

Mosaic Warfare and Human-Machine Symbiosis

Timothy P. Grayson,^a Samuele Lilliu^b

As the director of the Strategic Technology Office at the Defense Advanced Research Projects Agency (or DARPA), Timothy leads the office in development of breakthrough technologies to enable war fighters to field, operate, and adapt distributed, joint, multi-domain combat capabilities at continuous speed. He is also founder and president of Fortitude Mission Research LLC and spent several years as a senior intelligence officer with the CIA. Here he illustrates the concept of Mosaic Warfare, in which individual warfighting platforms, just like ceramic tiles in a mosaic, are placed together to make a larger picture. This philosophy can be applied to tackle a variety of human challenges including natural disasters, disruption of supply chains, climate change, pandemics, etc. He also discusses why super AI won't represent an existential threat in the foreseeable future, but rather an opportunity for an effective division of labour between humans and machines (or human-machine symbiosis). See video at https://youtu.be/_5MkXD_m6Qc.

Introduction

Samuele Lilliu (SL). Dr Grayson thank you very much for doing this.

Timothy Grayson (TG). It's a pleasure to be with you.

SL. I'd like to start with some introductions. What's DARPA, and what's your role at DARPA?

TG. Sure, absolutely. Well, thank you and thanks again for the opportunity to talk with you here today.

First of all, let's start with what DARPA is. DARPA stands for Defense Advanced Research Projects Agency. We are considered sort of the lead R&D arm of the US Department of Defense. We were created in response, way back in the late 1950s, by the Eisenhower administration, in response to the United States waking up to Sputnik.^{1,2} President Eisenhower in 1958, said "We don't ever want to be technically surprised". DARPA was actually [called] ARPA at the time; the "D" got added later, to distinguish it from other ARPA kinds of organizations around the government. It was created to make sure that we always at least keep abreast of, if not are leading in, technology. We like to say "The best way to avoid surprises is to create surprise". So we like to stay out there at the cutting edge of R&D. We can talk more as we go through the discussion, we've got a lot of different ways of how we do innovation.

I run what's called the Strategic Technology Office (STO) within DARPA. There are six technology offices [with] various different levels of research and types of technology maturity, from very basic research to more applied research. My office is at that end of the more applied types of research areas along with another office, the Tactical Technology Office (TTO) that mostly looks at weapons and



Figure 2 | Dr. Samuele Lilliu, Bullaki Ltd, director.

platforms, so the more physical material kinds of things. We work with what are called mission systems. So a lot of communications, sensors, things of that nature, but look at them again from a more mature research and a very systems type of view.

SL. For the YouTube watchers and the robots enthusiasts, DARPA collaborated with Boston Dynamics in the development of the famous BigDog and LittleDog robots.^{3,4} As you mentioned, DARPA was initially called ARPA, and in 1969, it came up with ARPANET, which was the first wide-area packet-switching network with distributed control and one of the first networks to implement the TCP/IP protocol suite.^{5,6} So basically, you guys invented the Internet. What are the top greatest DARPA inventions that changed the world?

TG. I certainly think the Internet is one of the things that gets raised there near the top. Another one is stealth technology. Most of the original prototypes of stealth technology came out of ARPA at the time. There may have been a myriad of other things from the M-16 rifle during [the] Vietnam [War] to self-driving cars in the DARPA Grand Challenge.⁷

What's interesting when you look across that array of different big breakthroughs, it highlights a couple of things about the Agency that are that are fairly unique. First of all, we're often contrarians, and stealth is a great example. The Air Force at the time was all about making planes go faster, supersonics, and just go faster. DARPA out there pushing prototypes for stealth was really trying to open up not just new technology but a new way of thinking about what the mission was and have conducted that mission. Oh, maybe you don't have to go so fast if it's really hard for a radar to see you.



Figure 1 | Dr. Timothy Grayson, DARPA's STO Office Director.

^a DARPA, USA, ^b Bullaki Ltd, UK

So that's one thing we do. We like to say "We don't do requirements." We do like to solve problems. We're focused, really mission-focused, on solving problems. But we don't wait for someone to tell us what to do. Otherwise we would have been sitting there in the 1970s, saying, "Let's figure out how to make an airplane go faster". Instead, it's like "Yeah, we understand what your fundamental problems are, Air Force, let's see if we can think of a way that technology can solve those problems entirely differently".

The other interesting thing about where our breakthroughs have been the Internet that you mentioned is certainly one of them. But talking about the Grand Challenge and the self-driving cars, a lot of times what we do, because we're out there being contrarians, we don't necessarily see the one for one immediate impact of what we create.

That original ARPANET, as you pointed out, was late 60s, it really didn't start getting converted to widespread and even used within the research community until the 80s. I had my own you UUNet account and some other things based upon early TCP/IP as a grad student, but it wasn't widely used until World Wide Web and the 90s. Then all of a sudden this boom of commercialization, never a thought when the ARPANET was created.

I think we've seen a similar thing with self-driving cars. I was one of the judges during that original Grand Challenge. No one there at that first Grand Challenge would have ever thought that there would be a whole industry of self-driving cars, but yet, when you look at a lot of the winners, the people who completed the grand challenge there, they're all now the teams at the forefront in the commercial and academic world, really advancing what will likely be a global commercial market.

SL. Yeah. And the grand challenge was the one... the long drive, right?

TG. That's right. The whole question was "Could a vehicle drive by itself on very rugged terrain across the desert?" It was fascinating. I was a judge for the first one and the farthest vehicle made it about seven miles. People kind of looked at it and chuckled a little bit and said "Wow, what are you crazy DARPA people doing?" But DARPA didn't give up and did a second Grand Challenge just a year later, without really the government providing the upfront funding. This was done as a challenge. People were building their own teams and raising their own money for it. I think it was five teams that finished this entire race. [There were] tremendous advances on incredibly complex terrain in the course of just one year.

Innovation at DARPA

SL. What's a DARPA-hard problem? Why do you need an organization like DARPA to tackle it?

TG. As I was pointing out, the first part of a so-called DARPA-hard problem, it's something that lends itself to this contrarian, alternative view. You know, if there's a nice clear technology roadmap "Here's where the research community is and here's the next logical step", we don't tend to get involved in those. We look at things first of all in very different ways "Is there a different way that technology can help solve this particular problem?"

But then the other part of it is [that] we're known as a risk-taking agency, but I like to characterize how we do things fairly uniquely as smart risk-taking.

I put it as two extremes. One is "I don't want any risk at all, I want you to sort of prove to me all the technology is going to work before I go try something and then I'll do little incremental baby steps". There's the other extreme, which I call "hope as a strategy". You know, someone comes forward with an innovative sounding idea and a pretty cartoon PowerPoint chart and say, "Wow, you're onto something. That's a really clever idea. So how are you going to do it?" and it's like, "I don't know, we'll hire smart people and they'll think about it".

We start with problems, where we understand enough about

the problem, that we know what the risks are. We have some way to rationalize that we have a chance of being successful. Then we can build programs around that, [which] are focused on retiring those biggest risks first. Then we can move on to building bigger systems into doing more exotic programs. That gets to one other thing I think is sort of magic about our model versus a lot of other R&D organizations. We do have latitude for a certain amount of curiosity-driven work and we give a lot of flexibility and latitude to our program managers. It really is bottoms-up based upon those program managers. But we're very problem-centric. We like to say we're mission-focused. So pretty much everything we do starts with "What problem are you trying to solve?" Then from there we can explore what are the different technology opportunities and have an extremely wide aperture.

I think it's that problem-focus that also creates a lot of ability for DARPA to move things quickly, especially when you mix it with the risk-acceptant culture.

SL. You did a PhD in physics, where you worked on quantum optics before it became quantum information. I guess people are now familiar with quantum computing. The possibility to build quantum computers can open new opportunities (e.g. cryptography or solving quantum systems). You then moved to DARPA where you worked as a scientist. I was wondering if you could tell us a bit about this journey and how that influenced your approach towards problems.

TG. Yeah, absolutely. Great, great questions. I like to joke around that I did quantum computing before it was cool. I was in a lab with a very prominent pioneer in the field of quantum optics, Professor Leonard Mandel, [who] unfortunately passed away a number of years ago.⁸ He was one of the pioneers of quantum optics and really laid the foundation. We were a lab group that was doing experimental realizations of quantum optics. It was interesting, it was a fascinating time for a PhD and I loved the research work. At the time, most of what we were doing was, frankly, experimental demonstrations of Copenhagen theory and various aspects of fundamental quantum mechanics.⁹⁻¹⁸ It was exciting stuff, it was really fun research. At the same time, there was sort of this practical side, nagging at me, that that said, you know, "Everyone already believes this, and accepts it, we just haven't demonstrated these particular principles".

