ECHODYNE

Highly-Adaptive Radar for Cognitive Imaging.

R

A
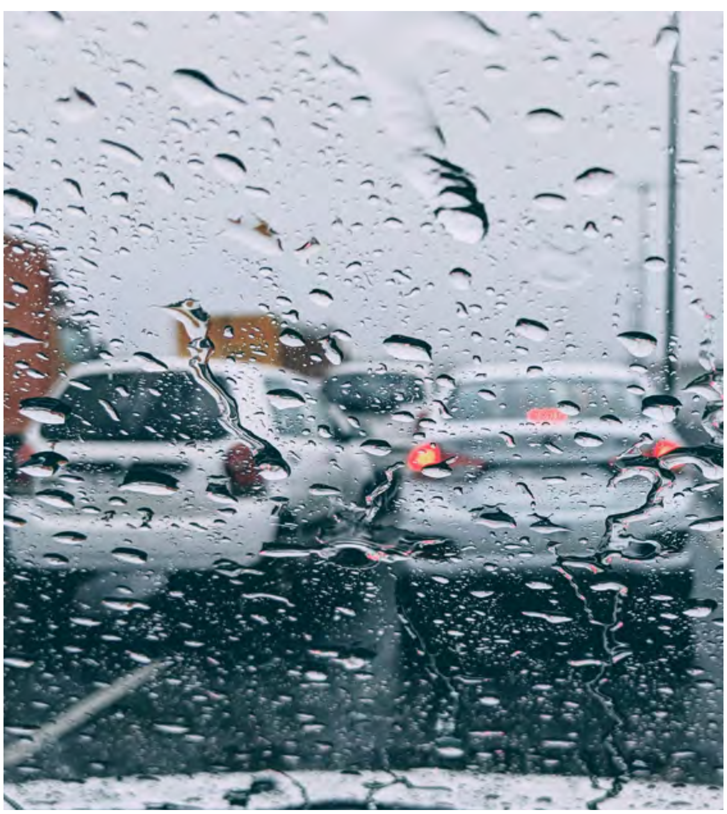
DA
R

Highly-Adaptive Radar for Cognitive Imaging

Tom Driscoll,* Robert Worl, John D Hunt, and Leo McCloskey†
Echodyne Corp. Kirkland, WA

In this white paper we look at how recent advances in knowledge-aided (KA) measurement techniques can enhance the capability of radar imaging in dense scenes and complicated environments. Following the thinking of the fully-adaptive cognitive radar systems pioneered for military radars, but with an eye towards commercial autonomous vehicles (AVs), we propose an architecture for a lightweight form of highly-adaptive cognitive radar (HACR). We discuss how dynamic reallocation of the limited time and bandwidth resources of a cascaded multiple input multiple output (MIMO) radar provides flexible, *adaptive*, measurement; which provides greater resolution and extent - in angle, range, and Doppler - than could be possible with a static equivalent.

I. INTRODUCTION TO COGNITIVE MEASUREMENT

The field of cognitive radar [1–3], and the broader topic of cognitive measurement [4], is built on principles of using prior information to design measurement to optimally exploit characteristics of both the channel and target(s). Cognitive measurement relies on knowing, forming, and refining a complete model for the complete multi-dimensional radar channel, environment, and targets, and using this information to improve acquisition of targets of interest. By designing the information transmitted (waveform(s)) and the channels in which it's transmitted (antenna gain pattern(s)) in context of prior information about target types and locations, significant improvements can be realized in the radar's measurement capacity [5, 6]. Such dynamic approaches are common in communications, with the channel-sounding used in wireless communication as the most ubiquitous example. Yet, such use is recent in radar.

Cognitive measurement techniques share a heritage with forms of adaptive digital signal processing (ADSP), especially techniques that utilize priors. However, in contrast to ADSP, which seeks to improve extraction of information from an existing measurement, cognitive radar uses prior information to *actively* design the interrogation signal(s), e.g the measurement matrix. The knowledge used often exists outside the radar data chain, and can include both endogenous and exogenous sources. Examples include historical/statistical data repositories, locational/positional information, and data streams from other sensors. The goal of this is improvement in the acquisition of the signal(s) of interest. In radar, those signal(s) of interest are reflections off target(s) of interest, measured against a background that obscures those signals - whether the hindrance comes from noise, interference, clutter, or other sources. (For convenience we'll use Signal-to-Clutter-plus-Interference-plus-Noise-Ratio SCINR, regardless of the source(s) of interference).

FIG.1 presents the basic cognitive radar approach, in context of a generalized radar system. Information about the channel and target(s) is contained in knowledge repositories - information such as maps, historical data, and real-time sensor data. This is consumed by a KA coprocessor, which uses that information to adjust free parameters within the synthesizer (modifying time-voltage spectral content) antennas (modifying the beam/null-forming spatial layer), and receiver processing.

