership) Award 2019, and the CEO of P3 Tech Consulting LLC.

Late to the Party?

In §1618 of the Fiscal Year 2017 National Defense Authorization Act (FY 17 NDAA), Congress required a timely, full, and joint DHS, DOT and the Department of the Defense (DOD) study “to assess and identify the technology-neutral requirements to backup and complement the positioning, navigation, and timing capabilities of the Global Positioning System for national security and critical infrastructure.” This was supposed to be due no more than one year after the date of enactment.


Three and a half years after the deadline, RAND, a nonprofit, nonpartisan, public-interest research organization, which runs the DHS’ Homeland Security Operational Analysis Center under contract, published this “partial response.”

While the DOT supported the effort, the DOD did not. As such, the report solely focused on domestic

non-military PNT issues associated with critical infrastructure.

Interestingly, the FY17 NDAA requirement was actually fulfilled, at least from the DOT perspective, by a DHS report submitted to Congress in April 2018: the Report on Positioning, Navigation, and Timing (PNT) Backup and Complementary Capabilities to the GPS National Defense Authorization Act Fiscal Year 2017 Report to Congress¹. That report found:

- **GPS IS NOT THE ONLY SOURCE OF PNT DATA.** Other sources are available for purchase, and include alternate space-based systems and constellations, terrestrial beaconing, time-over-fiber, cellular and wireless signals, and local terrestrial systems.
- **APPLY PNT PRINCIPLES.** Whatever the PNT source, users must apply the principles found in Executive Order 13905, Strengthening National Resilience Through Responsible Use of Positioning, Navigation, and Timing Services.
- **NO INCENTIVES FOR NON-GPS PNT.** Unless non-GPS PNT sources are free/low-cost or provide a unique benefit deemed valuable by the user, there is no reason to assume users will adopt new non-GPS PNT sources more widely than they have.
- **MANY SECTORS RELY ON PNT.** Disruption would cause significant costs, delays, or degradation of functions and service.
- **SECTOR NEEDS DIFFER.** Some sectors require very precise timing, while in others position and navigation precision is more important.
- **CHOICES WILL CAUSE ANY FAILS.** Critical infrastructure that ceases to operate due to GPS disruptions will do so because of design choices associated with a lack of information, cost, efficiency, and other considerations—not because of a lack of available options.
- **DECONFLICTION NECESSARY.** New non-GPS PNT systems that are designed without considering existing PNT systems may simply



Protect your position with true GRIT.

GNSS Resilience and Integrity Technology (GRIT) is a firmware suite of detection and protection technology developed for our OEM7 receivers to guard your position, navigation and timing. GRIT can detect spoofing and keep you on track, and the GNSS Interference Toolkit (ITK) identifies interference frequencies in your area, protecting you from unintentional or malicious interference. When you've got GRIT, your position is true.

Autonomy & Positioning – Assured | novatel.com/grit



compete with existing systems rather than fill perceived backup gaps.

- **NO ONE SOLUTION.** No single PNT system, including GPS, can fulfill all user requirements and applications.
- **PUBLIC-PRIVATE DEVELOPMENT NEEDED.** Position and navigation backups must be application-specific, developed in coordination with industry owners and operators.
- **TIMING IS THE EASIEST PROBLEM TO SOLVE.** Timing requirements are simple, with a minimal acceptable precision of anywhere between 65-240 nanoseconds. That same earlier FY17 Report made these four recommendations:
- **TEMPORARY GPS DISRUPTIONS:** End users should be responsible for mitigating temporary GPS disruptions. The Federal Government can facilitate mitigation for various critical infrastructure sectors, but should not be solely responsible for it.
- **PNT DIVERSITY AND SEGMENTATION:** Adoption of multiple PNT sources will diffuse the risk currently concentrated in wide-area PNT services such as GPS.
- **SYSTEM DESIGN:** PNT systems must be designed with inherent security and resilience features.
- **PURSUE INNOVATION THAT EMPHASIZES TRANSITION AND ADOPTION:** Incorporating PNT signal diversity should take into account factors such as business case considerations, financial costs, technical integration, and logistical deployment.

What's New?

This newest RAND report doesn't add much other than to cast a bit of doubt on the extent of monetary impacts that a GPS outage will actually have on critical infrastructure.

Real Risks. Overblown Impacts?

The RAND report validated that various threats could destroy, deny, or trick GPS signals. On the one end of the spectrum, acts of war can erupt into state-level nuclear exchanges.

Extreme space weather events can impact space systems. Both would have large scale to catastrophic effects. On the other end, localized GPS jamming or spoofing, whether negligently or intentionally caused, would likely have localized and short-lived effects. Other natural or intentional events of varying scope and duration, such as PNT terrorism, computer system failures, or simple communications breakdowns remain in the realm of

"Potential losses from GPS disruption

nationwide, on a daily basis, range from a low estimate of \$785 million to a high value of \$1,318 million."

2021 DHS/RAND Report

the possible. Similarly, PNT and GPS back-ups could also be vulnerable to these and other threats, such as terrestrial systems which seem more susceptible to physical attacks.

While acknowledging that insufficient data exists "to make defensible probability estimates," the report nevertheless concludes that, aside from unfathomable nuclear holocaust, virtually all other GPS disruption or loss scenarios would have small-scale impacts lasting only a few days. Only a geomagnetic storm implied economy-wide disruption. Even this, the report says, would last for only a few days. Major assumptions on the short-lived nature of these impacts include that law enforcement would catch perpetrators of intentional acts and that already-existing complementary tech, old fashioned pre-GPS practice, and/or steadfast American gumption would keep the economy grinding on.

Hundreds, Thousands of Millions in Losses

Estimated economic costs resulting from the varied threats to PNT from a nationwide outage, "range, on a daily basis, from a low estimate of \$785 million to a high value

of \$1,318 million." Based on 2018 rates, likely nation-wide aggregate losses by economy sector, in millions of dollars per day (unless otherwise noted), follow:

- Consumer Based-Location Services (cell phones and apps)—\$94M;
- Commercial Road Transport—\$141M
- Emergency Services—ranges from \$11M to \$72M with median at \$32M
- Agriculture—ranges from \$42M to \$514M based on time of year;
- Construction—\$24M;
- Surveying—\$7M
- Aviation—\$6M
- Railway—\$100,00
- Telecommunications—\$40M (day 1) up to \$456M (after 30 days);
- Electricity Generation/Transmission—\$9M
- Finance—\$9M;
- Port Operations—\$224M, aggregate 30-day outage \$6.7B
- Mining—\$32M
- Oil & Gas Exploration—30 day outage \$1.5B

Estimates drop by orders of magnitude when one makes reasonable assumptions about alternate available systems and the projected short duration of any PNT outages.

So Many Alternatives

According to the findings, many alternative sources and technologies already exist to increase national PNT resilience and robustness. No one-size-fits-all solution exists because users have different needs, such as:

- Precision Timing on Networks—modern communication networks, including financial networks and electrical power networks, need a common time standard for tightly synchronized operations. Examples include emergency services Land Mobile Radio systems and Financial Services high-frequency trading networks.
- Moderate-Accuracy Positioning of People and Road Vehicles—typically on smartphones and similar devices, used by the public and many commercial

emcore®

VISIT US AT JNC
Booth #211

Get 10X the Performance!

EMCORE Fiber Optic & Quartz MEMS IMUs

**ITAR
FREE**

THE EMCORE IMU ADVANTAGE

- Suitable for wide variety of high-precision Defense Industrial, Marine, and Commercial applications
- Industry's performance and CSWaP leading tactical grade FOG & MEMS IMUs:
 - EN-300 - 0.04°/hr in-run gyro bias, Ultra-low 0.015°/√hr ARW
 - SDI170 & SDI500 - 1°/hr gyro & 1 mg accel bias, Low 0.02°/√hr ARW
- Ruggedized, compact package designs suitable for shock and vibration environments
- Form, Fit and Function compatible upgrades for legacy FOG & RLG IMUs for easy integration and re-qualification (LN-200, HG1700-AG58)



EN-300
FOG IMU



SDI500
QMEMS IMU



SDI170 QMEMS IMU

FOR MORE INFORMATION

navigation-sales@emcore.com

+1 866 234 4976

www.emcore.com/nav



SYSTRON DONNER
INERTIAL

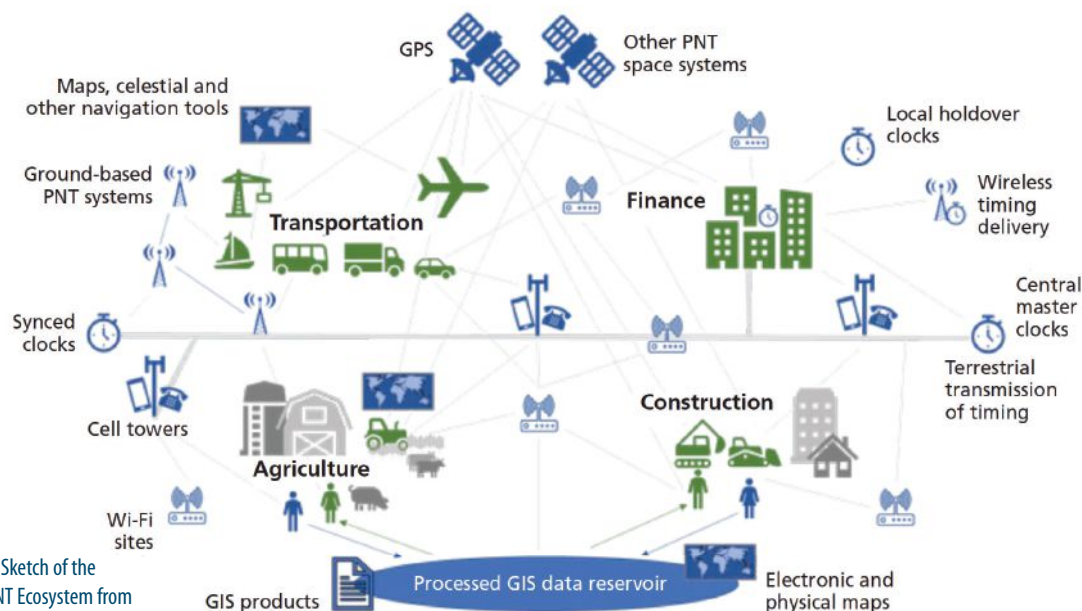


FIGURE 1: Sketch of the National PNT Ecosystem from the 2021 RAND/DHS Report.

Illustration courtesy of 2021 RAND/DHS Report.

industries for directional, logistical and similar applications.

- High-Accuracy Positioning of Equipment—for surveying and the guidance of expensive equipment for construction, mining, drilling, or agriculture.
- High-Reliability Positioning in Aircraft Landing—required in vertical position accuracy
- Low-Accuracy Positioning of Aircraft and Ships En Route—needed in wide-area coverage possible far from shore or land

Alternatives for these needs run the gambit of other wireless PNT signals, RF “signals of opportunity” (like Locata which is used over defined areas, such as ports and mines, where GPS reception might be unreliable or impossible), wireless time signals, wired time signals, user equipment-based solutions and PNT resilience technologies. These solutions can cross user-need categories. Many of these alternative PNT technologies or supplemental systems already exist but have not been widely implemented; others are in earlier stages of development (see Figure 1).

Of all of these, the report locks onto “one seamless backup for GPS”—other GNSS constellations that use similar signals in the same radio band, such as

Galileo, GLONASS and BeiDou. That said, the U.S. prefers not to depend on adversary sat systems. That leaves Galileo, which itself experienced a significant outage in December 2020.

Threats Don't Justify Backups

Cost drivers for complementary tech include new space infrastructure, new terrestrial infrastructure in covered areas, and individual user costs to utilize for PNT. Other than stating, “even costs of tens to hundreds of dollars per user could add up to substantial costs over relevant user bases,” the report outlines no other dollar figures for these “costly” systems (possibly because detailed cost estimates for potential alternative PNT systems are proprietary to various firms). Yet it concludes, “If the costs of a national GPS disruption are at most in the tens to hundreds of millions per day—and if any single backup system would mitigate only a portion of those costs—then the risk does not justify more than modest government investment in any single backup system.”

Take jamming in an urban area, as an example. To have more than a localized effect, someone would




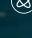
have to intentionally stage a large number of jammers of different levels of power and reach around the city to disrupt GPS. More than a hundred such local attacks per decade would be needed to match the costs of the cheaper alternative PNT systems to hedge against this possibility. “Given the complete absence of such attacks to date,” the report notes, “the threat of such attacks does not come even close to justifying investments in the backups or alternatives we examined.” Spoofing, on the other hand, can be readily detected in any sensitive applications through integrity monitoring, in RAND’s opinion.

So, what’s the plan? The report outlines four cost-effective measures the feds could implement:

- Information sharing and standards for hardening or quality monitoring for those systems now relying on GPS—the federal government should share information with the locals
- Expansion of law enforcement authorities focused on GPS jamming—make it painful for criminals to do this by allowing fines that could be recycled back into buying or

THE FUTURE OF FIBER OPTIC GYROSCOPES

BOREAS - GNSS / INS

-  0.005 ° Roll & Pitch
-  0.006 ° GNSS Heading
-  0.01 ° Gyrocompass Heading
-  0.001 ° /hr Gyroscope



BOREAS THE WORLD'S 1st DIGITAL FOG

Boreas takes FOG technology to the next level with new patent pending digital FOG (DFOG) technology. The revolutionary DFOG technology combines a specially designed closed-loop optical coil with advanced digital signal techniques that have been developed over the last 25 years. This ground-breaking gyroscope combined with our AI-based algorithm provides unparalleled performance and reliability, with the lowest SWaP-C (Size, Weight, Power and Cost) on the market.



**ADVANCED
NAVIGATION**

- maintaining resilient systems
- Timing-only backup through fiber/FirstNet, eLoran, or Satellite Time and Location (which uses the Iridium constellation)
- Expansions of economically successful but geographically limited systems offering high performance (whether 5G, MBS, or other) in neglected areas.

Additionally, maintaining “time-proven, robust fallbacks” to PNT, keeping diverse and dispersed PNT capabilities in the national PNT ecosystem, and designing non-GPS reliant systems are also essential.

Diverging Points of View

This latest RAND Report doesn’t deviate much from DHS’ FY17 Report, at least from a wave-top standpoint. However, when you compare it to DOT’s Complementary PNT Report, the conclusions diverge when it comes to toughening PNT. The breakdown, according to Karen Van Dyke, Director of DOT OST-R’s Office of PNT & Spectrum Management, who contributed information to the RAND report, when it comes to PNT resiliency, follows:

- **Protect:** GPS/GNSS Performance Monitoring and Interference Detection. DOT agrees with the RAND report on the need for an interference/threat detection capability.
- **Toughen:** GPS signal authentication and cyber-resilient user equipment. DOT considers this capability to be important, while the RAND report appears to minimize the likelihood that these are threats to GPS that need to be addressed.
- **Augment:** through complementary PNT service. The Office of the Assistant Secretary for Research and Technology at US DOT, through work authorized and funded under the FY 2018 National Defense Authorization Act (NDAA), has demonstrated mature PNT technologies that complement and could provide a backup to the GPS service

in case of a major disruption. Many of these technologies are already commercially available to owners and operators of critical infrastructure. During federally funded research demonstrations, these technologies and their vendors displayed significant potential to mitigate the risks faced by users who rely exclusively on GPS/GNSS services. This finding largely aligns with the RAND report.

VanDyke also points out that, in contrast to the RAND report, DOT

“The biggest challenge

is to build a dual-use open-systems architecture (OSA) with open, non-proprietary interfaces that sets the foundation for a PNT ecosystem. Folks are just starting to look at this now.”

Dr. Mikel Miller, Vice President for PNT Technologies at Integrated Solutions for Systems

believes there is an important role that the federal government, in particular DOT, needs to play with respect to:

- Development of safety-critical PNT requirements and standards for transportation services.
- Development of a PNT vulnerability and performance testing framework of demonstrated and suitable complementary technologies.
- Development of PNT performance monitoring capabilities to ensure PNT services provide operational resilience and achieve safety-critical standards.

Dr. Mikel Miller, Vice President for PNT Technologies at Integrated Solutions for Systems (IS4S), former Air Force (AF) Senior Scientist (ST-00) for PNT Technologies for the AF Research Laboratory, and current program manager for USAF’s Resilient Embedded GPS INS (R-EGI) project, agrees that no one-size-fits-all GPS alternative solution exists

and that the feds should play a role.


“While there are a lot of complementary PNT technologies out there,” says Miller, “many have not been employed because they are too difficult to integrate into existing proprietary systems. The biggest challenge is to build a dual-use open-systems architecture (OSA) with open, non-proprietary interfaces that sets the foundation for a PNT ecosystem. Folks are just starting to look at this now.”

This PNT ecosystem will enable 3rd-party developers to plug-and-play their applications and sensors into future PNT air, space, ground, or sea systems—which is directly in line with the Army’s Modular Open Systems Architecture (MOSA) approach.

According to Miller, who is also past President of the Institute of Navigation, the optimal solution may be for the government to own and maintain the integration architecture and allow industry to develop solutions to literally plug into it.

Not only has DOT distanced itself a bit from this RAND Report, incredibly, even DHS seems to have done so. Right up front, the RAND report states, “The results presented in this report do not necessarily reflect official DHS opinion or policy.”

As mentioned, the DOD did not contribute, nor has it not provided its own final assessment on PNT. Instead, the Government Accountability Office (GAO) has recently indicated that DOD needs to do more work to address PNT alternatives (reported online at [insidegnss.com/usg-accounting-office-faults-defense-dept-efforts-on-alternative-pnt/](https://www.insidegnss.com/usg-accounting-office-faults-defense-dept-efforts-on-alternative-pnt/)).

In the meantime, it would be great if DHS, DOT and DOD would actually get together and provide one true joint assessment as Congress directed. Only then can we truly know whether or not there is something to see, when it comes to PNT, in terms of vulnerability and the need for backup. 



Old Equipment? Time for Xidus!

Meet us in St. Louis!
ION GNSS+, Sep. 20-24
Booth no. 242



Xidus GNSS Simulator

Validate GNSS Systems with
German Precision Engineering.

WORK Microwave's Xidus is the obvious choice when validating GNSS receivers and other GNSS systems of systems.

All benefits of Xidus GNSS Simulator:

- World's best RF signal quality
- Multi-GNSS & Multi-RF in one box
- Perfect reproduction of received signals
- Most modern GUI technology & features



Find out more:

visit www.gnss-simulator.com or call +49 8024 6408 222

WORK Microwave GmbH, Holzkirchen, Germany • www.work-microwave.com

Concerned about your IMU's performance
after 30 years of storage or use?

Don't worry — we have taken care of it!

Predicting the future can be difficult. But providing predictability in inertial measurement units is our specialty. The STIM377H and STIM277H are hermetically sealed IMUs and gyro modules designed for applications that require the highest long-term reliability.

In our testing, we have simulated 30 years of operation or storage at high and low temperatures to give you the best insight into predicting their performance.

A 3,000-hour qualification program of HTOL, HTSL, LTOL and LTSL, including read-out at every 1,000 hours has confirmed performance at the end of the program.

By implementing the STIM377H or STIM277H, you reduce uncertainty in your design, whether it is planned for three decades of immediate use or for long-term storage towards future use.



Gyro modules and Inertial Measurement Units

Actual size

Q How to mitigate a spoofing signal while tracking it for intent analysis?

GNSS Solutions is a regular column featuring questions and answers about technical aspects of GNSS. Readers are invited to send their questions to the columnists, **Dr. Sam Pullen** and **Dr. Di Qiu**, who will answer them or select other experts to do so.

CHUN YANG AND ANDREY SOLOVIEV
QUNAV



DI QIU is a senior research engineer at Polaris Wireless. She has been involved in many aspects of positioning algorithms and location-based applications, including GNSS algorithm

development, indoor localization, sensor fusion, signals of opportunity, GIS-based map-matching, location-based security and location profiling using machine learning. She has a Ph.D. in aeronautics and astronautics from Stanford University. E-mail: dqiu@polariswireless.com.

In a GNSS spoofing attack, a terrestrial radio transmitter mimics GNSS signals at a greater signal strength than the actual system can transmit. In order to detect and mitigate a spoofing attack, it is important to know where a spoofing signal comes from and what it tries to accomplish for its ultimate neutralization.

A: A spoofing attack attempts to fool a GNSS receiver either by rebroadcasting genuine signals that are captured elsewhere or at a different time (a masking beacon or meaconing) or by broadcasting fake GNSS signals that are created to look like authentic GNSS signals (smart spoofing). A spoofer intends to mislead a GNSS receiver to use the spoofed signals as normal to produce a position fix at some place other than where it actually is (position-push) or to be where it is but at a different time (time-push) without knowing it.

A successful single-channel spoofing attack implemented is annotated in **Figure 1**, where the blue peak is the correlation with the authentic signal, the red peak is the correlation with the spoofing signal, the red dots represent the early, prompt, and late correlators, and the red arrows indicate how the spoofing signal gets close to and then moves away from the authentic signal.

What does a spoofer need to know to be successful? The spoofer has to know the target receiver well in time, position, and frequency in order to get into its capture window. The receiver has to be indifferent, that is, assume that the received signals are clean without questions asked.

