Canadian Maritime Security Network | iDEAS Series

A River Runs Through it

The River-Class as a option for the US Navy

Jonathan Cassels | February 2026

This paper is an updated and expanded version of an original published by Jonathan Cassels on December 14, 2025

Executive Summary

The cancellation of the US Navy's Constellation-class frigate exposed deep structural risks in American surface-combatant procurement, driven primarily by underestimated redesign complexity and industrial constraints. While the National Security Cutter offers a fast but limited stopgap, it cannot meaningfully replace the "mini-Burke" role Constellation was meant to fill. Canada's River-class destroyer, already redesigned, integrated with US combat systems, and in full-rate production, presents a uniquely low-risk alternative. Though costly, it offers full-spectrum capability, lower lifecycle costs, and growth potential, while stabilizing fleet structure and buying time amid mounting strategic and industrial uncertainty.

The announcement in late 2025 that the US Navy's Constellation-class frigate (Constellation-class) program would be cancelled, rather than rescued, was initially met with a mix of déjà vu, relief, and frustration. Déjà vu, because once again a supposedly affordable, risk-reduced shipbuilding program had spiralled into delays and cost overruns. Relief, because the fate of the Littoral Combat Ship, which continued to consume American treasure and time long after it was recognized as a failure, was avoided. Frustration, because the United States Navy has not commissioned a new, unambiguously successful, class of surface combatant since the first Arleigh Burke-class destroyer (Burke-class) entered service in 1991.

Even in the cancellation announcement, the Navy made clear that an alternative design was wanted quickly, and that it would not wait for a "blank-sheet" design to be developed. Less than a month later, a ship in service with the US Coast Guard, the National Security Cutter (NSC), was selected as the base design for a new class of frigate – one that will be significantly smaller and less capable than Constellation-class had been envisioned to be. Despite those limitations, it is not surprising that the NSC was chosen.

The United States' recent history of surface-combatant design is not an enviable one. Multiple projects have been launched only to be abandoned during the design process, including at least two cruiser concepts. If the Defiant-class "battleship" program proceeds as recently described, they may soon be joined by DDG(X) which, until late 2025, was intended to replace the Burke-class.

In fact, so far this century the US Navy has only had three new surface combatants actually touch water. One is the Zumwalt-class destroyer, which displaces roughly twice as much as Constellation-class was planned to. The other two are the Freedom- and Independence-class Littoral Combat Ships, each less than half Constellation-class's displacement. With size serving as an imperfect but useful proxy for capability, and all three classes having suffered from major structural, engineering, or mission-system issues which led to early retirements or abbreviated production runs, none of these offer a compelling alternative to Constellation-class.

Of course, they *could* not have done so. Constellation-class was conceived specifically to address the failures of recent US Navy shipbuilding programs. This means there's no US Navy ship design which could be used as a substitute for the Constellation-class, leaving only the NSC and foreign options. But the experience with the Constellation-class itself discourages most foreign options.

As developed, the Constellation-class was intended to be a full-spectrum surface combatant, capable of contributing meaningfully across the Navy's mission set. Where the full capacity of a Burke-class was unnecessary, Constellation-class ships were meant to be a less resource-intensive option that could operate independently, freeing Burke-class ships for more demanding tasks. When independent operations were not required, Constellation-class ships could augment task groups built around Burke-class ships and other high-end assets. It was, in effect, a "mini-Burke."

From the outset, the plan was to adapt an in-service parent design to meet the US Navy's needs. In 2020, the French-Italian FREMM design was selected based on its maturity and broadly aligned capabilities. It was initially expected that only 15-percent of the design would need to be altered, which promised reduced risk, cost, and development time.

Reality proved otherwise. By the time the program was cancelled, only 15-percent of the original FREMM design remained *un*altered - a reversal of the original ratio. As redesign cascaded, costs rose and schedules slipped, ultimately leaving the program at least three years behind schedule.

There's a popular narrative that Constellation-class's massive redesign was created by choice. A self-inflicted wound by meddlesome (read: stupid) admirals endlessly tinkering with a perfectly good foreign design. While this caricature contains a kernel of truth (some project requirements were changed later than best practice would suggest) it ignores the fact that rational, and sometimes unavoidable, constraints are part of large-scale redesign and poor outcomes do not require institutional or individual caprice. Ultimately, this narrative is misleading more than helpful as it reduces complex industrial and engineering problems to a morality play.

