Let's start by quickly recalling the steps of the proof. There are five stages:

(1) **Template**: We obtain a triangle decomposition of a constant fraction of G by taking a random injection into a "nice" Steiner triple system on a larger number of vertices (this number is somewhere between $2|V(G)|$ and $4|V(G)|$). We can show w.h.p. that (G, G^*) has some nice joint typicality property and $G \backslash G^*$ also has some nice typicality. This step is very straightforward, and in the interest of time, I will skip it. Check out Lemma 2.5 of Keevash's paper for details.

(2) **Nibble**: We start with $G \backslash G^*$ having the nice typicality probability (which is guaranteed by the previous step). Then, as we saw in Asaf's talk, we can use the Rodl nibble to obtain an "almost triangle decomposition" of $G \backslash G^*$ such that the "leave" is c_1-bounded. Recall that this boundedness condition not only quantifies the fact that we've been able to nibble away at most of the edges in $G \backslash G^*$, but also that this "sparsity" is spread throughout the vertices.

(3) **Cover**: Starting with a c_1-bounded leave L guaranteed by the above step, we find a set of edge disjoint triangles whose edges cover all edges in L, and possibly also spill over to G^*. This step is based on a "randomized greedy algorithm", and we will see how c_1-boundedness comes in here. Also, we will be able to arrange that the spill S is c_2-bounded, which will be important for the next steps.

(4) **Hole:** For a c_2-bounded spill S as above, we will find triangle decompositions M^O and M^I such that $\cup M^I, S$ is a partition of $\cup M^O$ and moreover, that $\cup M^O$ is c_3-bounded. This step will probably take up most of the time today.

(5) **Completion**: Finally, given the data above, we will "modify" M^O, M^I and M^C to obtain triangle decompositions M_1, M_2 such that $\cup M_2 = \cup M^O$ and $\cup M_1 = \cup M^O \cup L$, and also triangle decompositions M_3, M_4 such that $M_2 \subset M_4$, $\cup M_3 = \cup M_4$, $M_3 \subset T$. Then, $N \cup M_1 \cup (M_4 \backslash M_1) \cup (T \backslash M_3)$ gives the requisite triangle decomposition.

Due to lack of time, as well as for expositional clarity, I had to make some editorial decisions.

The first one is that I will not be focusing on the constants. In general, this is how you should think about them: $c_{i+1} \leq M \times d(G)^{-K} \times c_i$, where M and K are universal constants. Ultimately, we will go up to c_5 or so, and there will be some constraint on c_5 vs. $d(G)$, which we will meet by taking c_0 to be something like $d(G)^{-100{,}000}$ or so.

The second editorial decision is that I will skip details of concentration. Roughly speaking, most of the steps are randomized, and the analysis can be broken down into expectation and concentration.

Pretty much all of the interesting new ideas are already contained in the analysis of the expectation – the concentration is mostly a standard application of Azuma/Chernoff in some cases, and there are a couple of places where he uses an application of Freedman's inequality (see Lemma 2.9 of the paper).

Now, let's do the cover step. We will introduce "random greedy constructions" and their analysis here, so pay attention!

The algorithm for producing M^c given L which is c_1-bounded is simple. Arbitrarily enumerate the edges of L as $e_1, e_2, \dots, e_{|L|}$. Then, at the i^{th} strage, consider e_i and choose a random triangle containing e_i subject to the following constraints: both of its other edges are in M^*, and none of these two new edges was chosen in a previous step. Let S denote the union of all the new edges (this is the spillover). We want to claim that w.h.p., this process does not terminate before completion, and that the S is produces is c_2-bounded.