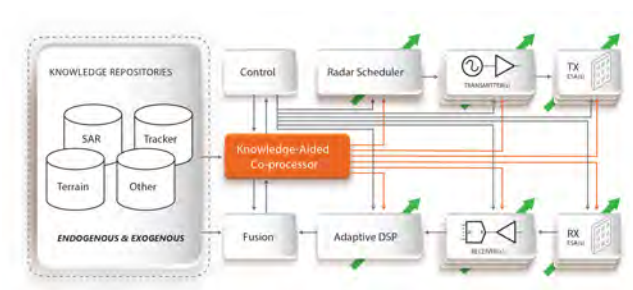

FIG. 1. Architecture of the cognitive measurement technique. Knowledge from repositories is used by a KA coprocessor to adjust the measurement matrix formed through the transmit-receive chain; controlling both the time/spectral content ((waveform(s)), and the spatial content at the physical layer (antenna patterns), as well as traditional ADSP throughout the receive chain.

The general form of the measurement equation is

$$\bar{g} = \bar{\bar{H}}\bar{f} + n. \tag{1}$$

where $\bar{f}$ is a vectorized scene, $\bar{\bar{H}}$ is a measurement matrix, n is an additive noise term $\bar{g}$ is the resulting measurement vector.

The greater the DoF's along all portions of this chain, the larger the parameter space available to exploit in optimization of H. Of course, no matter the quality/quantity of the knowledge repository, perfect knowledge of the channel and target are unobtainable, and in practice all KA approaches are a type of prior estimation. However, real-world testing shows dramatic improvement in the overall extraction of signal(s) of interest. FIG.2 evidences this point, presenting results for one radar operating in a series of A/B testing with and with-

* Also at ECE Department, Duke University.
† Corresponding author: Leo@Echodyne.com

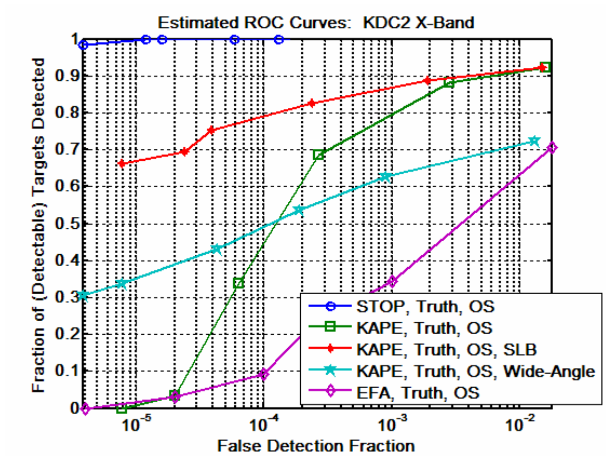

FIG. 2. Data from the DARPA KASSPER challenge, giving an example of the performance improvement possible using knowledge-aided techniques. Receiver operating characteristic (ROC) curves are shown for three KA variants developed by GRTI (prefix "KAPE") with a traditional post-Doppler STAP approach (prefix "EFA"). The clairvoyant unobtainable optimum is also plotted (prefix "STOP"). Reprinted with permission from the author[3]

out knowledge-aided techniques. The increase in the receiver operating characteristic (ROC) curve is dramatic.

It should be pointed out that these KA techniques improve the performance of a radar defined by a certain set of parameters/attributes that reflect its hardware realization. Attributes can include transmit power, receiver noise figure, number of channels and bandwidth (BW), and antenna sizes and efficiencies. Improvement in the SCINR could certainly be obtained through escalation of those performance-constraining hardware blocks; e.g. simply making a higher-power radar. However, in many cases these attributes have already been driven to the limits of what the application can realistically support, driven by bounds on cost, size weight, and/or power (C-SWAP). The KA approach, on the other hand, potentially only requires additional software (possibly supported by additional computation power). In analysis of specific cases in [7], KA design of $H_{tx}(\tilde{E}_i(t))$ resulted in between 7 to 24dB of improvement in SCINR (depending on waveform compression ratio). Achieving similar increase through hardware upgrades (ERP, G/T, etc), would be a significant undertaking, and dramatically increase the cost of the radar hardware. This points to the admission that all measurement systems are subject to some bound on the resources used to realize them. Cognitive measurement is a suite of techniques that get more performance out of the same resources through intelligent dynamic tasking.

As we'll see in the next section, design of high-resolution imaging radars face a daunting resource challenge, and solving the problem through brute-force seems incompatible with commercial C-SWAP constraints. The possibility that KA techniques could act as a force-multiplier on limited resources for this application is compelling. While there is certainly a cost to implementing the control and feedback loops required for cognitive measurement, that cost is small in comparison to the improvement they offer. After framing the imaging radar resource problem, we'll propose an architecture for a highly-adaptive cognitive radar that - while not as powerful as the fully-adaptive approaches in [1] - is aligned with the constraints and requirements of commercial AV's.