In hindsight, the notion that only 15-percent of the FREMM design would need modification should have been treated as an extraordinary claim requiring extraordinary evidence. A 15-percent redesign was an extremely optimistic floor for even the bare minimum level of modification the design required to be built efficiently by the American industrial base. By the time combat systems were modified, even to the extent needed to use US weapon stockpiles, the 15-percent claim was

extremely unlikely. Planned modifications always went deeper than this.

Redesigns have a way of cascading, with each forcing another. Once a project crosses a certain threshold of modification, the final scope of redesign cannot be known in advance. It can only be discovered – often painfully – in execution. It should have been realized that the Constellation-class project was above that threshold from the start, but the project was never given the resources, infrastructure, or time necessary to manage that scope of work. That is why the project struggled and was, ultimately, cancelled.

Given this experience, it is unsurprising that the US Navy was reluctant to pursue another design from overseas. While the correct lesson of Constellation-class is not "do not choose foreign designs," but rather "do not underestimate redesign – especially when time is scarce", any option from overseas – whether Spain's F-100 design, Japan's Mogami-class, or South Korea's Daegu or Chungnam-classes – would require significant time for redesign. After decades of failed programs, time is one resource the US Navy does not have.

Yet, just north of the border, outside the traditional American field of view, there is a ship that doesn't require redesign – either to begin construction in American shipyards, or to be compatible with US Navy systems – that comes remarkably close to delivering what the US Navy wanted from the Constellation-class. Canada's long-delayed but now accelerating surface-combatant replacement effort has reached full-rate production on the River-class destroyer (River-class).

The River-class is a derivative of the Royal Navy's Type 26 frigate, which was originally conceived as an anti-submarine-warfare-first, general combatant, to complement the Type 45 air-defence destroyer. Canada, however, used the design's generous growth margins to pursue a more ambitious goal: to create a surface combatant with credible capabilities across the full spectrum of naval warfare.

Image credit: Royal Canadian Navy Public Affairs, June 2024.

The result reads like an aspirational statement for the Constellation-class program. The River-class is an approximately 8,000-ton warship capable of over 28 knots, with advanced anti-submarine, anti-surface, and anti-air capabilities, including ballistic-missile defence and true area air defence. Comparable in displacement to early-flight Burke-class ships, it trades some missile magazine depth for strengths in areas where the Burke-class is weaker, while requiring a significantly smaller crew – a non-trivial advantage for a navy struggling to man and sustain a growing fleet.

Its combined diesel-electric and gas propulsion system also generates substantially more electrical power than the Burke-class can

sustain, while being more fuel-efficient. Together with the smaller crew, this will translate into appreciably lower operating costs over the ship's lifetime.

The River-class's advantages over the NSC as a Constellation-class replacement are substantial. That said, there are reasons the US Navy might prefer the NSC. From the beginning of what became the Constellation-class program, there was debate over whether the best complement to the Burke-class was the full combat spectrum approach the Constellation-class eventually followed (sometimes called the "mini-Burke-class" concept) or if a much smaller and even lower cost option was best, even if it could only assume a much more limited subset of the Burke-class's missions.

History gives some insight into this conflict. In the late 1980s, a period many today would identify as a time that the navy got its fleet mix "right," the US Navy was building three surface combatants simultaneously: First was the 9,800-ton Ticonderoga-class cruiser (Ticonderoga-class), the fleet's premier air-defence ship and a capable task-force flagship. Second was the 8,000-ton early-flight Arleigh Burke-class destroyer, a full-spectrum combatant with less capacity and cost than Ticonderoga, usable independently to free Ticonderoga-class ships for more demanding missions or in task groups alongside them. (This should sound familiar). Third was the 4,600-ton Oliver Hazard Perry-class frigate (Perry-class), a far smaller and cheaper ship with limited air-defence capability, optimised for specialized missions, like anti-submarine warfare, or for cases where capability beyond basic self-defence was unnecessary.

