DOI:10.36471/JCCM_September_2025_02

Thrifty self-assembly

1. **Magic sizes enable minimal-complexity, high-fidelity assembly of programmable shells**
   Authors: B. Tyukodi, F. Caballero, D. Hayakawa, D. M. Hall, W. B. Rogers, G. M. Grason, and Michael F. Hagan
   Phys. Rev. Lett. **135**, 118203, (2025)

2. **The polyhedral structure underlying programmable self-assembly**
   Authors: M. C. Hübl, T. E. Videbæk, D. Hayakawa, W. B. Rogers, and C. P. Goodrich
   arXiv:2501.16107

Recommended with a Commentary by Martin Lenz, CNRS, Université Paris-Saclay and ESPCI, France

Building machines at the nanoscale is unlike building a wind turbine. Whether the process is biological or artificial, it rarely involves an external operator grasping spare parts and adjusting them into place. Instead, the subunits making up the final assembly are typically delivered by Brownian motion. They come into contact through chance encounters, and stay bound only if the interaction between them is attractive.

This so-called self-assembly process presents challenges and opportunities. Because of the random nature of the encounters, any two particles that happen to attract may end up sticking together even if that impedes the intended assembly. On the upside, clever designs that disallow any spurious bonds can fully determine the final assembly simply by specifying the types of subunits and their interactions.

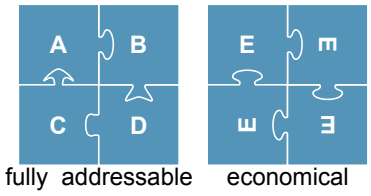

The simplest version of such a design is known as "fully addressable" self-assembly [1, 2]. To illustrate the concept, consider the task of engineering a square, two-dimensional assembly comprised of four identically-sized square subunits. In a fully addressable design, each of the four subunits is of a different type and equipped with completely specific interactions. For instance, the right side of subunit A in the illustration can only bind to the left side of subunit B, and its bottom side to the top of C. Its two remaining sides cannot bind to anyone. This strategy rules out spurious binding, even in very complicated designs. It is however difficult to implement using current experimental techniques – notably the assembly of nanoparticles, colloids or DNA origami subunits interacting through complementary strands of DNA. Indeed, producing many distinct subunits is costly, and implementing an even larger number of mutually exclusive, fully specific interactions is also

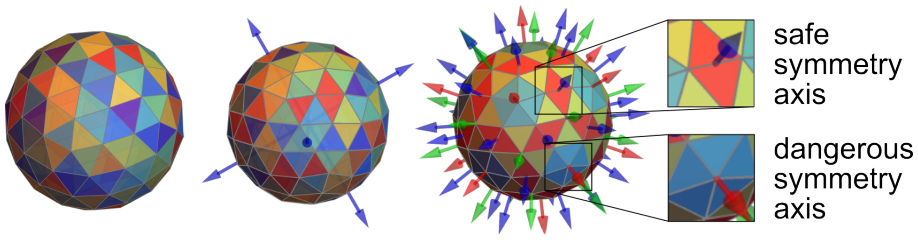

Figure 1: Three possible icosahedral shell designs. The leftmost shell is fully addressable: no two subunits are identical and so all have different colors. The righmost shell is the maximally economical design, which takes advantage of all 5-, 3- and 2-fold rotational axes of the icosahedron (*red, green* and *blue arrows*) to minimize the number of subunit types, and therefore of subunit colors. The middle shell shows an intermediate design which retains only "safe" axes, *i.e.*, those that do not go through a vertex.

challenging. These issues may be responsible for the rarity of fully addressable self-assembly in biology. Instead, living cells often assemble their proteins into highly symmetrical complexes [3]. This is analogous to making our big square out of four identical subunits of type E, each with two mutually attractive sticky sides. Applying this strategy to highly symmetric assemblies such as our big square is not only economical, but also speeds up self-assembly as subunits need a shorter time to find a suitable interaction partner.