II. HIGH-RESOLUTION RADAR IMAGING

Electronically scanned array (ESA)-based radars natively collect information in four basis vectors: range, Doppler, and spherical angle (azimuth and elevation). However, measurement capacity is rarely uniform across these. Range and Doppler extent and resolution are primarily defined by the waveform, which is limited by the system bandwidth of the receiver(s) and the time extent of the waveform (T). The product of these is the time-bandwidth-product (TBP), a key metric of radar performance. In modern radars the resulting 2D range-Doppler matrix (RDM) over the full coherent processing interval can easily be hundreds of thousands of points. Multiple receivers can be combined to increase the RDM area, although it's more typical that multiple receivers are employed as spatial sampling for angular resolution.

In the angle basis, achieving a desired resolution over a field-of-view (FOV) requires Nyquist-complete spatial sampling, with phase centers spaced at no less than the transverse wave-vector $1/(2k_\perp)$ [8]. In order not to degrade waveform capacity, each of these phase centers must receive the full TBP, resulting in a large number of RF channels on transmit, receive, or both. FIG.3 illustrates the situation, presenting a table of the required sampling phase centers to form a Nyquist-complete aperture with a given angular resolution in both azimuth and elevation over a desired FOV (where we assume the element array-factor has been designed to deal with grating lobes outside the FOV). To date, it has been impractical for commercial radars to realize these large RF-channel counts, though modern military active electronically scanned array (AESA) radars routinely have thousands of elements [9, 10]. Sparse or thinned antenna arrays with sub-Nyquist spatial sampling offer some respite, but thinning increases the array side-lobes, effectively reducing the dynamic range in the angular basis [11, 12]. Numerically optimized sparse arrays have been able to reduce channel count by 20% to 50%, while still maintaining better than 20dB of peak-side-lobe-level (PSLL) dynamic range over the majority of k-space [13, 14]. In all cases, the greater the instantaneous FOV of the array, the more difficult it becomes to design a sparse array with low PSLL. One approach some modern AESA's employ, and which a cognitive radar can leverage, is to design an array with high resolution over a narrow instantaneous field of view, and electronically scan this over a larger field-of-regard.

Angular Resolution (Az x El)		Field Of View (Azimuth x Elevation) 5°x10°	10°x10°	10°x20°	45°x20°	90°x20°	120°x20°
2D	20°x0°	2	1	1	2	4	6
	5°x0°	2	2	2	8	18	24
	2.5°x0°	3	4	4	18	36	48
3D	5°x5°	3	4	8	36	72	96
	2°x5°	5	10	20	90	180	240
	2°x2°	10	25	100	225	450	600
	1°x2°	25	50	100	450	900	1200
	0.5°x1°	100	200	400	1800	3600	4800

FIG. 3. Number of sampling phase centers (physical or virtual) required for a Nyquist-complete aperture with a given angular resolution and field-of-view. Highlighted cells show typical values for long-range and short-range ADAS radars (in gray), and what's often requested for Level 4+ AVs (in blue)

Recently, progress on the challenge of implementing large numbers of channels has been made by synthesizing a greater number of virtual phase centers from a given set of physical transceivers, using MIMO techniques [15, 16]. Given M transmitters and N receivers, MIMO enables the synthesis of up to MxN "virtual" phase centers, through the use of orthogonal transmitted waveforms. This approach offers a fairly dramatic reduction of the transceiver channels required to achieve a given angular span. However, all variants of MIMO fundamentally represent a "re-use" of the TBP of the radar, either time (TDMA) or bandwidth (FDMA or CDMA) - depending on the orthogonality basis - and this re-use of the TBP comes at a cost of degraded RDM. FIG.4 plots the total mean self-interference for a CDMA radar assuming ideal codes and perfectly incoherent cross-correlation (xC) (both of which are unachievable, and in practice the xC will be worse than this). Just as in communications, the more users of a given spectrum, the less un-polluted bandwidth is available to each. In radar, the larger the MIMO virtualization in terms of the number of simultaneous transmitters, the more cross-channel interference reduces the measurement dynamic range. While in many cases trading degraded RDM for improved angular resolution is a net-win, it can also obscure micro-Doppler or range-profile features critical for classification. Like so many other design options, MIMO is a trade-off.

In many historical applications of radar, limited angular resolution has not been an issue due to the reliance on super-resolution algorithms - with monopulse the best-known of these. When the radar scene to be measured is canonically sparse (e.g. small number of objects in an empty airspace) or has been sparsified through some assumption-set (e.g. only a small number of fast-moving targets) that information can be used to estimate the angle-of-arrival of targets with an accuracy greater than the antenna resolution. In well-calibrated systems, super-resolution can achieve accuracy/resolution