This three-tier structure served the US Navy well through the 1990s and into the 2010s. Since then, the Perry-class has left service, replaced inadequately by the smaller, less capable, and deservedly maligned LCSs – which have now also ended their production runs. Meanwhile, most of the Ticonderogas have been retired, with the few remaining expected to leave service by 2029. Only the Burke-class remains in production. Superficially, this suggests gaps at both the high and low ends of the fleet. But the Burke-class itself has changed dramatically. In its Flight III configuration, it now displaces roughly as much as the Ticonderoga-class and has assumed its role as the US Navy's premier air-defence platform.

This implies room for two designs "below" the Burke-class in capability. One to replace the Perry-class as a smaller, specialised combatant, and one to fill the niche the early Burkes once occupied, as a less resource-intensive complement to the premier air-defence ship. The NSC is a plausible starting point for the former; the River-class makes sense for the latter.

Needing both, and needing ships quickly, it might seem intuitive to prioritise the smaller, less capable, combatant as the lower risk program. All else being equal, those should be easier to build. However, the specific practical issues of these two designs invert this logic. The challenge with the NSC returns, again, to redesign. The NSC in service with the US Coast Guard is not a frigate. At the very least, it needs better sensors and more significant armament to fill the role even of a specialized combatant. This upgrade poses a significant challenge. Even though the NSC was originally marketed as being "designed for, not with" additional weapons and sensors, and there are even concept designs of it with more significant armament, in practice these claims and designs are based on very rudimentary examinations of the available growth margins in the ship's buoyancy, power, and physical dimensions. They do not indicate that the detailed engineering work needed to discover and solve the issues inherent in redesign has happened. In the case of the NSC, it's not even clear that the original margins still exist.

The original NSC design was structurally under-engineered. Post-commissioning analysis of the first ships suggested an expected service life of just three years. Extensive reinforcement was required, consuming buoyancy margins that might otherwise have supported combat systems. Even if the margin technically still exists to add combat systems, those systems will increase its weight further, potentially driving a need for additional reinforcement and cascading into more redesign.

Electrical margins present another challenge. Modern combat systems impose demands far beyond what was envisioned when the NSC was designed nearly two decades ago. The NSC has capacity for more systems than the US Coast Guard installed, but perhaps not for all those the US Navy would wish. Installing more electrical capacity is a major undertaking.

Survivability standards pose a third challenge. The Navy can waive these requirements, as it did for the LCSs, but it is understandably reluctant to field ships that cannot absorb damage without becoming combat-ineffective. These challenges create the temptation for, and some of them may require, deeper redesign than the US Navy's current timelines permit.

To credit the US Navy, it does not appear to assume the NSC redesign will be simple or quick. The apparent plan is to bolt on whatever systems can be added to the NSC with minimal modification and to begin building ships otherwise "as is" following the US Coast Guard design. More deeply modified, and capable, ships will only be built when redesign work is complete at an unknown point in the future. This plan may get hulls into the water quickly, but, the usefulness of these modified hulls, might be questionable.

In practice, this approach will initially deliver ships with capabilities similar to the in-service LCSs. Such NSC frigates may be sufficient to retire some additional LCSs, something the US Navy could welcome, but it is unlikely to meaningfully free up Burke-class ships from their present duties. The threat environment in most places where the Burke-class are used today is simply too significant for what LCSs or a similarly equipped NSCs can handle. As an example, only Burkes, not LCSs, played a role in the US response to the *Red Sea Crisis* of 2023-25 – precisely the sort of low-intensity littoral conflict LCSs were designed for – because the threat in that setting has evolved past their capabilities. More capable NSC variants are presumably possible, but their maximum potential capability is currently unknown, and they will only exist after a risk-prone redesign of uncertain duration.

Contrast this with the River-class in the mini-Burke role. Not only does it already have combat systems beyond what an NSC is likely to ever have, the systems that form its combat backbone are ones that are sourced from America. That includes an AEGIS combat system incorporating the US Navy's Cooperative Engagement Capability, a Lockheed Martin AN/SPY-7(V)3 radar, and Raytheon's AN/SLQ-32(V)6 electronic-warfare suite. No system changes are necessary.

This is not because the River-class avoided redesign entirely. Its Type 26 parent design includes none of the aforementioned systems. The difference is that the River-class has already passed through that redesign phase. It sits on the far side of the risks that doomed the Constellation-class and threaten to delay or skuttle the NSC and, indeed, any other options.