What Makes a GNSS Receiver Vulnerable to Spoofing?

The vulnerability of a conventional GNSS receiver to spoofing stems from its design premise to track a single signal per satellite. Typically, a GNSS receiver uses three 1-ms long correlators in tracking having a 3 dB resolution of $\pm\frac{1}{2}$ chips (150 m) in time

and ± 500 Hz in frequency. The code error discriminator in a delay lock loop (DLL) has a non-ideal responsiveness (not a Dirac delta function). Any signal that falls within a resolution cell, the capture window, as illustrated in **Figure 2(a)** can be “felt” by the receiver. As a result, the contribution to the total correlation values by a spoofing signal, just like multipath, biases the authentic signal estimates, leaving a door open to an encroaching spoofing signal.

Clearly, the early-prompt-late correlator architecture of a conventional GNSS receiver has a signal capturing window that is not narrow enough to reject any harmful signals. It is not wide enough, either, to spot any spurious signals that pop




FIGURE 1 An annotated spoofing attack process (from Reference [3]).



Live Remote GNSS Training with Real-Time Engagement

UPCOMING GNSS COURSES

November 2–3, 2021

Course 122: GPS Fundamentals and Enhancements

Instructor: Dr. Chris Hegarty, MITRE

November 2–5, 2021

Course 346: GPS/GNSS Operation for Engineers and Technical Professionals

Instructor: Dr. Chris Hegarty, MITRE

December 13–17, 2021

Course 557: Inertial Systems, Kalman filtering and GPS/INS Integration

Instructors: Dr. Alan Pue (Retired) APL/JHU and Michael Vaujin, Consultant

February 7–11, 2022

Course 551: Using Advanced GPS/GNSS Signals and Systems

Instructor: Dr. John Betz, MITRE Fellow Emeritus

All courses available for private remote group training.

See <https://www.navtechgps.com/gps-gnss-training/courses/>

“

The video quality was excellent. I don't feel as though going through the course remotely had any negative impact. It was still very personal, easy to ask questions, and I enjoyed the banter over coffee in the morning even if we were all scattered across the world. This was such a great experience. ”

—Shealyn Greer, Trident Research,
Course 346 (Remote Course, July 2020)



GNSS products, solutions, and training

QUESTIONS? Contact Trevor Boynton • tboynton@navtechgps.com • 800-628-0882 • +1-703-256-8900

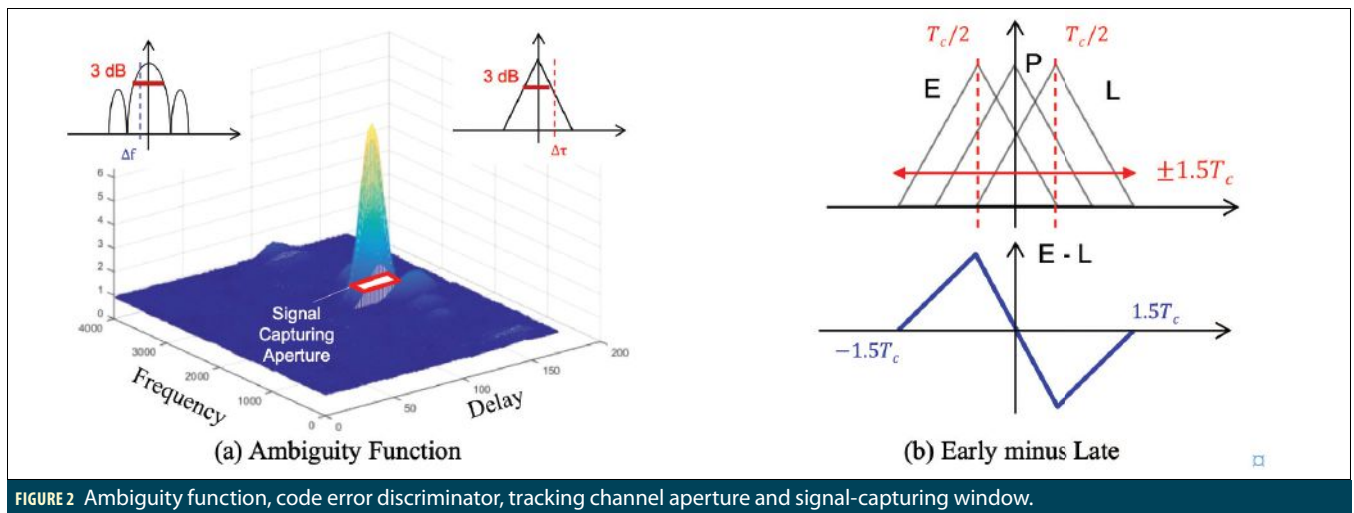


FIGURE 2 Ambiguity function, code error discriminator, tracking channel aperture and signal-capturing window.

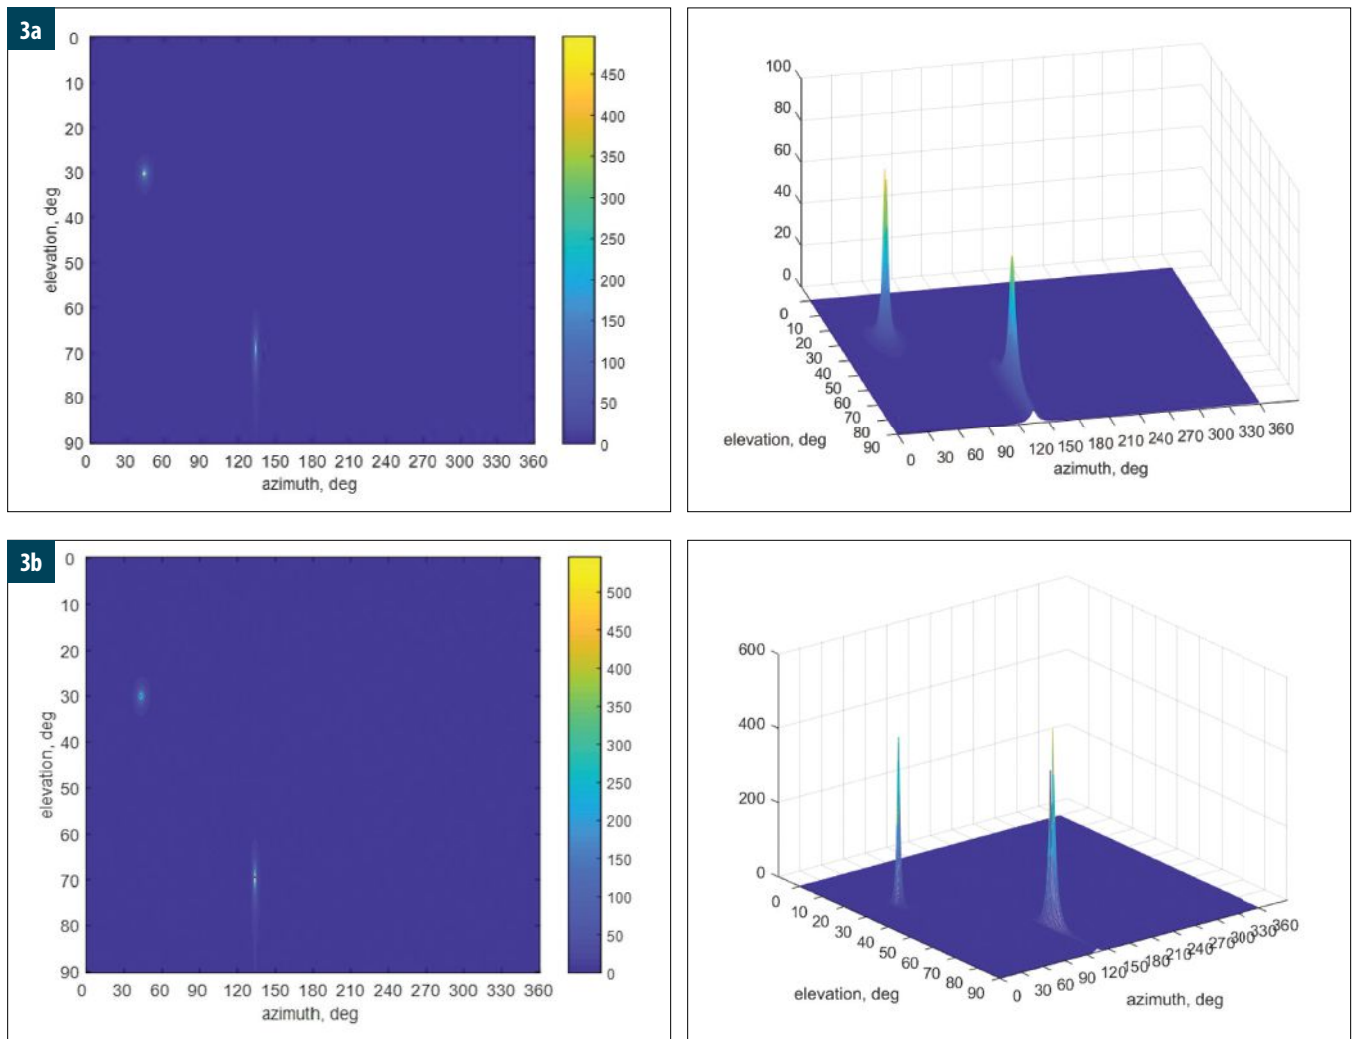


FIGURE 3 Effect of correlated signals on angular estimation with MUSIC vs. FBS-MUSIC. (a) $f_1 = f_2 = 1$ Hz, $\psi_1 = 0^\circ$, $\psi_2 = 90^\circ$. Both signals are sinusoidal with the same frequency of 1 Hz but different initial phases of 0° and 90° , respectively, sampled at 80 Hz over 1 sec. The two signals are uncorrelated and two sharp peaks appear at correct locations. (b) $f_1/f_2 = 1/10$ Hz, $\psi_1/\psi_2 = 0^\circ$. Both signals are sinusoidal with the same initial phase of 0° but different frequencies of 1 Hz and 10 Hz, respectively, sampled at 80 Hz over 1 sec. The two signals are uncorrelated and two sharp peaks appear at correct locations. (c) $f_1/f_2 = 10$ Hz, $\psi_1/\psi_2 = 0^\circ$. The two signals are identical sinusoidal with frequency of 10 Hz and initial phase of 0° , sampled at 80 Hz over 1 sec. The two signals are fully correlated (coherent). There is only one large peak between $(\phi_1, \theta_1) = (45^\circ, 30^\circ)$ and $(\phi_2, \theta_2) = (135^\circ, 70^\circ)$: a wrong angular estimate. (d) FBS-MUSIC: $f_1/f_2 = 10$ Hz, $\psi_1/\psi_2 = 0^\circ$. Two identical sinusoidal signals are fully correlated (coherent). With FBS-MUSIC, there are two distinct peaks. However, the peaks are widened, and their locations are biased.

up over the horizon. In other words, this “myopic” architecture cannot “see” beyond ± 1.5 chips as shown in **Figure 2(b)**. Except for in the initial acquisition mode, it has no way to detect the presence of a spoofing signal until it is too late, that is, the bad influence has already been exerted. This “myopic” architecture is the root cause of vulnerability of conventional receivers to spoofing.

Spoofing Detection, Mitigation, and Tracking

There are many techniques for post-factual detection of spoofing along the signal processing chain. Notably is the chip transient analysis that can detect the distortion in chip shape caused by overlapped spoofing and authentic signals with a protection level of a fraction of chip at the 30-m level.

The presence of a spoofing signal cannot eradicate the existence of an authentic signal even though it may overpower it. If a spoofing signal always aligns itself to the authentic signal, no harm is done. When a spoofing signal captures a receiver and pulls away, two correlation peaks will emerge in the time-frequency domain. This observation has been used for spoofing detection.





The use of a larger surveillance space (panoramic) eliminates the potential risk of conventional myopic tracking loops, thus detecting the presence of spoofing signals before they become dangerous. The idea was the basis of the Auxiliary Peak Tracking (APT) method and the All Signal Acquisition Processing (ASAP)

scheme. The latter goes one step further not only to mitigate the effect of spoofing on the authentic signal (electronic protection) but also to track the spoofing signals and use the spoofed solution for spoofing intent analysis (electronic support).

Spoofing Angular Estimation and Localization


The best way to mitigate spoofing is to prevent it from getting into the tracking and measurement generation processes of a GNSS receiver. From the time-domain waveforms, it is rather difficult to discern a spoofing signal from an authentic signal if they are off only slightly in time, frequency, or phase.

However, a spoofer cannot conceal its own physical presence.



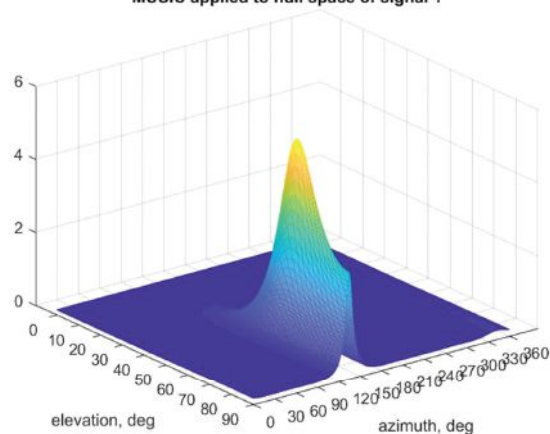
4a

forward and backward smooth MUSIC



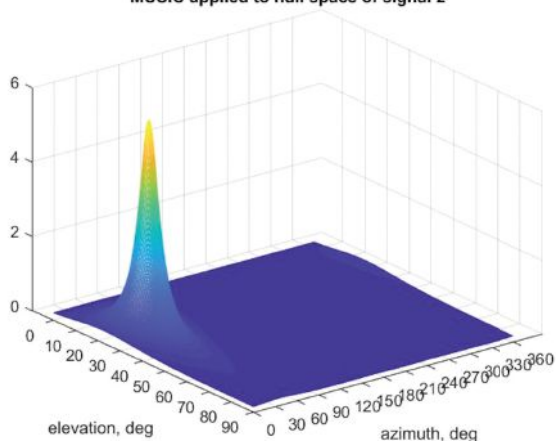
4b

MUSIC applied to null space of signal 1



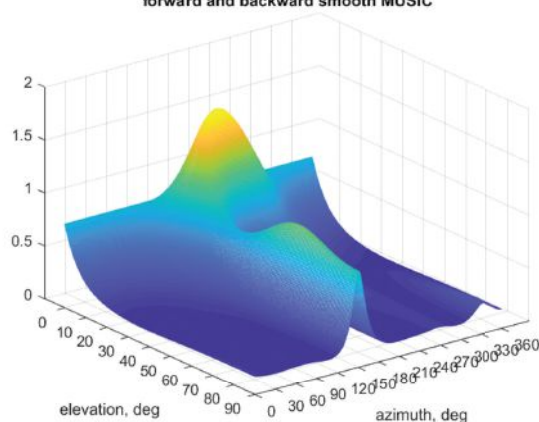
4c

MUSIC applied to null space of signal 2



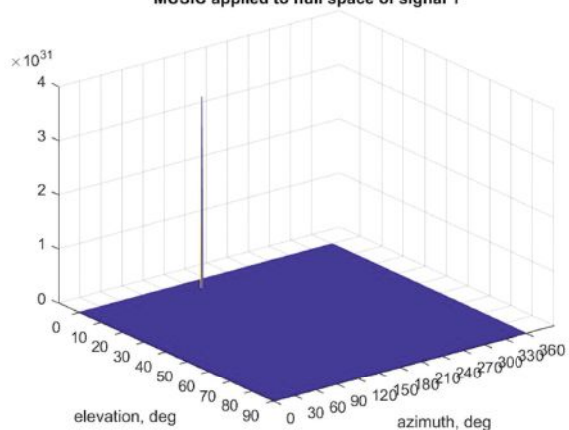
4d

forward and backward smooth MUSIC



4e

MUSIC applied to null space of signal 1



4f

MUSIC applied to null space of signal 2




FIGURE 4 Angular Estimation of Correlated Signals with FBS-MUSIC vs. Null-Space Projection. (a) Pseudospectrum of FBS-MUSIC Applied to Correlated Signals: Two Distinct Peaks at Biased Locations. (b) Pseudospectrum of MUSIC for 2nd Signal in 1st Signal Nullspace: More Accurate Angular Estimate. (c) Pseudospectrum of MUSIC for 1st Signal in 2nd Signal Nullspace: More Accurate Angular Estimate. (d) Pseudospectrum of FBS-MUSIC for Two Correlated Signals: Same Azimuth Different Elevation. (e) Pseudospectrum of MUSIC for 2nd Signal in 1st Signal Nullspace: More Accurate Angular Estimate. (f) Pseudospectrum of MUSIC for 1st Signal in 2nd Signal Nullspace: More Accurate Angular Estimate.

It has to counterfeit all satellites in view in order to push the position solution of a target receiver to the intended spoofed-to location. Yet, it is costly and even impractical for a spoofing system to place many transmitters around and above a target receiver to emulate a constellation. The spoofing signals are likely to be emanated from a limited number of directions.

A powerful way to discriminate against a spoofer is to exploit its spatial information, i.e., the angle of arrival (AOA). An antenna array can be used to steer a null in the direction of a spoofer so as to significantly reduce, if not completely eliminating, the influence of spoofing on authentic signals.

The AOA estimates to the spoofer used for spoofing detection and mitigation can be exploited advantageously for spoofer localization via triangulation from several distributed detectors or by one moving along a trajectory having a large angular extent over the spoofer. This provides an effective electronic support to ultimately neutralize the spoofer(s).

Correlated Spoofing: Forward and Backward Smoothing

Both GNSS authentic and spoofing signals are below the noise, which are “picked up” from the noise first by despreading correlation. An obvious benefit of such a post-correlation approach is its ability to exploit AOAs to differentiate the authentic signal from spoofing signal even when their energy peaks coincide in the delay-Doppler domain.

Instead of having a code and carrier tracking loop for each array element, the master-slave architecture only implements one tracking loop for a reference antenna (the master channel), and the code and carrier replicas from the master channel are used to drive the remaining antennas (the slave channels) so as to maintain the coherence across the array. The complex correlators contain the spatial information about both the authentic and spoofing signals that

are used for angular estimation. A spoofing signal is mitigated by generating a null in its direction while preserving the authentic signal in the main beam, thus closing the tracking loop for the master channel.

The multiple signal classification (MUSIC) algorithm is widely used for high-resolution angular estimation with an antenna array and has been applied for spoofing detection and mitigation. However, there is a technical complexity in that the MUSIC algorithm does not perform well with correlated signals, which is the case when a spoofing signal gets close to an authentic signal in time, frequency, and phase. A popular method to address this problem is to decorrelate the signals via forward and backward smoothing (FBS). Intuitively, two identical sinewaves are coherent. If they are off in phase by 90 degrees, their cross-correlation is zero and the two coherent signals are effectively de-correlated.

To illustrate the effect, consider the following example with a seven-element controlled reception pattern antenna (CRPA) receiving two signals from the two directions with azimuth and elevation $(\phi_1, \theta_1) = (45^\circ, 30^\circ)$ and $(\phi_2, \theta_2) = (135^\circ, 70^\circ)$, respectively, and both with SNR = 10 dB. The angular search spacing is 0.5° from 0° to 90° for elevation and 0° to 360° for azimuth.

Figure 3 shows the image (left) and surface (right) of the pseudospectrum of MUSIC or FBS-MUSIC where the peak location provides an angular estimate for uncorrelated signals (a-b) and correlated signals (c-d).

Correlated Spoofing: Nullspace Projection

An alternative way to deal with two correlated signals is to separate them via projection to each other's null space. Assume that the array output $\mathbf{x}(t)$ contains two signals $s_a(t)$ and $s_b(t)$ in directions \mathbf{a} and \mathbf{b} , respectively, buried in noise $\mathbf{n}(t)$ as:

$$\mathbf{x}(t) = s_a(t) \mathbf{a} + s_b(t) \mathbf{b} + \mathbf{n}(t) \quad (1)$$

By projecting $\mathbf{x}(t)$ onto the nullspace of $s_a(t)$ in direction \mathbf{a} via the

orthogonal projection matrix $\mathbf{P}_{\perp a} = \mathbf{I} - \mathbf{a}\mathbf{a}^H/(\mathbf{a}^H\mathbf{a})$ such that $\mathbf{P}_{\perp a} \mathbf{a} = 0$, we obtain the projected signal as:

$$\mathbf{z}^b(t) = \mathbf{P}_{\perp a} \mathbf{x}(t) = s_b(t) \mathbf{P}_{\perp a} \mathbf{b} + \mathbf{P}_{\perp a} \mathbf{n}(t) \quad (2)$$

which removes \mathbf{a} yet retains \mathbf{b} .

Similar operation can be carried out by projecting $\mathbf{x}(t)$ onto the nullspace of $s_b(t)$ in direction \mathbf{b} via $\mathbf{P}_{\perp b} = \mathbf{I} - \mathbf{b}\mathbf{b}^H/(\mathbf{b}^H\mathbf{b})$ ($\mathbf{P}_{\perp b} \mathbf{b} = 0$) as:

$$\mathbf{z}^a(t) = \mathbf{P}_{\perp b} \mathbf{x}(t) = s_a(t) \mathbf{P}_{\perp b} \mathbf{a} + \mathbf{P}_{\perp b} \mathbf{n}(t) \quad (3)$$

In this way, the two signals in $\mathbf{x}(t)$ are spatially separated into $\mathbf{z}_a(t)$ and $\mathbf{z}_b(t)$ without the influence of one on the other. We can then estimate the directions \mathbf{a} and \mathbf{b} and track the signals $s_a(t)$ and $s_b(t)$, respectively.