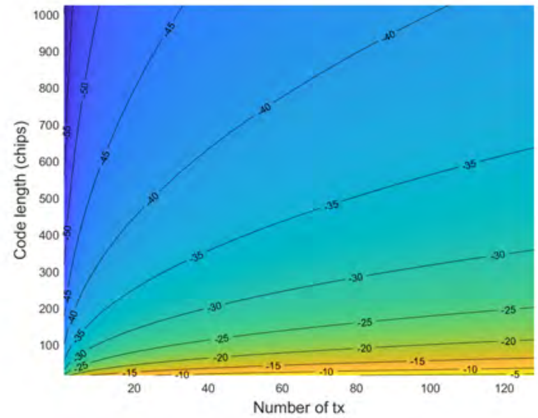

FIG. 4. Total average cross-correlation level for a coded MIMO radar using ideal codes, as a function of the code length and number of simultaneous inputs. This cross-correlation appears as interference in the basis of orthogonality (range-basis for fast-time-coded (e.g. PSK radar), or Doppler-basis for slow-time-coded (e.g. FMCW radar)).

improvement ratios of 100:1 or more [17]. However, these methods - often eigenspace estimation methods [18] - are strongly reliant on the validity of underlying assumptions about the sparsity of $\bar{f}$ and any regularization techniques applied. In the dense driving scenes for AV level 4+, the validity of these assumptions are highly suspect. Additionally, the computational cost of super-resolution algorithms can be significant, especially methods which attempt to handle poorly conditioned matrices. As a result of these factors, super-resolution approaches have a bound on the number of simultaneous detections that can be processed. Modern ADAS radars often only handle tens of detections [19], and imaging radars relying on super-resolution with scaled-up computation can handle hundreds to thousands[20]. Unfortunately, this number is *significantly* less than the number meaningful points which can exist in a typical scene $\bar{f}$ (the space-bandwidth product), which means information is being truncated/lost. Whether this information is critical to the AV depends on the exact details of the algorithm(s) and nature of the scene(s). However, in all cases, the takeaway is that super-resolution is a tool to be used judiciously, and antenna (Rayleigh) resolution is likely preferable when it can be achieved.

The above has presented some picture of the challenges of achieving a high-performance imaging radar, and the engineering trade-offs that are made between angular resolution and extent, transceiver channel count, and range-Doppler resolution and extent. Cognitive techniques help with the imaging problem by allowing trade-offs to be made in real time.

III. COGNITIVE RADAR FOR AUTONOMOUS VEHICLES

When seeking to apply these techniques developed by DARPA for high-end military radar systems to commercial applications, it is important to take a sober look at constraints and compatibility. While the fully-adaptive approaches outlined in [1] are impressive, they also impose serious requirements on system hardware and users. Though AV's are outfitted with high-end specialized processors, the total compute capability still pales in comparison to many battlefield systems. Radar processing techniques like space-time adaptive processing (STAP) are often executed in banks of top-of-the-line FPGA processors, and the addition of KA further escalates the required computing. Additionally, it is unrealistic to expect designers and operators of AV's to be experts in radar; and the DoFs available in fully-adaptive waveform design and ESA beamforming are a double-edge sword. To make these cognitive techniques commercially accessible, operational details need to be abstracted and the control simplified. While this has the unfortunate side-effect of reducing the total number of DoF's and truncating some KA pathways, one still gains significant advantage even in a greatly simplified architecture.

A first consideration is to examine the nature and location of knowledge repositories. AV's have tremendous quantities of information available to them: High-definition maps, historical/statistical databases, real-time vehicle-to-vehicle (V2V) communications, and an extensive suite of other sensors — all of which is prime for KA exploitation. Indeed, most AV Software Stacks (AVSS) already employ some KA approaches, even if semantics may differ, in their data processing and sensor fusion. Looking back at FIG.1, the classical cognitive architecture implies that both endogenous and exogenous data is made available *internal* to the radar, where a KA coprocessor works alongside the radar scheduler to define and direct measurement. In an AV system, this architecture would be incongruous. Given the quantity and nature of the information, passing all of it to the radar would require duplication of significant data bandwidth, storage, and compute. Instead, the more logical approach is to move the KA control block to where the information (and compute) already resides — outside the radar.

FIG.5 illustrates such an architecture, where control and analysis APIs are placed inside the AVSS. The differences are subtle but impactful. First, the existence of an external control loop is very different from the uni-directional flow of traditional (non-cognitive) sensors, which only permit slow-time changes of configuration/modes. A human analogy of "Brain" and "Eyes" (a common theme in cognitive measurement) helps explain both the utility and and design philosophy. Human eyes behave as dynamically taskable sensors, in much the same way as a cognitive radar does, but the decisions of *where* and *when* to look reside firmly in the brain. In an AV,

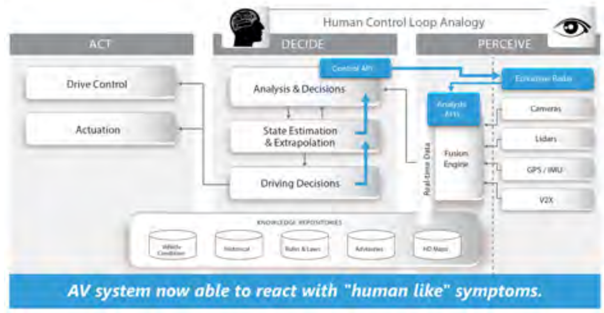

FIG. 5. Knowledge repositories and control flow for a cognitive AV radar. The control loop through the AVSS allows sensor feedback actions from all layers of the driving stack, many of which are reminiscent of human responses. Information about driving decisions (e.g. "turn left ahead") can be used to task the radar in ways that are reminiscent of *anticipation*. Tasking from state-estimation layers can be used as *reaction* (e.g. emergency breaking maneuver). Perception layers can task the radar for *confirmation* (e.g. tie-breaker on $Prob_{pedestrian}$).

the brain is the AVSS, and all decisions of how to modify $\bar{\bar{H}}$ should be made there. The radar sensor acts as the eyes, and its role is to execute measurement as tasked by the brain.