It was interesting that I was writing my dissertation the year that Peter Shor published his factoring algorithm,¹⁹ which, I would argue, sort of kicked off the whole quantum computing, quantum information area. But I had already made the decision that I was going to go do something more practical than quantum optics stuff. I went to work as a postdoc, for the Air Force using a lot of tools of what we had done in the lab. My dissertation was doing quantum optics based upon that two photon entanglement and to implement that we were doing a lot of work with nonlinear optics and spontaneous down conversion. When I started working for the Air Force lab, at Wright-Patterson (Dayton, Ohio), they were interested in laser sources based upon nonlinear optics. So we got into the business of developing different types of laser and other optical sources and doing research in nonlinear optics.²⁰⁻²⁵

But I think that's where again this this sort of practical side kind of tugged at me, because it was really interesting research going on there, the device and the system level. I kept saying, "What's this good for? How are we ultimately going to be using it?" and try to look at the full system level problem "If we could do this radical new laser source, how would it help us build a sensor?" and "Oh, what would I need to go along with it in terms of a new type of camera? Is there software and different types of algorithmic processing that might make it a more practical capability?" A lot of these LIDAR types of, you know, I'll call it beeps and squiggles, weren't intuitive to a human looking at a picture. So that led me to sort of this system level kind of thinking, I would call it. I guess to some basic researchers that sounds like boring or a little bit of a sellout. But I found it fascinating to think about how can we take really fundamental

science, but apply it in this problem-centric sort of way. And that really excited me.

I had the opportunity in the mid-90s, to move to the DC area and start to work for DARPA, first as a support contractor, to do a lot of the technical analysis to help with execution of programs. Then a couple years later, they yanked me over to the government side as a program manager. That was a big jump also because up until that point, I was doing very hands-on research. I had my own lab, my own lab group, some graduate students I was advising, publishing papers, all those good things that I'm sure for your viewers is their normal life. As a person at DARPA, I like to joke somewhat tongue in cheek, we don't actually do any real work. We just make PowerPoint and shovel money. There are no DARPA labs, despite what, you know, is in the Tom Clancy's video games and some of the movies. All of the actual research we do is extramural. We have our success based upon the teams of contractors, university people and some of the other government labs that actually do the real hands-on research. But at the same time, we demand that the people who come to DARPA be highly technical. They need to be leaders in their field, they need to have done hands-on research. So that they're not bureaucrats, they're not paper pushers. They drive the vision. They're helping to create the vision back to that smart risk-taking that I said. They know enough themselves that they know what the technical risks are, they know how to structure their programs around those risks. Then even though the work's being done extramurally by this group of contractors and university performers, they can provide the due diligence. Again, in a common sense way, in a lot of parts of the government that don't have the same level of technical expertise, their program managers have to resort to checklists and requirements and things like that. Our men and women are so experienced and acknowledged in their fields, that they intuitively know what the issues are, and can dig into them and ask the hard questions. And then based upon the response, then they're empowered to make decisions and pivot quickly, as their research evolves.

SL. Yeah, that's why I was going to ask you, how important is that leaders know and understand what's happening in the labs, because sometimes you have managers, people coming from business environments or managerial careers that have no idea about the technical sides, they don't have any [technical/scientific] background. I mean, the biggest example, probably is Elon Musk, right? He knows every single thing... probably... I'd imagine that's correct... and that's how he runs successfully his multiple companies. So how important is this aspect for you?

TG. It's very important, but it takes a certain breed of a researcher. You have to have been technical [when] coming in. Interestingly enough, we don't have any firm requirements or credentials. While the majority of people at DARPA have PhDs in a technical field, that's not a hard mandate. We do look for people who are technically accomplished, who have actually done their own research. The model I like to say, and even this is not hard and fast, we like to say that you've got to be a mile deep in some technical area, just so you have that, that experience-based, and you know what it's like to do research. But you also have to be inch deep, at least, across a very wide range. That's I think, one thing that differentiates a DARPA program manager from a lot of other very accomplished researchers. You could have someone who is one of the world leading researchers for their academic areas, and they work in it their entire career but you ask that person to step outside that lane, and they get very uncomfortable. DARPA Program Managers (PMs) have to also be very fast studies.

I think that may be the biggest characteristic, and I've never really thought about this, but it's academic curiosity. It's someone who is already proven that they can go deep, that they've got the technical chops, but then it's also augmented by that technical curiosity. Oh, a new challenge comes up, if someone presents them a new idea, I can be a quick study and I'm not going to be the expert who's going to go

toe to toe doing the research. But like you said with Elon Musk, he's not building SpaceX rockets or he's not building and designing batteries himself. But he's a quick enough study that he can ask the right questions and make informed decisions. That's kind of the model that we look at for our program managers.

SL. What kind of organizational structure do you have there? Is it flat or is it tall? Is the management command-oriented or is it flat?

TG. It's very flat. To my point about, we don't have DARPA labs, it pretty much begins and ends with the program managers. I mentioned we've got these six technical offices. They're populated by program managers, I don't know the exact number off the top my head, but somewhere around 100 or so program managers. Everything begins and ends with them. They generate the ideas, they execute their program activities, overseeing them. But, again, the actual research work is conducted extramurally. As a result, you know, we don't have a structure below them. Above them, it's basically folks at my level at the office management level and then the agency director, and deputy.

So within execution of an actual research portfolio, the PMs have more or less total autonomy and freedom on how to execute the programs. Now, if any of them listen to this podcast, some of them will probably grimace when I say this, but we at the office level have very little control or oversight of what they do or even what when they start the programs. Their ideas generate them. We do provide at the office level, what's called an Office Strategy and there's a similar strategy at the Agency level. But those are just general guidelines to lay down, these are the types of problems we're interested in. Then it's ultimately up to them to generate the ideas and then execute those ideas on their own.

SL. How do you protect intellectual property there to ensure that the things don't get stolen by someone... whatever agency, foreign agents, etc. It always happens in companies, it happens everywhere. What are the best strategies to protect intellectual property?

TG. I'll say there's, there's lowercase intellectual property and uppercase Intellectual Property. The thing you're referring to are our actual secrets and things of that nature. I can't go into a lot of that, but we are a government agency, we are a part of the Department of Defense (DoD). A good chunk of our work is classified and it's protected through all of the actual security and classification measures there.

On the capital, uppercase Intellectual Property, what's interesting about the way DARPA functions, we sort of stand with a foot in both camps. We do a lot of classified work that never sees the light of day but we understand that a lot of the technology is out there, in the commercial and in the academic world. If we do everything lurking in the shadows we can't engage with those communities, we risk just sort of being insular.

We do our best to, while still protecting all the rules and constraints of security, try to engage with those communities. The fact that I can be here talking to your podcast, is an example of how we try to be transparent and open. That does create a different kind of IP challenge, because we want to be able to engage with startups, we want to be able to engage with commercial companies that have some legal IP. So it's interesting that one of the things that also is nice about DARPA is we've got, even though we're a government agency, we have a lot of flexibility that other government agencies don't have. We've been able to get a lot of special authorities. A good example is on the contracting, because that's where a lot of the IP constraints pop up. In a traditional US government research contract, the default is what's called Government Purpose rights. That basically says "Hey, company or university, we're paying for your research". As a result, we want to be able to get access to it. We don't have to pay for it twice, once to fund your research and then once they have to buy a license back. But that being said, we understand we want to partner with people that have done a lot on their own and have put a lot of their own investment or a VC's investment or whatever. And

we respect that. So we will occasionally enter into special contracting relationships, where there's a policy within the US government called Other Transaction for Prototyping. DARPA makes a lot of use of that as an example, so that we could engage in more of a traditional type of almost business-to-business type of contract with people as opposed to a strict government contract. So lots of different ways that we try to engage and protect people's IP.

SL. What's the percentage of DARPA projects that pay off? How many develop unexpected spin-offs?

TG. That's a really tough question. Because like I was saying before, most of our technologies don't go immediately into use. We're usually working things that by their nature are contrarian and are thinking about a problem differently. I don't know the exact numbers, I would say the percentage that go directly into, say, a big military production program or something like that are actually really pretty small. I don't know, I'll make up a number, but probably somewhere around 10%, maybe even less, I don't know. But that's almost by design. Because if we were doing things that were so well aligned with the production programs, we're probably not out there taking enough risks, and we're probably not being contrarian enough. So I would say the majority of our efforts do transition, but transition indirectly. And I would say there are three big ways they transition. Probably the majority fall into this category, where someone is going to pick them up to do more research. It could be another government lab or it could be a company that chooses to do it on their own. A lot of our capability does come back and get used through these indirect paths.

We were talking before about ARPANET. ARPANET in 1968, or whenever it was, certainly didn't go anywhere right away. Twenty years later it started being used academically, another 10 years after that it changed the world. A lot of our technology, especially some of the more fundamental research goes into commercialization, where it might spin out, mature, and then shows up in all kinds of different products that come back and benefit the government and the military, but also the rest of the world. MEMS, micro-electromechanical system was something that a lot of the early research was done by DARPA and it didn't go directly into DoD products. The component level systems got mature, and then they start showing up in all kinds of products from ejection seats and military aircraft to your cell phone accelerometer. It's interesting, this is outside the field of my office, but a lot of the big current push for vaccines for the COVID-19 pandemic, DARPA did not directly fund any of that vaccine development, but some DARPA research about a decade ago, led to the original research into m-RNA-based vaccines. The fact that they were able to develop those vaccines so quickly, it wasn't just about governments throwing a lot of money at it, it was that there was this new technological foundation that enabled them to look at vaccine development in a different way. So we transition things in a lot of strange, different ways and I think a majority do have an impact in some regard, but just not in a way that people like to think in terms of that direct one for one transition.

SL. I'd like to change a little bit topic. I was wondering if you think there is any link between conflicts and the emergence of new ideas. So do you think extreme competition can lead to more innovation? I mean, we've seen plenty of innovation during the Second World War, even during the Cold War.