Consider the same example with two identical complex signals at $(\phi_1, \theta_1) = (45^\circ, 30^\circ)$ and $(\phi_2, \theta_2) = (135^\circ, 70^\circ)$, respectively, with lower SNR = 4.8 dB. **Figure 4(a)** shows the image and surface plots of the pseudospectrum of FBS-MUSIC. The FBS-MUSIC detects the two correlated signals, but their angular estimates are biased as $(\phi_1, \theta_1) = (32.00^\circ, 25.50^\circ)$ and $(\phi_2, \theta_2) = (133.50^\circ, 64.50^\circ)$.

Figure 4(b) shows the image and surface plots of the pseudospectrum of MUSIC applied in the nullspace of the first signal whose direction is known with an offset of 1° . After the projection nullifies the first signal, the second signal is detected with a more accurate angular estimate as $(\phi_2, \theta_2) = (133.00^\circ, 74.00^\circ)$.

Similarly, **Figure 4(c)** shows the pseudospectrum of MUSIC applied in the nullspace of the second signal using the direction estimate (ϕ_2, θ_2) from the previous step. After the projection nullifies the second signal, MUSIC detects the first signal again with a more accurate angular estimate as $(\phi_1, \theta_1) = (43.00^\circ, 30.50^\circ)$.

As shown in **Figs. 3(d)** and **4(a)**, FBS-MUSIC can work with correlated signals but may produce biased angular estimates. To further investigate, again consider two identical sinewaves with SNR = 6 dB, which have different elevations of 5° and 75° but the same azimuth of 155° .

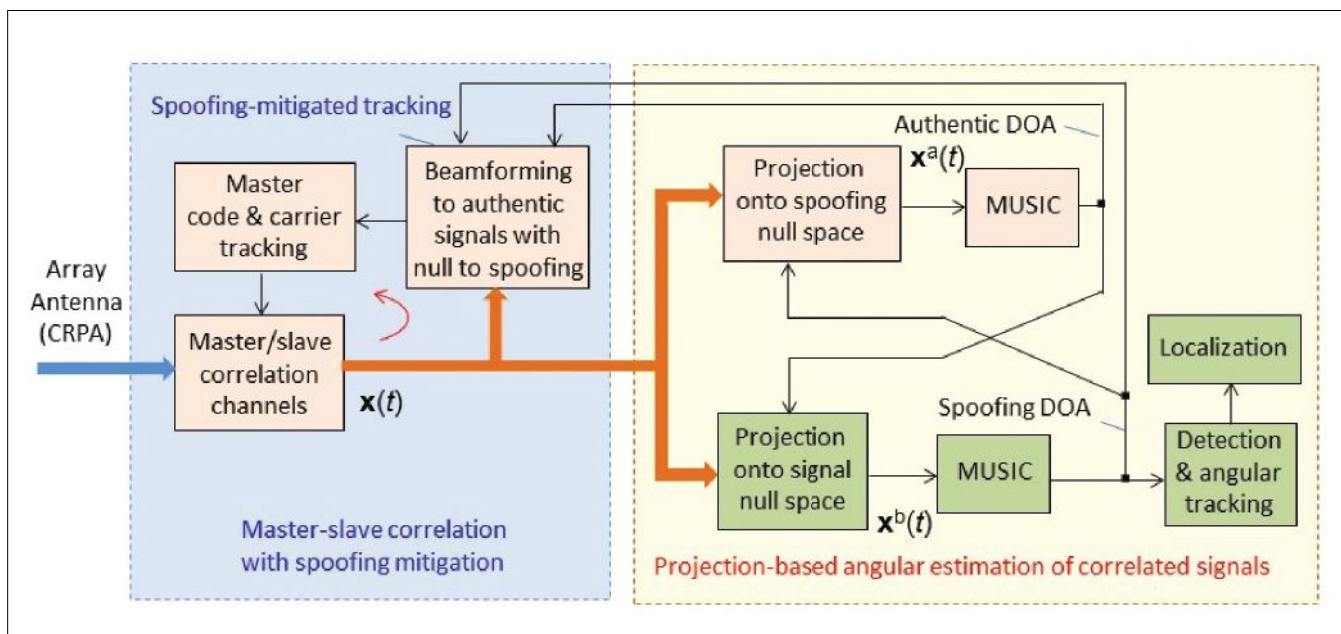


FIGURE 5 Projection-Based Method for Angular Estimation of Correlated Signals and Tracking.

Figure 4(d) shows the pseudospectrum of FBS-MUSIC, which has a dominant peak and a second elongated peak at (azimuth, elevation) = (194°, 9°) and (154°, 64°), respectively. Clearly, the angular estimates are biased.

Alternatively, we can project the array output to the nullspace of the first signal and then estimate the second signal and vice versa. An angular error of 1° is added to construct the nullspace projection matrix at (156°, 6°). The resulting pseudospectrum of MUSIC is shown in Figure 4(e), which leads to a more accurate angular estimate of (156°, 78°).


Similarly, construct another nullspace projection matrix at the angular estimate of (156°, 78°) of the second signal and then estimate the first signal. The resulting pseudospectrum of MUSIC is shown in Figure 4(f), which leads to an angular estimate of (150°, 6°). Clearly, the angular biases are largely removed.

Spoofing Signal Tracking and Angular Estimation for Localization

In GNSS applications, we start with no spoofing so a GNSS receiver can obtain good AOA estimates of authentic signals in view. The AOA of each authentic signal is

used to form a beam toward the GNSS satellite for tracking in a master-slave architecture on one hand, and to construct a projection matrix with a null toward the satellite on the other hand. In the satellite's nullspace, spoofing can be detected if present and its AOA is estimated, as illustrated in Figure 5.

In parallel, the AOA estimate of spoofing is used to steer a null toward the spoofer for tracking in the master channel on one hand, and to construct the projection matrix onto the nullspace of the spoofer for authentic signal tracking and AOA estimation on the other hand.

In this way, the spoofing signal is prevented from getting into the master tracking channel but left intact at the master-slave correlator outputs for its angular estimation and localization via triangulation. A tracking loop similar to that for the authentic signal in Figure 5 can be configured to track the spoofing signal and obtain a spoofed navigation solution for spoofing intent analysis. 

References

The online version of this article contains a list of reference papers.

Authors



Chun Yang is Chun Yang received his title of Docteur en Science from the Université Paris-Saclay (previously known as Université de Paris XI). After postdoctoral

research at the University of Connecticut, he has been working on adaptive array, synthetic aperture, and baseband signal processing for GNSS, radar, and communications signals and systems; nonlinear state estimation in target tracking, integrated inertial and distributed collaborative navigation; and optimization in resource management and information fusion. Dr. Yang is the recipient of 2007 ION Samuel M. Burka Award and the winner of 2009 IEEE NAECON Grand Challenge.



Andrey Soloviev is a principal at QuNav. His research and development interests focus on sensor-fusion and signal-processing implementations for

GNSS-degraded and GNSS-denied applications. He holds a Ph.D. in electrical engineering from Ohio University. He is a recipient of the Institute of Navigation (ION) Early Achievement Award and the RTCA William Jackson Award.



GPS Rack Mount Amplified Splitter

Ideally suited for timing and testing applications where the GPS carrier signal is required by up to 32 devices simultaneously.



● Test Labs

● Cellular Markets

● Public Safety

● Timing

*Available in 1x8 1x16 1x32
2x16 and 2x32

* Standard Splitter options available 1x2 1x4 and 1x8

Contact us for custom system design or more information
at 800-463-3063 or email salestech@gpsnetworking.com

WWW.GPSNETWORKING.COM



Real-Time Automated Aerial Refueling with Stereo Vision

Overcoming GNSS-Denied Environments In or Near Combat Areas

In-flight refueling requires sustained minimal separation between paired aircraft with little room for error. In or near combat zones, wide-area GPS-denial or spoofing means that an GPS-independent system must be available. Regardless of the selected sensor package, a common set of properties must be satisfied to facilitate mid-air docking: a high degree of accuracy, precision, and integrity.

At 0400 hours, a manned surveillance aircraft has been in the air for several hours. It approaches a refueling tanker, closing in for dock. Onboard the unmanned tanker, a computer automates the boom, maneuvering it to dock with the receiver—the trailing aircraft acquiring fuel from the tanker. Now fully refueled, the surveillance aircraft remains airborne and vigilant.

This scenario may be closer to reality than one may think. In June 2021, an unmanned MQ-25 Stingray successfully refueled a Navy Super Hornet. As simple as the scenario might seem, there are several factors required for successful automation of aerial refueling (AAR).

In the June Navy Stringray scenario, the autonomous MQ-25 tanker was equipped with a U.S. Navy drogue. This

drogue is essentially a basket extending from the tanker's rear via a fuel hose. The F/A 18 Super Hornet, a manned receiver, was responsible for approaching this basket and maneuvering its rigidly mounted fuel probe into the tanker's trailing drogue. Once docked, fuel transfer commenced. The MQ-25 is capable of flying via a GNSS-based flight path and generally holds a stable path while a receiver is attempting to dock with its drogue.

In-flight refueling requires sustained, minimal separation between paired aircraft with little room for error. In prior demonstrations, this has been achieved using differential GPS, but many different sensor packages could be employed with varying tradeoffs. Regardless of the selected sensor package, a common set of properties must be satisfied to facilitate

**JAMES ANDERSON, JOEL MILLER,
XIAOYANG WU, SCOTT NYKL, CLARK TAYLOR
AND WARREN WATKINSON**
U.S. AIR FORCE INSTITUTE OF TECHNOLOGY

midair docking: a high degree of accuracy, precision, and integrity.

Our research focuses on a tanker-centric stereo vision system for determining relative navigation between the tanker and receiver. Adding a vision-based system introduces several advantages. A second sensing system in addition to GNSS can be used for ensuring integrity. Fusing the GNSS and vision-based inputs may lead to improved performance over either individual system. Since refueling occurs near combat zones, if wide-area GPS-denial or spoofing occurs, employing an independent vision-based system allows refueling operations to continue.

A visual relative navigation approach also has some distinct advantages over a GNSS-based approach. To obtain cm-level accuracy or better, differential GNSS-based methods require constant communication and information sharing between the tanker and receiver. This implies modification of the receiver to enable AAR based on GNSS. A stereo vision technique modifies the tanker only; we assume no modification to the receivers. While the number of possible receiver aircraft is very large, the number of tanker aircraft is comparatively few. Therefore, the stereo vision approach may be easier to deploy and lower-cost to maintain than a GNSS-based technique. In addition, modern tankers are often already equipped with a stereo vision system to assist the human refueler, minimizing the quantity of new external equipment required for this approach.

Despite these advantages, the refueling receptacle on a receiver aircraft is only a few centimeters in diameter but approximately 25 meters from the ste-

reo cameras, requiring extremely high relative navigation accuracy to guide the boom. Additionally, we desire our approach to be completely independent of any GNSS system, adding redundancy and resiliency to the refueling system. Because refueling is a dynamic situation, any algorithm to automate mid-air docking must run in real-time.

Designing a system to meet these constraints requires innovations in several areas, including computer vision, machine learning, and parallel computing. In addition, verifying the system performance presents a unique challenge. Conducting real-world tests involving actual tanker and receiver aircraft is costly and time-prohibitive. As a starting point, virtual environments can be used to accurately simulate and verify the algorithm, as the truth data is known. Beyond virtual-only simulations, smaller scale real-world tests are also possible; however, knowing the

truth data becomes more challenging. As stated, any automated refueling algorithm must be accurate within a few centimeters, requiring a truth system with millimeter-level accuracy.

Approach

Our pipeline for estimating the receiver's pose (position and attitude) relative to the tanker while meeting these constraints, is shown in **Figure 1**. The pipeline consists of several components, some of which are standard image-processing steps (black solid lines), and some that have been customized to adapt to various aspects of AAR (red dashed lines). Before flight, preprocessing steps, including calibrating the stereo camera system and generating the reference point cloud, are performed. The first step of the algorithm is to acquire a pair of stereo images from 4k cameras in gray-scale. Stereo block matching (SBM) can be used to match similar pixels in the left and right

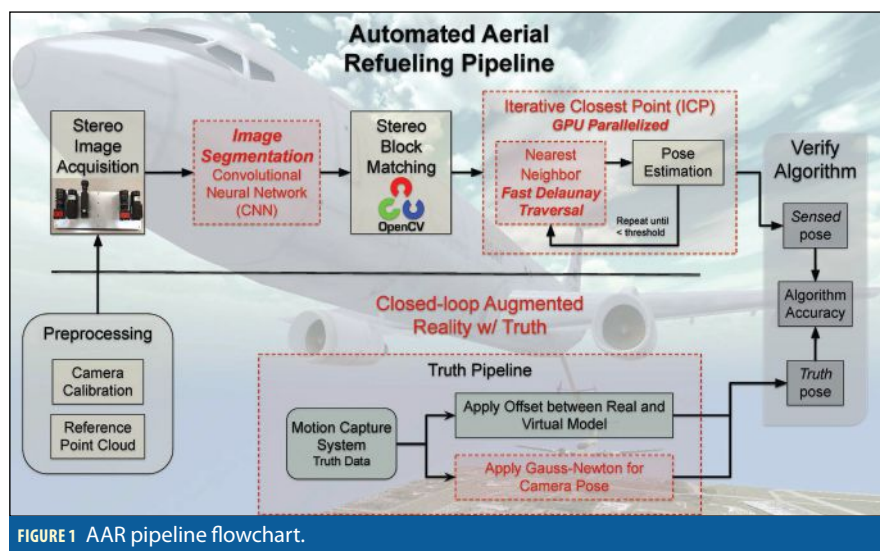


FIGURE 1 AAR pipeline flowchart.

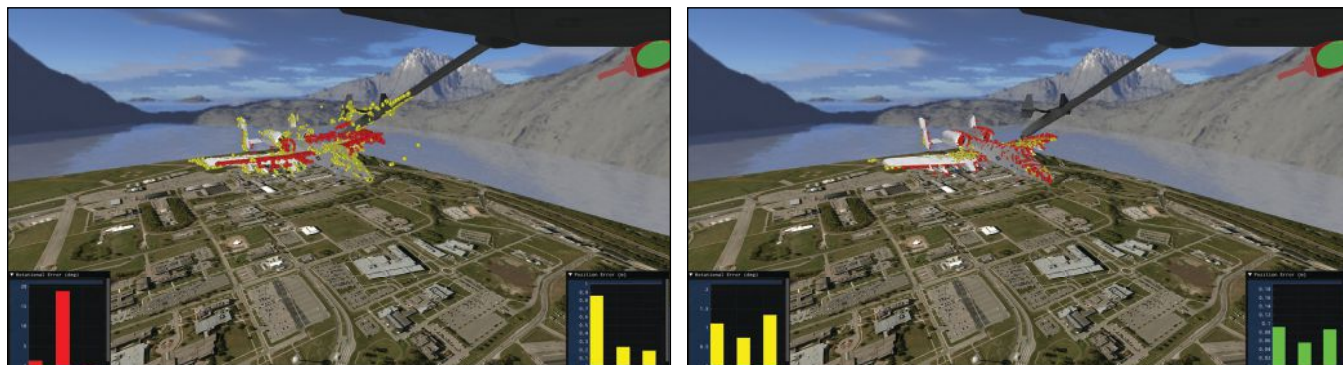


FIGURE 2 LEFT: An improperly converged red point cloud due to boom occlusions; note the spurious yellow point lying on the boom—these cause the errors. RIGHT: a properly converged red point cloud onto yellow sensed point cloud; the yellow sensed points on the boom have been filtered out.

images. The difference between the matches, or disparity, is used to reproject these 2D points into 3D space, forming a sensed point cloud of the receiver. (Figure 5 shows a sensed point cloud from an approaching receiver aircraft.) This sensed point cloud intrinsically encodes the pose of the receiver relative to the stereo camera system.

Optionally, the captured images can be fed into a convolutional neural network (CNN) trained to detect the cropping area of the receiver. This procedure allows the next step, SBM, to only evaluate pixels belonging to the aircraft; in our experiments this enhancement provides an 11x decrease in SBM computation when using 4k cameras.

However, one cannot directly extract the receiver's relative pose without further computations. To extract the relative pose from the sensed points, we employ point registration, specifically Iterative Closest Point (ICP). This process registers a reference point cloud onto the sensed

point cloud which computes the relative rotation and translation between the receiver and stereo cameras: $[R_{3 \times 3}, t_{3 \times 1}]$.

The reference point cloud is generated a priori from the known geometry of the receiver's airframe as shown in Figure 4. To automate selection of the correct reference point cloud, we employ CNNs to classify an approaching receiver and select its corresponding reference point cloud for the point registration process. Figure 6 shows a red reference point cloud registered onto a yellow sensed point cloud. Similarly, the right of Figure 2 shows a properly registered reference model lying on top of a sensed yellow point cloud. In turn, the sensed yellow point cloud is lying upon the surface of the receiver's actual airframe indicating a good reprojection of sensed 3D points. When all three align, our algorithm is functioning as designed. In the case of a simulation or virtual world, we know the pose of the actual receiver, so we can directly quantify the errors result-

ing from the stereo sensing and point registration.

To verify the algorithm produces an acceptable output, the algorithm's sensed pose is compared with the truth pose. To collect accurate truth data, we have developed a closed-loop augmented reality environment. First, a motion capture system (MCS) shown in Figure 3 provides the pose of the receiver. This pose is with reference to the MCS, so the virtual representation requires an offset to line up with the truth. This offset is determined via a novel use of the Gauss-Newton optimization technique.

Preprocessing Calibration

Possibly the most critical issue for obtaining accurate pose estimates is obtaining a precise intrinsic and extrinsic camera calibration. The intrinsic calibration matches the mathematical pinhole model to real camera attributes, accounting for properties such as focal length, lens distortion, optical center, and image resolu-

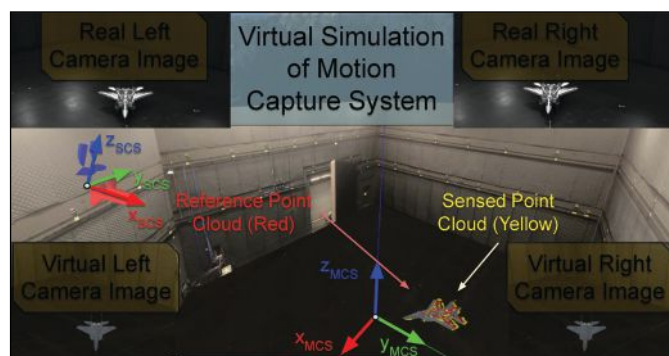


FIGURE 3 Virtual simulation of motion capture system. The top left inserts are images captured from real cameras in the MCS. The bottom images are virtual renderings using the same attributes of the real cameras (labeled on the left axes). The sensed points in yellow and the reference points in red depict the estimated pose. They lie on the virtual rendering of the aircraft, which is positioned using the truth pose obtained from the MCS.




FIGURE 4 Full reference model (left) and shelled reference (right).



FIGURE 5 A point cloud sensed by the tanker's stereo cameras.




FIGURE 6 A reference point cloud (red) corresponding to sensed point cloud (yellow).



SYNTONY
GNSS




The full power of **Constellator™** ready for CRPA Wavefront & Anechoic Testing

Achieve high GNSS dynamics testing, performance & scalability

- ▶ All known GNSS signals/bands and ready for future PNT signals
- ▶ 7+ Antenna Elements in conducted or over-the-air (OTA)
- ▶ Picosecond Antenna Synchro Phase
- ▶ 700+ of simultaneous & independent High Fidelity RF signals
- ▶ > 100 dB Jammer to GNSS signal ratio
- ▶ Advanced Jamming/Spoofing transmitters performances
- ▶ I/Q data stream

Request more info on **syntony-gnss.com/contact**

Or flash this
QR code:



TOULOUSE - PARIS - SAN FRANCISCO - NEW YORK - MONTREAL



FIGURE 7 Real image segmentation with a CNN. The green box indicates the truth bounding box and the purple depicts the CNN's prediction.




FIGURE 8 Delaunay triangulation of reference point cloud.

tion. The extrinsic calibration enables us to precisely determine the stereo camera baseline (horizontal distance between the cameras) and relative orientation between the left and right cameras. Once both the intrinsic and extrinsic parameters are known, we can un-distort and rectify image pairs and re-project matched features found in both the left and right images into 3D points. These 3D points form the yellow sensed point cloud shown in **Figure 3**. Note that performing calibration is a well-studied field and we utilize algorithms already implemented in OpenCV.

Reference Point Clouds

To estimate pose using a sensed point cloud, a reference point cloud of the receiver must be present to align the sensed point cloud with. Because we assume the type of aircraft being refueled is known, we assume a point cloud of the receiver can be created a-priori. In addition, because the receiver will essentially be observed from one perspective—above and in front of the receiver—we can reduce the reference model to include only points which can be seen from the refueling point of view. The result is a shelled reference model. In **Figure 4** the full set of points is shown next to the shelled reference model. As seen in the top pairs, the point clouds look very similar. However, in the bottom images, several points located on the side of the aircraft appear in the left image but not the right image. From the point of view of the tanker, these points would not be detected by the stereo cameras. Removing them from the model has the advantage of improving

accuracy as well as increasing the computation speed, as these points are not considered during the model registration steps in ICP.