This leads to an interesting point: taking an overall view of AV system control and data flow, "smart sensors" are the *opposite* of what's desired. It's no more logical for the radar to select $\bar{\bar{H}}$ than it is for your eyes to decide where to look. Smart sensors make sense when decision making *needs* to be done at the edge for reasons of upstream bandwidth; but modern AVs have high-speed interconnects that allow low-level data to be returned to the AVSS for processing [21]. Putting aside business considerations of value-chain ownership, the strongest systems-level solution is to perform as much analysis & control within the AVSS core as possible, and the HACR architecture pursues this objective.

Another serious consideration is cost. To date, cognitive radars have been expensive. Some of this is the compute required to perform KA processing, but a great deal is due to the hardware required to execute arbitrary control over tx & rx DoF's. The existing automotive market is extremely cost sensitive, and while the increased value of full autonomy is expected to dramatically shift some cost structures, pricing pressure will certainly remain [22]. ADAS radars achieve their breakthrough low cost-points thanks in large part to integrated MMICs produced in extremely high volumes by leading chipset manufacturing companies (Texas Instruments, NXP, Infineon). In order to meet realistic cost-points, a cognitive AV radar needs to be designed to leverage such MMICs; preferably using existing chipsets given the enormous costs to qualify a new design to automotive standards, and the fact that cost savings are not realized until volume-sales are successful. MMICs can further limit cognitive capabilities, especially as their cost-driven design strips away functionality not needed for ADAS radars.

Whatever MMICs are used, their operation needs to be synchronized with the rest of the system. Continuing down the architectural path started in FIG.5, FIG.6 presents block decomposition of how the control is implemented. A single "schedule execution" block internal to the radar abstracts all $\bar{\bar{H}}$ DoF's in the tx & rx chain, including MMIC functions and ESA beamforming. Valid DoF combinations are pre-computed and stored as a library of possible $\bar{\bar{H}}$. This approach has the advantage that, unlike in a fully-adaptive radar, the set of $\bar{\bar{H}}$ is bound and can be fully verified in production - a necessity for certain tiers of automotive QA. However, the obvious downside is that pre-computation does not allow for fine tuning of $\bar{\bar{H}}$ to exploit target/channel models. Of course, the production $\bar{\bar{H}}$ library can be augmented over time to fit known archetypes (possibly with over-the-air updates).

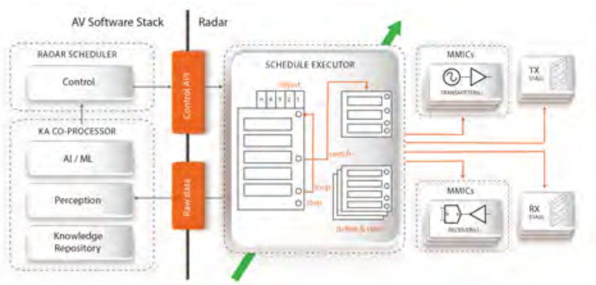

FIG. 6. HACR control architecture based on a schedule executor employing nested $\bar{\bar{H}}$ lists. DoFs in the tx and rx chains are abstracted, presenting a single "knob" interface.

To operate the radar, the external control API arranges $\bar{\bar{H}}$ entries into sequence lists, which are stored locally to the radar. The sole role of the sequence executor is to run these lists as a "dumb" state machine, and respond to any change-commands from the control API. Lists can be updated in real time, or created in advance (advantageous for scriptable events like "unprotected left turn"). Lists can be cascaded (linked finish-to-start) and timed (for start and/or stop); and at any time the control API can switch from any one list to any other. This control architecture follows four main design principles that reflect the needs of AVs: 1) Responsive: the radar should respond near-instantaneously, 2) Minimalist: making small adjustments to operation should require only small control effort, 3) Scriptable: measurement sequences should adhere to timelines, mirroring the behaviors of path-planning and state-extrapolation in the AVSS, 4) Self-sustaining: the radar should not require constant control input if no changes are required.

This section has presented many of the control-flow and hardware abstraction drivers behind the reduction from fully-adaptive to highly-adaptive cognitive radar for AVs. At this time, it's unlikely any *fully*-adaptive system could be realized that meets the needs and constraints of the automotive community, either in terms of C-SWAP or operational overhead. However, the highly-adaptive system proposed represents a significant leap-forward in terms of raw measurement capacity.