TG. I think the short answer is yes. I think the thing that is interesting to think about in the 21st century is what this competition means. I, surely to goodness, and this is the whole reason that sounds like an oxymoron, but the reason the Department of Defense exists, is to create peace and stability. We don't want to see a war. We certainly don't want to go create a war for purposes of creating innovation. Independent of whether there's war, there's always competition. I don't care if it occasionally is some military competition bumping up against each other or, unfortunately, when we see regional conflicts pop up because of natural instabilities there. There's

always some degree of global competition.

I think one of the things that really has changed the innovation landscape, and this is something that we think a lot about within DARPA, is [that] competition largely is even driven in the general economy and in commercial competition. You could think of that as a form of warfare. Look at the disruption that has happened over the last decade or two. All of these household name global industrial corporations have tumbled. That's a form of economic warfare, if you like. I was just reading an article this morning, over my breakfast on the predictions for 2021 and as the pandemic shakes itself out. That's not intentional warfare. That's a naturally occurring disruption. Nevertheless, pandemic is a form of disruption. One of the terms used in the startup world is creative destruction. In any kind of disruption it creates misfortune, it creates discomfort, but it opens the door for new opportunities that emerge.

I think there are two things that drive that. One, there's the old saying, "Necessity is the mother of invention". It goes back a little bit to what I mentioned, with DARPA's model about being mission centered. If you've got a really clear, tangible problem that motivates people and gets them focused. It's not technology or investment looking for a problem, [the problem] right there in front of you. I think that's one reason that conflict does accelerate innovation, that demand signal, that focus. The other is that it removes a lot of the barriers. Within government we love to complain about all of the bureaucracy and process and procedure. When you're faced with a serious conflict or any kind of disaster, all of a sudden, a lot of those various processes and checks and such become a lot less important. If people can go try something and realize, "Oh, it wasn't the end of the world, when we tried this thing a different way". Then after the conflict is over, people say, "Well, why are we doing it that way before?" Why can't we just keep doing it the way that worked during this conflict? So I think it's both that demand signal, but then also reducing barriers that competition and any kind of disruption does lead to innovation.

Model of the World

SL. How would you describe the world of warfare? When we try to describe a system, we have different parameters, we need to talk about domains, dimensions, rulers, and protractors or sensors, players, laws. So what are all these aspects? How do you accurately model this world? What do you need to take into account to make some good model of the world of warfare?

TG. So let me let me pull two threads here. The first one is sort of, I'll say, a bounding framework. There's something that's a sort of a driving trend right now within the US DoD. It's got a couple of different names, but for purposes of this discussion we'll use one of the names that goes under Joint All Domain Operations (JADO) or Joint All Domain Command and Control (JADC2). So we'll just call it JADO.

The key thing there is All Domain. So let's talk a little bit about what a domain is. Historically, you got an army that fights on the land, you got an air force that fights in the air, and you got a Navy that fights at sea. We talked about those physical domains. There's been a lot of press over the past year. We just passed the first anniversary of the US Space Force.²⁶ So that's an acknowledgement that space is a domain. You know, people talk a lot about cyber or the electromagnetic spectrum. Those are all domains. I've even heard of a domain referenced before when we start talking about things like information, I'll call it the cognitive domain. You know, people refer to hybrid warfare or gray zone warfare, hearts and minds, things of that nature. All of those are different domains. One of the big trends right now is a realization that you really put yourself at a disadvantage if you look at only one domain, or even if you look at multiple domains, but each individually, because the reality is that it is all very interdependent and codependent.

The second big factor that ends up popping up, and they're both

directly connected, and the punchline here, I want to come back to the question of complexity, and dimensionality and all this, the other one is speed.

SL. Yeah, time.

TG. Yep, Time. One of the terms that we've been starting to use around DARPA, I'll give credit to our Director for inspiring this term, Victoria Coleman, we're calling it Time Compression. It's "How can we make time speed up?" essentially. In simplistic terms, "How can we do things faster?"

There is this traditional determinism, particularly within DoD that likes to say "Okay, I want to study things and try to forecast what the future is going to be and in the process of doing that forecasting, come up with every possible contingency". Then I'm going to go build one heck of a powerful system, that either has a high enough performance level or is adaptable enough, that is going to address every single one of those contingencies. The reality is [that] we are now in such a complex multidimensional world, that all posit that this forecast model doesn't carry.

So instead, we have to be in a mode of rapid responsiveness. How could I acknowledge, not just assume, how can I acknowledge that I can't predict the future with any degree of accuracy, and then there are going to be contingencies that occur, that I just haven't forecasted? How can I be in a rapidly responsive mode? That's time compression. So a lot of what we're looking at right now is, both of those big themes pulled together, how can we be all domain and recognize all the interdependencies of this very high dimensionality space? How at the same time, can we do it with incredible speed? This is the notion of time compression.

I think there's an interesting aside. I love this story because it shines a light on this issue of time compression and it's about toilet paper. This was inspired by an article I read in Fortune magazine back early in the pandemic.²⁷ They were talking about the shock to the supply chain for consumer goods, whenever we had the run on stores and the shortage of toilet paper. They said [that] if you look in the business world, particularly in manufacturing, they're also a very forecast-centric type of world, they do a great job of data analytics. A company that makes the things will forecast out almost, I don't know how many decimal places, what they think their sales are going to be. It's all about managing that inventory, its efficiency. One of the things they said in this article is that one of the reasons the shelves ended up empty of toilet paper, is they were already running at something like ~93% capacity, just so they were making sure they had no inefficiencies. That's great from an efficiency and effectiveness standpoint, but it leaves no resilience. That's a very brittle system. All of a sudden, there's this contingency disruption of a pandemic and people doing a run on stores. It doesn't have the latitude to respond to that. You think, okay, I'll take a legacy forecast approach and says, well, I need to open up my error bars, I need to account for [the full range of possible demands]. Well, that's not practical because that would say "Okay, I need to have three times the quantity of milling machines". Well, if the typical toilet paper manufacturer went and bought three milling machines for everyone they have the day, they would go out of business. They can't afford that level over provisioning.

So the answer is, what are things I can do that can be rapidly adapted? How can I measure disruption? What are the knobs I can turn to get to good enough? It's not going to be optimal. But how can I get to good enough in a rapidly responsive manner? That's what I see is the big trend going forward. That's certainly where my organization has been focused.

SL. I think the problem with toilet paper was that toilet paper takes space. When people buy, let's say, five or six packs, the next customer will see an empty shelf. So they will think they're running out of toilet paper. So you could think about compressing it even more, so that it looks small, and you can put more and more so that the shelf doesn't get empty.

TG. That's it, you hit the nail on the head, in this article they talk about that. And that's exactly what one of these soft knobs is. They couldn't go whip up a new milling machine in a day. But what they could do is reprogram their production line, so that they could start packaging, you know, packages, smaller package and...

SL. Vacuum packed toilet paper.

TG. Not exactly vacuum packed but doing their packaging in smaller bundles with fewer roles per package. Yeah, it's the same idea. There were other things too, like, some of the shortages were caused by disruptions of the supply chain. So can I have what amounts to a vendor radar for new sources of pulp, and a new process for vetting and validating those providers in a faster manner?

So it's a great example of taking a very system architecture level thinking of these complex problems, as opposed to the obvious solution, which is build more milling machines. It may be the obvious solution but it is not very effective on the long run.

SL. I was reading about AlphaGo and all these AI systems, where they try to build super expert systems that can fight or play against humans. But when you model a game, for example, Chess, the sort of game tree complexity is 10^{120} , or if it's Go is 10^{700} . But when I think about the war game that must be like 1000s of orders of magnitude more complex than Chess and Go. So what are the approaches to achieve an acceptable level of modeling? Something that works, but that is not too complex. Because in theory, if you want to model things, you could start from the wavefunction and the Schrödinger equation. That's the most ridiculous thing that someone can do. So what's, what's the sort of modeling that you guys do?

TG. Yeah, no, that's a great, great question and something I personally think a lot about. I'll even toss it out to [a challenge to] your viewers, if anyone has any ideas on how to do this kind of framework analysis in a quantitative way, I'd be interested in hearing your ideas and, who knows, maybe you get a project out of it.

The fundamental way we're looking at it is managing, and in some ways this is borrowed straight out of network theory, it's looking at how do we manage complexity and dimensionality by breaking things apart into scale. If we think about some of the dimensions you mentioned, those are consistent with some of the [dimensionality] we see at the decision support level [just] inside an individual platform or a payload. One of the things I think we'll probably end up talking a little bit more about is our AlphaDogfight, you know, AI flying an aircraft. We have [another program with] an AI controlling a sensor payload. Those kinds of decision processes are on the order of what you're saying, maybe, you know, 10^{200} or so. Now imagine that platform or that payload is part of one mission unit and that, in turn is part of full squadrons or forces and then I'm talking All Domain and other types of things. Each of those is taking about that same dimensionality and now growing that, geometrically.

People who think that they can take that kind of endeavor and, as one of my bosses likes to say, sprinkle some AI pixie dust on it and just assume that some algorithmic approach is going to discover a way to manage that level of dimensionality, is just really not practical.

Instead, we say, "How can we break up those decision layers?" I'll think of it in terms of decisions that have to be made, "How can we partition those into a manageable degree of complexity, or a manageable dimensionality and then still create optionality by being able to abstract those interfaces in those boundaries?"