To further optimize computation time when performing ICP, the shelled reference model is further processed in two ways. First, a k-d tree typically is made that allows for searches of the data in approximate logarithmic time, $O(\log n)$ when no information about the aircraft pose is known. Second, we create a Delaunay triangulation of the point enabling us to run the nearest neighbor queries of ICP in an amortized constant time of approximately $O(1)$.

Stereo Image Acquisition

When attempting to refuel, the first step in our AAR pipeline is to capture a pair of stereo images. There are two items of particular import. First, the timing of the camera capture is extremely important. Both cameras need to capture the image at exactly the same time or else the sensed point cloud will have significant errors in the point cloud. Furthermore, when evaluating the performance of the AAR pipeline, the capture time and the timestamps for the truth data must be precisely aligned. Second, the resolution of the captured images significantly affects the final accuracy of the system. For the purposes of AAR, we began our experiments with 1280x960 resolutions, but we eventually transitioned to 4k cameras to achieve the accuracy required at the desired standoff distances. This increase in pixels, however, leads to a significant increase in computational requirements of the SBM step. Our approach to mitigating this issue is

described under Image Segmentation.

Stereo Block Matching

After images have been acquired, pose estimation based on that imagery is the key to accurate relative navigation. Many approaches to estimating the pose of known objects use features from the image and map them to the same features lying on the geometry's surface. These techniques require features on the object being tracked be visible in all imaging conditions (lighting, shadows, glint, etc.).

In addition, the fewer the number of features, the more accurately the features must be found in each image, leading to a sub-pixel accuracy requirement in many scenarios.

Alternatively, we chose to create a dense point cloud that generates thousands of points in 3d space (generally one point per pixel on the receiver aircraft), thereby averting sensitivity to improper matches. We found this approach to be more applicable to AAR for two reasons: (1) we assume no beacons exist nor can they be added to any aircraft, making the unique feature identification problem very difficult and (2) we believe that achieving accuracy through the use of 1000s of features is more robust to misidentification and image acquisition conditions than techniques based on fewer features. The creation of a dense point cloud is a well-studied problem and is solved utilizing SBM techniques found in the open source computer vision library OpenCV.

SBM employs two cameras with known separation geometry where each camera captures the scene from a slightly different perspective. A feature in one

image maps to an epipolar line in the other image. From these corresponding pair-wise matches, a disparity map is generated. The disparity map, in conjunction with the calibration parameters, can be used to reproject each pixel to a 3D location. This produces a 3D point cloud for each image pair. In this work each image pair produces about 40,000 points. Unfortunately, this many points makes the next step take too much computation time. Therefore, we chose to subsample the data to 5,000–6,000 points as this value appears to balance accuracy and computation time. These (Figure 5) sensed 3D points now serve as input to the pose estimation/registration algorithms.

Iterative Closest Point

In our algorithm, we have two point clouds to consider. The first is the yellow sensed point cloud generated from our stereo vision camera; this point cloud will be somewhat noisy as it is the output of our SBM process, see Figure 5. The second point cloud is the red reference model containing the ideal geometry of the receiver aircraft. This reference model's point distribution is chosen based on the stereo camera resolution and the expected receiver distance from the camera at the docking location.

Importantly, the sensed point cloud intrinsically encodes the relative pose between the stereo camera and the receiver. To extract this relative pose from the sensed point cloud, we estimate the rigid transformation (rotation

and translation) that best aligns the reference model onto the noisy, spurious sensed model, i.e., transform the red points onto yellow points as shown in Figure 2. Because we know the reference model is accurate, if we find this transformation we then arrive at a relative pose between the camera and receiver. Because we know the geometry of the reference model, this also estimates the specific location of the docking receptacle for AAR.

To estimate the aforementioned registration, we use a variant of the ICP algorithm. ICP works by first finding, for every point in the sensed point cloud, its “nearest neighbor” in the reference point cloud. An example correspondence is shown in Figure 6. After this, a summed outer product matrix is constructed from the outer product of each nearest neighbor pair. Decomposing the outer product matrix into orthonormal rotation matrices via methods such as singular value decomposition (SVD) yield a rigid rotation representing the change in orientation between the two point clouds. The vector between each cloud's center of mass yields a translation vector; thus, the rotation and translation produce the desired transform. After applying this rotation and translation, the nearest neighbor matches may have changed, so the process repeats. At the end of this iterative process, the rotation and translation that best maps the sensed point cloud onto the reference point cloud is output as the estimated pose of receiver.

Real-Time Execution

Because of the need for an AAR algorithm to run in real-time, the standard algorithms described above do not work “out of the box”. We have added a block between the image acquisition and SBM blocks, and modified the ICP algorithm to enable real-time performance.

Image Segmentation via CNN

When performing SBM, the expected error at a given range can be calculated

$$e_z = \frac{z^2}{bf} \varepsilon_d \quad (1)$$

where e_z is the depth error, z is the depth, b is the baseline, f is the focal length (in pixels), and ε_d is the matching error in pixels (disparity values, assumed to be one). While the field of view of the sensors, the depth, and the baseline are all fixed by the AAR scenario, we can increase the focal length f by increasing the number of pixels within the same field of view of the camera. AAR's long ranges and need for high accuracy require a relatively high resolution.

The downside of higher resolution images is the computation time required for SBM to generate the disparity map. However, a depth map is not required for everything in the capture images, just the receiver. At a distance of 25 meters, the receiver will generally not fill the entire image frame, causing SBM to waste significant computational resources. Therefore, we designed a method to efficiently find the receiver within the broader image and limit SBM

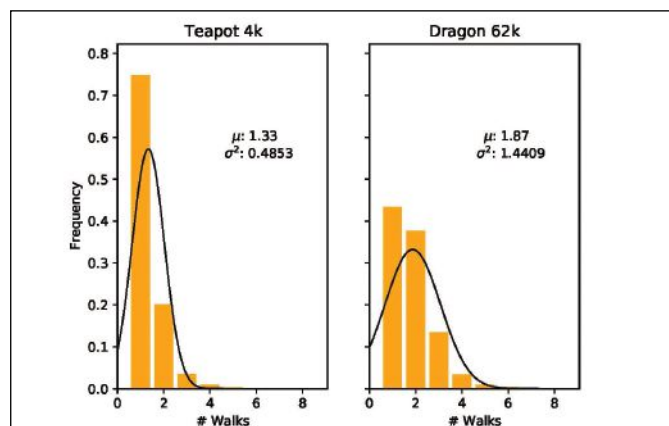


FIGURE 9 Average number of Delaunay walks to find the nearest neighbor. Increasing the number of points 10x only increases the number of walks needed by approximately 40%.



FIGURE 10 A virtualized refueling approach in the AftBurner engine.

to operate on that portion of the image.

To segment the image and reduce computation time required for SBM, a CNN was trained to output the bounding coordinates for the receiver. Generating images from actual refueling approaches is costly, and truth data is difficult to obtain. Therefore, the CNN was trained on images of real landscape scenery with a virtual aircraft rendered at a random but AAR-friendly pose. Once trained, we were able to utilize the CNN model on real images without adjustments. **Figure 7** shows the results of running a real image through the CNN. The green lines represent the truth bounding box while the purple lines denote the CNN's prediction. This technique led to an 11x decrease in SBM processing time for 4k images.

Parallel ICP on GPU

Aside from SBM, ICP also presents a barrier to real-time execution. To address this issue, we implemented a novel variant of the ICP algorithm using Nvidia's CUDA to enable a highly parallelized implementation on a graphics process-

ing unit (GPU). ICP's nearest neighbor (NN) matching step is, by far the slowest step, but also inherently parallel. This step iterates over the entire sensed yellow point cloud, and for each point, finds for the closest red point residing in the red reference point cloud. **Figure 6** shows these matched correspondences as the purple lines. Each of these NN searches can run in parallel on their own thread. By choosing a GPU with more Stream Processor Units than sensed yellow points, we effectively reduce a naive $O(n^2)$ algorithm to $O(n)$, at least with respect to time complexity. Beyond NN matching, the subsequent step, Rotation/Translation estimation, requires a summation of displacements from each pair of points, requiring all computation between points to be completed before decomposing the outer product matrix. Therefore, we leverage the parallel reduction: each task is split into sub-tasks with the ultimate goal of reducing all data into a single value. Executing ICP on a GPU reduced computation time by roughly 95x over serial versions. Although a vast

improvement of almost 2 orders of magnitude, ICP was still not fast enough to run in real time while registering point sets on the order of approximately 10,000 points. To overcome this hurdle we developed a novel Delaunay Triangulation algorithm.

Delaunay Triangulation

While parallelizing the ICP algorithm led to significant reductions in run time, our AAR algorithm required we register our sensed and reference point clouds (on the order of 10,000 points) in less than 20 msec. Analyzing the ICP algorithm, we found that ICP's nearest-neighbor search consumed the vast majority (>97%) of computation time. Fortunately, we were able to develop a novel algorithm that dramatically decreased the nearest neighbor search—amortized over time, this algorithm produced an $O(1)$ time complexity for any given NN query.

Traditional methods use tree-based approach such as, a k-d tree which partitions the data along the median value of each sequential axis. A CNN query using a k-d tree is on average is $O(\log n)$ with the worst case being $O(kn(1-))$ (for 3D points, k is 3) for points in particularly poor locations.

We investigated an alternative correspondence search utilizing a Delaunay triangulation of the reference points. A Delaunay triangulation turns a set of points into a graph with a specific structure that is useful for finding nearest neighbors. The graph is connected such that if a query point is closer to a node in the graph than all of its neighbors, then that node is the nearest neighbor of the




FIGURE 11 Inside the Motion Capture Chamber with 1:7 scale model aircraft and IR reflectors.




FIGURE 12 A virtual boom as presented in the virtual and the real world.

query point. Our Delaunay traversal algorithm walks along edges in the graph until there are no other points closer to the current candidate point. The main advantage to this approach is the ability to begin the search from any given point. In a k-d tree, searches must typically begin at the root node, leading to a best-case $O(\log n)$ run-time. In our Delaunay traversal, the search can begin from any arbitrary point. Because ICP is an iterative algorithm, we can start the NN search from the previous iteration's nearest neighbor. In general, this cache-friendly approach leads to a very small number of neighbors traversed per iteration.

In fact, our experiments have shown that on average, only 1-2 walks are required for each subsequent ICP iteration (Figure 9). Additionally, we have seen this performance regardless of the size of the data set. Thus, our Delaunay traversal approach reduces the $O(\log n)$ time complexity to an amortized $O(1)$ time complexity. Evidence of this behavior is seen when analyzing the average number of walks taken for point clouds of various sizes. In Figure 9, a 3D teapot model with ~4k points is shown to have taken 1.33 walks on average to find the nearest neighbor. In comparison, a 3D dragon model with ~62k points took 1.87 walks on average. Using this approach together with the parallelized ICP, we have been able to complete 30 iterations of the ICP algorithm in around 18ms for point clouds with ~10k point, making ICP computation a real-time possibility for AAR.

Low-cost Verification of AAR algorithms

As new algorithms are being developed and refined for AAR, the problem of verifying the performance of these algorithms becomes essential. Unfortunately, real flight tests are expensive, time-intensive, and may not provide perfect truth data needed for accuracy testing (differential GPS gives centimeter-level accuracy, but truth should really be closer to millimeter-level or better). Therefore, we have adopted two approaches to verifying our algorithms' performance.

First, we have constructed a virtual refueling scenario using the AftBurner (3D visualization/game) engine to test our vision pipeline. A geometrically

accurate refueling tanker is placed in the 3D virtual world with a pair of stereo cameras attached to the rear of the tanker. A receiver aircraft initiates a refueling approach and flies towards the tanker as shown in Figure 10. The virtualized stereo cameras capture image pairs and feed them to the AAR pipeline (Figure 1) to produce estimated poses for the receiver aircraft. This virtual environment is convenient as it provides absolute truth and the calibration, timing, and other issues are easy to resolve.

Second, we desired a test setup that utilized real stereo sensors so that the processing performance can be representative of real-world scenarios. In addition, we needed the ability to test closed-loop performance when movement of the boom (which will occlude key parts of the image) is added to the system. Therefore, we developed an augmented reality system.

Augmented Reality Setup

For our augmented reality framework, we utilized a motion capture chamber (MCS) that is approximately 15×20 meters in size (Figure 11). The motion capture system, through the use of 50 infrared cameras placed in the chamber, can track the position of markers (infrared reflective balls) anywhere in the room to about 1mm accuracy at a rate of 75Hz. By attaching multiple markers to an object in a known configuration, accurate estimates of attitude are also returned.

To simulate AAR, we mounted a pair of 4K stereo cameras approximately

8 meters high in the room to simulate a view akin to a tanker observing an approaching receiver aircraft. A 1:7 scale realistic model of a receiver aircraft is placed in the chamber with markers attached, see Figure 11.

To conduct an approach in the MCS room, we physically pull the real-life 1:7 scale aircraft towards the stereo vision cameras (Figure 13: Red Box). This simulates a receiver approaching the tanker; the MCS is large enough that we can evaluate camera-receiver ranges of 10-25 meters—these ranges are ideal for testing the most crucial phase of AAR docking and refueling. We capture the receiver's approach from two time-synchronized 4k cameras at 10Hz. We time-align and associate each captured image pair with the receiver's absolute truth pose provided by the MCS. This data is then streamed in real-time to the Augmented Reality AftBurner virtual world.

The physically acquired images (Figure 13: Blue Box) are fed into the aforementioned AAR vision pipeline, resulting in the yellow (sensed) dots that shown in Figure 13 in the yellow box. The pose obtained by the MCS is used to update the position of the virtual receiver aircraft in the Augmented Reality Virtual world. This is the gray virtual model that is somewhat hidden underneath the yellow points shown in Figure 13. Ideally, the sensed yellow points ought to lie perfectly on the skin of this virtual aircraft (the gray textured aircraft). Such behavior indicates the MCS and the stereo cameras are aligned and share a consistent

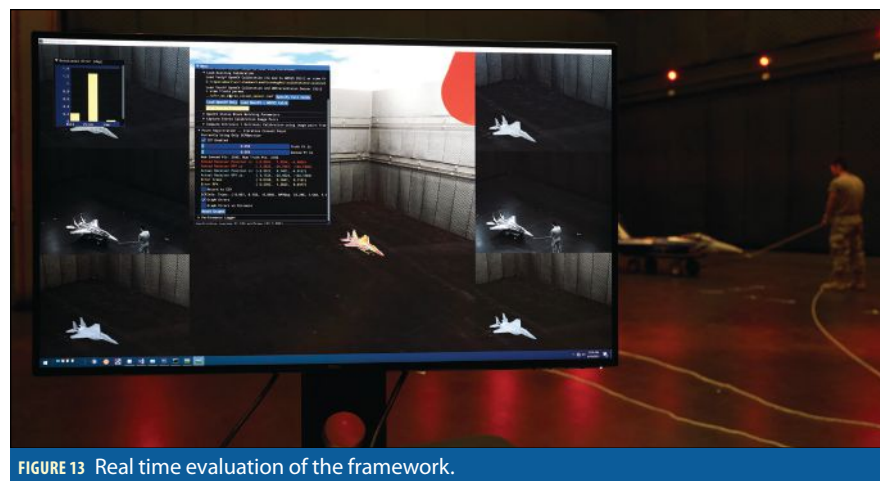


FIGURE 13 Real time evaluation of the framework.




FIGURE 14 Estimated corner coordinates (orange dots) before and after Gauss-Newton Optimization.

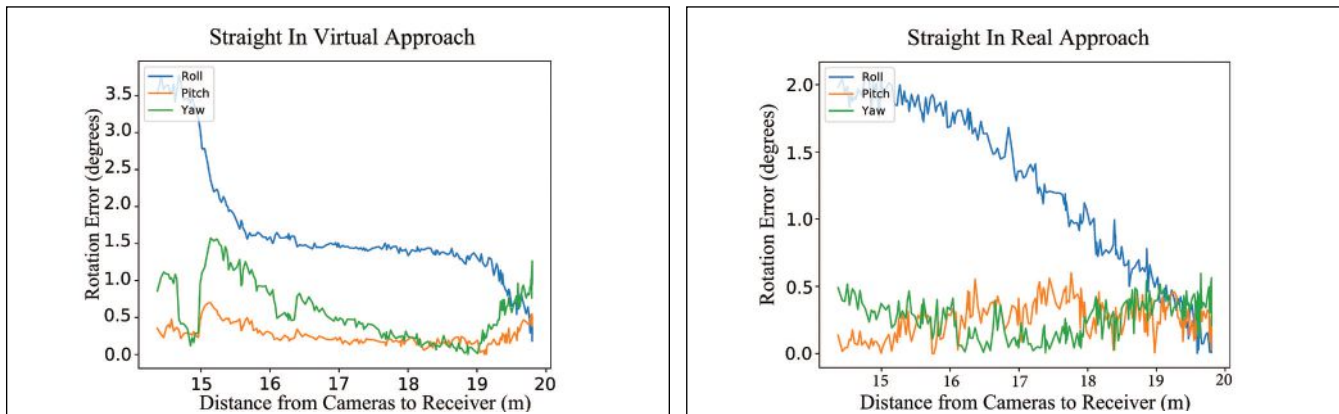


FIGURE 15 Rotational errors of the receiver aircraft in a straight approach.

coordinate frame. The final 2 minutes of a YouTube video at <https://youtu.be/Qw0W8qlOW8sc?t=213s> show this process unfolding in real-time. The final element is the red point cloud shown in the video and in Figure 13 in the yellow box. The red reference point cloud is registered onto the sensed yellow point cloud using our ICP variant. This computes the final pose sensed solely from the stereo cameras. This pose is compared against the pose from the mm-accurate MCS system to yield our total positional and rotational error. When the yellow point cloud lies nicely on the gray textured skin of the 1:7 scale receiver and the red points lie on the yellow points, the system is working well, typically with less than 3-5cm of error and less than 2 degrees of rotational error.

Beyond real-world tests, we can recreate the same path taken by the receiver inside the AftBurner engine directly using the MCS truth data. This enables us to verify the performance of our virtual simulation engine as we can directly compare the performance on both real and virtual imagery. Perhaps more important, though, is the potential this

system gives us for testing closed-loop performance. Rather than attempting to build a (large and unwieldy) 20m boom, we can virtually add the boom to real images captured to test performance of the system when the boom is being controlled using information from the visual processing pipeline. This gives rise to our augmented reality setup, as shown in Figure 12.

Gauss-Newton Optimization

Key to evaluating the performance of AAR in the MCS room is the ability to transform pose estimates from the stereo cameras into MCS-based truth coordinates. To find this transform, we employed a Gauss-Newton optimization technique designed to generate an accurate pose of the real stereo cameras in the MCS coordinate frame. To find this transform, we created a checkerboard that had several IR markers on it, allowing it to be precisely localized within the MCS room.

The Gauss-Newton approach minimizes the errors between the measured image pixel coordinates of corners on a chessboard in captured images and

where these corners are projected into the cameras based on the current transform from MCS to camera coordinates.

Before the optimization, the average distance between the estimated corner coordinates and the measured corner coordinates was about 280 pixels, see the left side of Figure 14. After optimization, the average offset was about 3 pixels, see the right side of Figure 14. The optimization allowed us to generate a good camera calibration for the real stereo cameras which minimizes reprojection errors and results in better convergence of the ICP algorithm.

With the transform output from the Gauss-Newton optimization, we achieved optimized alignment between the AftBurner generated imagery using truth data and the real-world captured imagery. Without this excellent alignment, the augmented reality would suffer biases, with this excellent alignment, objects virtually projected into the real world are within a few mm of where we expect them to reside—quite useful when quantifying features that are occluded by the motion of the virtual boom projected on to the real-world 1:7 scale receiver.

Testing Augmented Reality Framework

We conducted tests of the augmented reality environment with the physical 1:7 scale receiver approaching the stereo cameras in various patterns. For example, a straight-in approach involved pulling the physical 1:7 scale replica slowly towards the stereo cameras to simulate a real-world refueling approach (same approach shown in the aforementioned YouTube clip). The vision pipeline takes the imagery and produces pose estimates. The pose estimates are compared with the truth system poses to generate error statistics. These same truth values can also be used to generate virtual imagery, enabling the vision pipeline to produce pose estimates. In **Figure 15**, the error results for rotation estimates for both the virtual and real imagery are shown. Note that the overall magnitude and error characteristics of the two charts are very similar.

When analyzing several approaches collected via our augmented reality framework, we see a similar trend of pose errors between the real and virtual worlds. The virtual environment serves as an oracle: it predicts failure points of the real-world imagery and predicts accuracy within a small epsilon of the real-world measurements. When combined with the motion capture chamber and the truth data, the framework enables the approach collection to serve as an arbiter of truth and quantify any desired pose-estimation algorithm's efficacy.