So why does this work? Let S_{t-1} denote the new edges accumulated until stage $t-1$ (so that S_0 is the empty set). Assume inductively that S_{t-1} is c_2-bounded. Denote the edge e_t as (u, v). Then, by the typicality of G^*, we know that $|N_{G^*}(u) \cap N_{G^*}(v)| \approx d(G^*)^2 n$ (think of this as large). On the other hand, since S_{t-1} is c_2-bounded, the number of choices of vertices to complete the triangle which are excluded is at most $2 \times c_2 \times n$. So, if c_2 is much smaller than $d(G^*)^2$ (and indeed, this is the case), then we see that at least half of the choices that we would have had (if we had just chosen a random triangle, and not worried about not getting overlaps) are still there. In particular, if e' is some edge in G^*, then $Pr'_t(e')$, by which we mean the probability that e' is chosen in the t^{th} step of the process, conditioned on not having overlaps, is at most $2 \times Pr_t(e')$, where $Pr_t(e')$ is the probability of choosing e' in the t^{th} step when we don't worry about overlaps. The point is that we've only lost a constant factor, and the second probability is much easier to analyse.

Now, we use this control to show that we still expect S_t to be c_2-bounded. Indeed, fix any vertex v and consider $|S_t(v)| = \sum_{i \le t} X_i$, where $X_i = \sum_{v \in e' \in G^*} 1_{e' \subset T_i}$. Taking expectations, we get $E[|S_t(v)|] = \sum_{i \le t} Pr[X_i] = \sum_{i \le t} \sum_{v \in e' \in G^*} Pr'_i[e'] \le 2 \sum_{i \le t} \sum_{v \in e' \in G^*} Pr_i[e']$. This last sum we can bound – note that this is where the c_1-boundedness of L comes in!

Indeed, we write $\sum_{i \le t} \sum_{v \in e' \in G^*} Pr_i[e'] = \sum_{v \in e' \in G^*} \sum_{i \le t} Pr_i[e'] = \sum_{u' \colon (v,u') \in G^*} \sum_{i \le t} Pr_i[(v, u')]$.

For any fixed $e' = (u', v)$, $Pr_i[e']$ can be nonzero only in a step where e_i and e' share a vertex. In particular, it can be nonzero only in steps i such that e_i is either an edge in L incident to u' or to v'. However, since L is bounded, there are at most $2c_1n$ such values of i. Moreover, whenever this probability is not zero, it is equal to $1/|N_{G^*}(u_i) \cap N_{G^*}(v_i)| \approx \frac{1}{d(G)^2 n}$. Combining everything, we get that the inner sum is bounded by $2c_1 n \times \frac{1}{d(G)^2 n} = 2c_1/d(G)^2$. On the other hand, the outer sum is over at most n terms. Putting everything together, we get that the whole sum is bounded by $2c_1 n/d(G)^2 < \frac{c_2 n}{100}$, provided we define c_2 suitably. Now, we use concentration.

Let me emphasize what happened here. The c_2-boundedness was useful only to say that we had enough choices left, and to bound probabilities by some fixed constant times a probability we can control. On the other hand, the control on the latter probability is solely in terms of c_1 and $d(G)$ (in particular, it is independent of c_2 at the start of the step). So, it's not like one of those procedures where we have to worry about starting with some c_2 and getting a worse c_2'.

Before moving to Hole, we make the preliminary observation that the S we obtained from the previous step is tridivisible, in the sense that $|S|$ is a multiple of 3 and each vertex has even degree. Indeed, we can write $S = (\cup M^c) \backslash L$, where $L = G \backslash \{G^* \sqcup (\cup N)\}$.

Until now, we have covered every edge at least once, where the edges which are covered twice are precisely those in S. If we could obtain triangle decompositions M^O and M^I as in hole, then by adding all the triangles in M^I and removing all the triangles in M^O, each edge will be covered "once". Of course, this doesn't work as is, since there will be some overlapping edges between triangles of M^O, M^I and those coming from the template T. The Completion step will fix this part.

Roughly speaking, Hole lets us replace a bunch of edges with the "difference" of two triangle decompositions which is progress since these triangles live in the very nice part G^* of our graph. Then, we can use some sort of local "shuffling" to replace "bad" triangles one by one, and perform the completion step. More on this after the break.

The goal for the remaining time today is to do all of Hole except maybe for one analysis (which is conceptually not so different from the randomized greedy analysis we already saw for cover). Next time, we will do Completion, and this will finish the proof.

Now, Hole is carried out in two stages, each of which has multiple steps.