It's really about partitioning and abstraction.

If you look about the original inspiration for this, that this was before my time as office director, but a number of years ago, maybe 2015-2016, DARPA sponsored a conference called "Wait, What?"²⁸ And one of the speakers there was a chaired fellow or chaired professor from UC Berkeley [Prof. Alberto Luigi Sangiovanni-Vincentelli]. He was one of the creators of the original design tools for semiconductors and he really presented this notion better than anyone I've ever seen. It's intrinsic, you mentioned knowing the wave

function, well, you know, at one level, if you want to do a really good semiconductor, it would be great to know the wave function of every [individual transistor or gate, in an IC, but as you said, that would be totally impractical with the number of gates in today's devices.]

SL. Do you need that? Probably not.

TG. Yeah, so we've got Moore's Law,²⁹ because of this notion of being able to manage scale by partitioning and abstraction. Then once you've abstracted things, [you can build back up something big and complex by] being able to do composition, with those abstract elements, and that's [the same way] we're approaching this all domain warfare challenge [and addressing the dimensionality].

We're living right now in a very challenging, but also interesting intersection between technology and culture. Because to live with that model of abstraction and composition drives a certain amount of acceptance of uncertainty. You know, I don't know what's happening underneath, underneath the abstracted boundary of that next module but, trust me, it's going to satisfy some function. [This is the uncertainty a modern IC designer needs to live with.]

Well, usually, if you're talking to people in the military, they don't like an answer to just "Trust me". You know, they want determinism. They want a certain number of decimal places of certitude in it. But I argue that that is committing the statistics 101 fallacy of mistaking precision for accuracy. We're living in such a complex, dynamic world, that we have to be able to, back to the toilet paper, live with disruption, live with uncertainty. The only way you do that is by, you know, taking a more stochastic kind of model toward things.

Strategy

SL. So let's say that you have this model of the world of warfare. What can you use them for? How are these models deployed? Is it for training or actually fighting real wars?

TG. It's all across the board. Within my office, our guiding portfolio for these kinds of things we're calling Mosaic Warfare. And you can see our little logo here behind me. The whole metaphor of Mosaic Warfare is if you look, as a contrast, a jigsaw puzzle, that's a highly engineered architecture. I'm creating an overarching effect with a composition of existing pieces, but every single one of those pieces is very carefully engineered as to how it's going to fit into that broader picture. They're difficult to put together. They are very brittle and fragile once you've created them and they're not flexible at all. So the mosaic analogy is to say, I want to, again, a stochastic kind of model, I'm going to have a bag of tiles, so to speak, that are some arbitrary, perhaps even opportunistic distribution. I might have some control over the statistical distribution of those tiles but I'm not going to specify exactly what any one tile should be. But I'm going to have confidence that I can piece those tiles together in some way that can still produce an overall picture.

By the way it's not just completely arbitrarily, I've still got some kind of substrate and some kind of adhesive or mortar. I've still got a framework, I'm working with it. But I'm living with that, that uncertainty, this stochasticity of a bag of tiles, but now it gives me a credible ability to flex and adapt. In principle I can create that adaptation with much less difficulty than designing a new jigsaw puzzle.

By the way it's much more resilient to disruption and uncertainty because I can lose a tile. I can find one similar and throw back in.

So when we talk about modeling, you know, we have to do modeling across the board with those kind of things. Some of our modeling goes into just "How do we plan the Mosaic?" I can go back to the toilet paper example and say, "Okay, yeah, I know, I've got a problem, I can't keep the shelves stocked". Can I use modeling tools to figure out where within my process my supply chain is breaking, and be able to experiment with what are options to be able to create a more resilient supply chain?

We've got a program, for example, called PROTEUS [Prototype Resilient Operations Testbed for Expeditionary Urban Scenarios],³⁰ that on the surface looks almost like a video game. What it's actually

doing though is allowing people at things like military universities to explore new types of force structure, to create the notion of hybrid military units that could be designed much more finely, and specifically to a given mission need. So that's an example of using this kind of modeling in a planning kind of stage.

We've got a fair amount of modeling that goes on in how we design the networks we need, although that's a little bit less about modeling than it is about new networking constructs and communications, virtualization and interoperability kind of things. I'm happy to talk a little bit more about that later.

The third place where the modeling really becomes important, and this is where some of it actually gets used in operations, gets back to "How can we simplify the problems that human beings have to deal with in this kind of highly networked, very fluid, dynamic kind of architecture?". If you're responsible for being one of those tiles yourself, how do you know what your role is supposed to be in this Mosaic? That's where a lot of our work in AI has come in. But in the spirit of abstraction and composition, it's all up and down these various different degrees or levels of I'll call it a Decision-making Stack. So we've got some technology that functions at a very high level. Think of that as the Mosaic artist who's saying "I've got a certain function I want to have happen in the battlespace, and I want to use a new collection of tiles to go conduct that function". What are the best set of tiles to use at this moment in time? So it becomes automated modeling to make those high level decisions.

To your point about dimensionality and complexity, that decision maker using that model doesn't have to know anything about how to actually use that tile, or what other calculations might have to go in. It just sort of votes. The tile itself is working a lower level type of modeling. Again, managing dimensionality. Where it comes in, it says, "Okay, someone asked me to serve this role. I don't know why, I don't have understanding of the whole battlespace or commander's intent". Again, that level of complexity has been stripped away. I just know as a tile, I've been asked to do something. Can I do it? What calculations and modeling do I need? So that might be for example, for an aircraft. Just something as simple as route planning. Or we've got one program that's been looking at collections of air platforms. How do you deconflict the airspace in this very dynamic manner? Again, none of those functions needs to know that high level awareness.

Then - another layer of abstraction - the things that are doing things like that airspace planning, those modeling tools don't have to know how to actually actuate the platform. There's a different level of AI that can worry about actuating platforms, actuating payloads. This is where my challenge problem comes in. If someone has a great model who's an expert out there on network theory, I'd be really interesting to say, "How can we actually define a scaling law, you know, based upon a certain number of, you know, interfaces, divisions and boundaries and this notion of abstraction and composition?"

SL. So you came up with this idea of Mosaic Warfare as opposed to Monolithic Architectures, right? I mean, there is a huge rupture with the past. You gave a talk in 2018, where you spoke about the problem of Dominance and that we... the USA... needs to think about a different approach.³¹ You also spoke about Mosaic Warfare as a sort of planned strategy that can be laid into three phases.

TG. Let me talk about the Dominance one first, yeah, in some ways, I've really touched on it already if we just change the terms a little bit. Dominance, the way certainly the US military has thought about it, is back to that very deterministic, forecast-based model. I want to do lots of studies to try to predict what the future mission and the future threat is going to be and I'm going to go ahead and, by doing that right, just design something that is big enough and bad enough and high enough performance that it can accommodate all of those possible contingencies. That's the Dominance mental model.

What we're talking about instead of dominance is this notion of

what we really need to be focused on is “How do we achieve our objectives?” whatever those objectives might be. Frankly I think this mindset can apply again, outside the military. Back to “How can we make sure that people can buy toilet paper?”. It really is, how do I get away from this “I’ve got to forecast everything, I’ve got to provision for all possible contingencies?”.

Instead: time compression. How can I be rapidly responsive and adaptive, regardless of whatever the opportunity or the disruption is?

That’s really what we’re getting at with this notion of Mosaic. We’ve got to move away from the dominance mindset, because otherwise, you’re stuck in this classic cat and mouse problem that’s counter, counter, counter, counter. You know, as soon as you think you’ve built the biggest and baddest system, someone is going to now focus on how to either build the bigger and badder system or just some countermeasure that directly negates that capability.

Another analogy, I like metaphors: it’s like trying to pop a balloon with one finger. It’s sitting there bouncing on your finger, and you can’t get any leverage against it because however you push, it’s just gonna squeeze out somewhere else. That’s the Anti-dominance kind of approach. Dominance would say, I just want to make that as hard as a block of granite. Yeah, but then someone comes along with chisel and hammer and your granite is no good. So I’d rather be the balloon than the granite.

SL. How is this Mosaic philosophy being implemented in phases?

TG. I’ve recently been referring to what we’re calling three waves of Mosaic. The reason for laying it out this way, is that as I mentioned before, we’re right at the intersection between technology, but also culture and organization and process. Again, I don’t think what we’re doing is unique to the military or DoD. I’ve read a lot of articles about [that]. In fact, this is going on right now in the commercial world. There’s the Gartner hype curve³² that you’re probably familiar with or your viewers may have seen. A new technology comes along, and there’s first this huge excitement over its adoption. Then people look at how it’s being used, and all of a sudden, they’re not seeing quite the outcome that all of the hype seem to justify. Then you get the opposite response, they call the Trough of Despair. And then, ultimately people who stick with it slowly claw their way out of that Trough of Despair and you find out what really is this useful for and then you do see a little bit less hype, but adoption and the real impact.