IV. HACR RADAR: EXAMPLE AV SCENARIOS

In this last section, we put the entire picture together and walk through cognitive radar interrogation of two notional self-driving scenarios. The scenes concocted have a large number of factors happening simultaneously, as it helps to depict the range of measurement adaptation that's possible. In terms of beamforming and waveform design, the $\bar{\bar{H}}$ presented are notional, as a quantitative presentation is beyond the scope of this introductory paper; which instead seeks to show qualitatively how dynamic KA-tasking allows for a greater measurement extent, and how such tasking might operate.

To set context, it's worth giving basic constitutive relations governing any FMCW MIMO radar (for brevity, we assume basic proficiency with nomenclature, but refer to [23].)

$$
\begin{aligned}
\Delta R &\propto \frac{1}{BW} = \frac{1}{\alpha t_{PRI}} \\
R_{inst} &\propto \frac{1}{\alpha} \\
R_{MUR} &\propto t_{PRI} \\
R_{SCNR} &\propto N_p \Delta R \\
\Delta V &\propto t_{CPI} \\
V_{MUV} &\propto \frac{1}{t_{PRI}} \\
f_{FPS} &\propto \frac{1}{t_{CPI}} = \frac{1}{N_p t_{PRI}} \\
xC_{MIMO} &\propto \frac{\log N_{tx}}{\log N_p}
\end{aligned}
$$

Competing constraints are immediately obvious. Increasing chirp rate α helps range resolution (smaller ΔR) while hurting maximum detection range R_{inst}. Increasing t_{PRI} improves unambiguous range but hurts unambiguous velocity V_{MUV}. Increasing t_{CPI}, either through N_p or t_{PRI}, improves Doppler resolution ΔV but decreases frame-rate f_{FPS}. All of the trade-offs are tied to the receiver(s) bandwidth(s) and slow & fast waveform timelines. These relations are slightly modified for other modulation-modulation/waveforms schemes, such as PMCW or Pulse, but the fundamental challenge of resource constraints is unchanged. It is the job of the KA cognitive radar scheduler to decide how to deploy these limited resources in pursuit of optimal-utility measurement, at each moment in time and across the entire scene of interest.

In the first scenario presented in FIG.7, a passenger AV is depicted in an urban environment. We imagine notional cognitive tasking below:

- In gray: A fast-update background scan covers a wide FOV (possibly with a set of measurements). This covers the maximum azimuth extent the radar is capable of, and is continuously leveled/aligned in elevation with the road. In order to achieve high frame rate f_{FPS}, Doppler resolution is sacrificed. Long-range instrumentation is traded for high range-resolution. No KA/prior-information is considered to be used in this, but the short t_{CPI} means that this measurement can be repeated often without significantly impacting radar timelines. In this way, the gray scan serves as a KA input, identifying objects for additional/refined measurement. Data from different/sequential measurements should be associated/fused within the AVSS perception stack.
- In red: A long-range measurement, directed down the road. Instrumented range is emphasized, at the expense of range resolution. Only modest Doppler resolution is required, though a high unambiguous Doppler is desired in order to quickly assess closing speeds. HD maps are used to align the measurement with the road.
- In blue: A measurement designed to probe under bridges and lamp/sign posts for road-clearance. The scan FOV is expanded in elevation to cover both the road surface as well as the bridge/stoplight. Good range and Doppler resolution improve signal-to-clutter-ratio, and help with separation. HD maps provide prior-knowledge of the location of such objects, and accumulated historical/statistical data on the radar return signature can be used to strengthen analysis discrimination.
- In green: A classification & tracking measurement, focusing on pedestrians. Short range is acceptable, and this can be traded for improved range resolution. Long coherence times are used to provide fine micro-Doppler information, and slow-time coded MIMO modes are avoided to provide the maximum interference-free dynamic range. Extended FOV is not required for classification, and a single narrow look-direction may suffice. Once confidently classified, tracking can be performed with a modified measurement that gives up Doppler resolution in favor of a faster update rate (or frees-up timeline resource for other measurements.)
- In yellow: A measurement designed specifically to watch vehicles making unprotected left/U-turns. HD maps combined with historical/statistical traffic pattern data provide KA context that this is a critical area, and the gray short-range scan identifies geostationary object. This drives the tasking of a measurement with as-necessary range (dynamic decision) and ultra-high range resolution. Because the left turn maneuver will be nearly tangential to the ownship trajectory, geo-translated Doppler resolution may be unreliable as an indicator of motion, and instead maximum update rate on high range resolution is used. Combined with a FOV that eclipses the car body, time analysis of these clustered measurements provide a good indicator if/when the vehicle begins it's turning maneuver.