In the first stage, we construct an "integral" triangle decomposition (*of K_n; in particular, we use edges that are not in G)* such that the sum of edge multiplicities is one for edges in S and 0 elsewhere. Note that this needs the tridivisibility of S, which we know to hold. Actually, getting such integral triangle decompositions of tridivisible graphs was already known, except that Keevash needs to do it from scratch since he wants some boundedness property to hold (and here, he uses the c_2-boundedness of S). Once we have such an integral triangle decomposition, we will then modify it in the second stage to an integral decomposition valued in $\{\pm 1, 0\}$ and such that the $\{\pm 1\}$ values only occur on triangles which are actually in G^*. Then, taking M^O to be the triangles with 1 and M^I to be the triangles in -1 will complete the proof.

The first stage ends up being not so hard for the case of triangles (the proof Keevash gives in this paper doesn't generalise to designs). The idea is to construct the integral triangle decomposition as $\phi = \phi_0 + \phi_1 + \phi_2$ in three steps, where ϕ_0 has the right "number" (with multiplicity) of edges, $\phi_0 + \phi_1$ has the right "number" of edges, as well as vertex degrees (again, with signs) and finally, $\phi_0 + \phi_1 + \phi_2$ has the right everything.

For ϕ_0, we simply choose $|S|/3$ independent random triangles in K_n; note that the expected number of triangles containing a given vertex v is Binomial with mean $|S|/n < c_2 n/2$ (since $|S| \leq n \times c_2 n \times \frac{1}{2}$ by c_2-boundedness). Therefore, by standard Chernoff bounds, we know that w.h.p, the edge set of ϕ_0 is $1.1c_2$ bounded.

Given ϕ_0 as above, consider the signed collection of edges $S - \cup\phi_0$. Since both S and $\cup\phi_0$ are tridivisible, it follows that the induced signed degree at each vertex v is an even integer, and the sum of these signed degrees is 0. Let's break up the vector of signed degrees (denoted by J^*) into positive and negative parts $J^{*\pm}$. Note that since S and $\cup\phi_0$ are $O(c_2)$-bounded, it follows that each entry of $J^{*\pm}$ is bounded by $O(c_2 n)$.

Instead of introducing more notation, it's helpful to see the next step via an example. Say $n = 7$ and $J^{*+} = (2, 4, 0, 0, 0, 2, 0)$, $J^{*-} = (0, 0, -4, -4, 0, 0, 0)$. Then, we can pair up the vertices as follows: $(1, 3), (2, 3), (2, 4)(6, 4)$. Note that in general, there will be at most

$O(c_2 n \times n)$ many pairs of this form. The idea here is that we want to transfer degree from the first vertex of each pair to the second vertex of each pair, and we can do this by choosing two random independent vertices a_i, b_i for each pair x_i, y_i, and using the following construction (see figure on board).

This fixes the degrees, but we still want some boundedness property to hold. Namely, let $\cup\phi_1^{\pm}$ denote the new edges of each sign added at this step. Then, it is easy to see that $\cup\phi_1^{\pm}$ are both $O(c_2)$-bounded – a fixed vertex v can only appear $O(c_2 n)$ many times as some x_i or y_i, and in each such occurence, we add at most 2 new edges incident to this vertex. For every other round, the vertex is chosen with probability $1/n$, so we expect at most $O(1/n) \times O(c_2 n^2) = O(c_2 n)$ many edges to be added to a given vertex in any other round. Now, get concentration.

At this point, the only problem is that $J^1 = S - (\cup\phi_0 + \cup\phi_1)$ may not be equal to 0. In the STS case, there is an easy fix for this – since the (signed) collection of edges J^1 has signed degree 0 at each vertex, we can decompose it into a collection of cycles with edges alternating in sign (note that the cycles need not be proper). Each cycle can be further decomposed into 4-cycles (see figure on board), and the thing to note here is that in the process of decomposing into 4-cycles, the degree of each vertex only goes up by some universal constant factor (say 3). Moreover, since each of S, $\cup\phi_0$, $\cup\phi_1^{\pm}$ are $O(c_2)$ bounded, there are at most $O(c_2 n)$ cycles including any given vertex v in this decomposition.