A lot of what’s going on in the AI world is similar right now, you know, it’s like, okay, AI is going to change the future. Companies that have tried to adopt it, are like, okay, we’re spending a lot of money buying whatever this AI stuff is, where’s my return? They’re not seeing it. A reason for it is the exact reason why we’ve got three waves of Mosaic. For really disruptive technology you can get some marginal improvement if you just sprinkle in the technology, but you can’t get the orders of magnitude kinds of improvements if you aren’t simultaneously challenging your processes and your structures to go along with it. A business has to truly change its workflows and how it thinks about executing its business to get best advantage from automation. DoD not only is no different, DoD has an even bigger challenge, because it’s so locked in rigorously to doctrine, structure and tradition in a very disciplined manner.

What I’m seeing in these three waves, Wave One is really, to a large degree outside what DARPA is doing right now, although I’d like to say no different than the Internet had taken a couple decades to catch on. We’ve been working system of systems architectures at least going back to the late 90s, when I was a program manager. We’ve been pushing this for a long time. As recently as about five years ago, or so, my office was still working on system of systems. People would look at us like we were ogres with two heads, “What is this system of systems thing? Give me my next fighter aircraft”. The fact that there’s this whole Joint All Domain push within the Department is incredibly exciting. You know, I also use the term Mon-

olith Busting sometimes for Mosaic. System of systems is busting up monolithic platforms, where I’ve got to have the sensor, the weapon and the decider, all programmatically vertically-integrated and technology-integrated into one platform. So that’s good. Wave One is where the big military is right now, in starting to implement and experiment with systems of systems. You got to start somewhere. It’s exciting to see this happening. Put some markers down, try some pilot projects, provide some tangible, concrete examples of how you can get advantage by disaggregating capabilities, distributed capabilities. That’s Wave One and there’s goodness there. The challenge is, we risk replacing Monolithic Platforms, with Monolithic Architectures, in other words, jigsaw puzzles, as opposed to mosaics, and a lot of the wave one, Join All Domain activities are jigsaw puzzles. They want to study things. They want to figure out what’s the mission going to be, what exactly is the set of stuff that I want to go wire together to conduct that mission, and then how am I going to manually go integrate all of those. Again, I’m not knocking that you got to start somewhere, but if we stop at that point and say, “Okay, we’re just going to replace the whole DoD with these tailored architectures”, I would argue what we’ve done is replace Vertical Stovepipes, the platform-centric monoliths with Horizontal Stovepipes, the architecture-centered monoliths. That frightens me if we get to that point because, as anyone who’s tried to do system architecting knows, and you brought it up in your question about complexity, the more things we put together [the more] that complexity grows geometrically. Building systems of systems architectures is hard. If we try to build the whole DoD with system systems architectures, it has a real risk of just collapsing under its own complexity.

Wave Two is where we’re really focused right now, what we want to be able to do is to enable a military operator out in the field to say “I’ve got a bunch of stuff out there that by itself maybe has an existing function that it was designed to go do”. It in and of itself is a useful standalone capability. But now I’ve got a problem facing me and it could be a new mission, it could be a new adversary, it could be a new environmental problem or whatever. How can I take what I’ve got and take this architectural mindset and build a bespoke solution to this problem facing me today? How can I build an architecture that addresses that problem, basically, today, with whatever I’ve got? So this notion of it’s not everything wired together in a big mash, it’s looking for a federated approach. In this limited set of capabilities on a focused problem, there are these subsets of things that have to work together. Maybe they weren’t designed to work together, but I got a way to make them more interoperable on the spot. That’s what we see as Wave Two. I like to describe it as letting the warfighter do system architecting, without realizing they’re actually doing a technical act; they just think they’re doing mission planning. But there happens to be all these technical wiring diagrams going on behind them.

The challenge there, for those of your viewers who are old enough to remember the days of Windows 95. By the time we got to the 1990s, with a personal computer, it was pretty cool. Because you have a computer at home, you could configure it the way you want it; if you wanted a new printer you didn’t have to go buy a whole new computer just to get a new printer. So there were things that were tailorable architectures designed to whatever your need was. However, you remember those days adding a printer to your computer was not for the faint of heart. You had to tear open the box, pop in a card, and put in a floppy disk and do manual installation of a bunch of specialized software to make it work right. So that’s where we are with Wave Two. The technology that’s coming out of DARPA is, you know, think of as almost as Dell. We want to create the environment for the warfighter, where they can, again, focus on need. Being able to piece together an architecture is just like in 1990, we would have put together our home computer.

As an aside, one of the organizational things that I’ve been out there pounding the pavement about is we need a new function with-

in the military that for lack of a better term, I'm calling a Combat Support Geek Squad. For someone who didn't have the technical fortitude to rip open their computer to install their printer, they could call in Geek Squad, who is the support organization that can be out there 24/7 supporting operators. It's not like going back to one of the vendors or program office, but yet they're more technically skilled than your typical flight line mechanic. So that's Wave Two. That's the next logical step because we want to be able to demonstrate that we truly can flex to need.

Wave Three... and I'm not as much of an evangelist as I am about these things... I'm not willing to push too hard on Wave Three yet. But if we can prove that, we convince the warfighter that they really can operate in this very fluid, more stochastic, building a capability to need, it is going to change how we think about new systems. So think about how we're going to populate that palette with new tiles. The way we're thinking about Wave Two is the tiles, the platforms, the weapons, the sensors, are all going to be more or less the current things we've got today. When we buy the replacement for those in the future, if we're wildly successful with Mosaic, imagine I go buy a new sensor and I say, you know what, I'm not going to tell you what the performance requirements for this answer should be, I'm not even going to tell you quite what all the functions have to be, I'm not going to tell you how it's going to be used. But it's great. Trust me. We're trying to enable that kind of a future, where you could let innovators come in and say, "I think I've got this great new radar, I know it's got to be good for something," and you throw it into that palette of tiles.

One of the interesting things you asked about modeling earlier, you know, one of the things we're even exploring with our sort of game-theoretic models as a way to score [the value of] tiles. In this future vision we don't want to dictate how a tile is going to be used, how a particular capability is going to be used. DoD has a finite amount of money, it's not truly an open market. We can't just let the market decide the way a commercial environment would. So how do we figure out how we want to spend taxpayer money on these capabilities? We've contemplated using modeling tools against game theory to say, okay, someone brings me a new tile, it can be used in the course of this very fluid Mosaic, does it move the needle? Does this tile actually improve things? Does it make no difference, hopefully doesn't make things worse, but as a mechanism for deciding where to put future investment. That's way down the road. We're doing some more fundamental research, trying to understand how we would do that kind of modeling and evaluation in the future. But right now, we just want to prove that we can be good geeks and you know, help the warfighter build their Windows 95 computer.

SL. Now talking about technical challenges, I can think of what we do in the labs, for example, when we process multi-domain data. Sometimes you have a sample and then you get a microscope image, you get an X-ray fluorescence image, you get an X-ray diffraction pattern of that sample... well if I raster scan it, you get an image... and then you try to align the images, you try to register the images, and then you get all this data from different sensors, and then you need to transfer it to a computer. Then basically, you need to reduce the data, you need to generate knowledge, you need to synthesize that knowledge. Then you need to translate it in a way. You need to test hypotheses so that you can make informed decisions. All this stuff takes time and computing power. So how do you approach the data transfer between elements of the Mosaic and avoid time consuming sneakernet?

TG. Yeah, great question. So what you just described, I kind of real top level talked about these three thrust areas. We have planning, execution, the one in the middle we call interoperability, which is exactly this problem. The way we were tackling that is, to some degree, it's not exactly one to one. But to some degree, it parallels the ISO-OSI network stack model. We start at that physical layer. We've done research in the past, and we'll probably continue to do some

[more] research on how do we make sure that given nodes, given tiles if you like, or just different data sources, how can I make sure I've got a physical pathway that can move between two points? It's interesting, because I personally think one of the biggest problems is how can we make that process as, again, adaptable and flexible as possible, but at the same time, not making it overly complex?

People have, you know, the best, again, sorry, another metaphor, but the best analogy I've seen is comparing a Swiss Army knife to a Dremel tool. I can make a Swiss Army knife that does everything, I can make a radio gateway that speaks every part of the spectrum, and every possible protocol. The reality though is that that node is probably going to be insanely complex and insanely expensive and large and power-consuming to those platforms. The more Dremel tool model is maybe I want a node that only speaks two different waveforms, but I've got a bunch of those. Maybe I can build that smaller and cheaper. But I could take a bunch of those and scatter it, you know, throughout the environment. Again, statistically, if we are willing to live with a degree of uncertainty, statistically I have enough density of those sort of bilateral PHY-layer nodes, that I've got high confidence of being able to connect with each other.

Those are some of the things we're kind of looking at the physical layer right now. And there's the other interesting little side problems that gets to the software engineering side of the world. People certainly today know how to do software-defined radios. Gnu Radio is out there as an open source capability as an example. US DoD has invested a lot of money in software-defined radios in the past.

Back to the abstraction thing, there's a real challenge in the balance between efficiency and flexibility. I can have an expert build a computing platform, some radio card with a bunch of FPGAs and such. I can bring in an expert on that card, and do a really efficient job of writing the software and I can get a tremendous amount of options in signal processing power and the types of waveforms I can choose using that methodology. The problem is, if I ever want to change the card, or ever want to change the software, it's back to the drawing board. Oh, by the way, it's got to be that expert. That's another thing we're looking at is how do we create layers of abstraction, where we can break apart the process of doing the math and creating signal processing from the people that are building the hardware, so that both of them can evolve and be adapted, whether it's because of a different platform or different mission, different versioning or new technical opportunities, whatever the case might be, that the hardware and the software can evolve independently of each other. All of that is one big problem here.