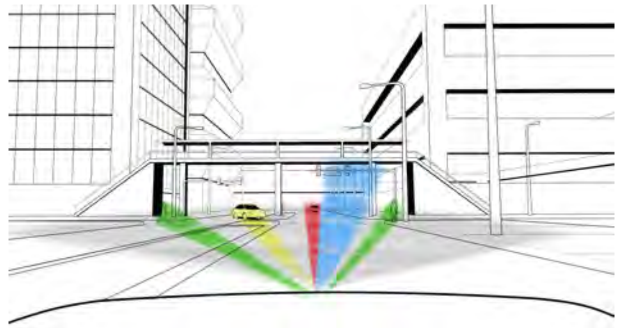

FIG. 7. Cognitive measurement through knowledge-aided beamsteering. A notional urban scene in which the AV interleaves long range measurement of, micro-Doppler measurement of, and fast-scan across a wide FOV.

In the second scenario presented in FIG.8, a transport AV is depicted in a highway environment. We imagine notional cognitive tasking as follow:

- In gray: A fast-update background scan (set of measurements), similar to the one utilized in FIG.7, but with slightly longer range due to the speeds/distances relevant to the highway scenario. In addition, the hills require that the fan-beam is continually adjusted in elevation to maintain coverage. Alternately, the scan could be broadened in elevation to fully overlap with the HD-maps adjusted road projection.
- In red: A pair of long-range measurements, directed down the road. As above, long instrumented range is emphasized. HD maps are again used to align the measurement with the road, this time in both azimuth and elevation. In this scenario the AV is intending to exit via the off-ramp, and so as the off-ramp is approached the AVSS path-planning layer begins tasking a second-long range measurement (our "anticipation" anthropomorphism). As the AV performs its exit, alignment of the off-ramp measurement is smoothly & continuously maintained. And at some point, the first down-road measurement can be stopped, freeing up resources.
- In blue: A measurement designed specifically to probe height for clearance. If under-sign lane obstruction is not a worry (or has been otherwise addressed), the FOR need only straddle the sign. This can be made a single-snapshot measurement, and once height is measured/confirmed it can be

added to repositories (local, and uploaded to fleet-wide/global) for future use.

- In yellow: A measurement monitoring an adjacent-lane car. Instrumented range is not required, but very high range and Doppler resolution are desired, along with low latency. Depending on the radar hardware, special considerations might also need to be made due to ultra-short range (e.g. provisioning for filtering or blind-ranges). In order to minimize latency, it may even be desirable to have fast-time data-vectors output directly to the AVSS rather than wait for slow-time Doppler FFT, which allows the AVSS to use of multiple banks of convolutional or rolling-STFT processing.

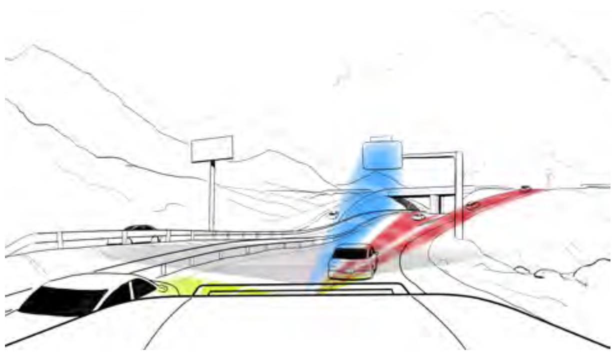

FIG. 8. Cognitive measurement through knowledge-aided beamsteering. A notional highway scene in which the AV interleaves long-range measurement of car-tracking with fine range and Doppler accuracy, elevation height finding.

V. SUMMARY

The primary goal of all AV systems is safety. A dramatic decrease in accidents is central to the expected economic impact that is driving investment, and until AVs are able to safely and independently operate in diverse conditions - environmental, geographic, and situational - they will remain an open-ended research field for academics and a cost-center for businesses. Achieving that safety will require significant improvements in many different areas, including but not limited to, Machine Learning perception, decision algorithms, hardware & software certification, vehicle maintenance & logistics, and policy. Amongst this litany of challenges, the performance of sensors is just one small piece of the puzzle; but it's a cornerstone piece. Sensors are the connection to the real-world, and troves of sensor data are the foundation of the ML approaches underwriting this age of autonomy. If sensors aren't up to the task, everything else is bound to fail; and its nearly tautological that "you cannot react to what you did not measure". Even when it's not the longest pole, higher quality and quantity of information flowing into the system from sensors can only make the situation better (though it could also point to the need for increased compute capability to consume that data).

The challenge for sensors - all sensors, not just radar - is that the total information relevant to driving exceeds the acquisition capacity of anything humankind has built. This is in large part due to the dynamic nature of the problem: everything is in motion and sensing requirements change from moment to moment. For example: it's possible to build excellent low-light cameras, and it's possible to build excellent day-light cameras, but covering the dynamic range of both simultaneously is exceedingly difficult. The obvious (and existing) solution is to extend total dynamic range by adding adaptability; e.g. an iris or adjustable filter. This dynamic range problem is not unique to man-made hardware, and biology has evolved a similar set of solutions: the true strength of the human eye is its adaptive dynamic range and taskability [24, 25].