But now, given a 4-cycle, we can choose a random vertex from the graph to serve as its "center", and we obtain four signed triangles which we can add to "cancel" the 4-cycle. So, the only thing that remains to be checked is boundedness. This follows since each vertex has $O(c_2 n)$ cycles including it, and we expect each vertex to serve as the "center" of a 4-cycle $O(1/n \times c_2 n^2) = O(c_2 n)$ many times. Again, concentration finishes the proof.

The last idea we will discuss today is that of "Octahedral Elimination". This will play a big role both in Hole as well as Completion. The idea here is that an octahedron has the following two desirable properties: (i) it has a signed triangle decomposition such that the sum on each edge is 0. Therefore, if we want to eliminate a triangle while keeping the signed sum of edges the same, we can replace it by the other seven triangles of an octahedron it is part of (ii) if we have two triangles of opposite signs which share a common edge (think of this common edge as "bad"), then we can replace them by the other six triangles of an octahedron in which they occur, and note now that we have also managed to eliminate the "bad" edge. This suggests the following two phase algorithm to finish Hole.

(1) In Phase I, we eliminate the triangles in $\Phi = \phi_0 + \phi_1 + \phi_2$: arbitrarily order the triangles as $f_1, f_2, \ldots, f_\ell$. Then, for each f_i, choose a random octahedron configuration Ω_{f_i} to replace f_i. The constraints on the choice of octahedron is that the nine new edges should all be in G^*, and moreover, that these new edges should be disjoint from all new edges at previous steps, as well as the edges coming from Φ.

Let Φ' denote the triangle decomposition at the end of Phase I. Note that Φ' and Φ have the same multiplicity on every edge. Note also that every triangle in Φ' has multiplicity in $\{\pm 1, 0\}$. The point here is that earlier, we had triangles which could have had three edges not in G^*, but every triangle in Φ' has at most one edge not in G^*. We've also moved from integer multiplicites to $\{\pm 1, 0\}$ multiplicities on triangles.

(2) In Phase 2, we get rid of these "bad" edges. Denote the set of new edges added in the last phase by Γ. Then, we have positive multiplicities on $\cup\Phi^+$ and Γ, and negative multiplicities on $\cup\Phi^-$ and Γ. When we add these multiplicites, we are left with multiplicity 1 on S. In particular, for any edge which is neither in S nor in Γ, we can find two triangles of opposite signs which share that edge (and only that edge). So, we fix a sequence of pairs of triangles doing this for all these "bad" edges (with appropriate multiplicities), and for each such pair, we choose a random octahedron to eliminate this pair (along with the edge). Again, this randomness is conditioned on all the new edges being in G^*, and being distinct from the edges in $\cup\Phi^+ \cup \Gamma$, as well as the edges chosen in previous steps. Denote all of the new edges added in this phase by Γ'.

Now, if we let Ψ denote the resulting triangle decomposition, then note that any triangle in Ψ only uses edges of G^*, and Ψ has the same edge multiplicity as Φ. Moreover, each triangle in Ψ is assigned multiplicity $\{\pm 0, 1\}$. So, if this algorithm works, we are done with Hole.

The analysis is very similar to the randomized greedy algorithm analysis we saw earlier today: we first show that Γ is c_2'-bounded. As before, the idea is that by the typicality of $G^*$and the boundedness of Φ, there are many octahedra available at every step, and if Γ is c_2'-bounded, at most half the choices are excluded. This allows us to replace the computation of the boundedness of Γ by a simpler random process (with independent steps), and since Φ is c_2-bounded and we are picking octahedra at random, Γ ends up remaining c_2'-bounded w.h.p. (by concentration).

Next, we want to show that Γ' is some c_2'' bounded. Again, the boundedness of Γ, Φ and the typicality of G^*, there are many available octahedra at every step. The boundedness assumption on Γ' implies that at most half the choices are excluded, and we are able to pass to a simpler random process. Then, since Φ is bounded, and we are picking octahedra at random, Γ' ends up remaining c_2'' bounded w.h.p. (by concentration). Depending on time, we might do details next time, but hopefully the idea is clear enough even now and we can all refer to the proof of Lemma 3.2 for details.