Virtual Boom

Another feature of our framework is the incorporation of virtual objects perspectively reprojected onto real-world imagery; ie, augmented reality. **Figure 12** showcases a virtual refueling boom perspectively reprojected on to real imagery of the approaching 1:7 scale receiver model. This allows us to quantify the degradation caused by boom occlusion without requiring a physical boom to be installed.

Because our augmented reality is designed with ease of adaptability in mind, we can replace the virtual Air Force refueling boom with other refueling tips such as the Navy's probe and drogue. This enhances our capability

to test various vision algorithms on other flying platforms. It also enables us to generate truth data for real approaches and augment that truth data with perspective-correct occlusions. This truth data can then be reprocessed by any number of vision algorithms to quantify the efficacy of a suite of relative vision-based navigation algorithms.


Conclusion

We employ stereo vision cameras as an alternative to traditional GNSS-based relative positioning algorithms for automated aerial refueling. While many portions of the solution consist of relatively standard computer vision techniques, several novel modifications have been made specific to the AAR problem. Specifically, we have 1) employed CNNs to identify the air frame of an approaching receiver. Subsequently, another CNN 2) computes a smaller region of interest around the receiver dramatically reducing the number of pixels stereo block matching processes.

We have developed a 3) novel Delaunay-based nearest-neighbor algorithm that transforms a time-intensive search to an amortized constant time operation, enabling real-time ICP on tens of thousands of points. We further enhance this algorithm via 4) GPU parallelization. This real-time functionality gives rise to a closed-loop augmented reality system that 5) enables virtually reprojected objects, such as a refueling boom or refueling drogue, to dynamically respond to motion of a real aircraft.

Similarly, the augmented reality combined with the high-accuracy truth system lets us 6) generate truth data sets for occluded approaches. This lets us evaluate different vision algorithms against each other as we have mm-level accurate truth. Finally, the virtual simulation is high enough fidelity that it 7) serves as an oracle able to reliably predict the outcome of real-world approaches.

Conclusion

The authors thank Daniel Schreiter and the Air Force Research Laboratory Aerospace Systems (AFRL/RQ) directorate for their support. 

Authors



James Anderson

received his Master's degree in computer engineering from Wright State University and is working on his Ph.D. there. His research work in collaboration with AFIT is on simulating and analyzing automated aerial refueling.



Captain Joel Miller

commissioned in 2017 from University of Colorado, with a B.S. in electrical engineering. At AFIT he is working on his Master's degree in computer engineering where his research is focused on automated aerial refueling.



Xiaoyang Wu

is a graduate student in software engineering at the AFIT. He received his B.S. in computer science from Manhattan College. His research involves augmented reality in conjunction with image processing and real-time visualization.



Scott Nykl

is an associate professor of computer science at AFIT. His areas of interest are real-time 3D computer graphics, computer vision, sensor fusion, parallel processing and interactive virtual worlds. He has a Ph.D. in computer science from Ohio University.



Clark Taylor

is assistant professor, computer engineering at AFIT. He has a Ph.D. in electrical and computer engineering from UC San Diego. He has published papers in video processing for UAVs, estimation theory, video comms and digital systems design.



Lt. Col. Warren

Watkinson

is Chief of the Sensor Plans and Advanced Programs Division at AFIT where he oversees advanced field and flight tests for air and space vehicles and a research portfolio of \$1.5 billion. He has an M.S. in computer science from Ohio State University and does Ph.D. research at Colorado School of Mines.

NOBODY'S FOOL

JAMMING AND SPOOFING DETECTION



GNSS professional users are becoming more and more aware that they have a problem—or, that they *may* have a problem. They suspect something is wrong, they're seeing anomalous behavior in their positioning domain. But how can they tell if they are being jammed or spoofed?

What's going on?!

Increasingly, professional and industrial users struggle to work around unexplained GNSS outages. Often, they don't realize an interferer has popped up in their midst. Situational awareness is now a key component whenever GNSS is employed. Am I being jammed? Am I being spoofed? From where, by whom?

These disruptions will only escalate over time. Fortunately, help is at hand in the form of an Interference Tool Kit and other countermeasures built into advanced receivers.

This free webinar describes a wide range of spoofing attack scenarios and their features, as well as the results in testing against these scenarios. It gives recent examples

of very dangerous, threatening behavior encountered in the field, and it offers a solution. Downloadable on demand, the webinar provides a solid technical grounding in the electronics of jamming and spoofing, and explores the tools available to detect, geolocate and mitigate such interference, whether intentional or unintentional. Actual field use cases are described, showing just how widespread such outages are becoming.

A spoofing detection unit already onboard a high-precision receiver collects metrics from the GNSS signal processing chain and provides a real-time indication if the receiver is under spoofing attack, at one-second intervals.

To ensure the resilience and integrity of your GNSS positioning, profit from the experience of our three expert speakers, each with years of experience combatting interference. Learn valuable lessons now to ensure GNSS resilience and integrity.

THE PANEL

Three experts provided diverse, complementary

perspectives on this dynamic and rapidly evolving applications area which spans all GNSS industry sectors.

LOGAN SCOTT is an expert consultant in systems/signal processing in advanced RF systems including GPS, RFID, navigation, communications, radar, and emitter location systems. A Fellow of the Institute of Navigation and holder of 45 US patents, he is the inventor of the Chips Message Robust Authentication (CHIMERA) signal concept for navigation-signal authentication.

NEIL GEREIN is Senior Director of Marketing at Hexagon's Autonomy & Positioning division, where he has worked for the past two decades, as aerospace & defense product manager, GPS systems engineer and other roles. He holds a Masters degree in electrical engineering from the University of Saskatchewan.

CHRIS MAYNE is Managing Director at Forsberg Services Ltd, a European positioning, navigation and timing component and systems provider, having been with the company since 2003. He holds Masters degrees in professional practice management and leadership and in mobile game design and M-commerce from Lancaster University.

WHY WOULD SOMEONE DO THAT?

Logan Scott opened the discussion by exploring the motivations for interference and the means of detecting it. The underlying message of his talk was, you may not be the target of the interference, but you will still be the victim: you and your work will still suffer, so you need to take appropriate countermeasures.

As an example, he played back an online video showing a teenager "going from zero to operational in about 10 minutes. He knows Linux but he's not an expert on GPS. He finds some routines, and if you watch this video it's really remarkable, 10 minutes in, his phone is saying it's in Cuba, and he really has no expertise in GPS. So, with software defined radios (SDRs), we are finding spoofing is available to basically anyone—people who are not expert in the subject."

Spoofing can cover a variety of criminal activities, and the motivation for most spoofing appears to be just that: theft or hijacking or illegal operations of other kinds.

THE PANELISTS



Alan Cameron
Editor in Chief
Inside GNSS



Logan Scott
GPS/GNSS expert consultant
Lonestar Consulting




Neil Gerein
Senior Director of Marketing
Hexagon Autonomy & Positioning



Chris Mayne
Managing Director
Forsberg Services Ltd.

PARTICIPANTS BY INDUSTRY

A diverse audience of professionals attended the webinar:



TO HEAR MORE, DOWNLOAD THE WEBINAR AT:

<https://novatel.com/tech-talk/webinars/nobodys-fool-jamming-and-spoofing-detection>




FIGURE 1 GRIT can detect a spoofing attack before it affects the position.

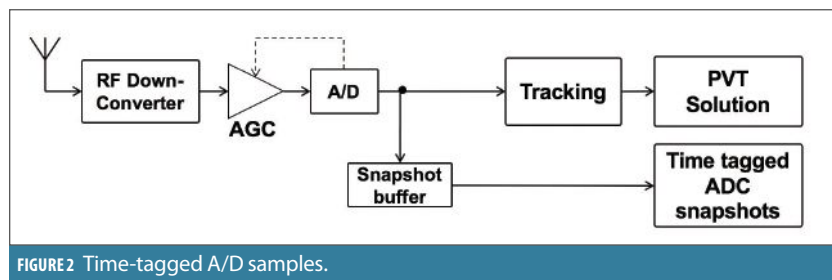


FIGURE 2 Time-tagged A/D samples.



FIGURE 3 Jamming detection array on highway by NovAtel headquarters.

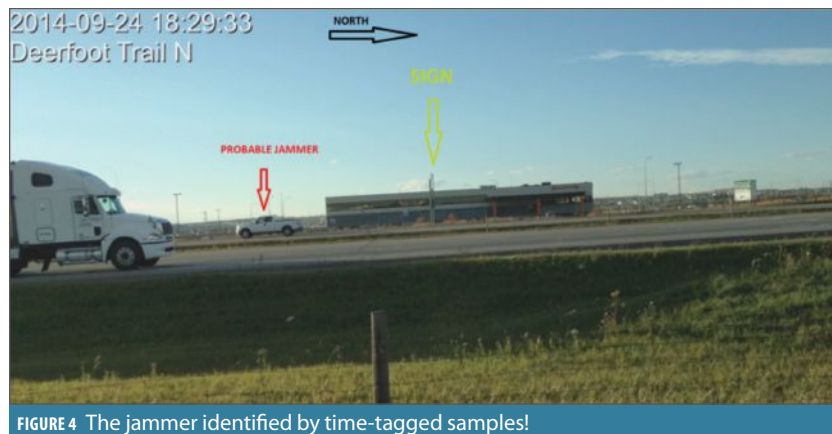


FIGURE 4 The jammer identified by time-tagged samples!

“One of the things that is really evolving right now is we’re seeing a lot more use of software defined radios,” continued Scott. “They have become inexpensive, but at the same time they’re capable of very sophisticated waveforms.”

The readily available HackRF gadget, open-source hardware for software-defined radios, serves as an instant jammer or spoofer. Essentially, a HackRF is a very cheap GNSS simulator.

He then covered the basics of jamming and spoofing detection. Basically, the receiver is your first line of defense.

RECEIVERS AT THE READY

At this point, Neil Gerein from Hexagon’s Autonomy & Positioning division took over, to explain the remarkable qualities and capabilities of the OEM7 receiver, with its on-board GNSS Resilience and Integrity Technology (GRIT, see Figure 1.)

GRIT jamming and spoofing detection and mitigation capability rides aboard every OEM7 receiver, and it can be accessed via firmware upgrades, in tailored configurations suited to the user’s needs.

Gerein then explained the jamming-detection methodology of timetagged snapshots of raw digital samples (see Figure 2). “Now you can look at your data not only in the frequency domain but in the combined frequency-time domain. This can give you a very exact timetagged and catalogued way to say what the particular interferer was.”

He illustrated this with a use case from his own back yard, so to speak.

“We had opportunity to do detective work along those lines at our

headquarters in Calgary,” he recalled. Staff had noticed some anomalous power measurements in the company’s reference network, happening every day for a few minutes in the early evening. They deployed a smaller reference network alongside the adjacent highway, where they suspected it was happening (**Figure 3**).

When they played the data back through some post-processing and employed the time-tagging method, along with a synchronized video camera that had also been monitoring that stretch of highway, they were able to pick out the jammer (**see Figure 4**): a white pickup truck that drove by everyday around six o’clock.

“This is right under the main flight path of the Calgary International Airport,” said Gerein. “This person may be just using the boss’s vehicle outside of regular work hours, but they’re potentially causing interference at a much larger level.

“It’s not just sophisticated jammers anymore. It’s everywhere.”

AWARENESS

Chris Mayne from Forsberg Services Ltd. in the UK then covered end-user awareness of GNSS interference, with more use cases from his company’s experience, in particular working with the British police in their efforts to track stolen vehicles, and in tracking down interference that was causing survey drones to crash.

He posed a question for thought to conclude the program. “We have experts on the ground demonstrating increased knowledge, and we have all sorts of publications and news reports [about jamming and spoofing]. But the question is, what combination of information, formats and alerts are


WHAT ATTENDEES WANTED TO KNOW

**Participants asked live questions during the webinar.
Here are a few of them:**


- Could you provide information as whether Chimera signals from NTS will be on offer to the civil community for test purposes?
- Could you comment on cybersecurity protections at the receiver level?
- Is an atomic clock necessary for the time tagging of A/D samples??
- How receptive is the police community to using interference reports as an alarm to possible criminal activity. Are they aware of its potential?
- Once a user is successfully alerted to an interference incident, what can they do to differentiate between malicious and unintentional jamming?
- Do the experts have any idea how the circle spoofing incidents affecting ships in the ocean were implemented ?
- Can you discuss spoofing and jamming as related to precise timing, used in the power grid, financial networks and other critical infrastructure?
- Could smart antennas be employed to perform plausibility tests with the satellites, an angle-of-arrival approach along the lines of the time-of-arrival method you discussed?

BY THE NUMBERS: PARTICIPANTS' VIEWS

Have you experienced jamming or spoofing in the past?



Is simple awareness enough?



most appropriate for the various target audiences? Keeping in mind that they’re technical and non-technical, in a vast range of different application areas.

“So who takes responsibility? Who’s going to tackle the problem, and who are the drivers for change?”

The program concluded with a lively question-and-answer session between the hundreds of attendees and the three panelists. See the box above for some of the questions posed.


To hear more, download the webinar at https://novatel.com/tech-talk/webinars/nobodys-fool-jamming-and-spoofing-detection?utm_source=insidegnss&utm_medium=article&utm_campaign=nobodys_fool. 

Photo courtesy of Daniele Borio.

GNSS Interference Mitigation

Modulations, Measurements and Position Impact

Interference mitigation techniques should protect GNSS receivers from interference and jamming without biasing their final position, velocity and timing solution. This column analyses five popular interference mitigation techniques, including the Adaptive Notch Filter (ANF) and Pulse Blanking (PB), evaluating their impact on pseudoranges and on the final position and timing solution. Several GNSS modulations are considered, showing the advantage of using GNSS signals with similar spectral characteristics.

GNSS receivers are now required to operate in complex radio frequency (RF) environments and to withstand significant levels of interference and jamming. For these reasons, receiver manufacturers are now implementing interference mitigation techniques able to improve receiver performance in the presence of jamming. Among the different techniques available in the literature, the most popular ones are probably pulse blanking (PB), to reduce the impact of pulsed interference, and notch filtering for continuous wave (CW) removal. While these techniques can significantly improve receiver robustness, they can also introduce biases and distortions. For example, the notch filter is known to introduce biases at the measurement level. When the fre-

quency of the filter notch is known and fixed, these biases can, however, be determined and compensated for during the process of measurement generation.

A review on interference mitigation can be found in the article published in the September 2017 issue of *Inside GNSS* listed in Additional Resources. The review also includes the framework of robust interference mitigation (RIM) that exploits principles from robust statistics. Techniques such as PB and frequency excision belong to the class of RIM techniques.

While significant work has characterized interference mitigation techniques at the signal-processing level, limited analysis has assessed their impact in the measurement and position domains. This article fills this gap and analyzes five popular interference mitigation techniques at the measurement and position levels, studying the potential introduction of biases. The impact on the timing solution has also been analyzed.

The five interference mitigation techniques considered include the Adaptive Notch Filter (ANF), a form of notch filter where the notch frequency is dynamically estimated, and four RIM techniques. The impact of the techniques is assessed considering several GNSS signals on two frequencies: on L1 (1575.42 MHz) the GPS L1 C/A signal and the binary offset carrier (BOC) modulation adopted by the Galileo E1B/C and the Beidou B1C signals are considered, while the wideband binary phase shift keying (BPSK) modulation adopted by the Galileo E5B

DANIELE BORIO AND CIRO GIOIA

EUROPEAN COMMISSION, JOINT RESEARCH CENTRE (JRC), DIRECTORATE FOR SPACE, SECURITY AND MIGRATION, ITALY



AUVSI
XPONENTIAL[™]
ALL THINGS **UNMANNED**

ATLANTA + ON-DEMAND | HYBRID EVENT SERIES
AUGUST 16 - 19, 2021 | ATLANTA
NOW - SEPT. 10, 2021 | ON-DEMAND

DEFINED

CHART THE PATH TO ASSURED AUTONOMY

Join us in Atlanta as our community of end users, technologists and policymakers come together to write the next chapter of autonomous innovation and assure its safe and seamless integration into everyday life. **XPONENTIAL 2021 is a reimagined hybrid experience** offering fresh insight across the full spectrum of unmanned innovation, from AI, to sensors, to cybersecurity. With on-demand sessions now available on our virtual platform and an in-person event happening August 16 – 19, XPONENTIAL 2021 offers you the best of both worlds.

MAP YOUR NEXT MOVE.

xponential.org

PATHWAYS

and the Beidou B2bI signals is analyzed on the 1207.14 MHz frequency.

The analysis highlights the complex interaction between the different GNSS modulations, interference mitigation techniques, measurement generation and position computation.

Interference Mitigation Techniques

The five interference mitigation techniques considered here can be classified according to the schematic representation provided in **Figure 1**. Interference mitigation techniques can be designed according to the interference cancellation (IC) principle: the mitigation algorithm estimates at first the interfering signal that is then removed (canceled) from the input samples. The ANF uses this principle and estimates the parameters of frequency

modulated signals. IC mitigation techniques usually adopt a signal model for the interfering signal. The ANF assumes that the interfering signal has a practically constant amplitude. Alternatively, mitigation techniques can be designed using principles from robust statistics. In this case, the interfering signal is at first projected into a transformed domain where it is expected to assume a sparse representation, i.e. to affect a limited number of samples that can be treated as outliers. A non-linearity is then used to mitigate the impact of these outliers. Finally, an inverse transform is used to bring the signal back into the time domain. Techniques implementing these three operations belong to the class of RIM. PB and frequency excision are obtained when the following non-linearity.

$$\tilde{Y}[k] = \begin{cases} Y[k] & \text{if } |Y[k]| < T_h \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

is adopted. $Y[k]$ denotes the signal samples brought into the transformed domain and T_h is a decision threshold. When processing is performed directly in the time domain, PB is obtained. If a Discrete Fourier Transform (DFT) is used to bring the samples in the frequency domain, then frequency excision is implemented. Incidentally, PB and frequency excision can also be interpreted as forms of IC where, in the first case, the interference term is modeled as a sequence of pulses and in the second as a combination of complex sinusoids.

In addition to non-linearity (1), we also considered the complex signum non-linearity:

$$\tilde{Y}[k] = \begin{cases} \frac{Y[k]}{|Y[k]|} & \text{if } Y[k] \neq 0 \\ 0 & \text{if } Y[k] = 0 \end{cases} \quad (2)$$

Combining the two processing domains and the two non-linearities a total of four RIM techniques are obtained: time domain pulse blanking (TDPB), time domain complex signum (TDCS), frequency domain pulse blanking (FDPB) and frequency Domain complex signum (FDCS).

IC and RIM techniques are effective when model assumptions are valid, i.e. when the interference signal has an almost constant amplitude for the ANF and when it has a sparse support in the processing domain of RIM techniques. We also consider here the case of pulsed interference. When the DFT is applied

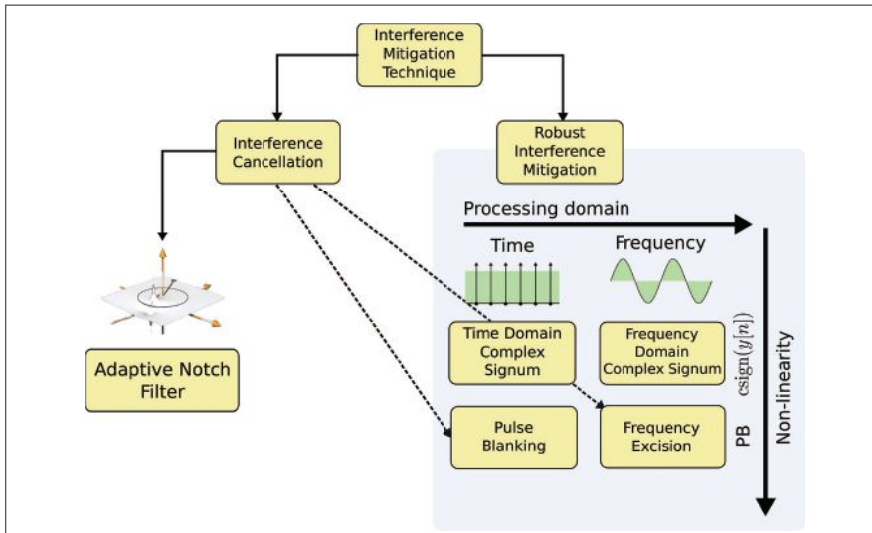


FIGURE 1 The five interference mitigation techniques analyzed in this article.