As we move up the stack, so to speak, we're doing a lot of work right now in software-defined networking. One of the big problems that DoD has, and again, but the military is not alone in this, is that a lot of these issues would be less technical challenges if we were a completely Greenfield kind of environment, I just want to throw away everything I've got and start from scratch. The reality is I don't know how many trillions of dollars of sunk capital equipment out there, there's a lot of legacy equipment. Some of it was fielded much before I was born. We can't just throw all of that away. So a lot of what we're doing is "How do we create layers of virtualization, on top of this very heterogeneous mix of legacy equipment?" And so that's been another big area of research for us. We're slowly adding more capabilities. A first step might be, I'm stuck with whatever's provisioned out there in terms of the actual network hardware, but how can I manage the data more smartly, in a virtualized network. Ultimately, as we have new options, like maybe a new software defined radio or something like that, I might have more knobs that I can turn. And I can do more to actually control and flex that networking environment, in addition to managing the flow of data.

As I move further up that stack, you touched on a really important point and talking about different sensor types and the different types of analysis you might want to do. Just because I can pass data between two points, doesn't mean it's useful to that endpoint, the

software can understand it, the people there can understand it, but the machines can actually talk to each other.

We've really been focusing a lot on that problem of data level interoperability. This is a little bit of heresy in the standards world. A lot of people are worried about this problem on a regulatory basis. But we are pretty against the notion of enforced global standards and the reason for that, there's nothing fundamentally bad about it, but the reason for that is that these require a lot of work to decide on. Two, they usually come with compromises either in functionality or performance. Three, after you've gone through all the pain of creating them, they're usually out of date [by the time you] start using it. Instead, we've actually built a technology that does auto generation of translators to move between different data types. What's really exciting about that, is it's not just a syntactic conversion of data formats, those examples you gave were, things have to be measured and characterized in different ways. That gets down into a very semantic description. What's cool about this, the software tool that we've got is some tortured acronym that I can't remember what it stands for, but the acronym spells very suitably STITCHES [System-of-systems Technology Integration Tool Chain for Heterogeneous Electronic Systems]. It will, if you lay down an architecture and say, "here's a whole set of different systems that have to talk to each other, and they've got all these different message types". They lay that out in the graph and hit compile, and it will generate a set of executable code that are all the message translators and get its message translation at an actual semantic level. Not just doing syntactically.

SL. Yeah. And the other thing I wanted to ask you is about the distributed computing and processing... parallel computing... instead of using a central base, where you gather all the data there, transfer everything in one single place, and then process everything there and distribute it to the operators or whoever, instead adopting a sort of distributed approach where processing happens in many nodes. Is that some sort of more acceptable approach?

TG. Yeah, we are really big on trying to push a very, you know, distributed kind of kick. Yeah, we don't like the idea of everything having to come back to a central spot. It creates brittleness, it also loads up your networks. The interesting thing, is all of that really ties back to that theme I raised earlier about abstraction and composition. You know, if I can be able to abstract individual elements into these partitioned chunks, if you like, collection of tiles, I can make very simple types of high level decisions that then are left to those individual elements to implement. It's how a lot of human military operators in the West function, has this notion of mission-command. We're trying to build the same thing into software and into architectures.

SL. How do you avoid overflow of information with operators?

TG. Well, once again, we're not centralizing information, we're partitioning. I start to be a broken record. But I come back to the abstraction and composition. Let's take, for example, in one of these distributed war game kinds of environments, you might have a small army unit with a missile battery and, let's say, they want to use their missiles as part of some targeting solution with the Air Force flying a radar sensor. To your point about information overload, if those poor guys in the army with the missile battery had to have every piece of information in the battlespace, and from every possible sensor, and they were distilling all of that themselves, how would they pick out how they should use their missiles? In our very distributed model, decisions are happening in different layers and in different pieces and in different locations. I might have a node someplace without knowing the details of that radar or the missile battery, saying "Hey, those two things would go well together". And then they give out tasking and then the guys in that army missile battery, say, "I don't really know the overall battlefield context, but I was told I need to launch a missile at these coordinates. I know how to do that function". It's the same kind of thing that gets to human management of information. It's the same problem you started out from a

control perspective and decision perspective. How do we manage dimensionality? They're all part of the same problem.

SL. How do you put together Mosaic Warfare and logistics? The pieces of the Mosaic need to be taken to the place where the war is being fought or whatever, if you don't talk about war, you talk about something else it's always the same thing. There is logistics involved.

TG. The way we're looking at that, and part of this work is organization again, I started out talking about DARPA, and we're not the office that does platforms. So rather than saying, "Hey, what's a new way to physically move things", we again look at most of the problem sets through the lens of information. We actually have a program underway right now, that is looking at awareness of logistics. It's almost like Uber for military logistics.

In fact, one of the things we're also really big on, this is another variant of Mosaic, but it's Mosaic of how we're actually doing the software engineering. We're big into a lot of the current modern software practices of micro services. Our approach to logistics is to build a micro services architecture, where we're building a lot of heavy computational types of algorithms for doing things like: how do I go find logistics related information? How do I do correlations? What are models for doing forecasting of various different kinds?

On the top of that, we're building a very app store like environment, where there are individual functions that can reach into that sea of heavy computational information and create certain products for different users' problems in that logistics chain. So for example, the organization that has to figure out, "should I go order more parts?" That's a very different problem than the organization that says, "How am I going to ship them off to a little island someplace?" And that's also a very different problem than the user who wants to know, "When's my stuff showing up?"

The analogy to Uber is [that] there's one big cloud analytics environment that Uber builds. But the app for the consumer versus the app for the driver versus the app back at the place that's doing billing are all very different lenses into that sea of data and computation. That's kind of how we're looking at the logistics problem.

SL. Recently, I saw this article about an attempt from here in Britain to build something similar to DARPA, they would call it BARPA (British Advanced Research Projects Agency). The UK, in terms of human resources is huge. The UK is second in terms of Nobel Prizes in the world, is also fourth in the Nature Index for high quality scientific publications output worldwide. Human capital is not a problem, but maybe budget is a problem. Do you think the DARPA model can be replicated in other countries?

TG. I absolutely do. The challenge is, it's not going to be one for one, it's not exactly cookie cutter. But it's interesting to me. It's less about budget and more about what do you want to get out of it. Are you willing to accept a couple of key attributes? So for example, this mission-centered focus that I mentioned, is a really big deal. Irrespective of how much top line budget you have, you need to have the notion of being able to be focused. I've seen a number of research organizations out there that do very good quality science, but it doesn't have the same impact as a DARPA because it might be doled out in small, incoherent allotments or just working on the next logical step in evolutionary tech roadmap. Being focused on solving problems gives us a critical mass, regardless of an overall agency budget. To put it into perspective, one of our program managers might manage a few 10s of millions of dollars and in other research places I've seen they are managing maybe \$100,000. It's difficult to get a critical mass, I would say, with something like that. The new UK version, regardless of how much money you have, it's more about getting the balance between the number of projects and the amount of money right.

Then there's the execution models. There's a model for how you hire people, and how you actually execute the work. One of the things I think is really important to note is that DARPA has zero full-time permanent technical employees, even to include myself,

the agency director, and all the program managers. All of us have expiration dates. The typical tenure at DARPA is about four years. What we ended up doing is, because we're very flat organizations, we talked about earlier, we recruit people who are already accomplished researchers who are mid-career, who can demonstrate to us they can think in this DARPA way. As my former boss, when I was a program manager liked to say, you know, "We bring them in, we squeeze all their good ideas out of them, then we toss them back out on the street". It's not like a great recruiting pitch, but it is sort of how it works. But the thing you get in exchange, first of all, I'd never heard of anyone leaving DARPA and being out on the street. But more importantly, it's a place where there is enough resource and enough freedom that someone can come in with this great grand idea they're trying to pursue, and they can't get anyone else to listen. Here's an environment where this is their chance now to go pursue their dreams. I think that's the most important thing that any new organization trying to recreate the DARPA model they've got to get, that human [aspect].

It's not about the pool of Nobel Prize winners or other accomplished PhDs; it's finding that right cultural mindset, and then being willing to cycle people through. No matter how good someone is, they get stale or they get locked into a certain area of research. When I worked for the Air Force [there were] some brilliant, very energetic, very enthused dedicated researchers, but their research specialty, say, was LIDAR. It didn't matter what the problem set was, it was like, "Hey, I can come up with a LIDAR solution for you". Well, what if tomorrow, the problem is a pandemic, you're not going to go solve a pandemic with LIDAR.

So DARPA is constantly retooling. And that gives an opportunity to focus on what is the problem at hand, and that's probably more than anything else the secret sauce.

AI and Human-Machine Symbiosis

SL. The last thing I wanted to talk about is the sort of rogue AI. I've got a few questions about that. A super AI would be some sort of AI able to learn and achieve certain goals faster than humans. And there are many discussions on cyber and physical existential threats. Probably, you know, people tend to get worried about AI when it's in the context of warfare, and people would associate this AI to what we've seen in Sci-Fi movies like Terminator... Skynet, and all these things. Some philosophers said that the issue with the super AI might be that when a human assigns a certain goal, for example, "get rid of email spam worldwide", then maybe the AI system finds its own inconvenient intermediate goals in order to achieve that primary goal, which is like "kill everyone". Do you think this represents an actual risk, or maybe the risk is that there would be more TED Talks, more books, and more movies on the subject?

