Why Non-PoW Non-PoS Scalable Permissonless Blockchains are Practically Unsafe?

Any type of Byzantine fault tolerant (BFT) system can ensure safety only when the strength of the adversary is limited to a certain fraction of the number of nodes or voting power or of the total computing power in the system. For example, traditional BFT solutions require that the number of Byzantine nodes would be limited to less than a 1/3 when either the environment is not fully synchronous or in the oral messages model. This assumption is also prevalent in permissioned blockchains, as the latter typically rely on traditional BFT solutions. Similarly, Proof of Work (PoW) based blockchains are only safe when the power of the adversary is restricted to less than 1/2 of the total compute power in the system, assuming that the network is synchronous with high probability. In fact, it is well known that an adversary that possesses more than 1/4 of the compute power can already gain an unfair share through selfish mining.

Yet, why should we accept such assumptions limiting the power of the adversary?

In the classical BFT model, it may be reasonable to assume that the servers executing the system are well protected, so that an attacker would need to invest significant resources in intruding each such server. Further, when different servers are placed at different administrative domains, as in the envisioned case for permissioned (non-private) blockchains, a successful intrusion into one server does not help intruding other servers. Consequently, one may be comfortable with the assumption that indeed less than a 1/3 of the servers are Byzantine.

In the case of PoW, it is assumed that obtaining 1/2 (or even 1/4) of the compute power of the system is not practical. In particular, at the scale at which Bitcoin currently runs, buying enough hardware to surpass these thresholds is considered too expensive. And while some downloaded software may covertly use its hosting machines to mine bitcoins on behalf of an attacker, it is still far from enough to tip the overall balance dramatically. Even in the case of a network of botnets, if the attacker tries to utilize them in a significant way for a long time, this will be detected by the respective machines' owners.

Yet, when eliminating the hard work associated with PoW, suddenly very little compute power is required in order to participate in a blockchain network. While this is an advantage from the point of view of electricity consumption and total transactions latency and throughput, it suddenly exposes the network to adversarial attacks. In this setting, a botnet network can easily be deployed to manipulate a blockchain network. Most such systems assume the use of Sybil resilient identity acquisition mechanisms to impose the Byzantine threshold assumption and thwart Sybil attacks. However, Sybil protection either relies on centralized authorities, or requires participants to perform heavy computations to acquire their identities. Relying on a centralized identity authority beats the purpose of permissionless blockchains. As for heavy computational approaches, if they are performed only once or very rarely, then botnets can gradually obtain enough identities to launch an attack. If these identity proving computations need to be executed frequently, we are back in the same inefficiency of PoW, only with a much more complex protocol built around it.