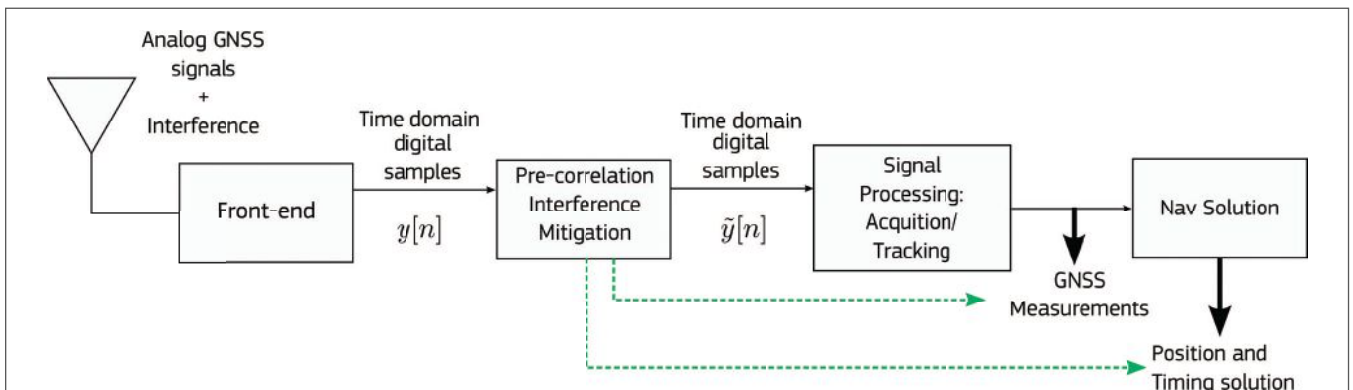


FIGURE 2 Interaction between pre-correlation interference mitigation techniques, GNSS measurements and position and timing solution.

to this type of signals and frequency domain RIM is used, the interference term is spread over several samples, the sparsity assumption is violated and receiver performance is actually degraded.

Measurement, Position and Timing

The five techniques considered are pre-correlation approaches that act directly on the samples provided by the receiver front-end. After pre-correlation interference mitigation, a new set of time domain samples is produced and passed to standard acquisition and tracking stages that process the GNSS signals and extract measurements such as pseudoranges, carrier phases and Doppler frequencies. Finally, GNSS measurements are used for computing the user position and clock parameters. In this respect, a potentially complex interaction occurs between pre-correlation interference mitigation techniques, measurements, the position and the clock solution. A schematic representation of the processing stages implemented in a GNSS receiver is provided in **Figure 2**. The complexity mainly arises from the relative distance, in terms of processing blocks, between interference mitigation, measurement generation and position solution.

For this reason, we adopted an experimental approach coupled with the software defined radio (SDR) paradigm to evaluate the impact of interference mitigation techniques. More specifically, a fully software GNSS receiver was used to process data collected in the presence of jamming and to assess the impact of interference mitigation techniques on the measurements and on the position and clock solution. The datasets collected according to the experimental setup described in the next section have been processed several times using the different interference mitigation techniques and the results obtained have been compared with those provided by standard processing in the absence of mitigation.

In the measurement domain, the analysis focused on pseudoranges and for each interference mitigation strategy, differences between pseudoranges computed with and without mitigation were formed:




FIGURE 3 Experimental setup adopted to test the impact of interference mitigation techniques a) Schematic representation of the setup. b) Actual view of one of the experiments involving Galileo E5B and Beidou B2bI signals.

$$PR_{diff}[m] = PR_{Mit}[m] - PR_{Standard}[m] \quad (3)$$

where $PR_{diff}[m]$ is the pseudorange difference at the epoch m and PR_{Mit} and $PR_{Standard}$ are the pseudoranges obtained with and without interference mitigation. This approach leads to the cancellation of all common errors preserving the differences introduced by the specific processing scheme. Ideally, interference mitigation techniques should not bias pseudoranges and pseudorange differences (3) should be zero mean. In this respect, we investigated the mean of pseudorange differences.

To analyze the impact of interference mitigation on the position solution, positioning errors were computed for all the processing strategies considered. Since the antenna of the receiver was carefully surveyed, it was possible to compute position errors with respect to the known antenna location. The errors were computed in a local East, North, Up (ENU) frame centered into the antenna reference position. In this way, it was possible to directly compare the position errors and analyse the impact of interference mitigation. The analysis focused on single point positioning (SPP) computed using a weighted least squares (WLS) approach where the weights were determined as a function of the satellite elevation.

Finally, we analyzed the clock bias. In particular, the antenna coordinates were fixed and a clock-only solution was computed. We considered this type of solution since it is usually adopted by GNSS timing receivers that fix at first the user position and compute only the clock bias. The clock time series were then compared.

Experimental Setup

Several experiments assessed the impact of interference mitigation in the measurement, position and clock domains. A SDR front-end collected in-phase/quadrature (I/Q) samples that were then processed using a custom Matlab software receiver.

A dedicated experimental setup was developed where a jammer was placed inside a shielding box whose output was connected to a variable attenuator, adopted to generate different jamming power levels. At the beginning of each test, the attenuation was set to a significant value and the amount of jamming power leaking through the attenuator was negligible. In this way, the first part of each test allowed the assessment of the impact of interference mitigation in the absence of interference.

The output of the variable attenuator was then combined with clean GNSS signals collected from a rooftop antenna. In this way, jammed GNSS signals were obtained. A schematic representation of the experimental setup adopted for testing the impact of interference mitigation and a view of one of the experiment conducted is provided in **Figure 3**.

Three modulations were considered: the BPSK(1) of GPS L1 C/A signals, the BOC(1, 1) modulation adopted by Galileo E1B/C and by the Beidou B1C signals and the wideband BPSK(10) modulation adopted by the Galileo E5B and Beidou B2bI component. For Galileo E1C, Beidou B1C and Galileo E5B signals, pilot processing was implemented. Signals on the L1 frequency were collected with a 10 MHz sampling




FIGURE 4 J/N profiles estimated for the three experiments considered in this paper. In all cases, front-end saturation occurs.




FIGURE 5 Pseudorange differences between standard measurements and observations from the five interference mitigation techniques. Test 1, GPS L1 C/A signals.

frequency, whereas the E5B components were collected using a 25 MHz sampling frequency. The Galileo E5B and the Beidou B2bI signals are considered wideband, and the main lobe of their spectrum occupies most of the frequencies captured by the SDR front-end. The custom software receiver generates Receiver INdependent EXchange Format (RINEX) files that were used for the measurement and position domain analysis. The parameters used for the different tests are summarized in **Table 1** that also describes the attenuation profile used for each experiment.

Three tests were considered:

- **Test 1:** performed on the L1 frequency band (1575.42 MHz centre frequency) with a 8 bit quantization. This test was used to analyse the impact of interference mitigation on the GPS BPSK(1) modulation and on the Galileo/Beidou BOC(1,1) component.
- **Test 2:** also performed on the L1 frequency band but with a 16 bit quantization. This test was conducted to analyse potential differences between 8 bit/16 bit quantization.
- **Test E5B:** performed on the E5B/B2bI frequency band (1207.14 MHz centre frequency). This test was conducted to analyse the impact of RIM on a wideband GNSS modulation.

The power profiles obtained for the three tests are analysed in **Figure 4** that shows the jamming to noise power ratio

(J/N) estimated from the samples collected using the SDR front-end. While a linear J/N profile was expected (see the corresponding parameters in **Table 1**), in all three cases the front-end experiences saturation. In this case, the jamming signal is so powerful to exceed the capabilities of the front-end quantization function that effectively clips the received samples to the minimum and maximum values allowed by 8/16 bits. Front-end saturation introduces significant signal distortions that further impact receiver operations. For the third test, the jammer attenuation was increased again after 550 seconds from the start of the test.

Additional details on the three tests and on the properties of the jamming signals used in each experiment can be found in the paper presented by the authors at the International Technical Meeting (ITM) of the Institute of Navigation (ION) and listed in Additional Resources.

Experimental Results

The pseudorange differences obtained for Test 1 and for GPS L1 C/A signals are shown in **Figure 5**. The four upper boxes correspond to the differences obtained for RIM techniques: in these cases, the time series are zero mean and no clear trend is introduced by the mitigation techniques. The first 200 seconds of the test are characterized by J/N values lower than 5 dB, which make the jamming component negligible. In this portion of the test, pseudorange differences oscillates around zeros with variations lower than 5 metres. The variance of the pseudorange differences increases with the J/N: this fact is expected and reflects the increased equivalent noise caused by the jamming component. The zero mean property observed for RIM techniques confirms the theoretical result obtained by Borio and Closas in their paper listed in Additional Resources.

RIM techniques do not bias the Cross-Ambiguity Function (CAF), the main quantity evaluated by the acquisition

	Test 1	Test 2	Test E5B
Parameter	Value	Value	Value
Sampling Frequency	10 MHz	10 MHz	25 MHz
Centre Frequency	1575.42 MHz	1575.42 MHz	1207.14 MHz
Sampling Type	Complex IQ	Complex IQ	Complex IQ
No. bits	8	16	8
Attenuation Step	1 dB	1 dB	0.25 dB
Attenuation Step duration	20 s	30 s	20 s

TABLE 1 Parameters used for the different tests.

and tracking blocks. The CAF is a generalized version of the correlation function and has a main peak in correspondence of the Doppler frequency and code delay of the signals acquired and tracked by the receiver. Since the CAF is not biased by RIM techniques, also the pseudoranges are not biased by this type of approaches. The results in the four upper boxes of **Figure 5** confirm this theoretical finding.

The impact of the ANF is analysed in the bottom left box of **Figure 5**: in this case, a significant bias can be observed. The ANF delays pseudoranges and introduces a time-varying bias. At the beginning of the test, the impact of jamming can be neglected and the only interfering term is a CW generated by the clock of the SDR front-end used for the data collection. This CW has a frequency equal to 1575 MHz, which is close to the GPS L1 C/A centre frequency. Thus, the ANF significantly impacts GPS L1 C/A signals. As the jamming power increases, the adaptive block of the ANF converges to a different solution leading to different biases on the pseudoranges. Indeed, the time-varying behaviour of the ANF is due to the changing jamming conditions occurring in Test 1.

While it is difficult to theoretically predict the bias introduced by the ANF, this delay is common to all measurements and a limited impact is found on the SPP solution, as better analysed in the following. Indeed, most of the bias introduced by the ANF is absorbed by the clock bias term.

The impact of interference mitigation techniques on Galileo E1C and Beidou B1C pseudoranges is analysed in **Figure 6**. In the figure, two RIM techniques and the ANF are considered. Also in this case, RIM techniques do not introduce biases on the measurements whereas the ANF delays the pseudoranges. The delay introduced on the Galileo E1C and Beidou B1C signals is however different from that observed for the GPS L1 C/A components. This difference is due to the different spectral interaction

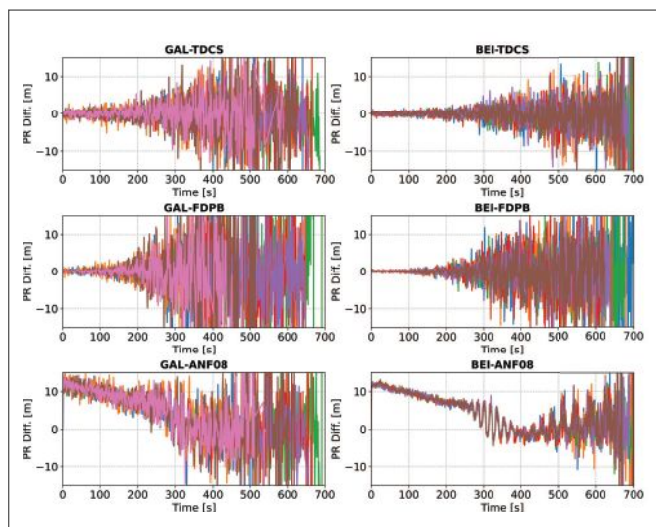


FIGURE 6 Pseudorange differences between standard measurements and observations from three of the five interference mitigation analysed, Test 1. comparison between Galileo E1C signals (left column) and Beidou B1C components.

NavtechGPS brings you ...


LabSat 3, YOUR Way!

Multiple frequencies and constellations



LabSat 3 Wideband

Portable, affordable, multiple L1 constellations, record-replay



LabSat 3

Replay in real-time: Chose GPS, GLONASS or BeiDou

LabSat
Real-Time



Ask us for a **FREE** demonstration unit.
People who try it, want to buy it.

https://www.navtechgps.com/brands/labSAT_by_racelogic/

NavtechGPS sells hundreds of GNSS products, including receivers, antennas, inertial systems, GPS jammer detectors, and more!
Contact us today.

NavtechGPS

+1-703-256-8900 • 800-628-0885
www.NavtechGPS.com

Your ONE source for GNSS products and solutions




FIGURE 7 Pseudorange differences between standard measurements and observations from three of the five interference.




FIGURE 8 Position errors in the ENU frame. Comparison between different interference mitigation techniques for GPS L1 C/A only positioning. Test 1.

between the ANF transfer function and the BOC and BPSK modulations adopted by the Galileo E1C/Beidou B1C and GPS L1 C/A signals, respectively.

At the beginning of the test, the delay introduced by the ANF on the Galileo/Beidou pseudoranges is about 10 metres against the 35 metres observed for the GPS L1 C/A components. The lower delay observed for Galileo/Beidou signals is due to the fact that the ANF notch is placed around 1575 MHz that it close to a spectral zero of the BOC modulation that is less distorted than the GPS L1 C/A signals. Thus, a lower delay is observed.

Far from its notch, the ANF transfer function is almost flat with a negligible group delay. For this reason, signals occupying spectral regions far from the ANF notch are less delayed. As for the previous GPS case, the biases introduced by the ANF are mostly common to all Galileo/Beidou pseudoranges and will be absorbed by the clock bias term in a Galileo- or Beidou-only solution. In a multi-constellation solution, these delays will be absorbed by the inter-system bias estimated by the receiver. **Figure 6** also provides a direct comparison between Galileo and Beidou pseudorange differences: the same delay is introduced by the ANF on the two components. In this case, only the BOC(1,1) component of the Beidou and Galileo

signals is processed and the ANF affects in a similar way modulations with the same spectral characteristics.

This result shows the advantages of using signals with the same spectral characteristics: the same delay is introduced by the ANF on interoperable modulations. In this way, the inter-system bias between Galileo and Beidou is not affected by the ANF allowing the use of advanced positioning algorithms exploiting the stability of inter-system biases. Beidou pseudorange differences are less noisy than the Galileo ones. This fact is due to the longer integration time (10 ms) adopted for the Beidou signals.

Results similar to those obtained for Test 1 were obtained in Test 2. These results can be found in our paper listed in Additional resources.

The results obtained for the test on the E5B/B2B frequency are shown in **Figure 7**: similar results are observed for the Galileo and Beidou signals. As for the previous cases, RIM techniques do not introduce biases in the pseudoranges and zero-mean differences are obtained. Moreover, since the Beidou and Galileo signals adopt the same modulation, the same delay is introduced by the ANF.

When comparing Test 1 and Test 2 E5B results, significantly lower differences are observed in the second case. More specifically, the ANF introduces a bias of about a metre, whereas RIM

techniques lead to average pseudorange differences in the millilitre level. These significantly lower differences are due to the wideband nature of the E5B/B2I signal and to the higher accuracy of the resulting pseudoranges.

Position solutions have been analysed for Test 1 in **Figure 8** that shows the position errors obtained using GPS L1 C/A pseudoranges. From the figure, it emerges that all five interference mitigation techniques do not bias the final SPP solution and, at least under low jamming conditions, all the time series are characterized by the same trend. The standard solution without mitigation provides reasonable position estimates for the first 550 seconds of test whereas time domain RIM techniques improve receiver performance leading to reduced errors for the first 700 seconds.

For Test 1, the jamming power was increased of 1 dB every 20 seconds. This implies that time domain techniques provided a margin of about 7.5 dB in terms of resilience to jamming power. In addition to this, the receiver front-end started saturating after about 400 seconds from the beginning of the experiment. Thus, time domain RIM allows receiver operations even in the presence of significant levels of front-end saturation. In this case, frequency domain RIM and the ANF provided limited benefits. Indeed, FDPB worsen the receiver performance. This

fact is clearly visible in Figure 8 where no position solution is available in the FDPB case after about 400 seconds from the start of the experiment. This result is due to the nature of the jamming signal used for Test 1 and to the selection of the FDPB threshold. Indeed the jamming signal used for Test 1 is pulsed in nature (see the papers from the authors listed in Additional resources). When the DFT is used to bring the samples in the frequency domain, the interference pulses are spread among all the frequencies and thus the signal does not admit a sparse representation, violating the main assumption of this type of processing. The second factor compromising the performance of FDPB is the selection of the decision threshold in (1). We used a threshold equal to 3 times the standard deviation of the samples in the absence of interference:

$$T_h = 3\sqrt{\text{Var}\{Y[k]\}} \quad (4)$$

where the variance $\text{Var}\{Y[k]\}$ has been estimated from the samples collected at the beginning of the experiment.

This choice has been dictated by the fact that, in the absence of interference, the samples, $Y[k]$, should approximately follow a Gaussian distribution and, under this hypothesis, only few of them should assume values greater than 3 times their standard deviation. This choice is however too conservati-

ve for the specific jamming signal used in Test 1. This fact is analyzed in Figure 9 that shows the square magnitude of the frequency domain samples collected at the beginning of the test and after 400 seconds. The corresponding FDPB threshold is also provided. At the beginning of the test, a very limited number of peaks passes the threshold. The two main peaks passing the threshold correspond to the clock CWs that are correctly excised. After 400 seconds, the jamming signal is so strong that most of the samples pass the threshold leading to the removal of a significant portion of the useful signals as well.

This result shows the importance of properly selecting the domain of operations and the parameter settings of interference mitigation techniques. Parametric approaches such PB and frequency excision may be significantly influenced by the selection of the decision threshold. The ANF and FDCS do not significantly improve the receiver performance but do not suffer the degradations of FDPB.

Results similar to those shown in Figure 8 were obtained for the Galileo and Beidou signals in a single-constellation SPP solution. Also Test 2 led to similar findings and the interested reader is referred to our paper published in NAVIGATION.

Position errors obtained for the test conducted on the Galileo E5B signals are shown in Figure 10: also in this case all the techniques lead to time series with a consistent trend. The advantages of using a wideband signal such as the Galileo E5B component are also evident: the position errors are significantly lower in magnitude than those observed for the GPS L1 C/A modulation.

The impact of interference mitigation on the clock bias is analysed in Figure 11 and Figure 12 that consider Test 1 and Test E5B, respectively. As discussed above, the clock bias absorbs most of the delays introduced by the ANF on the pseudoranges. From Figure 11, a delay up to 35 metres is introduced by the ANF on the GPS L1 C/A clock term whereas a bias of about 12 metres is observed for both Galileo E1C and Beidou B1C solutions. The delay on the GPS L1 C/A solution corresponds to about 120 nanoseconds. Since the clock bias is used by timing receivers to steer the local time to the GPS/Universal Time Coordinated (UTC) time scales, the ANF can render timing solutions unsuitable for high-end timing applications. On the contrary, RIM techniques do not introduce biases and the corresponding time series in Figure 11 and Figure 12 are zero mean. The results on the clock bias confirm the fact that the




FIGURE 9 Square magnitude of frequency domain samples, $Y[k]$, obtained at the beginning of the test and after 400 seconds. The samples are plotted in logarithmic scale for improved clarity. The FDPB threshold is also provided.

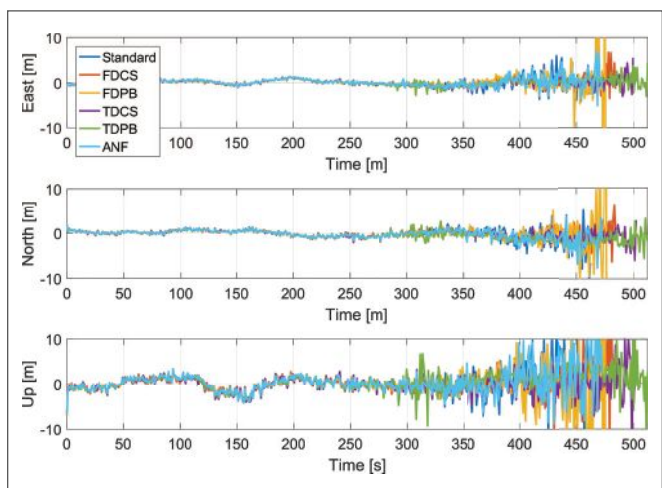


FIGURE 10 Position errors in the ENU frame. Comparison between different interference mitigation techniques for the Galileo E5B signals.




FIGURE 11 Clock bias differences between the standard solution without mitigation and the five interference mitigation techniques. GPS L1 C/A, Galileo E1C and Beidou B1C signals, Test 1.

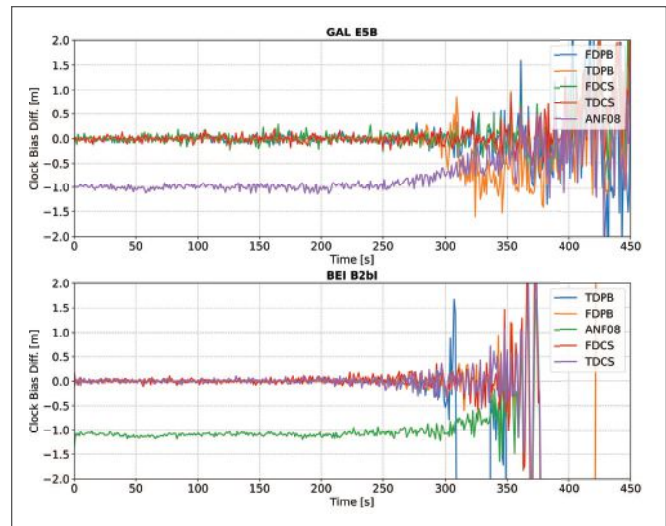


FIGURE 12 Clock bias differences between the standard solution without mitigation and the five interference mitigation techniques. Galileo E5B and Beidou B2b1 signals, Test E5B.