The River-class also benefits from its British heritage. The Royal Navy's survivability standards,

shaped by its Falkland War experience, are among the world's most demanding. While metrics differ, the outcomes align closely with US Navy priorities: ships that remain afloat and fightable after damage, and crews that come home. This removes another source of redesign pressure which exists in the NSC program, and would exist for most other foreign designs.

The River-class does include some non-American subsystems, but these are either UK systems for which Canada has already resolved supply-chain issues to build using them within the deeply integrated North American industrial base, or, Canadian systems which, while they might be new to the US Navy, are products of the joint US/Canadian manufacturing base. This means the River-class does not rely on electronic widgets or raw materials which are not available to US manufacturers. This gives the River-class a fundamentally different risk profile from other non-American designs.

Ships are designed around the materials, components, and fabrication practices of the industrial ecosystem in which they are built. When a design is transferred to a different industrial ecosystem, some of those solutions cease to be available, forcing redesign.

In some cases, these changes are manageable. Substituting a slightly different steel plate thickness may require analysis but, assuming the thickness provides the needed strength and is within the ship's buoyancy allowances, little else.

In other cases, the consequences cascade. A specialized component made from an unavailable alloy may present a tremendous challenge. Importing the component, or the raw materials to make it, might be an option if America is willing to accept the higher costs and supply chain risk. But it also might not be an option at all, if the country of origin doesn't have the slack manufacturing capacity to meet the need, and no manufacturer in America has the experience needed working with the material to create a component to the necessary tolerances. This can force redesign, altering dimensions and interfering with adjacent systems, which in turn require redesign themselves. Many mechanisms can cause this and it is how the scope of redesign can grow unpredictably.

It runs counter to the 'meddling (read: stupid) Admiral' trope, which is blamed for the failure of the Constellation-class. But many of Constellation-class's redesigns began this way, with additional tinkering only happening once some level of redesign was required. Any non-American design originating outside the US industrial base must successfully navigate this process before construction can begin in American yards. And that takes time.

Canada is the exception. It's a somewhat shallow demonstration, but Canadians buy two-by-fours, not 50-by-100 millimetre lumber - even though Canada officially uses metric measurements. Why? Because doing otherwise would make it harder for Canada to trade goods to and from the United States.

This mundane example reflects a deeper reality: For all practical purposes, the Canadian and American industrial bases are one and the same. There are effectively no raw materials or industrial components available in Canada that are not also available in the United States, and vice versa. This means American industry can freely "copy Canadian homework" in ways that do not apply to designs from elsewhere because a design which can be built in Canada can skip the redesign phase when building in America too.

And that's what the River-class represents. Like the Constellation-class, it derives from a foreign parent design. It required redesign. However, it's already in full-rate production, a milestone the Constellation-class never reached, which means that the scope of the River-class's required redesigns is known, and largely complete.

This does not quite mean the US Navy could adopt the design tomorrow and begin construction in Wisconsin the next day. Even transferring a proven design between American shipyards requires some engineering work. But that work is fundamentally different – faster, cheaper, and far less risky than redesigning the ship itself.

Some have proposed bypassing redesign altogether by waiving the requirement that US Navy ships be built domestically, allowing designs like the Mogami-class to be constructed overseas. Two issues exist with this: One is that this merely moves the problem. Unless the United States is prepared to rely on foreign shipyards for maintenance and modernization throughout a ship's service life, the redesign work still has to be done in order to maintain the ships. Second, it only reduces the amount of redesign, it doesn't eliminate it. Even just changing the voltage and shape of crew-accessible electrical outlets to work with US electronics is a redesign project that can cascade and grow in scope.

While all of this means that the River-class bypasses some of the issues the Constellation-class encountered, and other foreign designs would encounter as well, none of this should allow us to ignore the River-class's weaknesses. Cost being the most obvious. Canada's Parliamentary Budget Officer estimates the acquisition of 15 River-class ships at over C$80 billion: roughly C$4.6 billion per ship. Understanding this figure requires knowledge of its accounting assumptions, for example, it includes taxes paid back to the government – something government estimates frequently omit. However, adjusting for this, and converting to US dollars, still gives a cost estimate of over US$3 billion per ship.