One might naively assume that identical particles should be used whenever two subunits within the intended assembly are related by a symmetry operation, and only in those cases. However, the two articles discussed in this commentary show that the potential of this economical approach is both narrower and greater than this argument suggests. In the first article, Tyukodi *et al.* demonstrate that not all symmetries are created equal, and that while some can safely be taken advantage of, exploiting others makes for an error-prone assembly process and should be avoided. By contrast, in the second article Hübl *et al.* show that even subunits that are not related by symmetries can sometimes be safely identified without decreasing the assembly yield.

Tyukodi *et al.* consider a task reminiscent of the assembly of a viral capsid, namely the self-assembly of a closed shell out of triangular subunits. To perform it, viruses often use a highly symmetric icosahedral design to maximize subunit economy [4]. One can in principle design icosahedral shells comprising an arbitrarily large number N of subunits; taking full advantage of their symmetries allows the use of as few as $N/60$ unique subunits [5]. As shown in Fig. 1, in this design identical subunits are related by a rotation. Running numerical simulations of the assembly process, the authors however show that two types of rotations must be distinguished. Rotations whose axis runs through a vertex of the assembly are indeed associated with a risk of off-target assembly. The reason why is shown at the bottom right of Fig. 1, where five identically-colored triangular subunits join at their tips. As subunits come together during the assembly process, such configurations run the risk of incorporating one too few or one too many subunits. This results in a defective vertex with coordination 4 or 6, and ends up spoiling the rest of the assembly process. By contrast, rotation axes that run through an edge (see Fig. 1) or the center of a triangle (not shown here) are incompatible with such mistakes, and are thus safe to use. The authors then consider shells of increasing

sizes, and ask in each case whether the rotation axes associated with icosahedral symmetry are of the dangerous or of the safe type. They then compute how many distinct subunits and interactions are required to assemble the shell based only on the safe rotations. They find that the answer depends very sensitively and nontrivially on N. For instance, assembling shells of size $N = 960$ or $N = 1040$ respectively requires 2880 and 1040 distinct interactions between subunits, whereas 245 interactions are sufficient to successfully assemble a shell with $N = 980$. This huge difference in self-assembly cost is likely to apply to more complicated designs as well, and thus implies a broadly applicable, easily intelligible design guideline.

Unlike this first article, the second does not attempt to specify a single target assembly by defining a set of subunits and interactions – a characterization in what Hübl *et al.* term the "primary design space". Instead, it discusses the magnitudes of the associated subunit chemical potentials and bond strengths – coordinates in the "secondary design space". They show that setting these values can in some cases force an ambiguous design, *i.e.*, one that is *a priori* compatible with many target assemblies, to form a single final product with a very high yield. To identify these cases, they denote by μ_α the chemical potential of the reservoir of subunit α and by ϵ_i the energy of bond i. They then note that according to Boltzmann statistics, the equilibrium concentration of an assembly s containing N_α copies of subunit α and n_i instances of bond i goes as

$$\rho_s \propto \exp\left(\sum_\alpha \beta\mu_\alpha N_\alpha + \sum_i \beta\epsilon_i n_i\right) = e^{\beta \boldsymbol{M}_s\cdot\boldsymbol{\xi}}, \tag{1}$$

where $\beta^{-1} = k_B T$ denotes the thermal energy. Here the vector $\boldsymbol{M}_s = (N_1, N_2, \ldots, n_1, n_2, \ldots)$ characterizes the structure of assembly s, while $\boldsymbol{\xi} = (\mu_1, \mu_2, \ldots, \epsilon_1, \epsilon_2, \ldots)$ is a point in secondary design space. The authors consider the equilibrium state of the system in the presence of strong interactions, *i.e.*, in the formal $\beta \to \infty$ limit, and leave aside any question of kinetic trapping or slow diffusion that might arise under such conditions. In that limit, the concentration ρ_s can only remain finite if $\boldsymbol{M}_s \cdot \boldsymbol{\xi} \leq 0$. In practice, any assembly that violates this condition will form at a high rate, depleting the reservoirs of its constitutive monomers and forcing their chemical potentials to decrease until the condition is met. The intersection of all conditions $\boldsymbol{M}_s \cdot \boldsymbol{\xi} \leq 0$ associated will all possible assemblies s defines a polyhedron in secondary design space (Fig. 2). Among the points of this polyhedron, those located on its outer faces are the most interesting. Indeed, at these points $\boldsymbol{M}_s \cdot \boldsymbol{\xi} = 0$ for only one species. This implies that this species will be produced with a 100% yield in the $\beta \to \infty$ limit. Not all possible assemblies are however associated with a face of the polyhedron, and those that are not can never be obtained with a high yield – they are not "designable". This reframing of the initial problem thus provides a geometric criterion to identify designable assemblies, and the authors confirm the validity of this formulation in DNA origami experiments. This result enables an intuitive, if abstract, understanding of subunit design, and an extreme speed-up of its numerical treatment.