ANF affects Galileo and Beidou signals, which share the same spectral characteristics, in the same way introducing a common delay.

Finally, **Figure 12** further highlights the advantages of using wideband signals such as the Galileo E5B and Beidou B2b1 components. The time series are less noisy than the ones obtained in Test 1. Moreover, the bias caused by the ANF is reduced to about a metre.

Conclusions

This article experimentally evaluated the impact of five interference mitigation techniques on pseudoranges and on the final position and timing solutions. The analysis included the ANF, time domain PB, frequency excision and two additional RIM techniques. Moreover, several GNSS modulations were considered from three GNSSs: GPS, Galileo and Beidou. From the analysis, it emerged that RIM techniques, including time domain PB and frequency excision, do not introduce biases at both measurement and position/clock level whereas the ANF delays pseudoranges. The delays are however common to all the measurements and only the clock component is affected in the final position solution. In this respect, delays of several tens of nanoseconds were observed. These delays, which are difficult to predict and depend

on the jamming signal tracked by the ANF, are modulation dependent: signals sharing similar spectral characteristics are delayed in a similar way.

The analysis focused on a SPP solution and additional work is required to determine if a more evident effect could emerge in a Precise Point Positioning (PPP) framework. The analysis also confirmed that the domain of operation of RIM has to be carefully selected along with the parameters of the RIM non-linearity. In this respect, parameters improperly set can lead to performance degradations. 

Manufacturers

The SDR front-end used for the experiments is an Ettus Universal Software Radio Platform 2 (USRP 2) device from Ettus Research, a National Instrument Brand, Santa Clara, California.

Additional Resources

- (1) D. Borio and C. Gioia "Robust Interference Mitigation: a Measurement and Position Domain Assessment," Proceedings of the 2020 International Technical Meeting of The Institute of Navigation, San Diego, California, January 2020, pp. 274-288.
- (2) D. Borio and C. Gioia "GNSS Interference Mitigation: a Measurement and Position Domain Assessment" NAVIGATION, the Journal of the Institute

of Navigation. Spring 2021; Vol. 68, Issue 1, pp. 93-114, <https://onlinelibrary.wiley.com/doi/full/10.1002/navi.391>.

(3) D. Borio and P. Closas "A Fresh Look at GNSS Anti-Jamming", Inside GNSS, Volume 12, Number 6, September/October 2017, pp. 54-61.

(4) D. Borio and P. Closas "Robust Transform Domain Signal Processing for GNSS". NAVIGATION, the Journal of the Institute of Navigation. Summer 2019; Vol. 66, Issue 2, pp. 305-323, <https://onlinelibrary.wiley.com/doi/full/10.1002/navi.300>.

Authors



Daniele Borio holds doctoral degree in electrical engineering from Politecnico di Torino in April 2008. He worked a senior research associate in the PLAN group of the University of Calgary and is currently a scientific policy officer at the Joint Research Centre of the European Commission (EC) in the fields of digital and wireless communications, location and navigation.



Ciro Gioia received a Ph.D. in navigation sciences from Parthenope University. After working as an external consultant at JRC, he is now scientific project officer at the JRC focusing on location and navigation.



ION

GNSS+

2021

September 20-24, 2021

Exhibit Hall: September 22 and 23

St. Louis Union Station Hotel

St. Louis, Missouri

Meet me in St. Louis.



**The 34th
International
Technical Meeting
of the Satellite Division of
the Institute of Navigation**

REGISTER TODAY

ion.org

Resilient PNT for Critical Applications

The Need for Modular Open System Architectures—and Much More

There has been much discussion of the need for resilient PNT over the past few years as dependencies have grown and an evolving threat matrix has become more active. As a nation, we need a measured and cost-effective response commensurate with the level of threats and the possible consequences.

LOGAN SCOTT

LOGAN SCOTT CONSULTING

Do we even know what resilience is? Mostly it looks like extra cost when you don't need it. Then, when it is needed and you don't have it, it looks like failure—sort of like the Texas power grid back in February. Resilience has costs, and budgets are bounded. My working definition for resilience is that it is about building sufficiently secure and reliable systems out of insecure and unreliable components operating in an indeterminate and evolving environment.

In Richard Cook's 3-page long masterwork "How Complex Systems Fail," he observes: "complex systems run as broken systems. The system continues to function because it contains so many redundancies and because people can make it function, despite the presence of many flaws." Furthermore, he notes that resilience can be improved by "establishing means for early detection of changed system performance in order to allow graceful cutbacks." Situational awareness is the key in discovering where problems might be developing and then, taking corrective action before catastrophic failure.

How do you measure resilience? One way is to try to break it and then decide whether the protection is adequate for the domain of use. Any system will break under sufficient stress. Determining what is sufficient is a

hard question but it is a key question. Resilience could end up being quantified using a series of tests like UL standards for safes. You expose the safe to a skilled safecracker and see how long it takes them to break in. Interestingly, the highest security rating, TXTL-60, only guarantees protection for 60 minutes. After that, you need another plan.

"The state of safety in any system is always dynamic; continuous systemic change ensures that hazard and its management are constantly changing."

Richard Cook, *How Complex Systems Fail*

Finally, how do you maintain resilience? Again, from Richard Cook: "The state of safety in any system is always dynamic; continuous systemic change ensures that hazard and its management are constantly changing." As systems and threats evolve, new failure modes and attack vectors develop, and so, the challenge is to respond with commensurate protections. Operator training and controlled exposure to threats is essential. Experience builds confidence and speeds reaction times.

The Need for Standards and Mandates

A core question we have not addressed at a national policy level is how to incentivize resilience. When seat belts were invented, they were made available as an option by Ford and others. It was not a popular option. Less than 2% of buyers elected to get them. Then, through a series of federal mandates, they became required equipment and later, we saw requirements to use them. The point being that safety standards are needed, they have costs, and, if they are left optional, they may not be implemented or used.

I like that DHS is addressing resiliency as a risk management question, but leaving the implementation of protections entirely up to the user community strikes me as unworkable. Our user communities are rarely aware of the potential risks and possible consequences, much less how to address them. Even when they are aware of the risks, industry is often more driven by cost considerations under nominal conditions, and they fail to prepare.

I'll pick on the Texas power grid in February again, but I could also have picked on "just in time" manufacturing systems vulnerable to supply chain disruptions, or the pilots who let the Ever Given onto the Suez Canal. By establishing standards and exposure-based testing procedures, vendors and buyers in critical infrastructure domains can avoid the more egregious outcomes in a cost-effective manner.

The Need for Modular Open Systems Architectures (MOSA)

In prior discussions I noted that building a resilient architecture is not just about having the right parts; they must be integrated correctly (and tested). MOSA is about effectively leveraging the capabilities of diverse system com-

ponents and maintaining currency as new innovations and technologies become available. MOSA is a platform, it is an operating system, it is an enabler. It is not a point solution.

A cell phone's positioning process is a great example of MOSA. Android phones come in diverse flavors and have a rapid innovation cycle based on a rich and constantly evolving ecosystem of parts. Yet, they all manage to integrate sensors together to establish position with good accuracy both indoors and outdoors. That said, cellphones performed abysmally when exposed to inadvertent spoofing at the ION GNSS+ conference in 2017.

In many ways, integrity and resilience are highly intertwined problems and so, there are opportunities within MOSA constructs to approach the problem of safely integrating less than 100% trusted, 100% reliable components. Adding trust modules using, for instance, Bayesian inference approaches can substantially harden systems operating with uncertainty. When you hear that rattling sound in your car, you may not know what it is, but you do know to investigate. Experience with cybersecurity shows the need for a rapid update and response cycle—MOSA will help.

The Need for Cybersecurity and Authentication

Modern PNT systems are computers, often running a full operating system. For good or ill, they will almost certainly connect with a network in some manner. Authentication is about knowing where your data comes from, knowing where your software comes from, and establishing a chain of evidence to establish provenance. Any software updates and/or data ingested need to be authenticated to establish that they come from a trusted source. This includes not only ephemeris but also maps, databases, reference station data, and any cryptographic key material essential to operation. The

Chips Message Robust Authentication (CHIMERA) concept adds to this a means for authenticating pseudorange measurements that form the basis for position and time estimates.

In connected applications, it is important to recognize that when a receiver reports its position, there are diverse “man-in-the-middle” attacks that can corrupt or alter data before it gets to its destination. When a ship's autopilot receives position and velocity reports from a GNSS receiver, how does it know that there is not a data-altering dongle between it and the receiver? Spoofing is an effect, not a method, and cyberspoofing is often a much easier (and more powerful) method compared with RF spoofing.

A detailed discussion is out of scope for this article but much of the public key infrastructure (PKI) protocols and


procedures used in Internet communications have direct applicability. Trusted platform modules (TPM) and subscriber identity modules (SIM) such as those used in computers and cell phones can further enhance security. Updates analogous to antivirus protections can maintain a receiver's ability to recognize threats as they evolve.


In the MOSA paradigm, if subsystems can report problems due to spoofing, jamming, cyberattack, hardware failure, software corruption etc. and, there are performance and security monitors in place to aggregate information and watch for discrepancies, a more effective and resilient response can be mounted.

The Need for Trustable Multi GNSS

In the quest for resilience and augmentations, I am not convinced that we as a nation have fully explored how to safely integrate foreign navigation systems into critical applications. Access to more signals and more systems offers considerable resilience potential. Instead, the FCC has unilaterally restricted their use under part 25 rules with limited justification. Perversely, this has led to some US companies flying their satellites under foreign flags so as to gain legal access to FCC-proscribed navigation signals. These same restrictions limit the performance and utility of precision positioning systems, receiver autonomous integrity algorithms, positive train control systems, snapshot RTK processing and spoofing detection processes.

One of the things I found fascinating about the Galileo failure in July 2019 was that the satellites all continued to produce good ranging signals. If you could provide your own ephemeris, say from JPL, NASA, NSWC, and/or other sources, you still got great performance. Treating foreign navigation satellites as “signals of opportunity” and using curated and signed US generated ephemeris strikes me as a powerful and inexpensive augmentation.





Much less trust is placed in the foreign state, yet you get a lot of augmentation benefit for minimal cost. Additionally, you limit the impact of a global system's outage. If Galileo had been the only game in town, its one-week outage would have been catastrophic. As it was, its absence was noted but had almost no effect on GNSS dependent operations.

The Need For An Honest Evaluation of Ligado's Impact

GNSS really is different from communications. The FCC, by setting a standard where the mechanism of harm is to place GNSS receivers in deep and uncontrolled saturation, ignores the possibility of normally harmless signals mixing and causing harm. None of the testing to date has explored this issue and so, our national policy might be the RF equivalent of mixing alcohol and fentanyl and hoping for the best. Furthermore, the FCC showed almost no cognizance of the importance of GNSS-based remote sensing in monitoring climate change. The RF smog that Ligado's signals will create restricts our ability to develop a clearer picture of what is happening to the planet. The FCC's decision needs to be revisited using sound engineering as a basis.

The Need for Situational Awareness

Smoke alarms do not extinguish fires, but by providing an early warning, they allow for a more timely and effective response. Intelligent receivers

provide warnings that interference is taking place so users can take corrective actions. With unambiguous statements like "I am jammed and do not know where I am / what time it is," less time is spent debugging GNSS dependent systems. Multi-sensor systems can avoid ingesting hazardously misleading information (HMI) into their PNT estimation processes and so, can avoid corrupting estimates. Fleet managers will know when an employee is jamming his work vehicle and can have a quiet conversation to correct this behavior.

Building a modicum of intelligence into a receiver is not hard nor is it costly. As a minimum, sudden changes in automatic gain control (AGC) settings are highly indicative of interference. If the AGC is telling you there is a lot of power coming into the receiver's front end and your C/N_0 meters are saying SNR is high, you probably have a spoofer. Add to this a few other simple techniques (see chapter 24 of Position, Navigation, and Timing Technologies in the 21st Century), and the receiver can provide a feature-rich description of the RF environment it is experiencing.

Earlier, I noted that most receivers connect to networks. Aggregating reports from multiple receivers, it then becomes possible to geolocate interference sources and characterize their behavior using crowdsourcing methods. With this more global level of situational awareness, we would now have a much better picture of where the interference

hot spots are and a better understanding of the motivations driving their use. Again, early warning of trends provides a basis for a commensurate and less costly response.

Finally, providing law enforcement with accurate jammer locations would allow for enforcement actions, not so much because they want to catch jammers, but because of the criminal activity motivating the jamming.

Augmentations and the Role of Markets

One of the most insidious things about GNSS is its price to the user: free. That, combined with its worldwide coverage and superb accuracy in both time and positioning creates significant barriers to entry for new offerings. In the commercial arena, a new entry that provides only the same services as GNSS, maybe a little better, seems doomed. A successful entrant will need to have a value-added proposition with features that cannot be met using GNSS.

Communications facilities, indoor operation, proof of integrity and location, and uninterruptable service would all be on my short list. In large measure, these capabilities can be provided by combining GNSS with other sensors and systems, especially if we use the full constellation of 125 navigation satellites on orbit and healthy now. That said, I do expect new entrants.

5G NR and 802.11 both have strong potential to meet the requirements of my short list, especially as they move towards higher frequencies. Yes, the ranges there will be short, but infrastructure densities will be high. Both technologies have strong and active initiatives within their standards-setting process oriented towards providing accurate, high-integrity positioning. Also, because they are extant systems, there is less pressure to offer ubiquitous service at inception, a daunting challenge for a brand-new entrant.

I expect LEO satellite systems will also have a role. Because they have high angular rates across the sky, you can get nearly instant-on cm-level positioning. Operating at higher frequencies, eg. X, Ku or even V-band, they can simultaneously provide strong communications

capabilities when outdoors and so, might play very well in the autonomous vehicle markets. Yes, the antenna issues are challenging, but they ride a wave of actual deployments.


The Role of Government

So, what is the proper role for government? Well first, stably fund, maintain, and operate GPS. GPS is foundational critical infrastructure, not easily supplanted.

Recognize that GPS signals are extremely weak. Like fish in a river when the river becomes polluted, expect GPS receivers stressed by interference to become less resilient and more prone to unexpected anomalies. Monitoring the health of our spectrum on a continuous basis is of paramount importance, so we can act early and with resolve. Crowd-sensing approaches using intelligent receivers offer an important means to do so, but it will require coordination and open sharing of results.

Providing secure ephemeris and integrity data to support safe use of multi GNSS should be funded. Most of the parts are already in place, and it is mostly a matter of setting up a service offering. Of course, removing FCC roadblocks to its use is also essential. PPP data services should be considered as part of the package to promote rapid adoption of “safe ephemeris.” The same public-facing servers could also provide key materials for civil signal authentication, e.g. fast CHIMERA keys.

Beyond that, government’s role should be one of, dare I say it, leadership. Defining what we want for resilience and what standards of performance are needed in critical applications is only part of the solution. We also need to take action to ensure these standards are met by introducing clear requirements and ensuring necessary infrastructure is available. Developing an integrated infrastructure plan that

uses GEO, MEO, LEO and terrestrial components to our best advantage is a necessary step. Government needs to influence approaches not only as a provider of public infrastructure but also as a customer for private infrastructure. Resilience is best achieved as a cooperative undertaking with industry, but absent leadership, nothing is going to happen. Until it does. 

Author



Logan Scott is an expert consultant in systems/signal processing in advanced RF systems including GPS, RFID, navigation, communications, radar, and emitter location systems. A Fellow of the Institute of Navigation and holder of 45 U.S. patents, he is the inventor of the Chips Message Robust Authentication (Chimera) concept for navigation signal authentication.



NEW ELLIPSE-D



0.05°
ATTITUDE

0.02°
HEADING

1 cm
POSITION

The Smallest Dual Frequency & Dual Antenna INS/GNSS

- » RTK Centimetric Position
- » Quad Constellations
- » Post-processing Software

www.sbg-systems.com

Evaluation of Sensor-Agnostic All-Source Residual Monitoring for Navigation

Photo courtesy of National Robotics Engineering Center, Carnegie Mellon University.



The addition of alternative sensors such as cameras, magnetometers, and small ranging radios increases the likelihood of a mismodeled and/or faulty sensor, affecting the accuracy and performance of the overall navigation solution. Unlike two-sensor systems such as GPS-inertial integration, systems of three or more sensors present the problem of ambiguity as to which sensor is adversely affecting the solution. This presents the need for a robust framework that can maintain navigation integrity despite the additional sensor modalities.

Much time and effort has been invested into alternative navigation to operate independently of GNSS using sensors such as vision-aided navigation, navigation by very low-frequency radio, and more recently, through map-matching of magnetic anomalies. The primary reason behind these efforts is the need for navigation when GNSS signals are unavailable due to occasional outages, signal obscuration or in contested environments due to jamming or spoofing. The increased possibility of GNSS outages brings the need for error characterization and integration of a multitude of alternative sensors into a single platform

with a robust integrity framework. The integrity framework must be capable of monitoring all sensor measurements to perform fault detection and exclusion (FDE) of sensors experiencing adverse mismodeled effects. Additionally, the framework must be able to solve for and apply corrections to the stochastic model to account for these effects.

There has been extensive research into integrating two sensors on the same platform and assessing sensor measurements. However, when more than two sensors are incorporated on the same platform, identifying the culprit sensor becomes challenging. One method of handling sensor faults is through solution separation, using a main filter and several subfilters, each excluding one sensor. However, this is not as effective in detecting faults that do not result in a significant difference in state estimation between the uncorrupted and remaining (corrupted) subfilters. Faults producing such differences resulted in undetected faults when compared with an FDE with a more rigorous statistical algorithm.

A new algorithm, Sensor-Agnostic All-Source Residual Monitoring (SAARM), uses a sum of squared residual covariance Mahalanobis distances as a moving average χ^2 -test. Such a technique is part of a larger framework known as Autonomous and Resilient Management of All-Source Sensors (ARMAS).

Background

The ARMAS framework draws from several research papers and includes such concepts as FDE, re-calibration, and remodeling of faulty sensors and combines them into one robust framework. The ARMAS framework collects input measurements from each sensor and categorizes them into five major modes of operation: monitoring, validation, calibration, remodeling, and failing. **Figure 1** displays the modes of operation. Of these five modes, only the sensors in monitoring mode are able

ANDREW APPEGET, ROBERT C. LEISHMAN
AND MAJ JONATHAN GIPSON
AIR FORCE INSTITUTE OF TECHNOLOGY

to affect the navigation solution. The remaining modes allow the sensors to be re-calibrated and/or remodeled so that they can be validated and reused in the navigation solution. Sensors that have been placed into failure mode are allowed to be recovered via a process known as Resilient Sensor Recovery (RSR), whereby the sensor may re-attempt validation after a user-specified time period and potentially be placed back into monitoring mode.

The SAARM algorithm considers sensors operating in monitoring mode, and in the event of a fault, determines which sensor should be excluded for possible re-calibration and/or remodeling. SAARM employs multiple filters and calculates the Mahalanobis distances for each sensor/subfilter pair. The framework leverages past research of modeling and stochastic error estimation for one- and two- sensor systems to derive an overarching heterogeneous sensor-independent algorithm. This algorithm seeks to, first, evaluate if there is a potential inconsistency in the navigation solution and second, to isolate and exclude the sensor causing

such inconsistency for follow-on validation, calibration, and/or remodeling. SAARM operates within the ARMAS framework which allows additional modeling and stochastic error estimation techniques to be added and evaluated in parallel, depending upon the application. As such, the approach is a scalable, modular framework that can be added or modified as sensors are added or removed from a system, or as the modeling of particular sensors matures.

This article focuses on sensors that are in monitoring mode and are directly affecting the navigation solution. The monitoring mode contains the FDE layer that uses the SAARM algorithm to determine if there is a fault in the navigation solution, and to determine which sensor is the culprit causing the fault. Once the culprit sensor is determined, the sensor is removed from the monitoring mode and placed into validation mode at which time the sensor can no longer affect the navigation solution. The ARMAS/SAARM framework presents the navigation solution to the user as a single Bayesian filter that is supplied by all measurements vetted by this

FDE subfilter. Previously, the SAARM algorithm had been tested with only synthesized data and compared with conventional filtering techniques. This article incorporates real-world measurement data in a 3D locally tangential Cartesian plane. The data was collected from a flight test conducted by the Autonomy Navigation and Technology (ANT) Center of the Air Force Institute of Technology (AFIT). SAARM is evaluated by processing individual pseudorange measurements from multiple GNSS signals from both GPS and GLONASS satellites. A linearly growing range bias is added to a single satellite measurement to test the SAARM algorithm. The navigation solution is compared against truth data to determine the overall accuracy.

SAARM Algorithm

The SAARM algorithm is based upon a likelihood ratio test comparing the covariance of the measured residuals with the expected covariance based upon their estimated stochastic distributions. The residual measurement r is based upon the extended Kalman Filter's (EKF) predicted states.

The Chi-Squared test is derived by taking the expected value of the covariance of the measurement residuals, combining that with the Gaussian measurement equation, and applying the formula for the Mahalanobis distance and summing.

$$\chi^2_{[i,j]} = \sum_{k=1}^k \mathbf{r}^{[i,j]}(t_k) \left[\mathbf{P}_{rr}^{[i,j]}(t_k) \right]^{-1} \mathbf{r}^{[i,j]}(t_k) \quad (1)$$

The value $P[i,j]$ is the expected covariance of the residuals pertaining to sensor i , subfilter j . The value M is the number of samples within the time window chosen by the user, k , the current time index, i and j the appropriate sensor/subfilter pair corresponding to the residual $r[i,j]$, and t , the current time in seconds.