Yet, looking a little deeper, there are costs built into that number which wouldn't be reflected in a US program. Prior to the River-class, Canada hadn't laid down a true warship in over 30 years. This lack of experience has direct impacts on cost that won't be a factor in US shipyards, which have significant and much more recent experience. This gap in production also meant no Canadian yards could begin construction on the River-class before being largely rebuilt. Rebuilding yards is incredibly expensive and much of that cost, which also won't be duplicated in the US, was built into the River-class project's budget.

Accounting for all that, however, a River-class built in America would still likely have an acquisition cost similar to a Burke-class: around US$2.5 billion. That raises a legitimate question: why build a mini-Burke for the price of a real one?

Lifecycle cost is part of the answer. Fuel and crew dominate lifetime expenses. Acquisition cost is a relatively small fraction of the total. The River-class's propulsion system is more efficient, and its crew of roughly 200 is far smaller than the Burke-class's 300-350. Over decades, the savings from just those two factors will be significant.

But even that defence of the River-class's price tag is shallow. To begin to explore it more fully, we need to acknowledge another one of the River-class's perceived shortcomings: The fact that it has an underwhelming 24 VLS cells. This statistic is sometimes used in a misleading way when

comparing to European designs. The River-class's Mk 41 VLS cells can each carry one large missile *or* four small missiles but are frequently compared on a 1:1 basis with VLS systems that can only carry one small missile. Many ships also rely on their VLS cells to hold their point defence and anti-ship missiles, both of which the River-class has dedicated launchers for. However, compared to the Constellation-class's planned 32 Mk 41 cells, with dedicated point-defence and anti-ship launchers, it is a reduction.

Yet the Constellation-class failed precisely because it tried to do too much within too tight tolerances without having the time, resources, or infrastructure to make it work. Expecting an alternative to achieve everything the Constellation-class sought isn't realistic.

All the same, the Royal Canadian Navy has acknowledged this shortcoming in the River-class, and stated that it too would prefer more VLS cells on the platform. However, the time crunch created by the looming retirement of the Halifax-class, Canada's only in-service and rapidly aging surface combatant, has convinced the Royal Canadian Navy to proceed with building its first River-class "as-is" and invest the time spent building them to do the design work properly for a "Flight II" version with more VLS cells.

With its own time-crunch, this is an approach that the US Navy could emulate. Begin construction now on an available design which meets most of the force's perceived needs – much more so than the base NSC design – and use the cushion of time created by those ships to properly complete the redesign that's required to create an ideal variant for the future.

There is ample room within the River-class to support that evolution. While it is too early to know for sure, Canada's Flight II River-class will likely have between 30 and 36 Mk 41 VLS cells. This relatively modest increase over the 24 VLS cells in the base configuration isn't an inherent limitation of the River-class's design. The Royal Canadian Navy is extremely enamoured with the River-class's "multi-mission bay", which is a reconfigurable space directly forward of the helicopter hangar which can be used to increase the ship's capacity for helicopters or unmanned aerial vehicles, deploy unmanned underwater vehicles or large boats, or carry cargo and support humanitarian and relief efforts, among other things. Keeping this space puts a limit on how many VLS cells can be added to the design.

There's an argument that this is an excellent feature for the US Navy as well. Especially given the rapid changes in drone technology, and the impact they're having on warfare, as well as the expectation that this ship will be used for taskings the Burke-class is overkill for, like humanitarian work. However, if VLS cells are the sole priority for the future, there are concept designs (based on the Type 26 parent design, not the River-class specifically) which suggest that replacing the multi-mission bay could allow for as many as 96 Mk 41 cells (matching the Burke-class), or even 128 cells if the main gun were sacrificed.

These upgunned versions of the River-class are concepts. But their existence shows that the Type 26 has the basic space, buoyancy, and power requirements for this to work. Realizing these concepts would require extensive detailed design work, exactly the type of work this piece cautions against relying on going smoothly in the near-term. It is also possible that engineering challenges would be discovered that make the practical limit for VLS cells somewhat lower than this theoretical maximum. But the fact that such growth is plausible underscores an important point.