Although both articles discussed here are admittedly concerned with very simplified theoretical models, they address a crucial challenge that has emerged from the rapid progress of self-assembly mediated by DNA interactions. Indeed, and despite the development of some impressive platforms [6], producing a large enough number of distinct subunits and interactions to enable fully addressable self-assembly remains impractical except in simple cases.

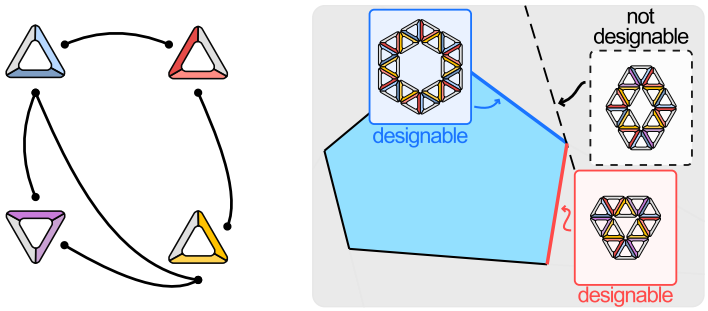

Figure 2: Designable assemblies are associated with geometrical structures in secondary design space. *Left:* A set of rules in primary design space specifying 4 subunit types (*colored triangles*) and 5 interactions (*black lines*) between their sides. This specific set implies a 9-dimensional secondary design space composed of 4 chemical potentials and 5 interaction energies. It allows for the formation of 283 distinct assemblies. *Right:* Cartoon two-dimensional cut of the secondary design space showing the polyhedral region (*light blue*) that satisfies all $\boldsymbol{M}_s \cdot \boldsymbol{\xi} \leq 0$ constraints. Choosing the assembly parameters so that $\boldsymbol{\xi}$ lies on the *red face* produces the corresponding structure with 100% yield in the $\beta \to \infty$ limit. By contrast, the assembly associated with the dashed line can never be assembled with 100% yield.

While many other studies propose algorithmic recipes to produce simplified, more manageable designs, the work of Tyukodi *et al.* and Hübl *et al.* stand out by putting forward simple geometrical rules that govern these simplifications. The resulting designs remain challenging to implement, meaning that they are probably not yet ripe for practical applications and still need to be combined with other innovations. This, in my opinion, is what makes the authors' new insights so valuable. Indeed, their simplicity and interpretability implies that they can be communicated to and appropriated by the rest of the community in a way that no black box algorithm or machine learning procedure can. That gives them the potential to be incorporated into new self-assembly frameworks in ways that likely neither the authors nor I can foresee at this point in time.

References

[1] P. W. K. Rothemund, E. Winfree, Proceedings of the thirty-second annual ACM symposium on theory of computing, 459 (2000).

[2] S. Hormoz, M. P. Brenner, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 5193 (2011).

[3] S. E. Ahnert, J. A. Marsh, H. Hernández, C. V. Robinson, S. A. Teichmann, *Science* **350**, aaa2245 (2015).

[4] F. H. Crick, J. D. Watson, *Nature* **177**, 473 (1956).

[5] D. L. Caspar, A. Klug, *Cold Spring Harb. Symp. Quant. Biol.* **27**, 1 (1962).

[6] R. Saha, D. Hayakawa, T. E. Videbaek, M. Price, W.-S. Wei, J. Pombo, D. Duke, G. Arya, G. M. Grason, W. B. Rogers, S. Fraden, arXiv:2502.05388.