As we seek to follow the lead of biology, cognitive measurement provides an overarching philosophy and mathematical framework for the path. As with many new approaches, it will take some exploration find the "right-fits" for early applications, and the hardware and compute required for fully-adaptive approaches may be beyond the reach of commercial systems for yet some time. The HACR architecture presented in this paper is an initial effort to reduce and simplify cognitive enough to bring it within reach of AVs operating at level 4 and 5; those with the greatest need and highest value proposition. However, the radar hardware is only one-half of the HACR system, and ultimately the degree of performance improvement realized in-field will depend on the breadth and depth of knowledge mined, and the extent to which the AV software layers actively use this knowledge to refine measurement. When sensor hardware and AVSS work together closely, exploiting a fast feedback and control loop, the combined system is capable of dramatically expanding the scene information which is made accessible to perception.

[1] Joseph R Guerci, *Cognitive Radar: The Knowledge-Aided Fully Adaptive Approach* (Artech House, 2010).

[2] S Haykin, "Cognitive radar: a way of the future," Signal Processing Magaine, IEEE **23**, 30–40 (2006).

[3] J. Guerci and E. Baranoski, "Knowledge-aided adaptive radar at darpa: an overview," Signal Processing Magaine, IEEE **23**, 41–50 (2006).

[4] Simon Haykin, *Cognitive Dynamic Systems: Perception-Action Cycle, Radar, and Radio* (Cambridge University Press, 2012).

[5] Maria S Greco, Fulvio Gini, Pietro Stinco, and Kristine Bell, "Cognitive Radars: A Reality?" (2018),

arXiv:1803.01000.

[6] Martin Drahansky, M T Paridah, Amin Moradbak, A Z Mohamed, Taiwo Abdulwahab, Folahan Owolabi, Mustapha Asniza, Abdul Khalid, and H P Shawkataly, "Sense Smart, not hard: A layered cognitive radar architecture," IntechOpen **i**, 13 (2016).

[7] J R Guerci, "Cognitive radar: The next radar wave?" Microwave Journal **January 10** (2011).

[8] B D Steinberg, *Principles of Aperture of Array System Design* (Wiley, 1976).

[9] "4 game-changing underlying technologies for advanced radar," https://www.ni.com/en-us/innovations/white-papers/18/4-game-changing-underlying-technologies-for-advanced-radar.html (Published: August 27, 2019).

[10] "AN/FPS-108 "Cobra Dane" ," https://www.radartutorial.eu/19.kartei/01.oth/karte003.en.html (Accessed: Dec 2019).

[11] Randy L Haupt, "Interleaved Thinned Linear Arrays," IEEE Transactions on Antennas and Propagation **53**, 2858–2864 (2005).

[12] Robert J Mailloux, *Phased Array Antenna Handbook* (Artech House, 2018).

[13] V S Gangwar, A K Singh, Hemant Patidar, and S P Singh, "Optimistic Design of Thinned Planar Antenna Array for Radar Operating Scenarios," 2016 International Conference on Microelectronics, Computing and Communications (MicroCom) **1**, 1–4 (2016).

[14] M G Bray, D H Werner, D W Boeringer, and D W Machuga, "Optimization of Thinned Aperiodic Linear Phased Arrays Using Genetic Algorithms to Reduce Grating Lobes During Scanning," IEEE Transactions on Antennas and Propagation **50**, 1732–1742 (2002).

[15] Sun Shunqiao, Petropulu Athina, and Poor H Vincent, "MIMO Radar for ADAS and Autonomous Driving : Advantages MIMO Radar for ADAS and Autonomous Driving : Advantages and Challenges," IEEE Signal Processing Magazine (2019).

[16] Vito Giannini, Manju Hegde, and Curtis Davis, "Digital Code Modulation MIMO Radar Improves Automotive Safety," Microwave Journal (2019).

[17] Samuel M. Sherman, *Monopulse Principles and Techniques* (Artech House, 1984).

[18] A. Paulraj, B. Ottersten, R. Roy, A. Swindlehurst, G. Xu, and T. Kailath, "Subspace methods for directions-of-arrival estimation," Handbook of Statistics **10**, 693–739 (1993).

[19] BOSCH, "Fourth generation long-range radar sensor (lrr4)," (2019).

[20] Michael Meyer and Georg Kuschk, "Automotive Radar Dataset for Deep Learning Based 3D Object Detection," 2019 16th European Radar Conference (EuRAD) , 129–132 (2019).

[21] J Yoshida, "Autonomous vehicles driving new data interfaces," EE Times (2018).

[22] Todd Litman, "Autonomous vehicle implementation predictions," (2019), Victoria Transport Policy Institute.

[23] M Jankiraman, *Design of Multi-Frequency CW Radars* (Scitech, 2007).

[24] "XKCD : Visual Field," https://xkcd.com/1080/.

[25] "Notes on the Resolution of the Human Eye," https://clarkvision.com/imagedetail/eye-resolution.html.

ECHODYNE

The radar platform company.

12112 115th Ave NE
Kirkland, WA
98034

echodyne.com