The River-class isn't limited to being a "mini-Burke". The flight I design the Royal Canadian Navy is building can be built as-is by the US Navy in order to fill the "mini-Burke" role, saving money compared to operating the same number of Burke-class ships, but that's not the limit of the design. It is really more accurate to think of the River-class as a Burke-class with different priorities, and with a modern systems architecture that leaves room for a broad range of configurations and whatever advancements the next 30 years of naval warfare bring.

This is important, because the Burke-class is at the absolute limit of what can be integrated into its 1980s architecture and it too needs replacement in the relatively near future. In mid-2025, the Burke-class's replacement plan appeared to be to bring in the Constellation-class to take on some of the Burke-class's lower-demand taskings, and to aggressively pursue a project for the DDG(X) destroyer, which could take on the "premier air-defence" taskings from the Flight III Burke-class in the future.

After the cancellation of Constellation-class, and the selection of the NSC, a significant gulf was left in the US Navy's future fleet mix. The NSC is a 4,600-ton ship which, as a frigate, may grow to around 5,500 tons. The Constellation-class was going to be a 7,200-ton ship. The Burke-class is currently being built as a 9,800 ton-ship. The DDG(X) is set to be a 13,000-ton ship. Without Constellation-class, and with Burke-class nearing the end of its economically viable production life, the US Navy may soon be building an (at most) 5,500 ton frigate and a 13,000 ton destroyer. Nothing in between.

The DDG(X) program itself, as a blank-sheet design, comes with many risks. But, to complicate this situation further, the White House and US Navy recently announced a plan to build a 35,000 ton "battleship" known as the Defiant-class (Defiant). While Congress will have a say before this plan can go ahead, it is notable that the plans put forward for the Defiant quietly proposed cancelling the DDG(X).

To prevent an outcome where the US Navy is building just a 5,500 frigate and a 35,000 ton "battleship", the stated intention behind the battleship plan is to keep building Burke-class ships into the indefinite future. But this is simply not feasible. There is a limit to what new systems can be patched onto old frameworks, and changes in industrial practices will continue to inflate the Burke-class's production costs. Eventually the Navy will have no choice but to stop building the Burke-class and, without something else, will be left with an enormous unfilled gulf in its construction plans. The River-class can be that "something else".

This returns this discussion to the topic of cost. The River-class is a more modern, more adaptable ship with room for future growth. It also has the latent capacity to become a direct successor to the Burke-class. As such, it should not be surprising that its cost is similar to the Burke-class. That is simply the price of a survivable, full-spectrum surface combatant in an era of drones, hypersonic weapons, and dense sensor networks. These high capital costs should not be viewed as an impediment to pursuing such ships, where there is a need, especially for the United States whose President has recently foreshadowed up to a 50% increase in US Military spending, most of which will, initially, be available for capital projects.

That is not to denigrate smaller specialized combatants with even lower costs, like what the NSC may become after it has been redesigned. Those ships are useful. But they require full-spectrum combatants to complement them, while a fleet of just the most premium combatants is difficult to

build at scale.

More to the point, all the shipbuilding plans the US Navy has for the future have significant risk. Whether discussing the risk inherent in the NSC's redesign, or the possibility that the "threat floor" will become higher than its "capability ceiling"; the likelihood of the Burke-class's age creating an industrial cost spiral, or the chance of new threats appearing that its systems can't be adapted to counter; or enormous capability gulfs being left in the US Navy's force structure by decisions and developments in the DDG(X) or Defiant programs – the River-class provides an alternative which derisks all of those programs.

The case for the River-class is not that it is perfect, cheap, or tailor-made for the US Navy. It is that it can be built now, meeting urgent fleet needs, while preserving the option to pursue greater ambition later. It offers the US Navy a way to satisfy many priorities in one low-risk project instead of having to pursue multiple high-risk no-fail projects at once. This case for the River-class is an appealing and unique proposition. And it is enough to ask whether the best future for the US Navy is one where a River runs through it.

Jonathan Cassels is a communications professional in the non-profit sector where, among other things, he works on programs to support and find homes for unhoused Canadian Veterans. He has past experience working in the Canadian finance sector, doing democratic advocacy as the past Chair of Fair Vote Canada's Board of Directors, and working in the office of Ontario's Shadow Minister for Finance. Jonathan applies his knowledge of politics and finance to his defence analysis, which is frequently published under the moniker "CDNPolicyHawk".