A vector of summed values can be collected by summing down the subfilter columns of the test matrix via (7):

$$\mathbf{s}(j) = \sum_{i=1}^I \mathbf{T}(i, j) \quad (2)$$

To identify the culprit sensor, assuming only one faulty sensor, the vector \mathbf{s} will consist of all values of 1, while one value will be equal to 0, indicating the subfilter which excludes the faulty sensor. The faulty sensor has now been identified and can be excluded from the monitoring mode, placed into validation mode, and routed to the proper settings in the ARMAS framework where this sensor can no longer corrupt the navigation solution. SAARM is also able to detect multiple simultaneous sensor faults by adding additional layers consisting of J subfilters where J consists of all of the combinations of excluded sensors.

The SAARM algorithm will calculate a 2D Guaranteed Position Zone (GPZ) based upon the estimated covariance of each subfilter's horizontal position states (East and North). This data will be used to calculate and display a χ^2 error ellipse

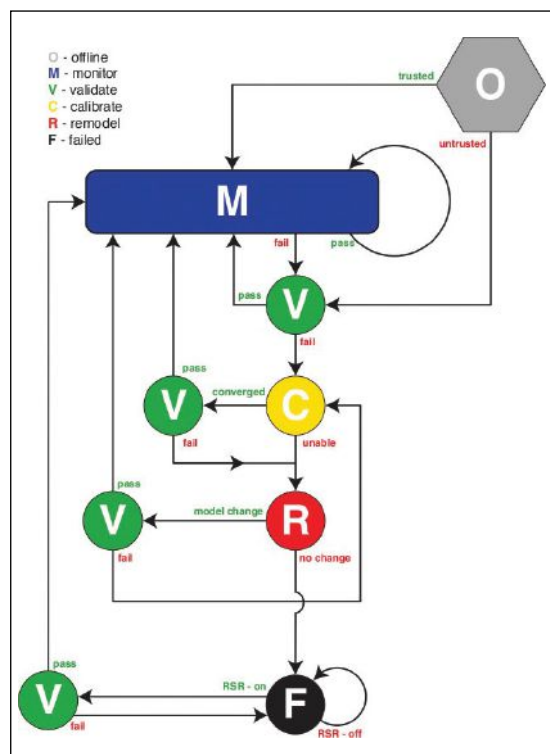


FIGURE 1 ARMAS Diagram detailing the five modes of operation: Monitoring, Validation, Calibration, Remodeling, and Failing.




FIGURE 2 Truth Flight Path of UAS. The coordinate system is based upon a local-level plane on the Earth, whose origin is the initial position of the UAS.




FIGURE 3 SAARM algorithm with no corrupted sensors (just before fault is applied). Each dot represents the position solution of each subfilter in the locally tangential plane (where x = Latitude and y = Longitude), whose origin is the initial position of the UAS. The altitude is the altitude above the initial position. Each filter is named after the sensor that is excluded, by constellation and PRN number.

with confidence level $100(1-\alpha)\%$ for each subfilter. The GPZ is defined as the union of all of the subfilter's 2D error ellipses for a given confidence level. For altitude, the integrity will be based upon the union of each subfilter's 2σ covariance bound. One important assumption of SAARM is that there is full state observability across all subfilters. Another is the assumption that at least one of the filters is informed entirely by properly modeled, uncorrupted sensors, and at least one filter contains consistent estimation error statistics.

State Dynamics

A bank of EKFs are required for the SAARM integrity filter, where one subfilter is required for each of the combinations of ways to exclude one or more sensors, depending upon how many integrity layers the user chooses to implement. Each of these subfilters are modeled the same, with the only difference being which sensors are excluded from measurement updates. The states to be estimated correspond to the position, velocity and acceleration of the vehicle and the clock errors for the GPS and GLONASS constellations.

The state dynamics assume standard EKF dynamics formulas, with 11 states: 3 each for position, velocity, and acceleration with clock offset states for the GPS and GLONASS constellations. The

acceleration and clock states are modeled as FOGM dynamics models. The GPS and GLONASS measurements are modeled as a typical pseudorange with Gaussian measurement error distributions. The states are modeled with:

$$\mathbf{x} = [E_{pos}, N_{GPS}, U_{GPS}, E_{vel}, N_{vel}, U_{vel}, E_{acc}, N_{acc}, U_{acc}, b_{u,GPS}, b_{u,GLO}]^T \quad (3)$$

where $b_{u,GPS}$ and $b_{u,GLO}$ are the FOGM clock states. It is assumed that the location of each SV is known to an accuracy of 1–2m given each constellation's broadcast ephemeris message.

Here we process real measurement data from 6 GPS satellites and 4 GLONASS satellites. For the purposes of SAARM, we treat each satellite as an independent sensor.

The position is calculated in the local tangential frame where the origin is set to the initial truth position of the platform. The position of the satellites is calculated using the broadcast ephemeris message, using Keplerian orbital parameters for GPS satellites, and using an ODE solver for GLONASS satellites. The GLONASS satellite locations were transformed to account for a translation and rotation correction due to GLONASS's use of a different ellipsoid.

UAS Data Collection

A dataset was collected during an unmanned aerial system (UAS) flight

test conducted by the ANT Center at Camp Atterbury, IN. Measurement data from approximately 27 minutes of one flight scenario is used here. The aircraft flew at predominately two different altitudes: 250 and 100 meters above the local surface. The truth dataset was derived using the ANT Center's Scorpion framework to process measurements from several sensors, including inertial measurements, to arrive at a standard EKF solution at a given timestamp. This solution was used as the truth trajectory. The timestamp was pulled from the UNIX system time on the computer running the Scorpion software. The plot in **Figure 2** depicts the flight trajectory of the truth dataset. The origin of the coordinate system was set to the initial position of the UAS at the beginning of the data collection.

GNSS range data was collected from a GNSS receiver every 0.2 seconds in GPS time ($dt = 0.2$). To synchronize the truth data with the range data that was collected, the truth position data was converted to the Cartesian (Earth-Centered Earth-Fixed, ECEF) coordinate system and linearly interpolated in each dimension to match the times at which the GNSS measurements were received. To account for leap seconds, 18 seconds were added to the UNIX time to match up with the time of reception of

the pseudorange measurements, which were reported in GPS time. To take advantage of a GNSS sensor's ability to make ionosphere range corrections, the ionosphere bias was removed from the measurement prior to being processed by the SAARM subfilters. The bias was removed via the L1 Klobuchar parameters broadcast with the GPS ephemeris message. These parameters were applied to the GLONASS pseudoranges as well.

Parameters and Initializations

Table 1 defines the parameters for the system dynamics and measurement covariance. These values were replicated in every subfilter as they were initialized. Every subfilter was initialized with the same initial prediction error covariance (P_0) data.

The initial state estimation was set to zero, where the origin of the position coordinate system is the initial truth position, although unknown to the filter. The filter processed pseudorange

measurements with the ionosphere removed, and with the satellite positions precalculated and converted to the local-level frame. The ARMAS integrity filter requires the selection of two key parameters by the user: the monitoring period, M in seconds; and probability of false alarm. For this dataset,

was set to 30 s; the α for each individual filter was set to 6.67×10^{-6} .

Flight Test Results

The SAARM algorithm was post-processed from data collected from a GNSS receiver's raw pseudorange measurements during the ANT Center's flight test. It is assumed that no previous bias, with the exception of minor multipath, was incorporated into the receiver's measurements. A growing range bias of 1 m every second (i.e. 0.2 m every 0.2 seconds) was added to one SV's measurement to test the SAARM algorithm to detect and exclude this measurement. The 2D and Altitude GPZ, as well as the detection and exclusion of one culprit SV measurement, are illustrated in **Figures 3 through 5**. The names of the subfilters are based upon the particular constellation and PRN of the excluded measurement. The values of x and y are based upon the distance from the initial truth position of the UAS, x for Latitude, y for Longitude. The horizontal dilution of precision (HDOP) values are also depicted, along with the faulty sensor,

the distance of the bias added to the faulty sensor, as well as the 2D, 3D, and altitude errors from the truth.

Figure 3 depicts SAARM just before the fault is applied, depicted by the small red triangle in the truth path. **Figure 4** depicts the detection of a faulty sensor, but SAARM has not yet identified the culprit. **Figure 5** depicts SAARM after the culprit has been identified and excluded. The observed images demonstrate that the SAARM algorithm is able to work with real-world measurements.



The faulty sensor as well as the biases that are applied, the time elapsed since the beginning of the data set, and 2D and 3D errors from the truth are depicted in the middle-right window. HDOP values for each subfilter are depicted in the lower right, providing useful data to assist with one particular phenomenon observed throughout the running of the ARMAS filter.

Analysis

HDOP data provides insight into the observability of the system, although it is not exactly observability. The correlation between geometry and observability is especially strong in this experiment, as the only measurements in the filter are GNSS measurements. When there is weak or only partial observability, the subfilter solutions can vary dramatically and can seem to "swing" wildly from one solution to another. Loss of observability

Parameter	Value
R_{GPS}	$(1)^2 \cdot \frac{I}{6 \times 6}$ [m ²]
R_{GLO}	$(4)^2 \cdot \frac{I}{4 \times 4}$ [m ²]
τ_a	$100 \cdot \frac{I}{4 \times 4}$
σ_a	100 m/s ²
$\tau_{b,GPS}, \tau_{b,GLO}$	3600 s
$\sigma_{b,GPS}, \sigma_{b,GLO}$	8000 m

TABLE 1 System Dynamics and Measurement Statistics.



also removes the power of this integrity approach to detect and remove sensor problems. An example scenario of when this might happen is when there is low availability of GNSS satellites and, due to the relative geometry, two or more satellites can appear close together, effectively reducing the number of satellites. Effects of this phenomenon can be observed when running the SAARM algorithm with only GPS satellite measurements.

To illustrate, the SAARM algorithm compares two different scenarios near the same time: one incorporating only 6 GPS SV measurements, and the other that includes 4 GLONASS SV measurements for a total of 10 SV measurements. **Figure 6** shows the first scenario with bad geometry while **Figure 7** shows the latter scenario with better geometry.

The measurements from satellites below 20 degrees elevation were not used due to greater multipath and increased measurement dropouts. The 20-degree cutoff line is indicated in **Figure 6**. As illustrated in **Figure 6**, the subfilter excluding GPS PRN 29 suffers from bad geometry, as the HDOP is greater than 10. The bad geometry causes intermittent integrity warnings due to false alarm errors, although the algorithm did not outright exclude any sensors.

Careful observation of the skyplot in **Figure 8** details why this occurs: Excluding PRN 29 removes the middle

satellite from the overall geometry. The remaining satellites are not well spread out from the perspective of the GNSS receiver on the UAS. Therefore, for the geometry depicted in this dataset, when only incorporating the six GPS satellites above 20 degrees elevation, SV 29 is the most critical to maintaining a good geometry. This SV's critical importance is underscored by the low availability of GNSS-only measurements above 20 degrees elevation. The remaining subfilters that excluded other satellites did not suffer from this issue. In general, this is a violation of SAARM's position state observability assumption for all subfilters.

Observability

The ARMAS framework with SAARM was originally conceived and simulated with basic 2D sensors and assumed fully overlapping position observability in every subfilter within the FDE subfilter layer. Pseudoranges are extracted from six SVs for processing in ARMAS as individual pseudorange sensors. During the analysis of the 27-minute flight test dataset, it is clear that the latter half of the flight test dataset contains numerous sensor dropouts. Initial analysis reveals that a sensor dropout causes spurious behavior in ARMAS. SAARM is unable to provide a subfilter consensus to identify a failed sensor when the FDE

subfilter layer loses overlapping position state observability. In other words, SAARM can detect a sensor fault but cannot exclude the faulty sensor if even a single FDE layer subfilter loses position state observability due to dropout, poor geometry, etc.

Since the initial simulation of ARMAS was performed with fully overlapping position state observability in the FDE layer, this problem was not identified in development. **Figure 9** shows a local covariance analysis of the observability layer subfilters. For the observability analysis, a second integrity layer of subfilters is added, similar to a case when one would expect two faulty sensors. There are now 2 integrity layers consisting of 15 subfilters in the observability layer for 6 pseudorange sensors. Each subfilter in the observability layer is informed by 4 unique pseudorange sensors, the minimal case for a stable 3D solution.

In other words, a single pseudorange sensor dropout results in the loss of position state observability and is evidenced by an increase in the trace(Pxyz). At approximately $t = 1175$ sec, SV 5 experiences a transient dropout which affects all but 5 of the observability layer sub-subfilters (No GPS 5 GPS 13, No GPS 5 GPS 15, No GPS 5 GPS 20, No GPS 5 GPS 29). As a side note, these observability layer sub-subfilters would

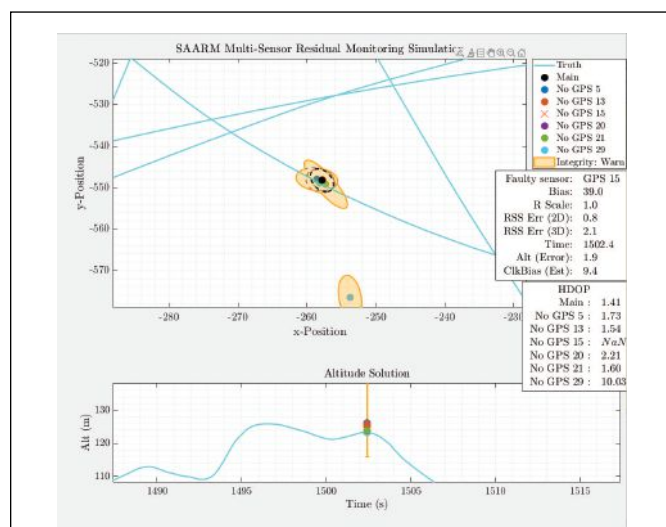


FIGURE 6 Subfilter with bad geometry deviates from main solution. The filter excluding GPS SV 29 has an HDOP over 10. The approach is aware that the integrity is compromised but does not identify and exclude a fault sensor.




FIGURE 7 Subfilter with better geometry. The filter excluding GPS SV 29 now has an HDOP below 2. The solution of this subfilter does not deviate by as much and false positive detections do not occur as often.


form the new FDE layer if GPS 5 were excluded from the monitoring mode. This analysis forms the genesis of local observability research at the subfilter level to preserve the consistency of the FDE layer and the integrity functions of ARMAS in the event of a single simultaneous sensor failure.

Conclusions and Future Work

Processing data from a flight test, the SAARM algorithm processed pseudorange measurements with one measurement corrupted with a growing range bias of 1 meter per second. The algorithm successfully detected and excluded this corrupted satellite measurement with no false exclusions before or after this time. Observing the geometry of certain subfilters gives further insight into the factors leading to possible false

alarm errors in the event of low availability due to jamming, spoofing, or operating in urban or other obscured environments. Incorporating additional GNSS signals through adding the GLONASS constellation alleviated the issue of bad geometry. Insight into observability can be acquired through adding another integrity layer.

The ARMAS framework and SAARM algorithm can perform fault detection and exclusion with real-world data fused from multiple GNSS constellations and support further research and applications. Future work will incorporate sensors such as ranging radios, inertial measurement units, lidar and velocity sensors into ARMAS. These sensors will update at different rates and will thus test ARMAS with real-world asynchronous measurements. Research will also

explore estimating the attitude and utilization of error states within SAARM. The ARMAS framework will also incorporate all five modes of operation to test the re-calibration or resilient sensor recovery of failed sensors. One such method to test ARMAS is a bias that would be present and then disappear so that the sensor would return to monitoring mode. Follow-on applications would then be to apply the ARMAS framework in real-time to detect bad measurements to improve navigation integrity. 

Acknowledgements

The authors wish to thank Lt. Col Juan Jurado of the U.S. Air Force Test Pilot School for developing the ARMAS framework and Dr. John Raquet for providing guidance and reviewing the material.

Authors



Andrew Appleget is a research engineer at the Autonomous Navigation and Technology (ANT) Center at the Air Force Institute of Technology (AFIT) at Wright-Patterson Air Force Base (WPAFB). He received his MSEE from Ohio University.



Robert C. Leishman is Director of the ANT Center at AFIT and research assistant professor in the Dept. of Electrical & Computer Engineering. His research spans autonomous vehicles, robust alternative navigation, image processing, sensor fusion, and control. He has a Ph.D. in mechanical engineering from Brigham Young University.



Maj. Jonathon Gipson is a Ph.D. candidate at the Air Force Institute of Technology and an experienced flight test engineer in the Air Force. He is a graduate of the Air Force Test Pilot School.




FIGURE 8 Skyplot of GPS satellites. GPS SV 29 is centered and is critical to geometric diversity toward the end of the data run (NOTE: Satellites below 20 deg. were not used, as noted by the black cutoff line)




FIGURE 9 Observability Layer Covariance Analysis of SV 5 Sensor Dropout, Flight Test Data, 12 Oct 2018, Camp Atterbury, IN.

Consult conference websites
periodically for changes.

See additional listings at
www.insidegnss.com/events

August 2021

AUGUST 16–19

AUVSI XPONENTIAL
Atlanta, Georgia

www.xponential.org/xponential2020/public/enter.aspx

AUGUST 24–27

JOINT NAVIGATION CONFERENCE (JNC)
Cincinnati, Ohio

www.ion.org/jnc

September 2021

SEPTEMBER 21–24

ION GNSS+
St. Louis, Missouri
www.ion.org/gnss



Photo courtesy of Daniel Schwen, via Wikimedia Commons.

SEPTEMBER 21–23

INTERGEO
Hanover, Germany
www.intergeo.de/en/news/intergeo-2021-hanover-germany

SEPTEMBER 27–OCTOBER 1
**INTERNATIONAL COMMITTEE ON
GLOBAL NAVIGATION SATELLITE
SYSTEMS (ICG): ANNUAL MEETING**
Vienna, Austria

unoosa.org/oosa/en/ourwork/icg/meetings/ICG-2021.html

October 2021

OCTOBER 4–7

**INDOOR POSITIONING AND
INDOOR NAVIGATION (IPIN 21)**

Barcelona, Spain

ipin-conference.org/2021/



Photo courtesy of Jorge Franganillo, via Wikimedia Commons.

OCTOBER 11–13

**ASSOCIATION OF THE U.S. ARMY (AUSA)
ANNUAL MEETING & CONVENTION**

Washington, D.C.

meetings.ausa.org/annual/2021/

November 2021

NOVEMBER 15–18

**NAVIGATION 2021:
THE EUROPEAN NAVIGATION
CONFERENCE AND THE
INTERNATIONAL NAVIGATION
CONFERENCE**

Virtual

rin.org.uk/mpage/Navigation2021

January 2022

JANUARY 24–27

**ION INTERNATIONAL TECHNICAL
MEETING AND PRECISE TIME AND
TIME INTERVAL MEETING**

Long Beach, California

www.ion.org/itm/index.cfm



Photo courtesy of Ed g2s to English Wikipedia, Public domain, via Wikimedia Commons.

April 2022

APRIL 11–14

ION PACIFIC PNT 2022
Honolulu, Hawaii

www.ion.org/pnt/past-meetings.cfm

ADVERTISERS INDEX

Company	Page Number
Advanced Navigation	19
Analog Devices	7
CAST Navigation	67
Emcore	17
GPS Networking	31
Institute of Navigation	55
Jackson Labs	9
L3Harris	2
NavtechGPS	25, 51
NovAtel	15, 68
Racelogic/LabSat	13
SBG	59
Sensoror	22–23
Silicon Sensing	11
Spirent	3
Syntony	35
TeleOrbit	6
VectorNav	5
Work Microwave	21
XPonential	47

GNSS Simulation | IMU Simulation | Wavefront Simulation | Interference Simulation

ONE SIZE **DOES NOT** FIT ALL.



Whether you need to test a single navigation system or simulate an entire squadron, CAST Navigation will configure the capabilities you need today for **GNSS, Wavefront, IMU, and Interference Modeling Simulations**. Over time, our modular, synchronizable systems can expand your capabilities - when you need them - to meet future requirements.

Through 40 years of GNSS/INS simulation innovation and a resolute commitment to our customers, CAST Navigation earned the trust of military and commercial clients around the world. Contact CAST Navigation's GNSS/INS simulation specialists to learn how we can simplify your testing and integration programs.

CAST
NAVIGATION
www.castnav.com



RTK From the Sky™ will change the world

RTK From the Sky is bringing instant GNSS accuracy to the world. Ultra-fast PPP convergence in minutes for centimetre-level accuracy on land, air and marine applications. RTK From the Sky combines benefits of multi-GNSS station reference networks with advanced receiver positioning algorithms and technology. That means the highest reliability data for assured, reliable positioning anywhere, anytime. Our solutions provide you the surest path to success — a path followed by countless leaders worldwide in the fields of transportation, agriculture, defense, surveying, mining, marine and construction. We can help you, too.

Autonomy & Positioning – Assured | hexagonpositioning.com/rtk-from-the-sky