TG. Well, I tend I tend to lean more toward your latter there. You know, never say never. But I'm not worried about it, certainly within our lifetimes, or for that matter, my grandchildren's lifetimes. As AI advances, we should watch AI advances, keep an eye on what you're talking about. I don't think we're in any danger of that happening right now. You know, first of all, I'll say, AI does a great job, and can do some of the things you talked about, you know, process much, much faster than a human can, in so called closed-world problems. But for more open-world kind of problems, things that take more context, more intuition, more inference, it's still extremely, extremely early days for AI. I don't think we're going to get anywhere close to that existential kinds of risks that you're referencing in Hollywood. Until we get to that so called really competent third wave AI.

One of the reasons people point to DoD is because it is about competition and conflict and things like that is being the driver of the killer AI. Both from my own experience, as well as a lot of things I've read, there are so many checks and balances within the military. And frankly, culturally, as I was talking about before, the military is so conservative. If the killer AI were going to emerge, I don't think

it would be in the military. I've actually heard it postulated that the finance world is a better place for it to emerge.

But back to my comment about closed-world, I think what's a very realistic thing in constrained cases, and we've even seen some of it, are the unintended consequences of the AI running off, so called rogue, toward the equivalent of a local maximum and you get some unintended consequences for what it was built for. Now an AI based [stock] trading app, I don't think is likely to all of a sudden somehow mutate and become Terminator. But it can, if left unchecked, wreak havoc in financial systems.

A lot of the way that the military is looking at it even as greater and greater amount of automation is done, that the term that's used is human-on-the-loop. I hate to be a broken record, but it comes back to our Mosaic model of abstraction and composition. If we're breaking apart decision making and control into these layers, there can be some layers that the AI is responsible for, and some layers that the human is responsible for. There will be natural boundaries then on what the AI can run off and do even if it decides to go stupid.

A great example is something that got a lot of media attention a couple months ago, that was run out of my office, it was called the AlphaDogfight Trials.^{33,34} It's being conducted under the Air Combat Evolution (ACE) program. It got tons of attention. For those who are watching this, you can go on to DARPA's YouTube channel and look up AlphaDogfight, and you can see the whole thing.³⁵ What got everyone's attention is, and it relates to sort of your question, and it creates that situation "So maybe this is the first step to Terminator". The final event was the winning AI agent. It was a contest with eight different AI agents that competed against each other tournament style. Then the winning AI agent flew against a human pilot. This was a real accomplished fighter ace, active duty Air Force sitting in a simulator. Sadly for the poor pilot, he lost five-nothing. It was pretty dramatic. There were lots of things that weren't completely realistic or whatnot. But frankly, I think in the balance of things, the number of not realistic things was probably equally split in terms of who would favor. But despite how eye opening and sort of titillating it is that the AI beat the human so handily, it misses the point of the program.

The real thing we're trying to do in the ACE program, is figure out how AI and humans work together. In the subsequent follow on program, it's going to be focused on how we train and how we create a protocol to train trust in AI. The analogy my program manager likes to use is, the first time he got in a car that had adaptive cruise control, and his car's speeding down the road, and there's the sea of red lights in front of them, it took a moment of panic: "do I trust the AI to stop? Or do I stomp on the brake?" That's the real push in the ACE program, is "How to get a human fighter pilot comfortable with the plane flying itself?"

But even more fundamentally, it gets back to your question and back to this notion of human-on-the-loop. What we're really trying to do in the long run is human-machine symbiosis. How do we create a division of labor? Such things like flying a plane are incredibly dynamic, very, very difficult, require a lot of hand eye coordination, actuation coordination, but flying an aircraft, if given a specific objective, is actually a very closed-world problem. It doesn't require a lot of inference. That's exactly the kind of thing that a computer can really excel at. It can even process more data than a human can. It's not just about speed. One of the things we saw is that the machines can think in as many dimensions as you want, you can give it a six degree of freedom or a nine degree of freedom state vector for the opposing aircraft. Humans can't think that way.

SL. Just think in 4 dimensions.

TG. [Think of] the highway, someone cuts you off, you know, they don't think in barely even three dimensions. But yet at the same time, the computer does a really lousy job of the higher level strategic things. That's the kind of division of labor that we're working on. That's why I don't think that we're anywhere near a Terminator

nor are we really on a path for Terminator, but yet, where the real promise of an AI is gonna be.

SL. So then the problem there was media that didn't get the point of this exercise, right?

TG. Human machine symbiosis is boring. [They prefer to focus on] fighter pilots loose into a computer...

SL. That's actually more interesting... it depends who's thinking about it. Also in this way, if the pilot can forget about all these technical details of like piloting the airplane, can focus on other things, like maybe political considerations, or maybe tactical considerations and strategy and all these things, right?

TG. That's right! The way we describe it, it allows the fighter pilot to become a battle manager thinking higher level. And also, from a training perspective. One of the terms I heard used by a former general who ran a training range, he says, "I've got to stop spending so much time training fingers, and more time training brains".

Think about video games. I'm not a gamer... I'll say, I'll tell those young whippersnappers, those kids out there. But you watch someone who's a gamer and they can move from game console to game console and game to game and fairly intuitively pick up a new gaming system. If we had a bunch of AI at that lower level, could we make learning the use of a system as intuitive? Now, the human is more transportable. Maybe that's another layer of abstraction, I'll come back to one of my favorite themes: "Abstract away the systems from the human". And this is something that military does a horrible job in. But the commercial world, I think, is getting there. That's behind a lot of the principles of UI, UX user experience, in the design thinking, and I hope we see more of that within the military.

SL. Interfacing humans and machines. DARPA has been interested in Brain Machine Interfaces (BMIs) since the 70s and now we have Elon Musk trying to achieve this with Neuralink. So is the way humans acquire and output information a bottleneck in warfare? Or maybe this is not a problem? What do you think?

TG. Well, so, first of all, I can't comment too authoritatively about that, because most of that kind of work goes on in another office at DARPA, called the Biosystems Technology Office (BTO). I don't have a lot of insight into those specific programs. I'll come back to that human-machine symbiosis problem. Personally, I'd love to have something like [a Neuralink]; I could read faster and get through more quickly.

There have been a couple really exciting things that have gone on out of those BTO programs, where it has created more fine level control of prosthetics or for people who are quadriplegic, totally paralyzed and bedridden, and yet with one of those brain-machine interfaces, being able to control all kinds of things in the physical world or experience things that they wouldn't normally. All of that is, I think, a very exciting and a big opportunity.

But to me, it doesn't really change the fundamentals in this human-machine symbiosis. To me the bigger, or at least the more interesting, thing from my office's and my personal perspective is not how do we bring the two together more closely through one of those interfaces, but rather, how do we understand where it makes sense to split them apart. Where are the natural boundaries for division of labor? And that really gets more to what are the type of the information, what are the ways to control information, dimensionality or complexity. And that, to me, is where the big breakthroughs with AI will occur. I'm not saying there's anything wrong at all with the other research, but to me, I find it more exciting to think about this information centric problem.

SL. How do you think AI or advanced AI, super AI, will help scientific research in the future? Do you think we're going to see more symbiosis between AI and scientists, maybe for things like repetitive tasks and tedious things like you know, when you try to test different materials...

TG. I would certainly think so. Just off the top of my head, two big opportunities pop out. One is exactly what you described, you

know, the tedious tasks. I'll go back to my experience of doing quantum optics as a grad student. I don't know how many hours I spent in a pitch black room twiddling with aligning mirrors and such, with a beam I could barely see. I used to think to myself, "Wouldn't it be great if I could have robotic mirror mounts that arranged and aligned themselves on the optics table automatically." I was ready to actually go patent this at one point in time. I thought, if you had these little robotic mirror mounts they could drive themselves around an optics table, and then had a big feedback loop that they could align this interferometer on your own. Boy, would that be great? Well, I say that half-jokingly, but you could imagine all kinds of different types of self-configuring scientific apparatus. That's the equivalent to the fighter doing tactical maneuvers. Then that would free up the human researcher to think the bigger thoughts if you weren't having to spend days upon days twiddling mirrors. So I think that's one big area.

The other big area is to help develop or transfer intuition. Again, I think we're a long way away from machines actually having intuition, but I think they can help humans with their own intuition. I think they can do a certain amount of transference of experience. Do it in such a way that allows researchers to explore and challenge hypotheses. The open-world problem is the challenge of coming up with a hypothesis in the first place. And again, I think that's something that's going to be the domain of humans for a long, long time to come. But once you've got a hypothesis, you know, humans are notorious for getting, you know, locked into tunnel vision, "I've got my hypothesis, now, I'm going to build my experiments and do my data analysis", that in some cases, not all, but in some cases can actually be self-confirming. I've seen AI based tools that allow you explore other competing hypotheses that maybe are going to be wrong. In fact, they very well might be wrong. Because again, machines aren't great at intuitive leaps. [But then again, those leaps may open up a new line of thinking for the human to get out of that tunnel vision.] There's interesting research going on right now in what's called self-aware AI. So it's not fully understanding why it came up with something, but it can at least say, "Oh, I got to this particular neural net", or "I got to this particular outcome, based upon some particular model that was provided to me as input, or some particular data set that was provided as input". Tools like that will allow researchers to go say, "Oh, I'm stuck in tunnel vision, I'm actually sitting on a local maximum someplace and the AI just provided me a pointer to another alternative hypothesis, based upon my initial input". So I think exploring the hypothesis space is, and maybe part of it is just from a literature search perspective. Maybe there's a set of esoteric journal articles someplace. Only if you knew that, there might be something there that might change your hypothesis. Something to allow you to go down a different path, and I think the AI will provide a really good opportunity for those types of things.

SL. Yeah, I think what you said about providing the "why" is a very important thing, because as far as I knew usually AI systems just function as a black box, you don't know what's happening inside and if the thing gives you an insight about why it's giving you that output that's great. I also heard about what you just said about scientific papers analysis. These systems can find emerging patterns, emerging things, especially for medicine and drugs discovery. Okay, so shall we close it here? We've done almost two hours.

Conclusion

TG. Oh, wow. Okay. Yeah. Time flies.

SL. Yeah. Is there anything else you would like to add?

TG. Yeah, I don't believe so. We covered a lot of water. But I'll just come back to you know, my big themes. Maybe there is one last thought I was going to do just highlight. (My boss says I have to put a penny in the jar every time I say Mosaic.) I really am a believer in this notion of Mosaic.

In the context of those big themes we hit, a federated approach, not a common approach, being able to be adaptive and resilient to

disruptions as opposed to forecasting and everything pre-planned, things that are by exception, as opposed to so called boil the ocean and do everything top down, abstraction and composition as a way to do this in a flexible way and manage complexity.

Those are to me such fundamentally powerful themes that I've been out there on a quest for what other endeavors do they apply to. We talked about toilet paper earlier, somewhat tongue in cheek, but as a simple example. I actually am very interested in how we can apply these principles of Mosaic to things like renewable energy and climate change, efficient [distributed] manufacturing, ... global supply chains, medical distribution. So whatever the endeavor might be, anything that's a system problem or an architectural problem, I think these principles that we're exploring right now very much apply.

From the perspective of your viewers or anyone else that I'm talking to out there, I've actually been asking, what are the dual use applications? You know, can we find ways to apply this type of technology to other human endeavors outside of only the warfighting aspect? And that's a good way to close the loop back to your question about conflict driving innovation.

SL. Okay, thank you very much, and have a great holiday.

TG. Yeah, thank you. I really appreciate the opportunity to talk and you have a wonderful holiday too.

Conflicts of Interest

There are no conflicts to declare.

Acknowledgements

SL thanks DARPA's communication team for coordinating and filming this interview at DARPA.

Author Contributions

TG was interviewed by SL. TG and SL wrote this manuscript.

Keywords

This article is licensed by T.G. et al. under a Creative Commons Attribution 4.0 International License ([CC BY 4.0](https://creativecommons.org/licenses/by/4.0/)).

References

- Wikipedia. DARPA, <https://en.wikipedia.org/wiki/DARPA#:~:text=Originally%20known%20as%20the%20Advanced,1958%20by%20President%20Dwight%20D.&text=The%20name%20of%20the%20organization,to%20DARPA%20in%20March%201996.> (2020).
- C. Herzfeld. How the change agent has changed. *Nature* **451**, 403-404, (2008).
- R. Playter, M. Buehler & M. Raibert. BigDog. *Proc.SPIE* **6230**, (2006).
- M. Raibert, K. Blankespoor, G. Nelson & R. Playter. BigDog, the Rough-Terrain Quadruped Robot. *IFAC Proceedings Volumes* **41**, 10822-10825, (2008).
- S. Lukasik. Why the Arpanet Was Built. *IEEE Annals of the History of Computing* **33**, 4-21, (2011).
- Wikipedia. ARPANET, [https://en.wikipedia.org/wiki/ARPANET#:~:text=The%20ARPANET%20\(an%20acronym%20for,technical%20foundation%20of%20the%20Internet.](https://en.wikipedia.org/wiki/ARPANET#:~:text=The%20ARPANET%20(an%20acronym%20for,technical%20foundation%20of%20the%20Internet.) (2020).
- Wikipedia. DARPA Grande Challenge, https://en.wikipedia.org/wiki/DARPA_Grand_Challenge, (2020).
- U. o. Rochester. Physicist Leonard Mandel, a Founder of Quantum Optics, Dies, <https://www.rochester.edu/news/show.php?id=778>, (2001).
- P. P. Yaney, T. P. Grayson & J. W. Parish. Measurements of temporal and spatial scales in gas flowfields using a two-pulsed-laser correlation scheme. *Symposium (International) on Combustion* **23**, 1877-1883, (1991).
- X. Y. Zou, T. Grayson, L. J. Wang & L. Mandel. Can an empty de Broglie pilot wave induce coherence? *Physical Review Letters* **68**, 3667-3669, (1992).
- X. Y. Zou, T. P. Grayson & L. Mandel. Observation of quantum interference effects in the frequency domain. *Physical Review Letters* **69**, 3041-3044, (1992).
- X. Y. Zou, L. J. Wang, T. P. Grayson & L. Mandel. New technique for controlling the degree of coherence of two light beams. *Optics and Laser Technology* **24**, 289-291, (1992).
- T. P. Grayson & L. J. Wang. 400-ps time resolution with a passively quenched avalanche photodiode. *Applied Optics* **32**, 2907-2910, (1993).
- T. P. Grayson, X. Y. Zou, D. Branning, J. R. Torgerson & L. Mandel. Interference and indistinguishability governed by time delays in a low-Q cavity. *Physical Review A* **48**, 4793-4796, (1993).
- X. Y. Zou, T. Grayson, G. A. Barbosa & L. Mandel. Control of visibility in the interference of signal photons by delays imposed on the idler photons. *Physical Review A* **47**, 2293-2295, (1993).
- A. Fougères, J. W. Noh, T. P. Grayson & L. Mandel. Measurement of phase differences between two partially coherent fields. *Physical Review A* **49**, 530-534, (1994).
- T. P. Grayson & G. A. Barbosa. Spatial properties of spontaneous parametric down-conversion and their effect on induced coherence without induced emission. *Physical Review A* **49**, 2948-2961, (1994).
- T. P. Grayson, J. R. Torgerson & G. A. Barbosa. Observation of a nonlocal Pancharatnam phase shift in the process of induced coherence without induced emission. *Physical Review A* **49**, 626-628, (1994).
- P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 124-134, (1994).
- B. L. Johnson & T. P. Grayson. Moving target exploitation. *Proceedings of SPIE - The International Society for Optical Engineering* **3393**, 172-183, (1998).
- D. A. Garren, T. P. Grayson, R. O. Johnson & T. M. Strat. Theoretical analysis of a continuous tracking system. *Proceedings of SPIE - The International Society for Optical Engineering* **3709**, 184-195, (1999).
- C. Y. Chong, D. Garren & T. P. Grayson. Ground target tracking - A historical perspective. *IEEE Aerospace Conference Proceedings* **3**, 433-448, (2000).
- D. R. Kirk, T. Grayson, D. Garren & C. Y. Chong. AMSTE precision fire control tracking overview. *IEEE Aerospace Conference Proceedings* **3**, 465-472, (2000).
- R. H. Giles, H. J. Neilson, D. M. Healy, T. Grayson, R. Williams & W. E. Nixon. Acquisition and analysis of X-band moving target signature data using a 160 GHz compact range. *Proceedings of SPIE - The International Society for Optical Engineering* **4379**, 289-299, (2001).
- T. P. Grayson. Curved focal plane wide field of view telescope design. *Proceedings of SPIE - The International Society for Optical Engineering* **4849**, 269-274, (2002).
- U. S. Force. About, <https://www.spaceforce.mil/About-Us/About-Space-Force/>, (2020).
- J. Wiczcner. The case of the missing toilet paper: How the coronavirus exposed U.S. supply chain flaws, <https://fortune.com/2020/05/18/toilet-paper-sales-surge-shortage-coronavirus-pandemic-supply-chain-cpg-panic-buying/>, (2020).
- DARPA. About "Wait, What?", <https://www.darpa.mil/work-with-us/about-wait-what#:~:text=Wait%2C%20What%3F%20was%20a%20fast,idea%20further%20into%20the%20future.> (2015).
- G. E. Moore. Cramming more components onto integrated circuits. (1965).
- DARPA. Prototype Resilient Operations Testbed for Expeditionary Urban Scenarios (PROTEUS), <https://www.darpa.mil/program/prototype-resilient-operations-testbed-for-expeditionary-urban-scenarios>, (2020).
- DARPA. Mosaic Warfare and Multi-Domain Battle, <https://www.youtube.com/watch?v=33VAnIEJgk&t=81s>, (2018).
- Wikipedia. Hype Cycle, https://en.wikipedia.org/wiki/Hype_cycle, (2020).
- DARPA. AlphaDogfight Trials Go Virtual for Final Event, <https://www.darpa.mil/news-events/2020-08-07>, (2020).
- DARPA. AlphaDogfight Trials Foreshadow Future of Human-Machine Symbiosis, <https://www.darpa.mil/news-events/2020-08-26>, (2020).
- DARPA. AlphaDogfight Trials Final Event, <https://youtu.be/NzdhIA2S35w>, (2020).